Gap Analysis, Problem Statement and Requirements for Source Address Validation in Inter-domain Networks

Jianping Wu, Dan Li, Libin Liu, Mingqing Huang, Lancheng Qin, Nan Geng

March 29, 2023

Background

- Goals
  - Perform the gap analysis of existing inter-domain SAV mechanisms
  - Summarize the fundamental problems of existing inter-domain SAV mechanisms
  - Describe the requirements for the new inter-domain SAV mechanism
- Historical versions
  - draft-wu-savnet-inter-domain-problem-statement-00, IETF 114 SAVNET WG
  - draft-wu-savnet-inter-domain-problem-statement-01, Sep. 25, 2022
  - draft-wu-savnet-inter-domain-problem-statement-02, Oct. 22, 2022
  - **draft-wu-savnet-inter-domain-problem-statement-03, IETF 115 SAVNET WG**
  - draft-wu-savnet-inter-domain-problem-statement-04, Nov. 29, 2022
  - draft-wu-savnet-inter-domain-problem-statement-05, Dec. 15, 2022
  - draft-wu-savnet-inter-domain-problem-statement-06, Mar. 4, 2023
  - **draft-wu-savnet-inter-domain-problem-statement-07, IETF 116 SAVNET WG**

Comments on Version-03

- Michael Richardson: Your document needs to visit the RIRs/NOGs for feedback (NANOG, RIPE, APRICOT).
  - Response: We have discussed a lot with China Telecom, China Mobile, NANOG, MANRS and APNIC community. Their suggestions are fully considered when we wrote the draft.
- Alvaro: Slide #12, if we build solutions based on these requirements, these requirements are not specific enough. E.g., the small overhead requirement, we need to quantify.
  - Response: Some more descriptions about the concrete requirements have been added to the draft, and they present some baselines in the requirements section.
- Roland Dobbins: Reflection amplification is over emphasized. Direct path boost attack is under emphasized.
  - Response: Reflection amplification is the most important attacks caused by source address spoofing, as emphasized by MANRS. But we also include other attacks.

Comments on Version-03

- Barry Greene: Remove "EFP-uRPF" from the SAVNET work: Don't use hypothetical protocols like EFP-uRPF as a foundation for the SAVNET Work. EFP-uRPF is a theory. There has never been an effort to code EFP-uRPF.
  - Response: RFC8704 takes a gap analysis on existing SAV mechanisms and proposes EFP-uRPF to narrow the gaps. Since EFP-uRPF is an improved mechanism and is also standardized, our draft does an analysis on EFP-uRPF. In the section of existing SAV mechanisms, the draft was revised to point out that EFP-uRPF has not been implemented in practical networks so that people can learn about the current status of EFP-uRPF.
- Barry Greene: Replace "Misaligned incentive" : "Misaligned incentive" is disrespectful to the operations who have the bear the deployment, operational, and capital SAV cost. I would change "misaligned incentives" to "deployment & operational incentives." .
  - Response: The draft has been updated to make it clearer, and the descriptions related to misaligned incentive have been removed.

Comments on Version-03

- Barry Greene: Should park the existing work and focus on a new joint authored "state of SAV" today.
  - Response: The draft gives analyses on existing inter-domain SAV mechanisms and finds the technical gaps/problems that can be fully or partially to be solved, and describes the motivations of SAVNET architecture and BAR-SAV.
- Li Chen: Requirements section could be reframed to "requirements of a solution in the pursuit of these goals", but pithier.
  - Response: In requirements section, the draft has been updated by clarifying which kind of requirements we are talking about and making the listed requirements more practical.
- Yuanyuan Zhang: Should low operational overhead also be a requirement for the new inter-domain SAV mechanism?
  - Response: The operational overhead was also analyzed for inter-domain SAV. The draft has been updated to give more descriptions on the operational overhead.

Main Updates Compared to Version-03

- Updates in Introduction section
  - Boundary between intra-domain and inter-domain SAV mechanisms
  - Goals of inter-domain SAV mechanisms
- One new Existing Inter-domain SAV Mechanism section
- Updates in Gap Analysis section
- Updates in Problem Statement section
- Updates in Requirements section

Boundary between Intra-domain and Inter-domain SAV Mechanisms

- Intra-domain SAV mechanisms
  - An AS X defends against source address spoofing without the collaboration of other ASes
    - Goal 1: Prevent a subnet of AS X from spoofing the addresses of other subnets (either within the AS or other ASes)
    - Goal 2: Prevent the incoming traffic to AS X from spoofing the addresses of AS X
- Inter-domain SAV mechanisms
  - Multiple ASes collaborate with each other for defending against source address spoofing
    - AS X helps defend against spoofing traffic originated from AS A which spoofs the addresses of AS B

Goals of Inter-domain SAV Mechanisms

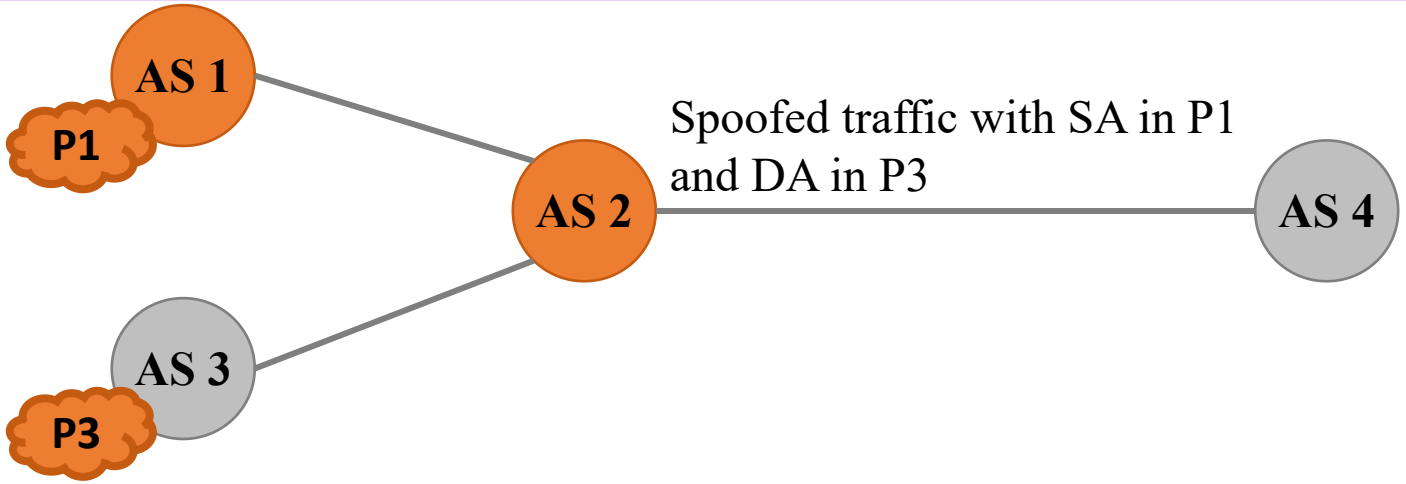

- An example to illustrate inter-domain SAV
  - P1 is the source prefix of AS 1, and P3 is the source prefix of AS 3
  - Both AS 1 and AS 2 deploy intra-domain SAV and inter-domain SAV
  - AS 3 and AS 4 deploy neither intra-domain SAV nor inter-domain SAV
  - AS 4 sends spoofed traffic with SA in P1 and DA in P3 to AS 3 through AS 2

Goals of Inter-domain SAV Mechanisms

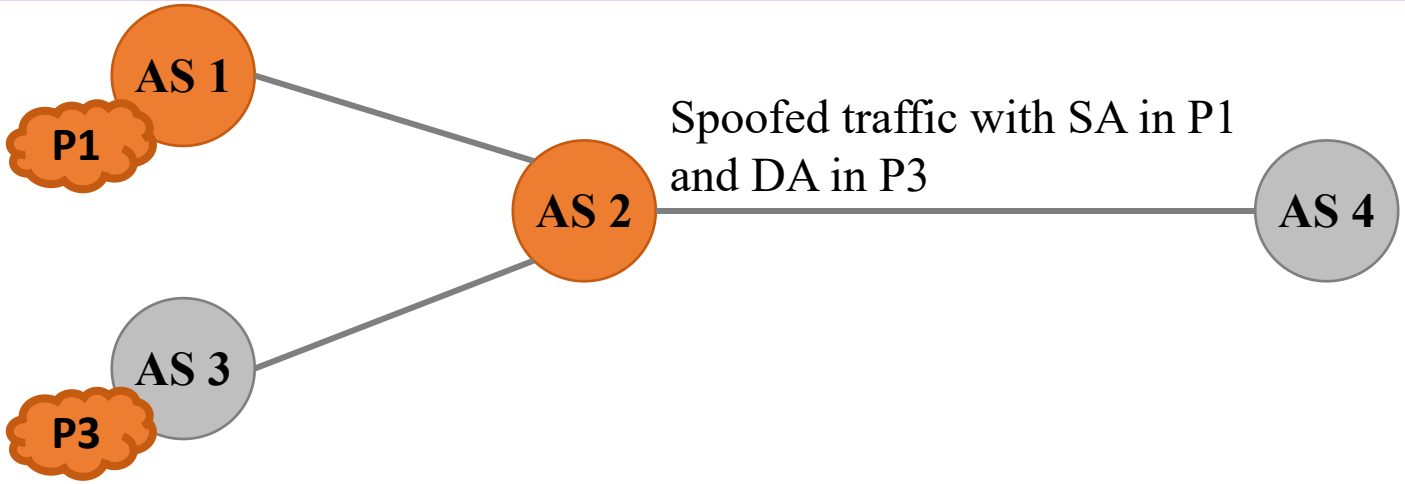

- Intra-domain SAV cannot help in this scenario
  - Although AS 1 deploys intra-domain SAV, the spoofing traffic from AS 4 to AS 3 do not go through AS 1, **they cannot be blocked by AS 1**
- **Inter-domain SAV can help in this scenario**
  - Since AS 1 and AS 2 deploy inter-domain SAV, AS 2 knows the correct incoming interface of packets with P1 as source addresses, and the spoofing packets can thus be blocked by AS 2 since they come from the incorrect interface

Main Updates Compared to Version-03

- Updates in Introduction section
- One new Existing Inter-domain SAV Mechanism section
  - Review existing inter-domain SAV mechanisms
- Updates in Gap Analysis section
- Updates in Problem Statement section
- Updates in Requirements section

Existing Inter-domain SAV mechanisms

Ingress filtering [RFC2827, RFC3704, RFC8704] is the current practice of inter-domain SAV

- ACL-based SAV [RFC2827, RFC3704]
- Strict uRPF [RFC3704]
- Loose uRPF [RFC3704]
- FP-uRPF [RFC3704]
- VRF uRPF [https://www.cisco.com/c/dam/en_us/about/security/intelligence/urpf.pdf]
- EFP-uRPF [RFC8704]
- Source-based Remote Triggered Black Hole (RTBH) Filtering [RFC5635]
- Carrier Grade NAT (CGN)
- BGP Origin Validation (BGP-OV) [RFC6811]

ACL-based SAV [RFC2827, RFC3704]

- Working principle
  - ACL rules can be configured to filter packets based on their source addresses
- Deployed at provider interfaces, customer interfaces, or peer interfaces, and recommended to deploy at provider interfaces or customer interfaces
  - At provider interfaces, block source prefixes that are clearly invalid in the inter-domain routing context, such as suballocated or internal-only prefixes of customer ASes
  - At customer interfaces, prevent customer ASes from spoofing source addresses of other ASes that are not reachable via the provider AS
  - Implemented at border routers or aggregation routers if border ACLs are not feasible
- ACL rules need to be updated in a timely manner when prefixes or routing of ASes change, which relies on manual configuration

Strict uRPF [RFC3704]

- Working principle
  - The most stringent mode of uRPF-based mechanism
  - Only permit packets that have a source address that is covered by a prefix in the FIB, and that the next hop for that prefix is the same as the incoming interface
- Recommended for deployment at customer interfaces that directly connect to an AS with suballocated address space
  - Used to prevent spoofing attacks from that AS or its downstream ASes

Loose uRPF [RFC3704]

- Working principle
  - Match the source address of the packet with one or more prefixes in the FIB, regardless of which interface the packet arrives at. If the source address is not routable, discarding the packet
- Deployed at provider interfaces of an AS
  - Used to block packets with source addresses that are obviously disallowed, such as non-global prefixes (e.g., private addresses, multicast addresses, etc.) or the prefixes that belong to the customer AS itself

FP-uRPF [RFC3704]

- Working principle
  - Maintain a reverse path forwarding (RPF) list, which contains the prefixes and all their permissible routes including the optimal and alternative ones
  - Permit a packet only if the packet's source address is encompassed in the prefixes of the RPF list and its incoming interface is included in the permissible routes of the corresponding prefix
- Recommended to be deployed at customer interfaces or peer interfaces, especially those that are connected to multi-homed customer ASes

VRF uRPF

- Working principle
  - Use a separate VRF table for each external BGP peer
  - A VRF table is a table that contains the prefixes and the routes that are advertised by a specific peer
  - Check the source address of an incoming packet from an external BGP peer against the VRF table for that peer. If the source address matches one of the prefixes in the VRF table, allowing the packet to pass. Otherwise, it drops the packet
- VRF uRPF can be used as a way to implement BCP38
  - The operational feasibility of VRF uRPF as BCP38 has not been proven

EFP-uRPF [RFC8704]

- Working principle
  - Consist of two algorithms, algorithm A and algorithm B
  - Based on the idea that an AS can receive BGP updates for multiple prefixes that have the same origin AS at different interfaces.
  - For example, this can happen when the origin AS is multi-homed and advertises the same prefixes to different providers. In this case, EFP-uRPF allows a packet with a source address in any of those prefixes to pass on any of those interfaces
- Deployed at customer interfaces or peer interfaces of an AS
- Not implemented in practical networks yet, but BCP84 suggests using EFP-uRPF with algorithm B at customer interfaces

Source-based RTBH filtering [RFC5635]

- Working principle
  - Enable the targeted dropping of traffic by specifying particular source addresses or address ranges
  - Use uRPF, usually loose uRPF, to check the source address of an incoming packet against the FIB. If the source address of the packet does not match or is not covered by any prefix in the FIB, or if the route for that prefix points to a black hole (i.e., Null0), discarding the packet
- Filter out attack traffic at specific devices (e.g., ASBR) in an AS based on source addresses

Carrier Grade NAT (CGN)

- Working principle
  - Used by service providers to translate between private and public IPv4 addresses within their network
  - Enable service providers to assign private IPv4 addresses to their customer ASes instead of public, globally unique IPv4 addresses
- Cannot determine whether the source address of an incoming packet is spoofed or not, additional SAV mechanisms need to be implemented

BGP Origin Validation (BGP-OV) [RFC6811]

- Background
  - Attackers can bypass uRPF-based SAV mechanisms by using prefix hijacking in combination with source address spoofing. By announcing a less-specific prefix that does not have a legitimate announcement, the attacker can deceive existing uRPF-based SAV mechanisms and successfully perform address spoofing
  - To protect against this type of attack, a combination of BGP-OV and uRPF-based mechanisms like FP-uRPF or EFP-uRPF is recommended
- Working principle
  - BGP routers can use ROA information, which is a validated list of {prefix, maximum length, origin AS}, to mitigate the risk of prefix hijacks in advertised routes

Main Updates Compared to Version-03

- Updates in Introduction section
- One new Existing Inter-domain SAV Mechanism section
- Updates in Gap Analysis section
  - SAV at provider interfaces
  - SAV at customer interfaces
  - SAV at peer interfaces
- Updates in Problem Statement section
- Updates in Requirements section

SAV at Provider Interfaces

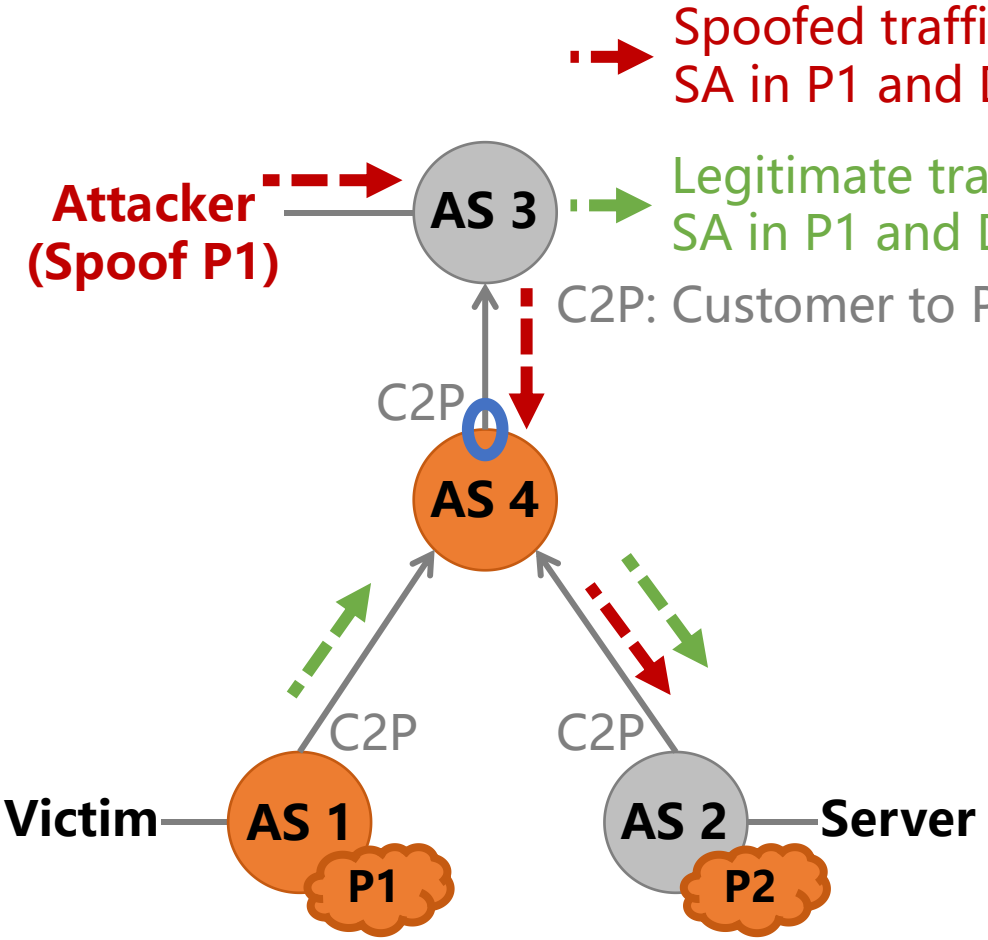

A Scenario of The Reflection Attack from Provider AS

- SAV mechanisms
  - **ACL-based SAV**, **source-based RTBH filtering**, or **loose uRPF**
- Both AS 1 and AS 4 have deployed SAV
- SAV at AS 4 is considered facing AS 3
- An attacker which is inside of AS 3 or connected to AS 3 through other ASes spoofs traffic with SA in P1 and DA in P2

SAV at Provider Interfaces

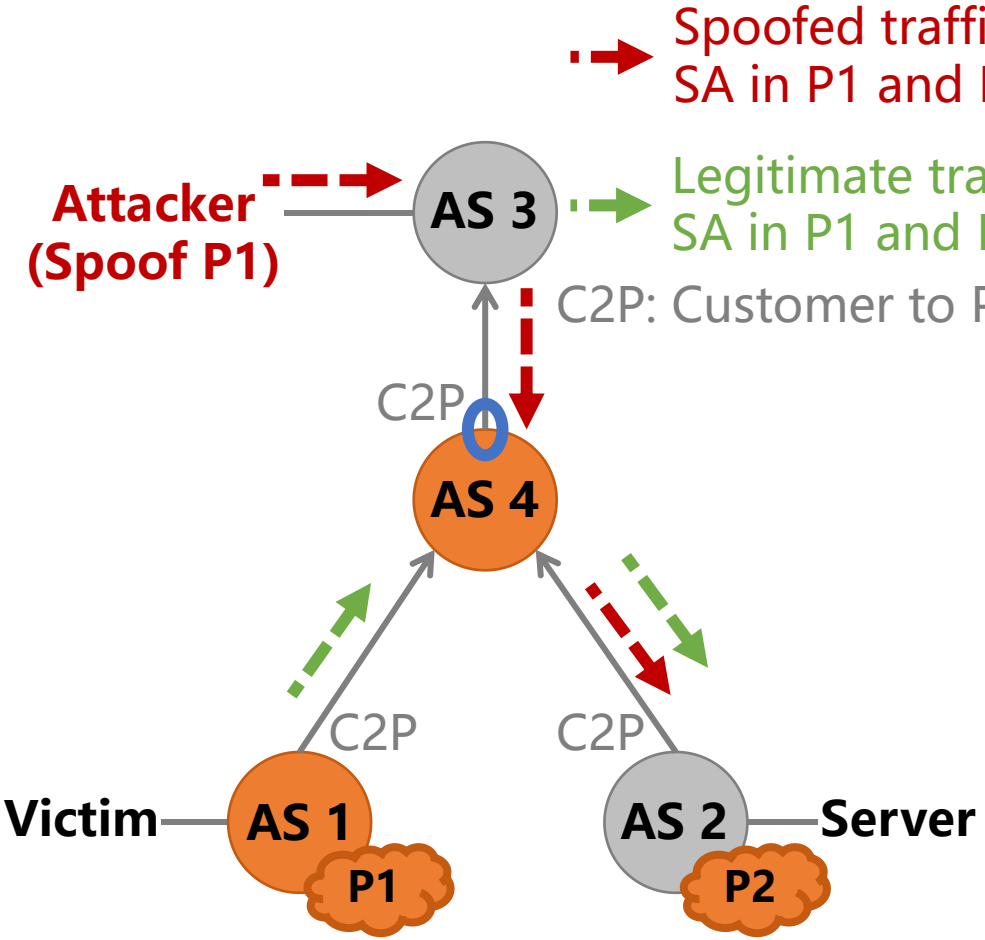

A Scenario of The Reflection Attack from Provider AS

- When deploying **ACL-based SAV** at AS 4
  - To avoid improper block or improper permit, operators need to perform timely update of ACL rules based on the prefix or topology changes of AS 1 and AS 2, which incurs **high operational overhead**
- When deploying **source-based RTBH filtering** at AS 4
  - To avoid improper block or improper permit, operators need to update specified source addresses in a timely manner, which incurs **additional operational overhead**

SAV at Provider Interfaces

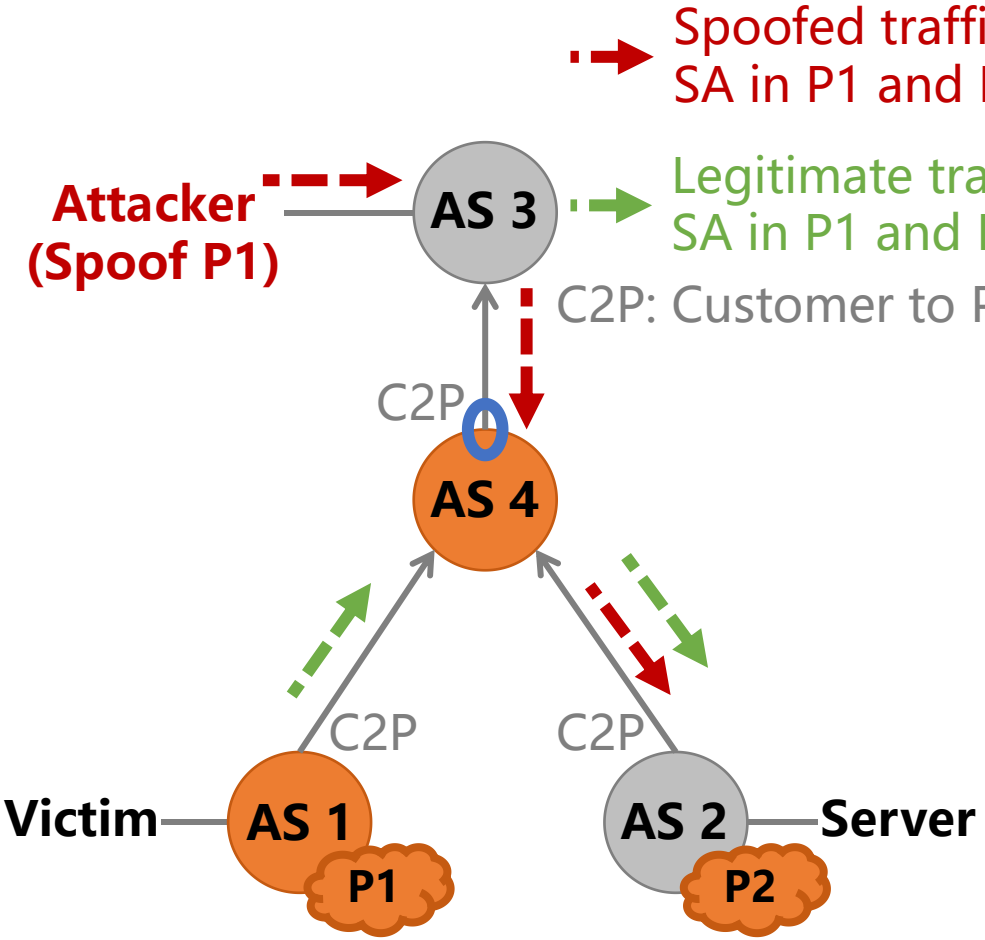

A Scenario of The Reflection Attack from Provider AS

- When deploying **loose uRPF** at AS 4
  - AS 4 would **improperly permit** the spoofed traffic since AS 4 with loose uRPF cannot determine whether legitimate traffic with SA in P1 will come from AS 1 or AS 3

SAV at Customer Interfaces

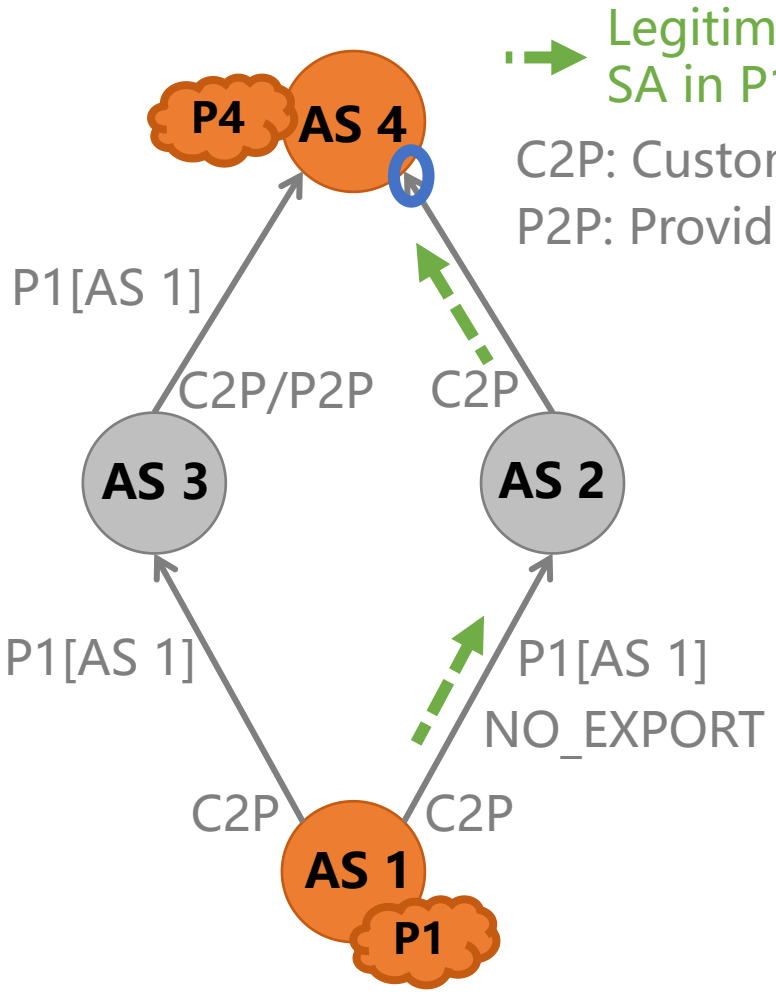

Limited Propagation of Prefixes Caused by NO_EXPORT

- SAV mechanisms
  - **ACL-based SAV**, **source-based RTBH filtering**, **strict uRPF**, **FP-uRPF**, **VRF uRPF**, or **EFP-uRPF**
- Both AS 1 and AS 4 have deployed SAV
- SAV at AS 4 is considered facing AS 2
- AS 1 adds "NO_EXPORT" in BGP update to AS 2
- AS 1 selects AS 1→AS2→AS4 as the best path to P4

SAV at Customer Interfaces

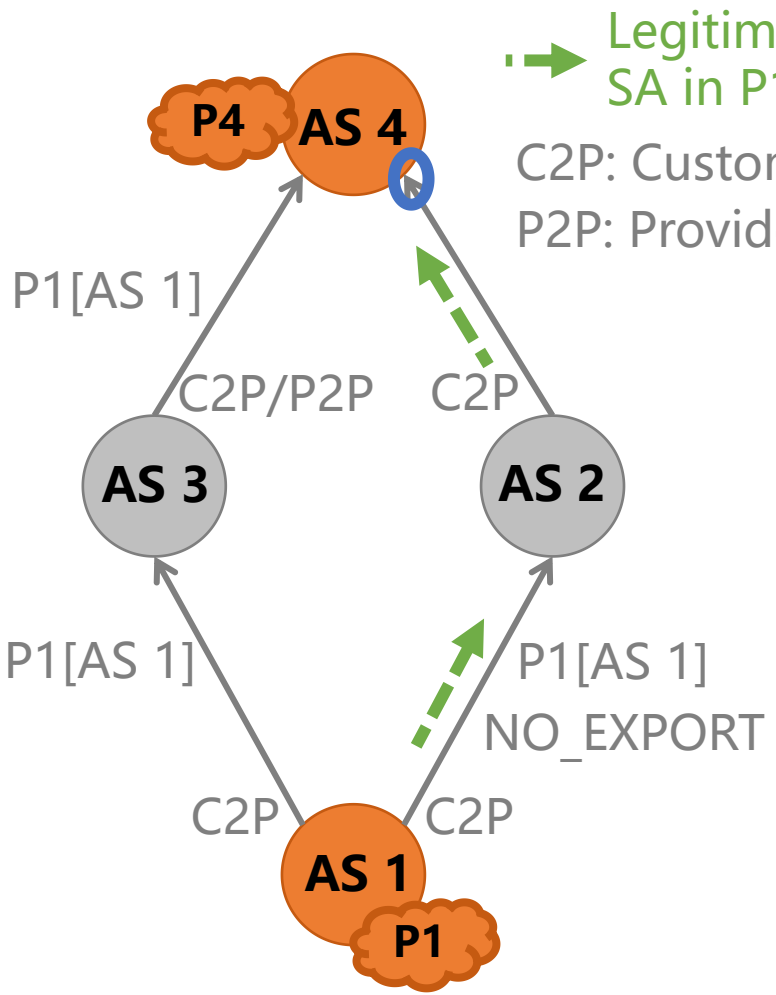

Limited Propagation of Prefixes Caused by NO_EXPORT

- Both **ACL-based SAV** and **source-based RTBH filtering** have **operational overhead** like performing SAV at provider interfaces
- Assuming AS3 is the customer of AS4
  - **Strict uRPF**, **FP-uRPF**, **VRF uRPF**, and **EFP-uRPF** with algorithm A would **improperly block** the legitimate traffic with SA in P1
  - **EFP-uRPF** with algorithm B works well
- Assuming AS3 is the lateral peer of AS 4
  - **Strict uRPF**, **FP-uRPF**, **VRF uRPF**, and **EFP-uRPF** with algorithm A/B would **improperly block** the legitimate traffic with SA in P1

SAV at Customer Interfaces

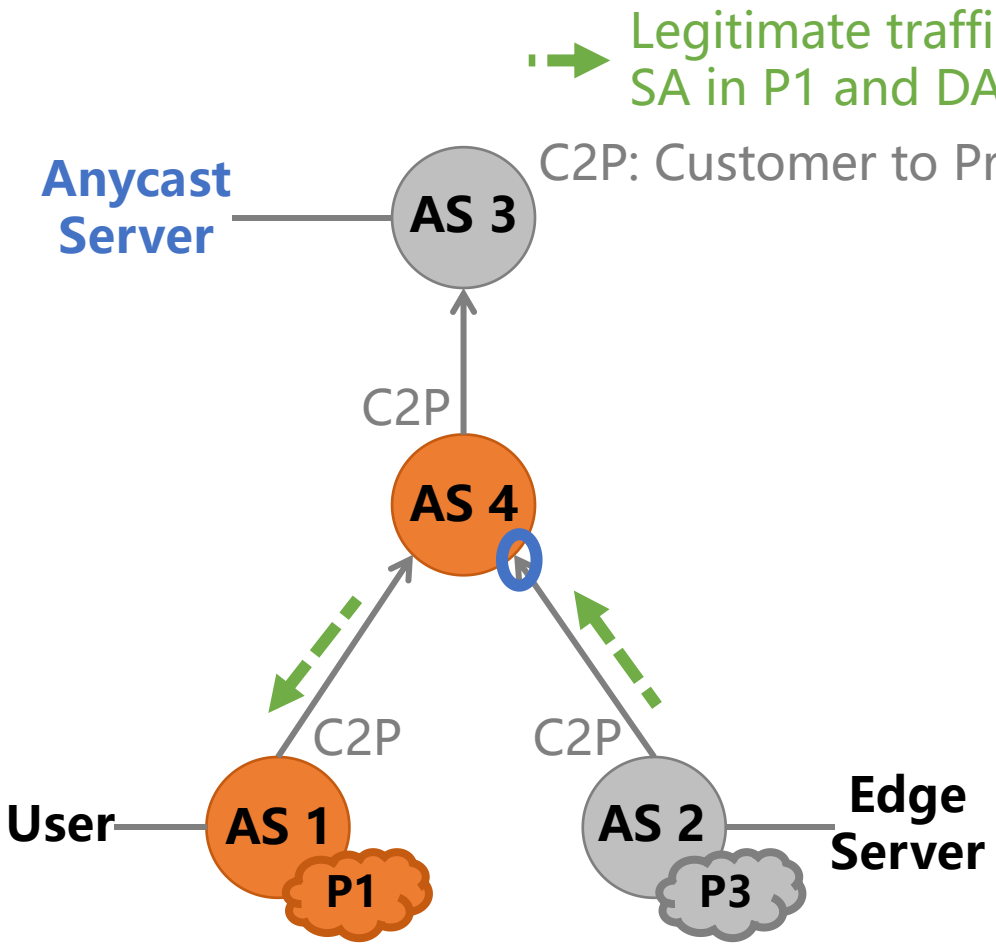

A Direct Server Return (DSR) Scenario

- Both AS 1 and AS 4 have deployed SAV
- SAV at AS 4 is considered facing AS 2
- P3 is not advertised by AS2 through BGP (e.g., P3 is only used in DSR scenario)
- When deploying **strict uRPF**, **FP-uRPF**, **VRF uRPF**, or **EFP-uRPF** at AS 4
  - AS 4 would **improperly block** the legitimate traffic with SA in P3

SAV at Peer Interfaces

- Both **ACL-based SAV** and **source-based RTBH filtering** have the same **operational overhead** as performing SAV at provider interfaces or customer interfaces
- **FP-uRPF**, **VRF uRPF**, or **EFP-uRPF** may **improperly block** the legitimate traffic in the cases of limited propagation of prefixes or hidden prefixes, e.g., DSR, like performing SAV at customer interfaces

SAV at Peer Interfaces

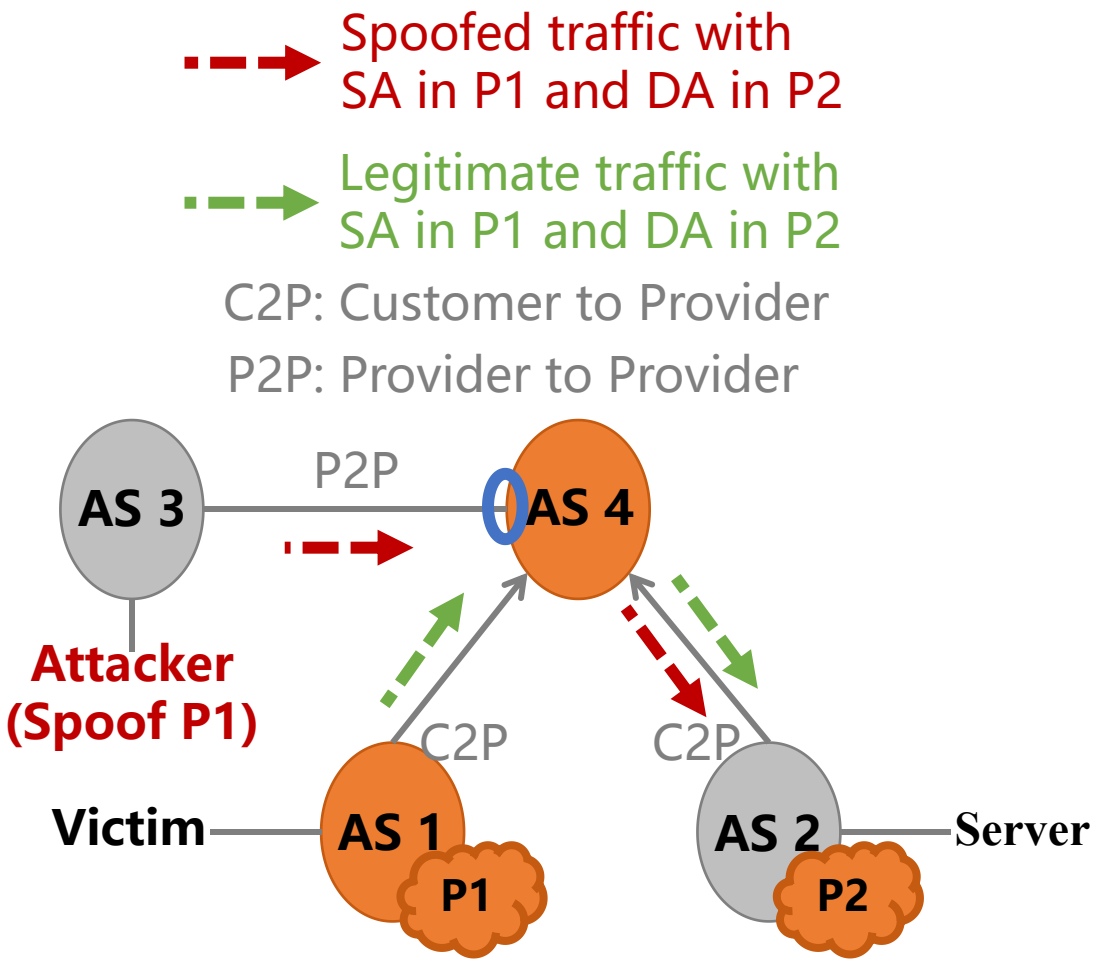

A Scenario of The Reflection Attack from Peer AS

- SAV mechanisms
  - **ACL-based SAV**, **source-based RTBH filtering**, **FP-uRPF**, **VRF uRPF**, or **EFP-uRPF**
- Both AS 1 and AS 4 have deployed SAV
- SAV at AS 4 is considered facing AS 3
- When deploying **EFP-uRPF** with algorithm B at AS 4
  - AS 4 would **improperly permit** the spoofed traffic with SA in P1

Main Updates Compared to Version-03

- Updates in Introduction section
- One new Existing Inter-domain SAV Mechanism section
- Updates in Gap Analysis section
- Updates in Problem Statement section
  - Remove the subsection of “Misaligned Incentive”
  - Summarize the problem for each inter-domain SAV mechanism
- Updates in Requirements section

Problem Statement

- ACL-based SAV
  - Problem: high operational overhead
  - Reason: need to manually update ACL rules to adapt to network changes
- Source-based RTBH filtering
  - Problem: high operational overhead
  - Reason: need to manually update the specified source addresses
- Strict uRPF
  - Problem: improper block when AS is multi-homed and has asymmetric routes
  - Reason: perform SAV only based on the local FIB which may not include the asymmetric routes of the legitimate traffic

Problem Statement

- Loose uRPF
  - Problem: improper permit
  - Reason: oblivious to the incoming interfaces of packets
- FP-uRPF and VRF uRPF
  - Problem: improper block in asymmetric routing scenarios, e.g., limited propagation of prefixes
  - Reason: perform SAV based on the local RIB which may not have the prefixes with limited propagation and their permissible incoming interfaces
- EFP-uRPF
  - Problem: improper block in the cases of hidden prefixes, e.g., DSR
  - Reason: not learn the hidden prefixes, which are legitimate source prefixes

Main Updates Compared to Version-03

- Updates in Introduction section
- One new Existing Inter-domain SAV Mechanism section
- Updates in Gap Analysis section
- Updates in Problem Statement section
- Updates in Requirements section
  - Remove the subsections of "Direct Incentive" and "Acceptable Overhead"
  - Add the subsection of "Automatic Update"
  - Revise the description of the requirements

Requirements for New Inter-domain SAV Mechanisms

- Automatic update
  - The new inter-domain SAV mechanism MUST be able to adapt to dynamic networks and asymmetric routing scenarios automatically, instead of relying on manual update
- Accurate validation
  - The new inter-domain SAV mechanism SHOULD improve the validation accuracy in all directions of ASes over existing mechanisms
    - Avoid improper block and minimize improper permit by learning the real forwarding paths or the minimal set of acceptable paths that cover the real forwarding paths
    - Multiple sources of SAV-related information can help reduce the set of acceptable paths and improve the validation accuracy
- Working in incremental/partial deployment
  - The new inter-domain SAV mechanism SHOULD provide effective protection for source addresses when it is partially deployed in the Internet

Requirements for New Inter-domain SAV Mechanisms

- Automatic update
  - The new inter-domain SAV mechanism MUST be able to adapt to dynamic networks and asymmetric routing scenarios automatically, instead of relying on manual update
- Accurate validation
  - The new inter-domain SAV mechanism MUST be able to accurately validate packets in all
    - Avoid improper block and minimize improper permit by learning the real forwarding paths or the minimal set of acceptable paths that cover the real forwarding paths
    - Multiple sources of SAV-related information can help reduce the set of acceptable paths and improve the validation accuracy
- Working in incremental/partial deployment
  - The new inter-domain SAV mechanism SHOULD provide effective protection for source addresses when it is partially deployed in the Internet

These requirements serve as practical guidelines that can be met, in part or in full, by proposing new techniques

Acknowledgements

- Many thanks to Jared Mauch, Barry Greene, Fang Gao, Anthony Somerset, Kotikalapudi Sriram, Yuanyuan Zhang, Igor Lubashev, Alvaro Retana, Joel Halpern, Aijun Wang, Michael Richardson, Li Chen, Gert Doering, Mingxing Liu, John O'Brien, Roland Dobbins, etc. for their valuable comments on this document

Thanks!