National Institute of Standards and Technology

INFORMATION TECHNOLOGY LABORATORY

Enhanced Feasible-Path Unicast Reverse Path Filtering

draft-sriram-opsec-urpf-improvements-02

K. Sriram, D. Montgomery, and J. Haas

IEPG, Singapore
November 2017

Acknowledgements: The authors are grateful to many folks who offered feedback and suggestions in the OPSEC WG meeting at IETF-99 and earlier on the GROW mailing list.

Difficulties with Adoption of uRPF Solutions

- Strict uRPF is usable in very limited scenarios
- Loose uRPF is not very effective for denying traffic with IPv4 address spoofing (except bogons, etc.)
- Feasible path uRPF is a refinement but ISPs apprehensive that they might deny traffic with legitimate customer source IP addresses
  - When faced with multi-homing and asymmetric routing
- Is there a way to make feasible-path more generalized and accurate?
- Goal: Encourage wider deployment of uRPF

Reverse Path Filter (RPF) List

The list of permissible source address prefixes for incoming data packets on a given interface.

Key Principles of Enhanced Feasible Path uRPF Version-01 Algorithm

Algorithm for customer facing ISP eBGP router:

1. Set A = {AS1, AS2, …, ASn} is the list of all unique origin ASes in all received routes
2. Set X_1 is the list of unique prefixes that have a common origin AS1
   - Those routes have potentially been received on different customer/ peer/ provider interfaces
3. Include X_1 in Reverse Path Filter (RPF) list on all interfaces on which one or more of the prefixes in X_1 were received
4. Repeat Step 2 and 3 for all ASes in set A

Basic Scenario A

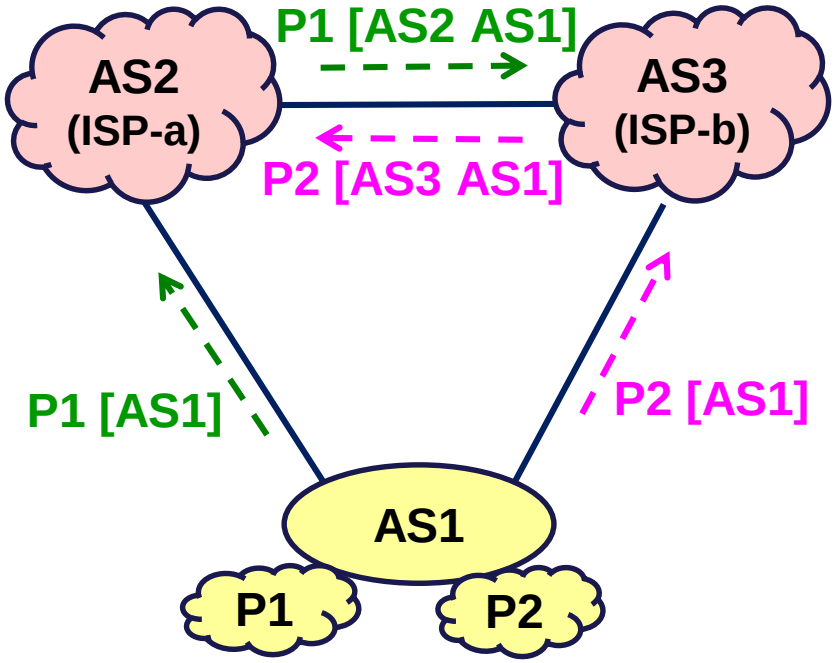

Consider data packets received at AS2 with source address in P1 or P2:

- X Strict uRPF fails
- X Feasible-path uRPF fails (since routes for P1, P2 are selectively announced to different upstream ISPs)
- ✔ Loose uRPF works (but not desirable)
- ✔ Enhanced Feasible-path uRPF works best

Basic Scenario B

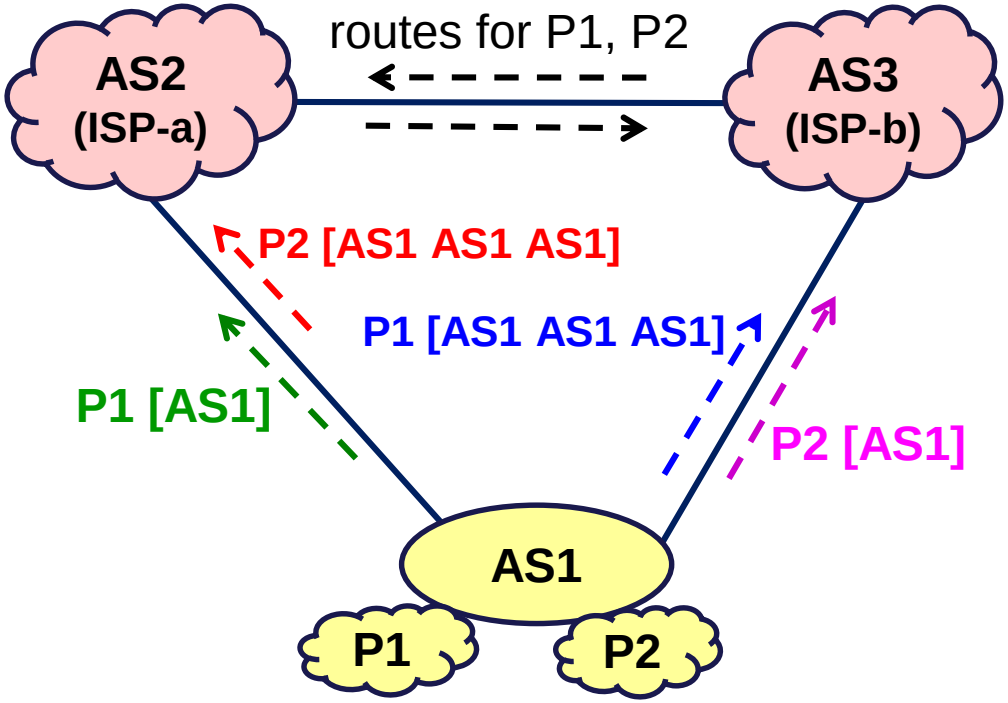

Consider data packets received at AS2 with source address in P1 or P2:

- ✔ Feasible-path uRPF works (if customer route preferred at AS3 over shorter path)
- X Feasible-path uRPF fails (if shorter path preferred at AS3 over customer route)
- ✔ Loose uRPF works (but not desirable)
- ✔ Enhanced Feasible-path uRPF works best

Scenario 1

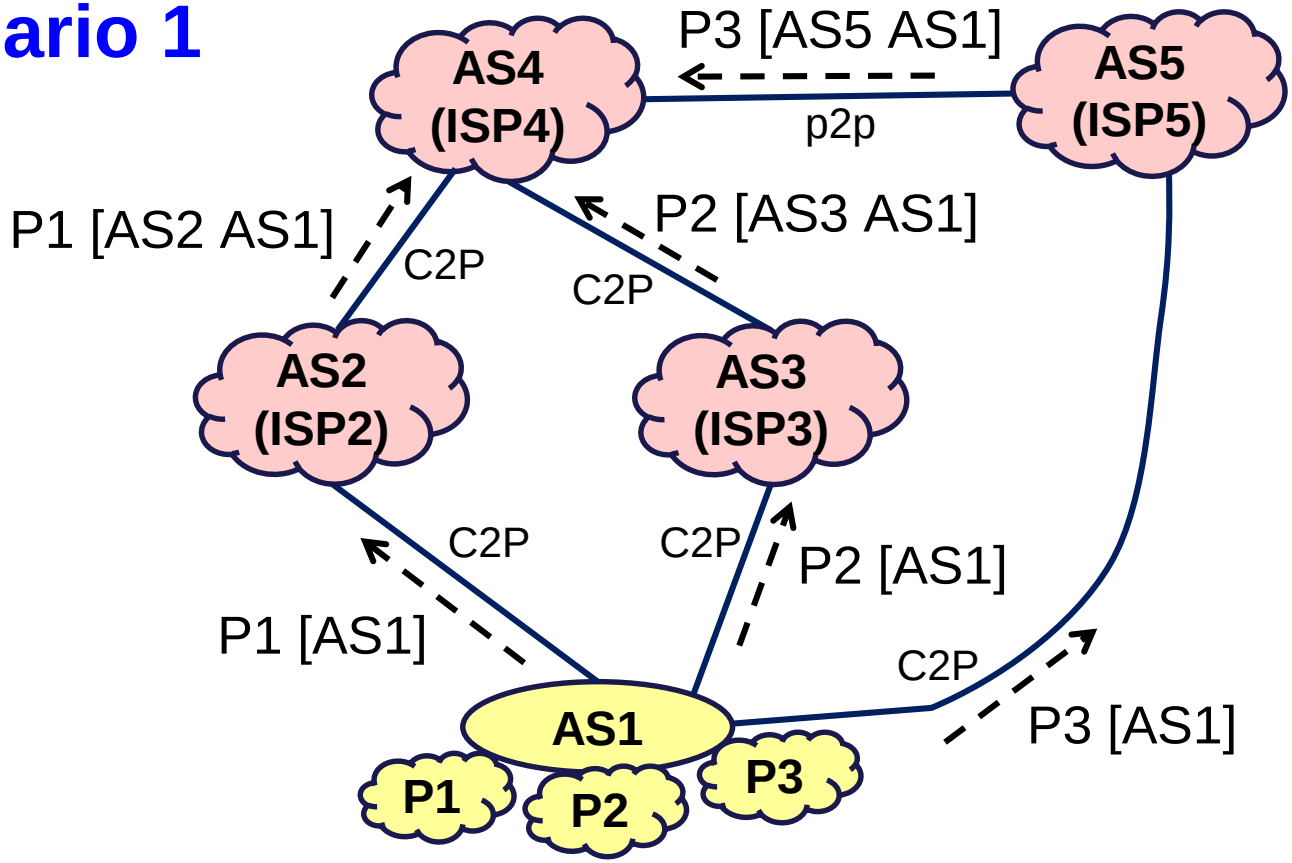

Consider that data packets (sourced from AS1) may be received at AS4 with source address in P1 or P2 from any of the neighbors (AS2, AS3, AS5):

- X Feasible-Path uRPF fails (since routes for P1, P2 are selectively announced to different upstream ISPs)
- ✔ Loose uRPF works (but not desirable)
- ✔ Enhanced Feasible-Path uRPF works best

Scenario 2: Example of a Challenging Scenario (from OPSEC & GROW WG discussions)

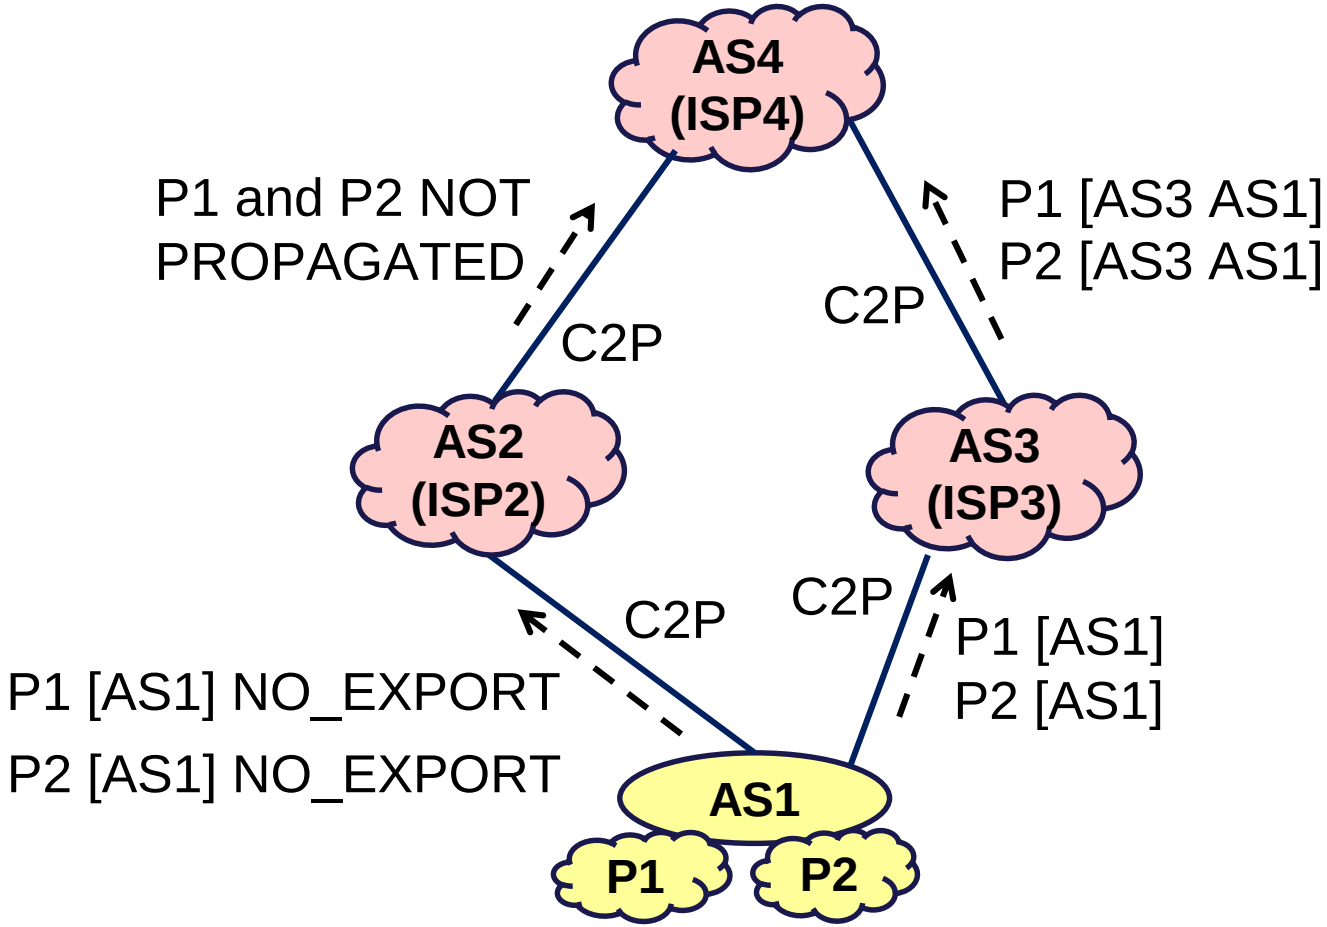

Adding More Flexibility to Enhanced Feasible Path uRPF

Updated Algorithm (meets with the challenge)

- Let I = {I1, I2, ..., In} represent the set of all directly-connected customer interfaces at customer-facing edge routers in a transit provider's AS.
- Let P = {P1, P2, ..., Pm} represent the set of all unique prefixes for which routes were received over the interfaces in Set I.
- Let A = {AS1, AS2, ..., ASk} represent the set of all unique origin ASes seen in the routes that were received over the interfaces in Set I.
- Let Q = {Q1, Q2, ..., Qj} represent the set of all unique prefixes for which routes were received over peer or provider interfaces such that each of the routes has its origin AS belonging in Set A.
- Then, Z is the RPF list for each of the interfaces in Set I.

MAY use the enhanced FP uRPF as described on Slide 4 or the Loose uRPF for Peer & Provider interfaces.

Scenario 3: Example of a Challenging / Complex Scenario (and it works)

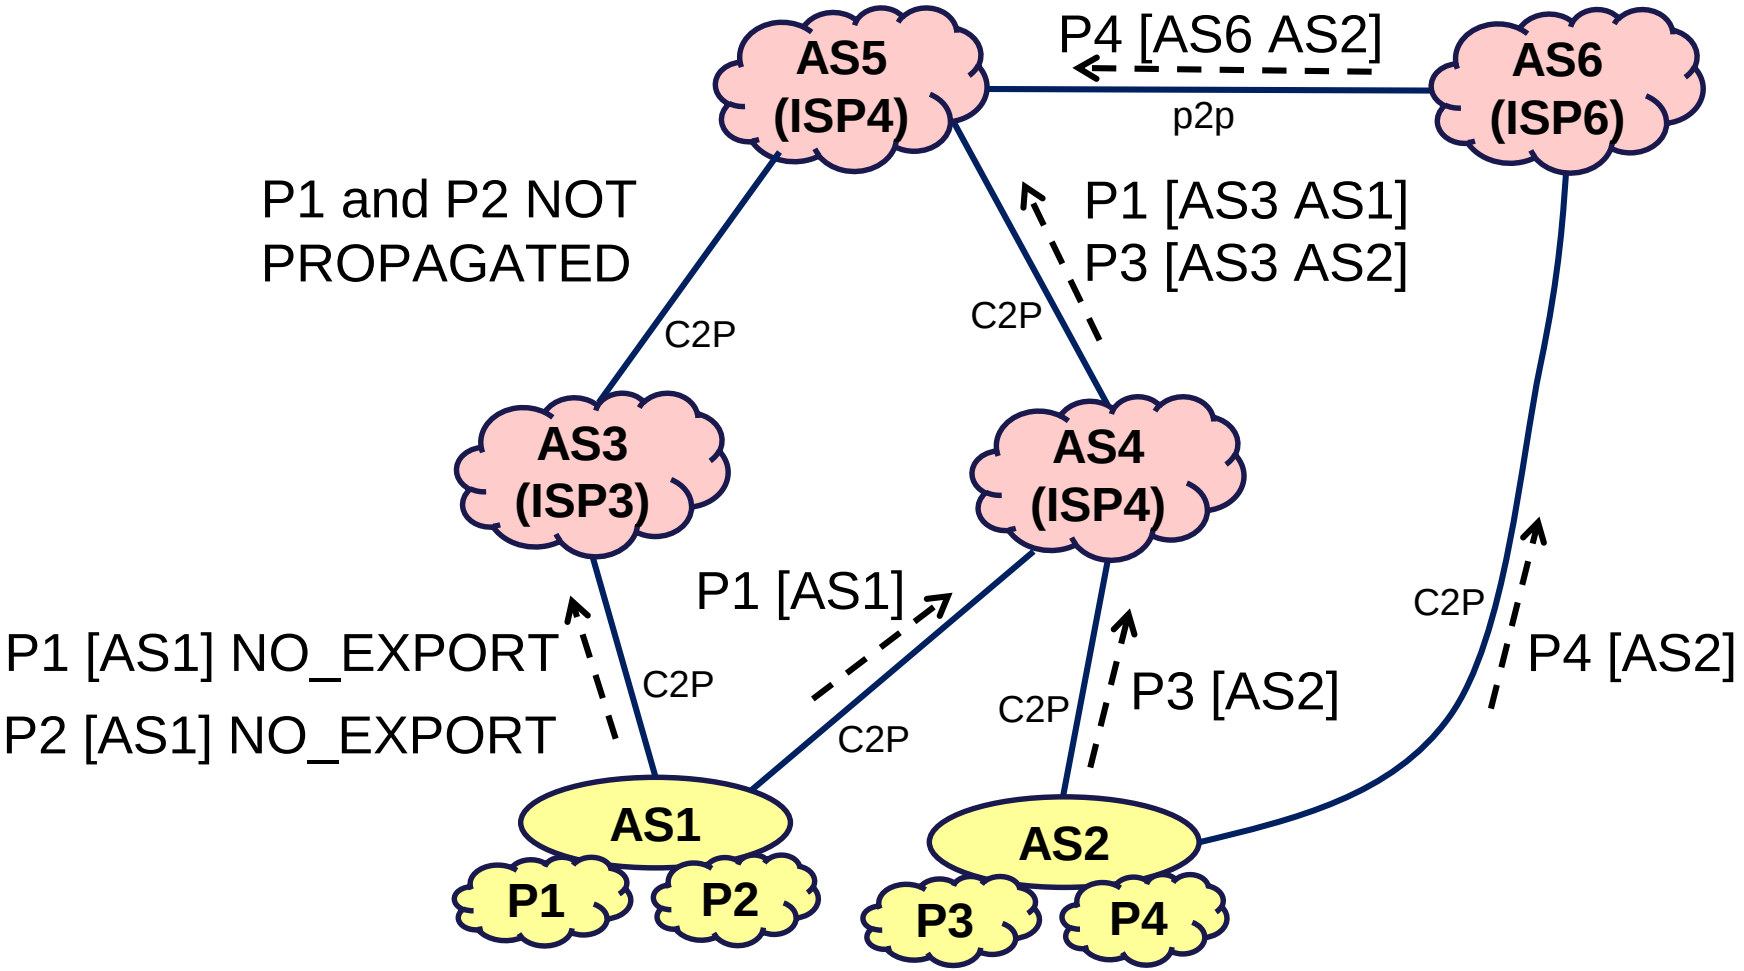

Customer Cone Size (# Prefixes) = RPF List Size

Type of ISP	Measured Customer Cone Size in # Prefixes (in turn this is an estimate for RPF list size on line card)
Very Large Global ISP	32392
Very Large Global ISP	29528
Large Global ISP	20038
Mid-size Global ISP	8661
Regional ISP (in Asia)	1101

References:
1. K. Sriram and R. Bush, "Estimating CPU Cost of BGPSEC on a Router", Presented at RIPE-63; also at IETF-83 SIDR WG Meeting, March 2012.
2. CAIDA AS ranking, http://as-rank.caida.org/

Available FIB Sizes in Router Line Cards

Type of ISP	Guesstimated Line Card FIB Memory Size (#prefixes) [cisco1][cisco2]
Very Large Global ISP	2M to 6M
Large Global ISP	1M
Mid-size Global ISP	0.5M
Regional ISP (in Asia)	100K

- RPF list sizes (slide 11) seem very small compared to the corresponding Line Card FIB sizes – correct?

[cisco1] https://www.cisco.com/c/en/us/support/docs/routers/asr-9000-series-aggregation-services-routers/116999-problem-line-card-00.html

[cisco2] https://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/5_x/nx-os/unicast/configuration/guide/l3_cli_nxos/l3_manage-routes.html#22859

Summary

- The proposal adds better logic to feasible path uRPF
- Performs well under various challenging scenarios
- We have given consideration to implementation feasibility
- Proposed method should help alleviate ISPs' concern about customer service disruption