Methods for Detection and Mitigation of BGP Route Leaks

ietf-idr-route-leak-detection-mitigation-00

(Route leak definition: draft-ietf-grow-route-leak-problem-definition)

K. Sriram, D. Montgomery, and B. Dickson

SIDR WG Meeting, July 24, 2015

Acknowledgements: The authors are grateful to many folks in various IETF WGs for commenting, critiquing, and offering very help suggestions (see acknowledgements section in the draft.)

Illustration of Basic Notion of a Route Leak

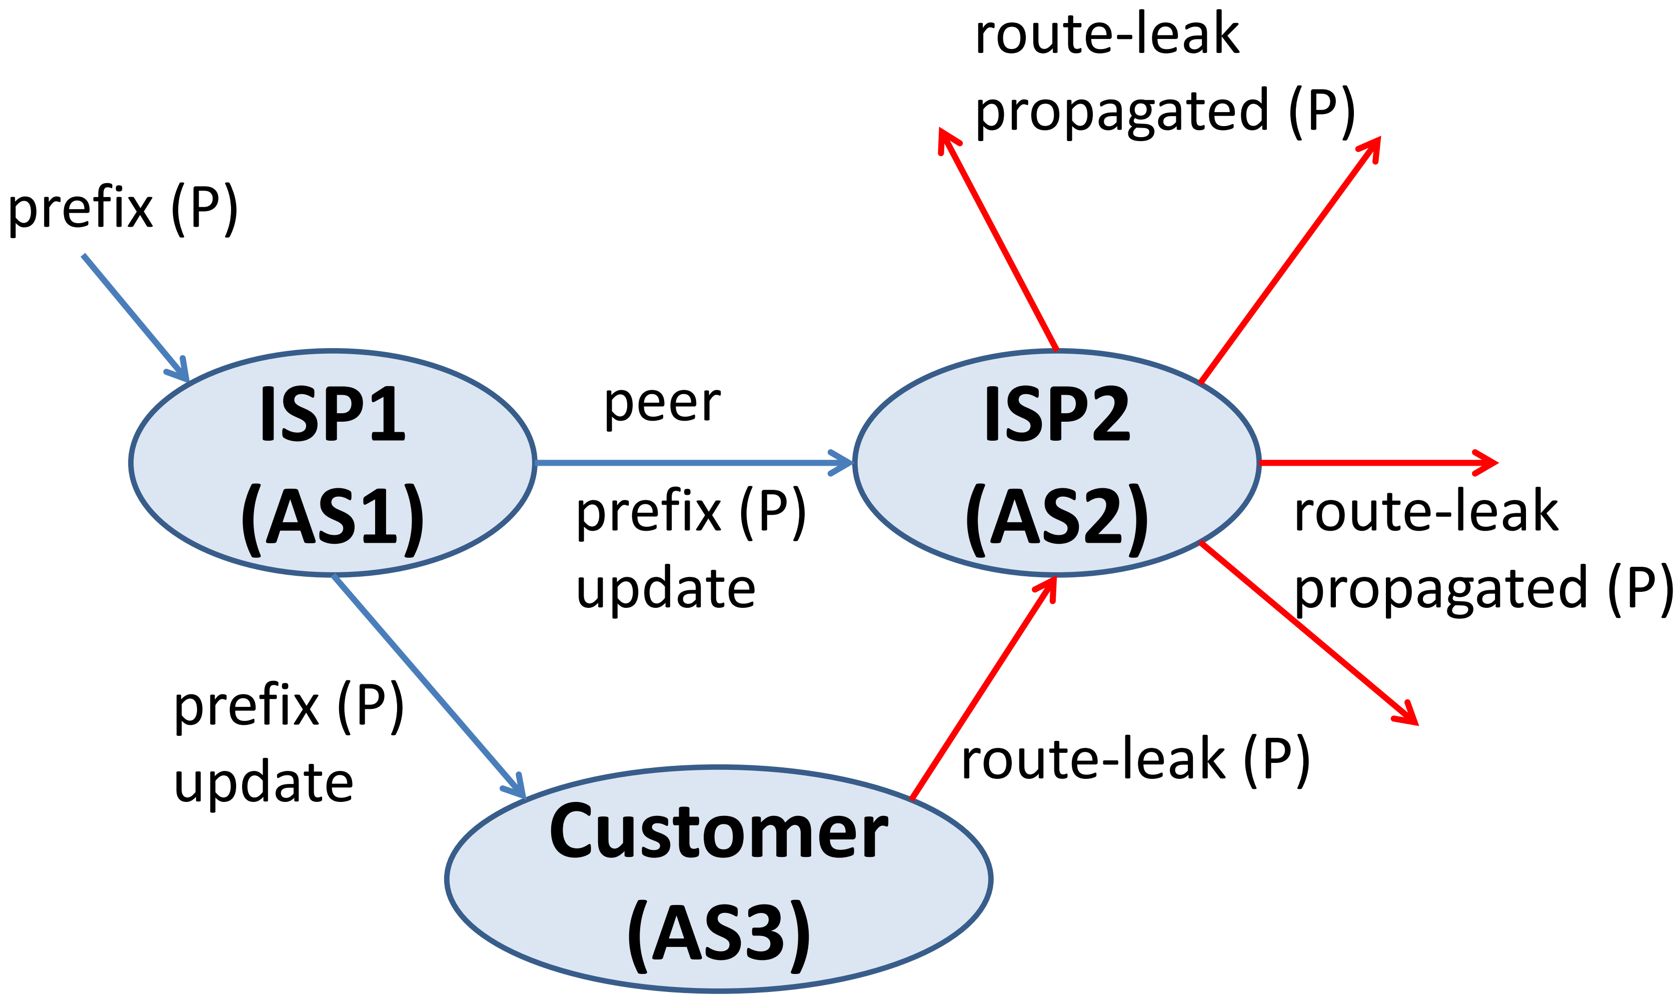

prefix (P)

ISP1 (AS1)

peer

prefix (P) update

ISP2 (AS2)

route-leak propagated (P)

route-leak propagated (P)

prefix (P) update

Customer (AS3)

route-leak (P)

In general, ISPs prefer customer route announcements over those from others.

Hathway / Airtel Route Leaks of Google Prefixes

March 12, 2015

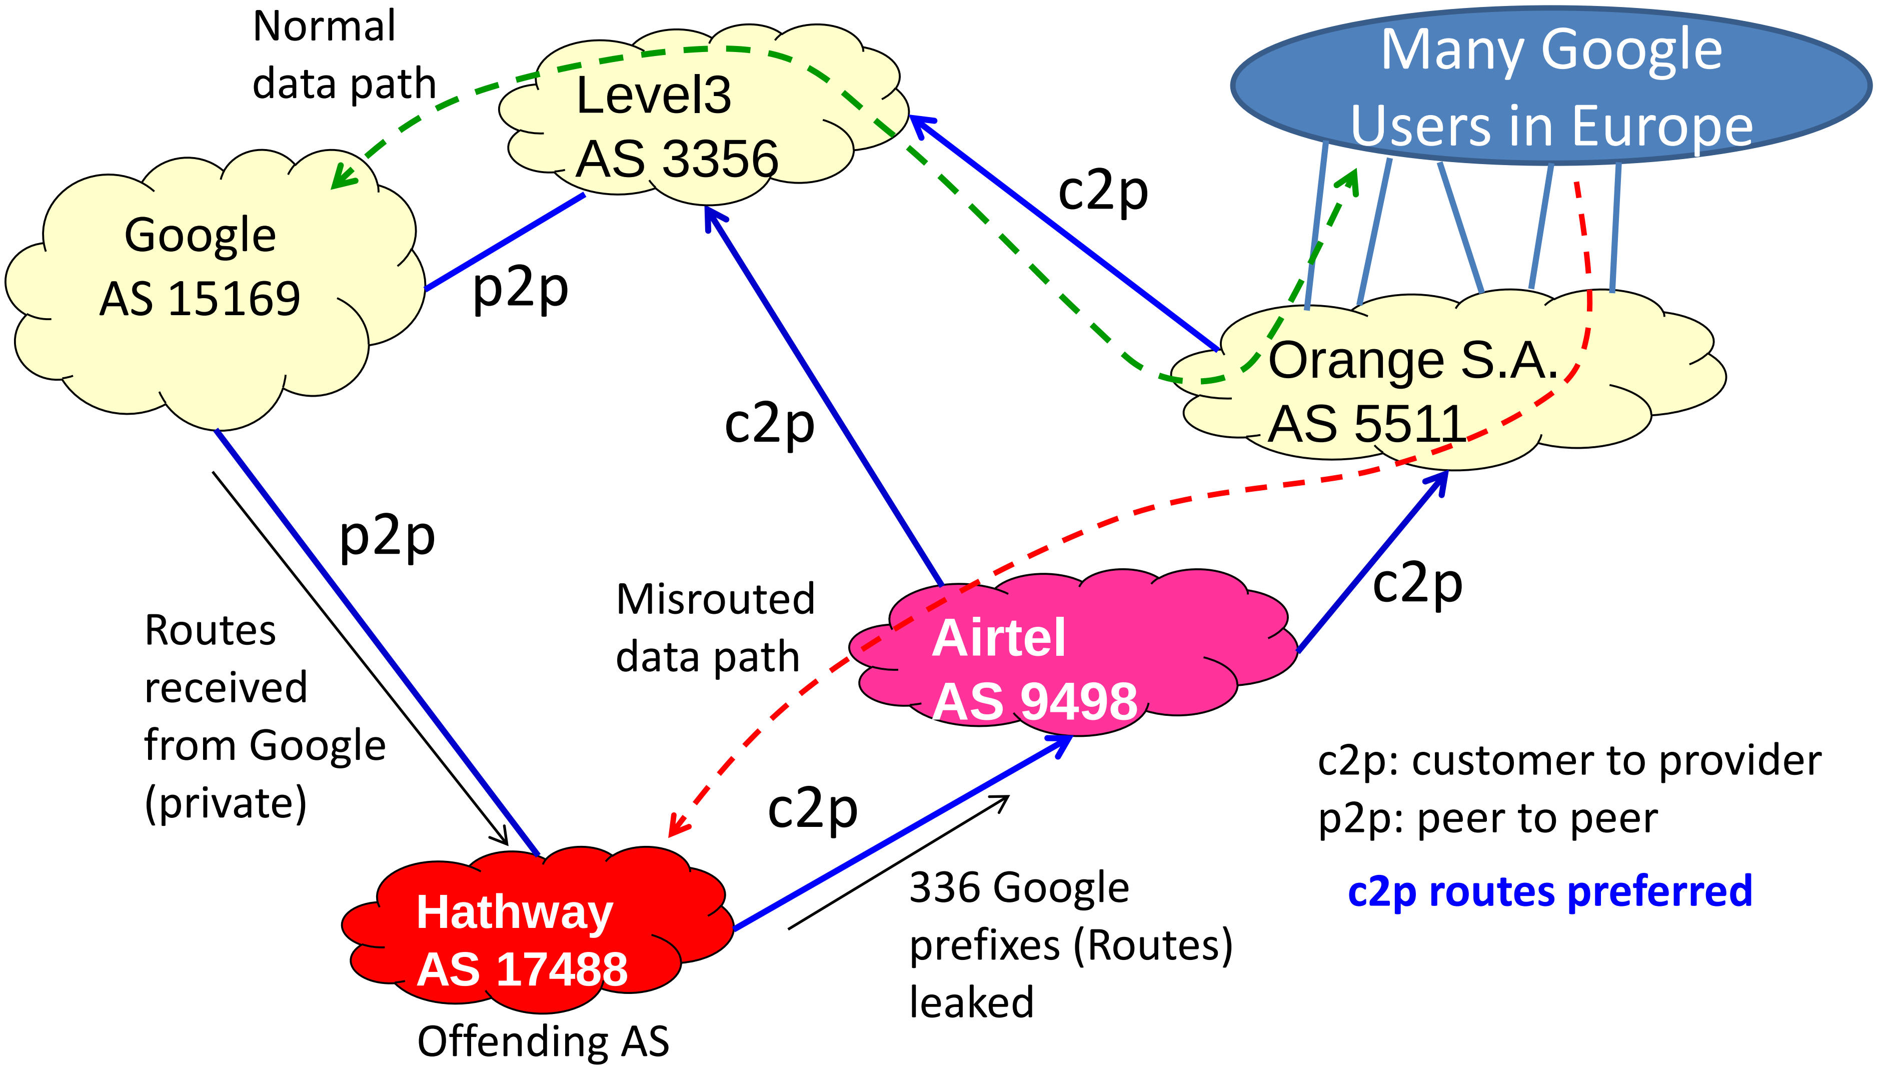

c2p: customer to provider
p2p: peer to peer

c2p routes preferred

Incident analysis: http://research.dyn.com/2015/03/routing-leak-briefly-takes-google/

Anatomy of a Route Leak: Seven Types

Type 1: U-Turn with Full Prefix

Type 2: U-Turn with More Specific Prefix

Type 3: Prefix Reorigination with Data Path to Legitimate Origin

Type 4: Leak of Internal Prefixes and Accidental Deaggregation

Type 5: Lateral ISP-ISP-ISP Leak

Type 6: Leak of Provider Prefixes to Peer

Type 7: Leak of Peer Prefixes to Provider

Details and example incidents provided in:
draft-ietf-grow-route-leak-problem-definition-02

Route Leak Detection/Mitigation in Origin Validation and BGPsec

Type of Route Leak	Detection Coverage
Type 1: U-Turn with Full Prefix	None
Type 2: U-Turn with More Specific Prefix	Origin Validation (partial); BGPsec (100% detection)
Type 3: Prefix Reorigination with Data Path to Legitimate Origin	Origin Validation (100% detection); BGPsec does not detect
Type 4: Leak of Internal Prefixes and Accidental Deaggregation	Origin Validation (partial); BGPsec does not detect
Type 5: Lateral ISP-ISP-ISP Leak	None
Type 6: Leak of Provider Prefixes to Peer	None
Type 7: Leak of Peer Prefixes to Provider	None

Solution lies in RPKI/OV

Data still flows to the legitimate AS (no detour)

Basic Idea and Mechanism – Date back to 1980's

- "Information flow rules" described in "Proceedings of the April 22-24, 1987 Internet Engineering Task Force"
- "Link Type" described in RFC 1105 (obsolete), June 1989
- "Hierarchical Recording" described in "Inter-Domain Routing Protocol (IDRP)", IETF Internet Draft (expired), November 1994.
- BGPsec based solution to detect accidental and malicious route leaks
  - ➢ Discussed in the SIDR WG since 2011
  - ➢ Documented by Brian Dickson in 2012:

    https://tools.ietf.org/html/draft-dickson-sidr-route-leak-def-03 (expired)

    http://tools.ietf.org/html/draft-dickson-sidr-route-leak-reqts-02 (expired)

    https://tools.ietf.org/html/draft-dickson-sidr-route-leak-solns-01 (expired)

Basic Design Principle for Route Leak Detection

RLP is Route Leak Protection field; Transitive per hop attribute

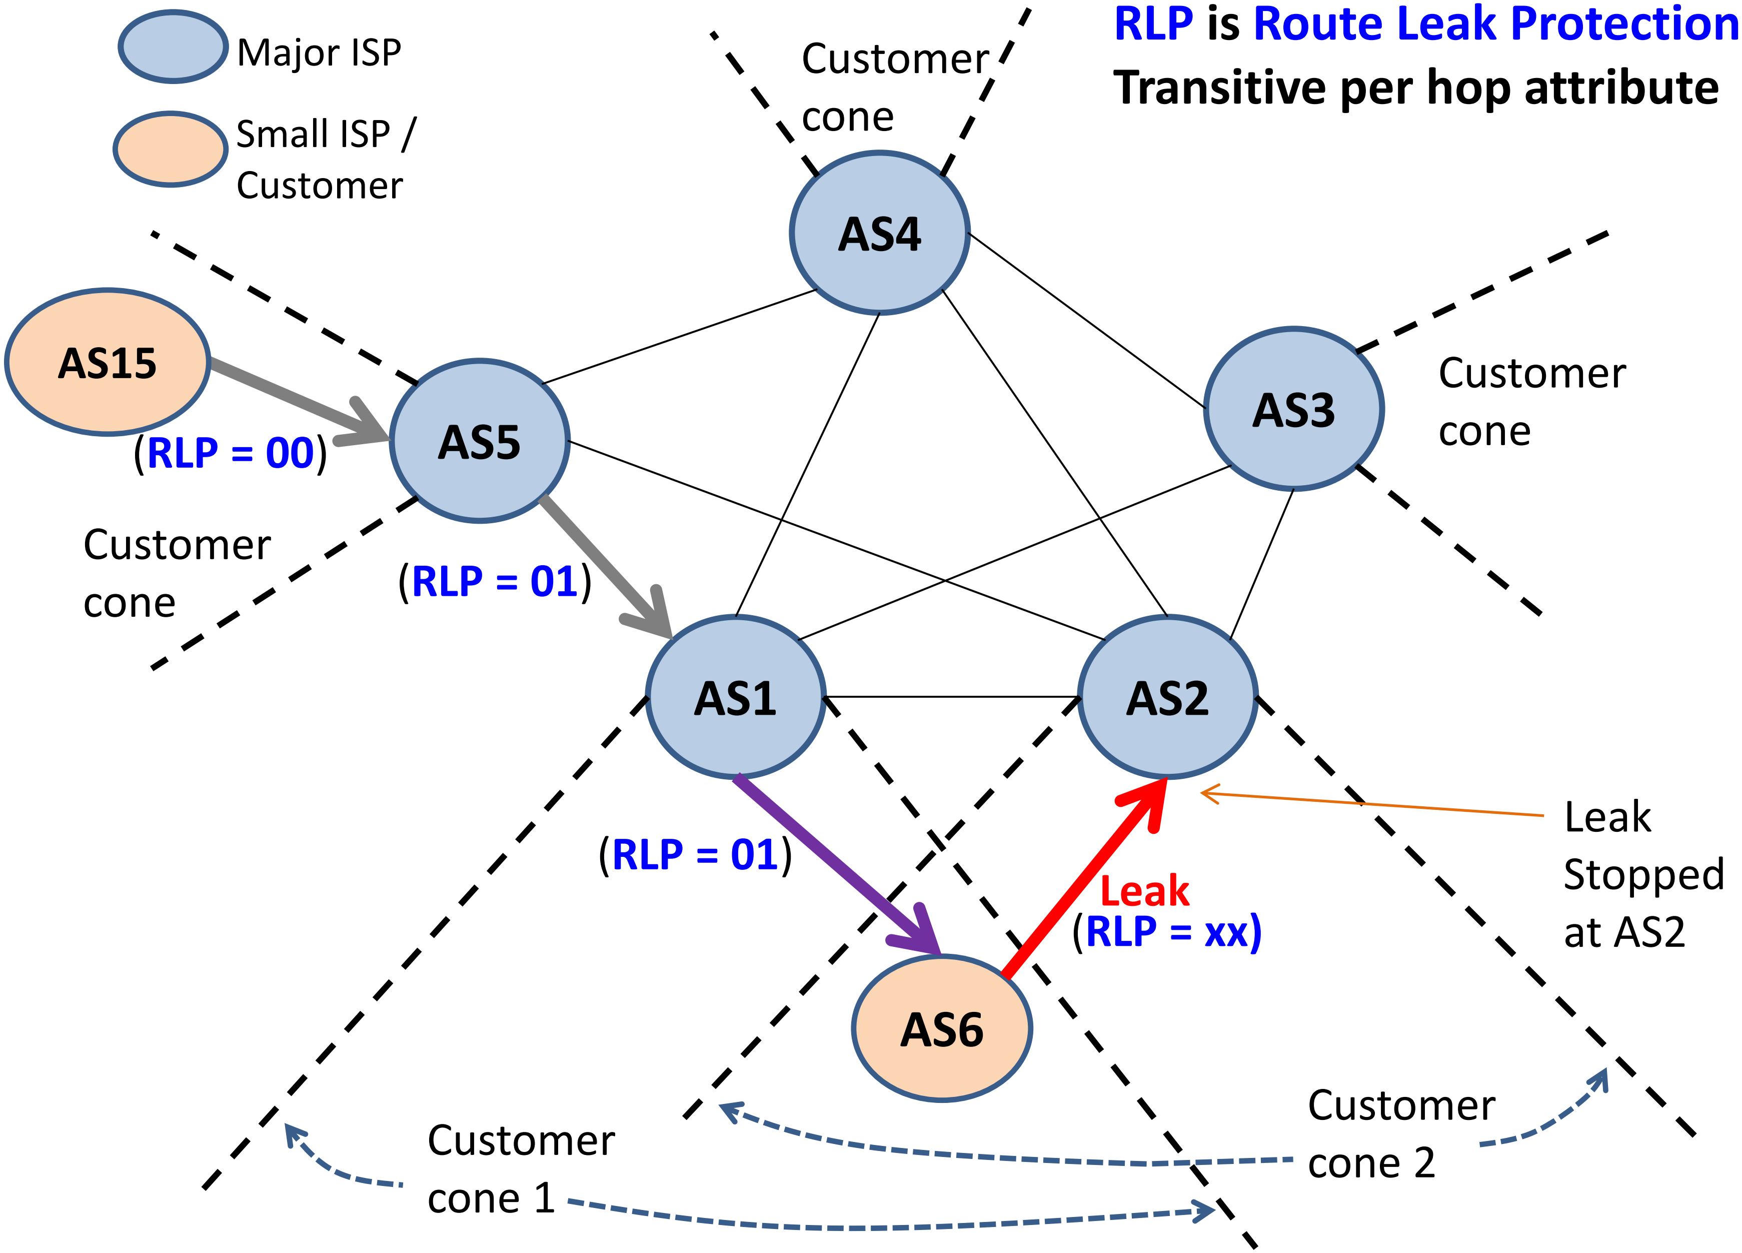

Route Leak Protection (RLP) Field Encoding by Sending Router

- RLP is proposed to be a 2-bit field set by each AS along the path
- Can be carried as a transitive per hop attribute in BGP or in the existing Flags field in BGPsec (TBD)
- The RLP field value SHOULD be set to one of two values as follows:
  - **00:** This is the **default value** (i.e. "nothing specified"),
  - **01:** This is the **'Do not Propagate Up or Lateral'** indication; sender indicating that the prefix-update SHOULD NOT be subsequently forwarded 'Up‘ towards a provider or to a ‘Lateral’ peer
  - 10 and 11 values are for possible future use.

Sending Router’s Intent

- Note: There is no explicit disclosure about the nature of a peering relationship.
- By setting RLP indication to **01**, merely asserting that this prefix-update that I’ve forwarded to my neighbor **SHOULD NOT** be propagated ‘Up’ (i.e. on a c2p link) or ‘Lateral’ (i.e. on a p2p link) by said neighbor or any subsequent AS in the path of update propagation.

Recommended Receiver Action for Detection of Route Leaks of Types 1, 2, 5, 6 and 7

Receiving router SHOULD mark an update a Route-Leak if ALL of the following conditions hold true:

a) The update is received from a customer or lateral-peer AS

b) The update is ‘Valid’ per RPKI-OV and BGPsec (BGPsec path validation not applicable if update not signed)

c) The update has the RLP field set to '01' indication for one or more hops (excluding the most recent) in the AS path.

Note: Reason for “excluding the most recent” – an ISP should look at RLP values set by ASes preceding the customer AS in order to ascertain a leak .

An Example Receiver Action for Mitigation of Route Leaks

- If a prefix-route from a customer AS or a peer AS is detected and marked as a “Route-Leak”, then the receiving router SHOULD prefer an alternate unmarked prefix-route if available
- If no alternate unmarked prefix-route is available, then the prefix-route marked as a “Route-Leak” MAY be accepted

This in only an example. We do not specify receiver action for mitigation as it may vary based on operator policy.

Adoption and Path for Success

- Mid and large size ISPs can participate early, and be the key detection/mitigation points for route leaks.
- More the ISPs that adopt, greater the success (benefits accrue incrementally).

Note: In a case like that of Hathway/Airtel leak of Google prefixes (see Slide 3), the attack is mitigated if Google would set its RLP field value to 01 in its prefix update announcement to Hathway, and Airtel would in turn use the receiver action recommended on Slide 12 to detect the leak from Hathway.

Accidental vs. Intentional (Malicious) Route Leaks & Solution Steps

Today: Current BGP (without route leak solution; assuming prefix filters aren’t doing job adequately)

- Vulnerable to accidental (99%) and malicious (1%) route leaks

⬇

Step 1: BGP with proposed route leak solution (with RPKI/OV but without BGPsec)

- Detects/mitigates accidental (99%) but not malicious (1%)

⬇

Step 2: BGP with proposed route leak solution (with RPKI/OV and BGPsec)

- Detects/mitigates accidental (99%) as well as malicious (1%)

Is there a new attack vector in using RLP bits without security (BGPsec)? (1 of 2)

Upgrade Attack: RLP ‘01’ → RLP ‘00’ to avoid route leak detection

- For a prefix-route that keeps propagating in the ‘Down” (p2c) direction, this poses no problem
- When propagated ‘Up’ (c2p) or ‘Lateral’ (p2p), the worst that can happen is that a route leak goes undetected
- No worse than BGP today
- Less than 1% of all route leaks may go undetected in this manner (malicious intent or faulty implementation)

Is there a new attack vector in using RLP bits without security (BGPsec)? (2 of 2)

Downgrade Attack: RLP ‘00’ → RLP ‘01’

Result: A prefix-route is mis-detected as a route leak, but …

- Default is RLP set to ‘00’ – that helps reduce errors of this kind
- Every AS or ISP wants reachability for prefixes it originates; so it is not likely to set RLP ‘01’ intentionally
- Receiver would prefer an alternate ‘clean’ prefix-route from a provider or peer over a ‘marked’ prefix-route from a customer; may end up with a suboptimal path
- In order to have reachability, receiver would accept ‘marked’ prefix-route if there is no alternative that is clean
- Low probability that all received routes for a prefix are detected as ‘route leaks’. If it happens, need a tie breaker policy to prefer one (up to the operator to choose a mitigation algorithm)

Summary and Conclusion

- Identified categories of route leaks
- Some of these are already mitigated in OV or basic BGPsec
- Presented an enhancement of BGP that detects and mitigates all route leaks (when combined with Origin Validation)
- Analyzed whether RLP bits (without BGPsec) could become a new attack vector and extent of damage
- RLP field should be protected in order to detect and mitigate malicious route leaks
    - ✓ RLP field can be placed in existing Flags field and protected under path signatures in BGPsec

Backup Slides

Questions at the mike at IDR WG Mtg. in Dallas

- Wes George: Have you considered if techniques in RFC 7454 “BGP Operations and Security” may be adequate to address route leaks?
  - ➢ Answered in section 5.2 in the draft
- Keyur asked about combining the proposed RLP solution with AS path filtering and ORF techniques.
  - ➢ Also answered in section 5.2 in the draft

Discussion & Examples – How it works!

Example 1: Multi-homed Customer Leak

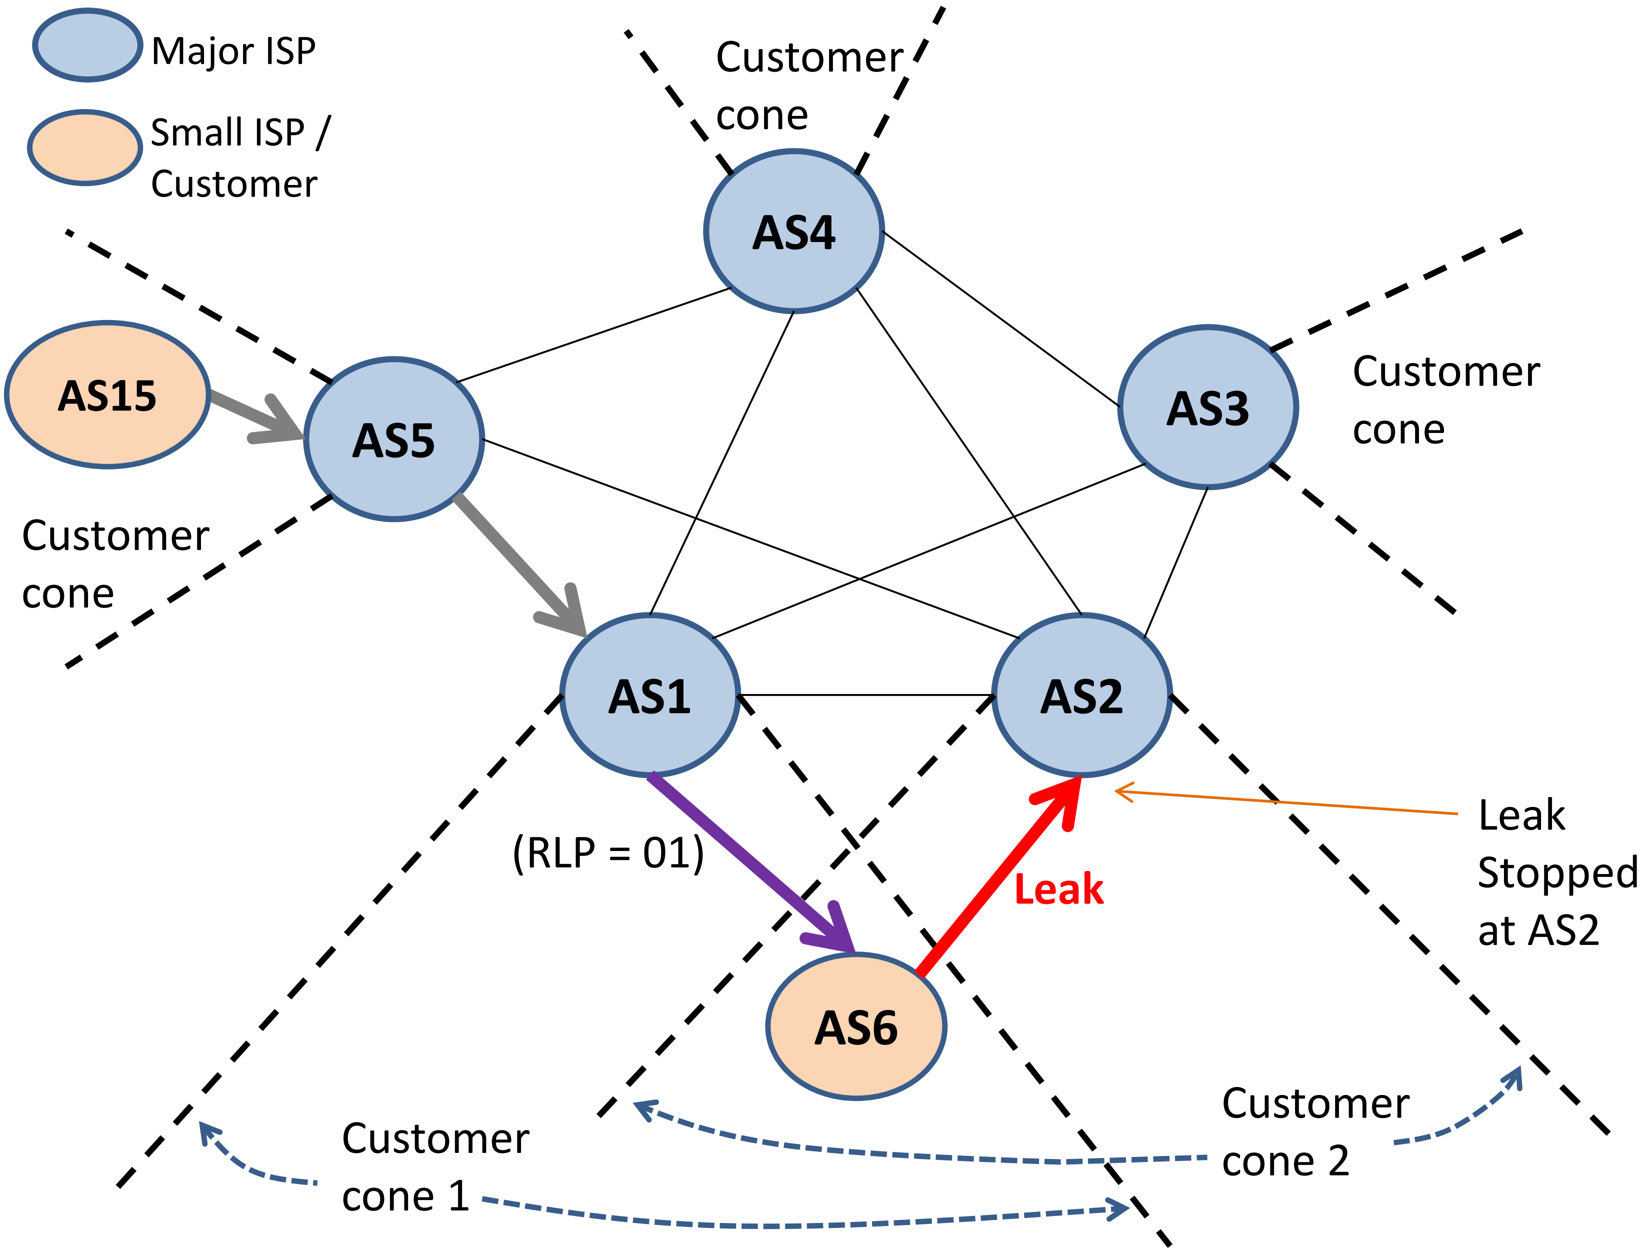

Example 2: Lateral Across Customer Cones and Then Leaked Up to Other ISP

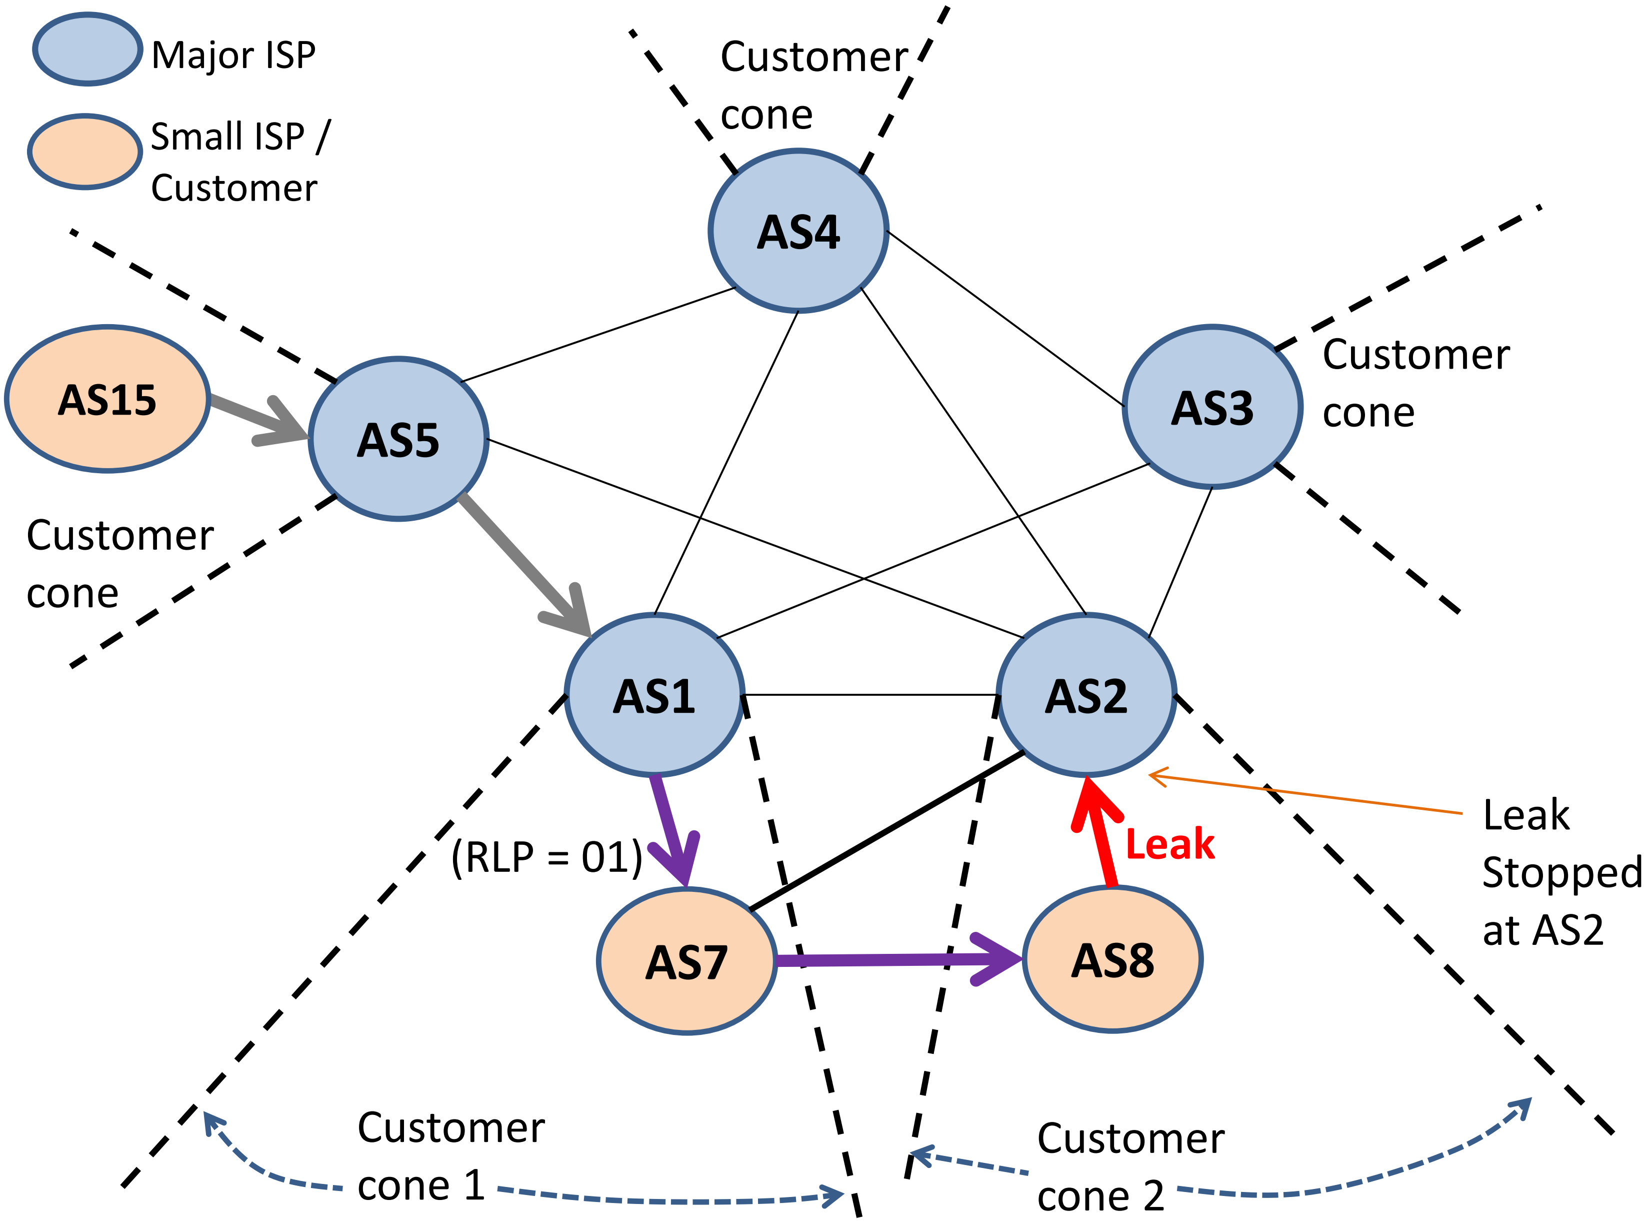

Example 3: Customer's Customer is Multi-homed and Leaks

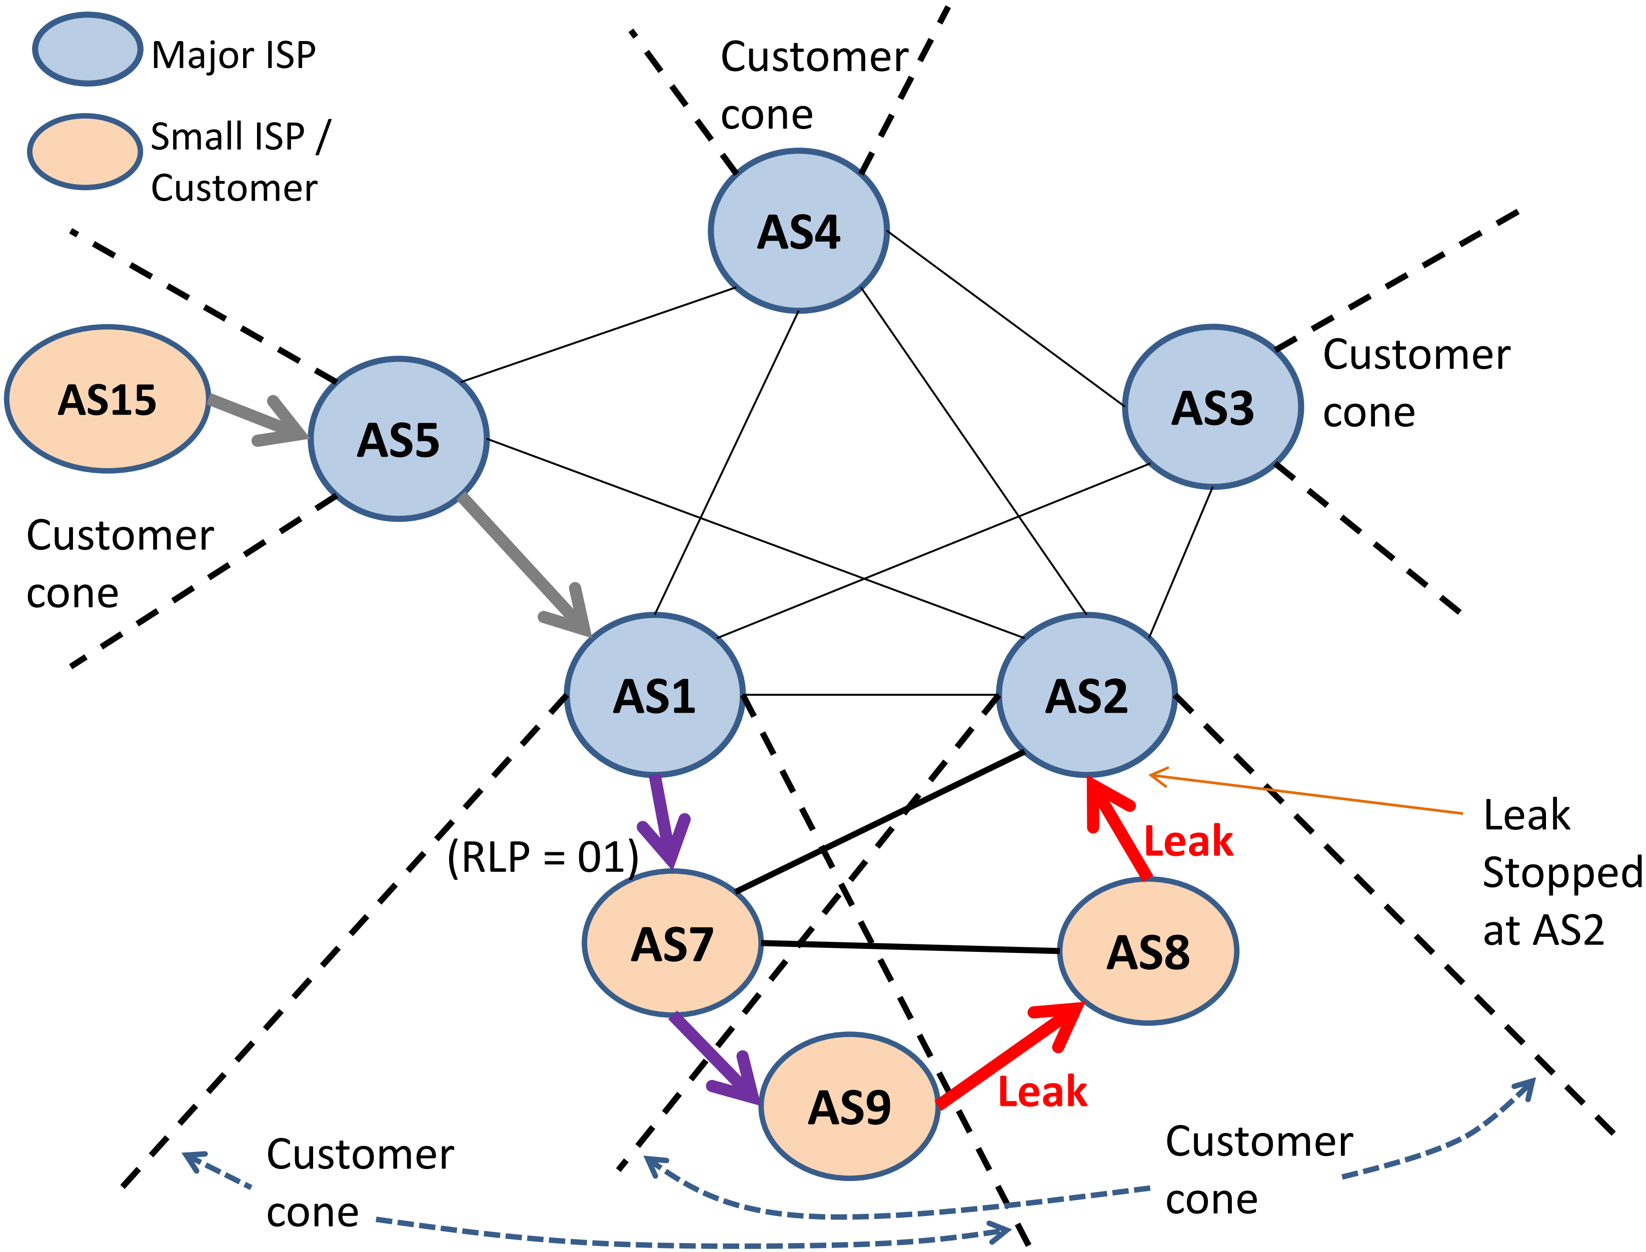

Consideration of DDoS Mitigation Service Provider

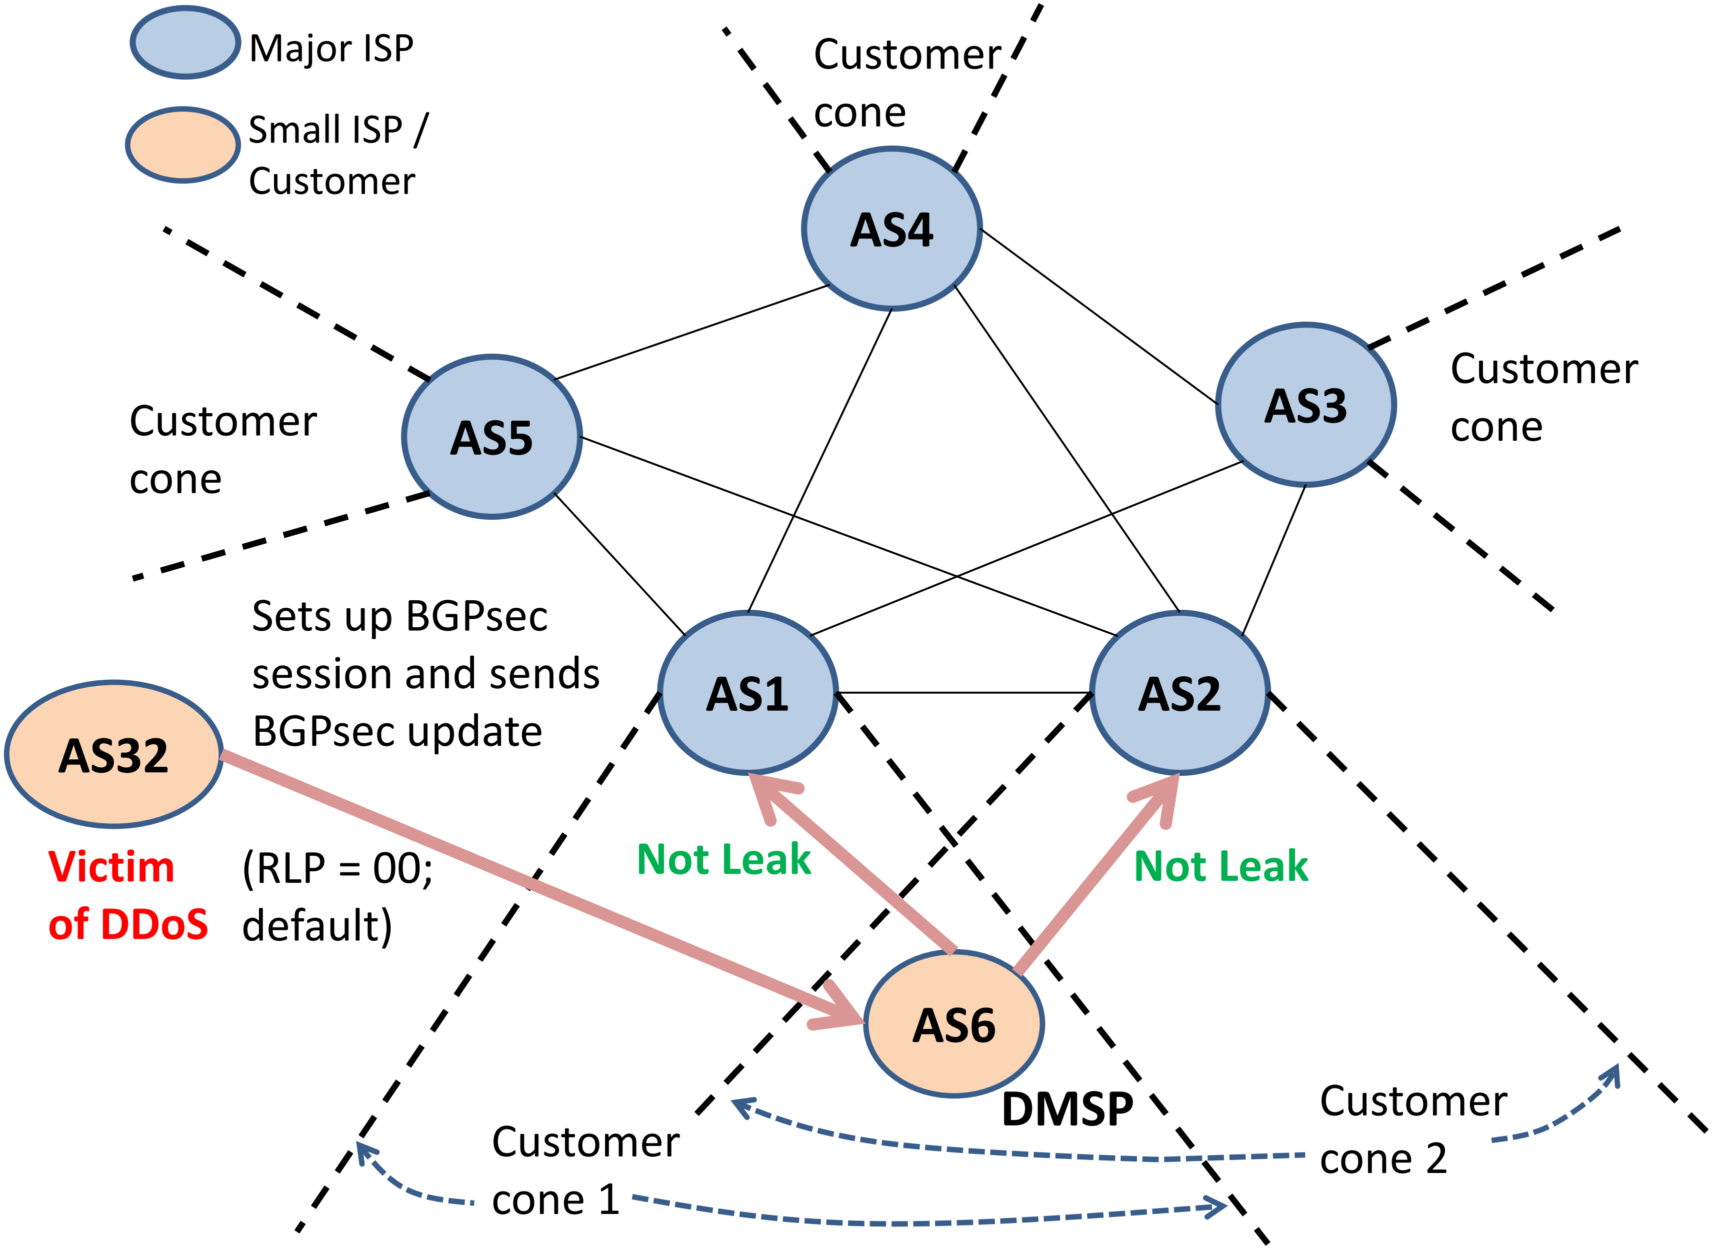

Stopgap Solution when Only Origin Validation is Deployed

Construction of Prefix Filter List from ROAs

1. ISP makes a list of all the ASes (Cust_AS_List) that are in its customer cone (ISP’s own AS is also included in the list)
2. ISP downloads from the RPKI repositories a complete list (Cust_ROA_List) of valid ROAs that contain any of the ASes in Cust_AS_List
3. ISP creates a list of all the prefixes (Cust_Prfx_List) that are contained in any of the ROAs in Cust_ROA_List
4. Cust_Prfx_List is the allowed list of prefixes that are permitted by the ISP's AS, and will be forwarded by the ISP to upstream ISPs, customers, and peers
5. Any prefix not in Cust_Prfx_List but announced by any of the ISP’s direct customers is not permitted to be propagated upstream

Exception to the Rule in Case of DDoS Mitigation

- DDoS Mitigation Service Provider (DMSP) requires exemption from the rule of Cust_Prfx_List described in the previous slide
- ISP and the DMSP make a prior arrangement on this
- DMSP can propagate upstream to the ISP any prefix-update it receives from its DDoS'ed customer (in emergency), and the ISP will not treat it as a route leak
- This helps prevent any disruption or delay in the DMSP's mitigation services under emergency scenarios