

Methods for Detection and Mitigation of BGP Route Leaks

draft-sriram-idr-route-leak-detection-mitigation-00

K. Sriram and D. Montgomery

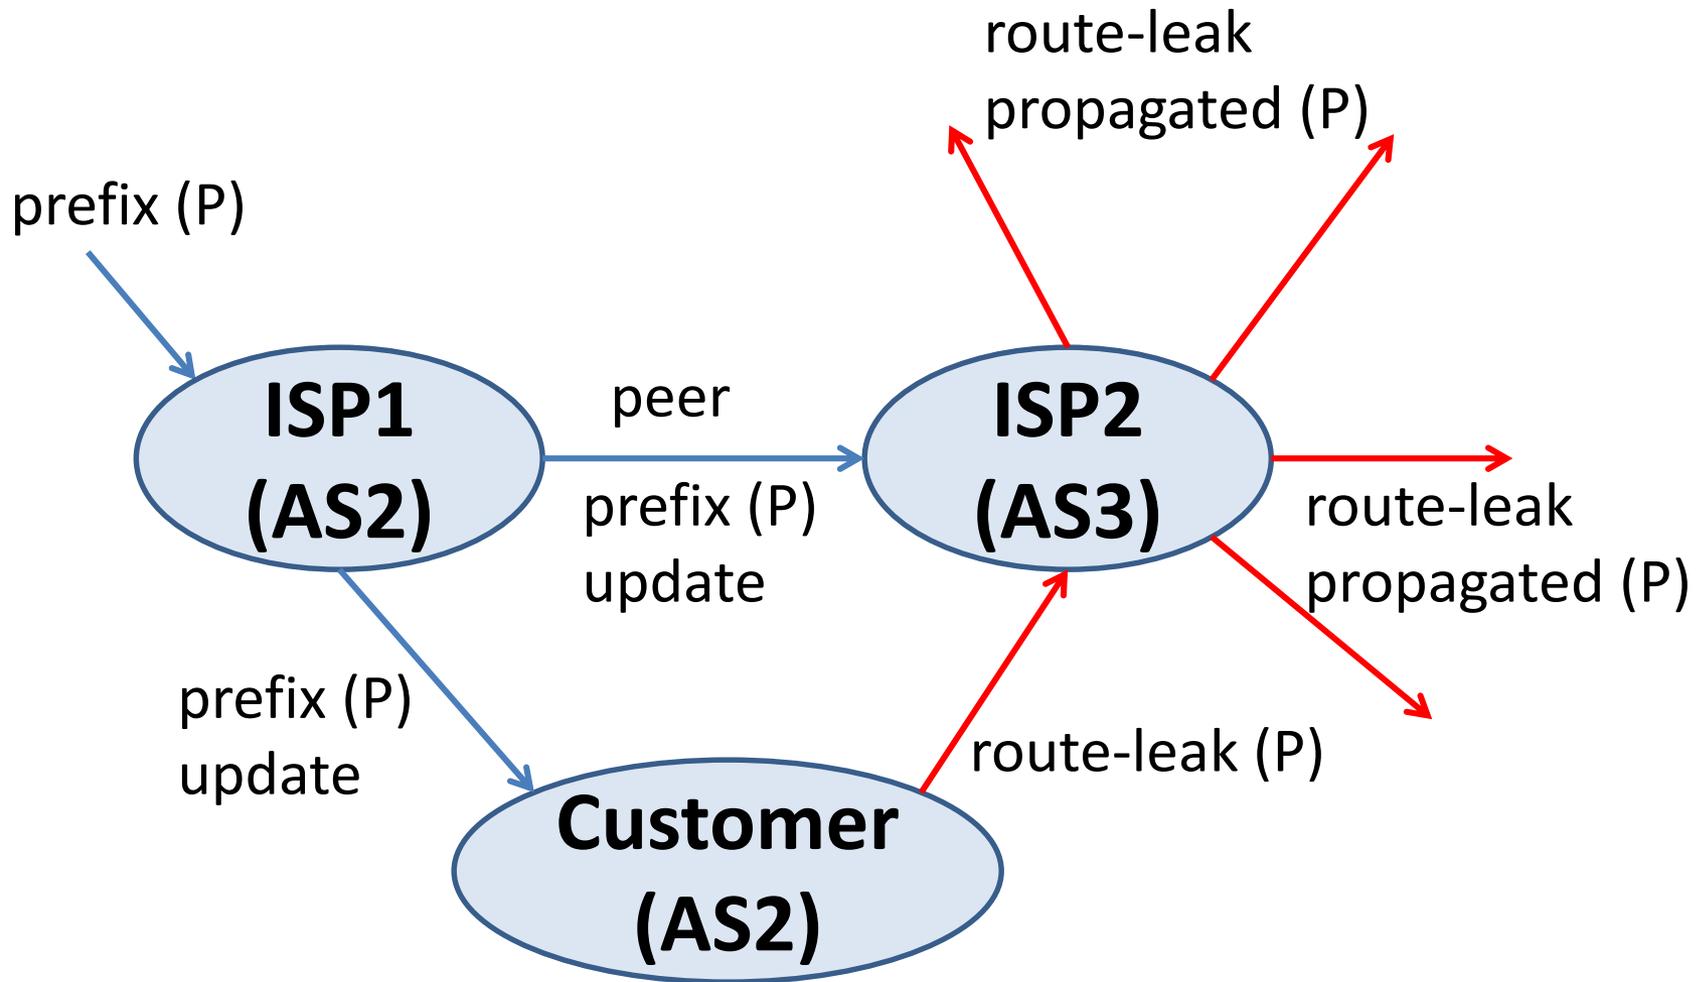
NIST

IDR WG Meeting, IETF 92, Dallas, Texas

March 24, 2015

Acknowledgements: The authors would like to thank Danny McPherson, Eric Osterweil, Jeff Haas, Warren Kumari, Jared Mauch, Amogh Dhamdhere, Andrei Robachevsky, Brian Dickson, Randy Bush, Chris Morrow, and Sandy Murphy for comments and suggestions.

Illustration of Basic Notion of a Route Leak



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

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-01**

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

*BGPSEC protocol performs path validation only, and does not include OV (spec version 11)

Begin Sender Specification

(Simple Enhancement to Existing BGP or BGPSEC)

Route Leak Protection (RLP) Field Encoding by Sending Router (Method 1)

- RLP is proposed to be a 2-bit field set by each AS along the path
- Can be carried in a Transitive Community attribute in BGP or in the 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' indication; sender indicating that the prefix-update SHOULD NOT be subsequently forwarded 'Up' towards a provider AS,
 - 10 and 11 values are for possible future use.

Route Leak Protection (RLP) Field Encoding by Sending Router (Method 2)

Only the following is different w.r.t. Method 1:

- 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: "Do not Propagate Up" indication
 - 10: "Propagate to Customers Only" indication
 - 11: "Do not Propagate" (i.e. NO_EXPORT)

Agreeing on the semantics of these indications is important.

End of Sender Specification

Sending Router's Intent

- Note: There is no explicit disclosure about the nature of a peering relationship.
- (In Method 1) 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) by said neighbor or any subsequent AS in the path of update propagation.

Recommendation for Receiver Action for Detection of Route Leaks of Types 1, 2 and 7 (When Sender is using Method 1)

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 AS.
- b) The update has the RLP field set to '01' (i.e. 'Do not Propagate Up') 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 .

Recommendation for Receiver Action for Detection of Route Leaks of Types 5 and 6 (When Sender is using Method 1)

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

- a) The update is received from a peer AS.
- b) The update has the RLP field set to '01' indication for one or more hops (excluding the most recent) in the AS path.

Note: In this case, the RLP indication of '01' is more strictly interpreted to mean that the update should not be propagated on a lateral peer link either.

An Example Receiver Action for Mitigation of Route Leaks

- If an update from a customer AS or a peer AS is detected and marked as a “Route-Leak”, then the receiving router SHOULD prefer an unmarked update from another neighbor AS, if available.

Path for Success

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

Note: In a case like that of Moratel's leak (in November 2012) of Google's prefixes, the attack is mitigated if Google would set its RLP field value to 01 in its prefix update announcement to Moratel, and PCCW would in turn use the receiver action recommended on Slide 11 to identify the update from Moratel as a Route Leak.

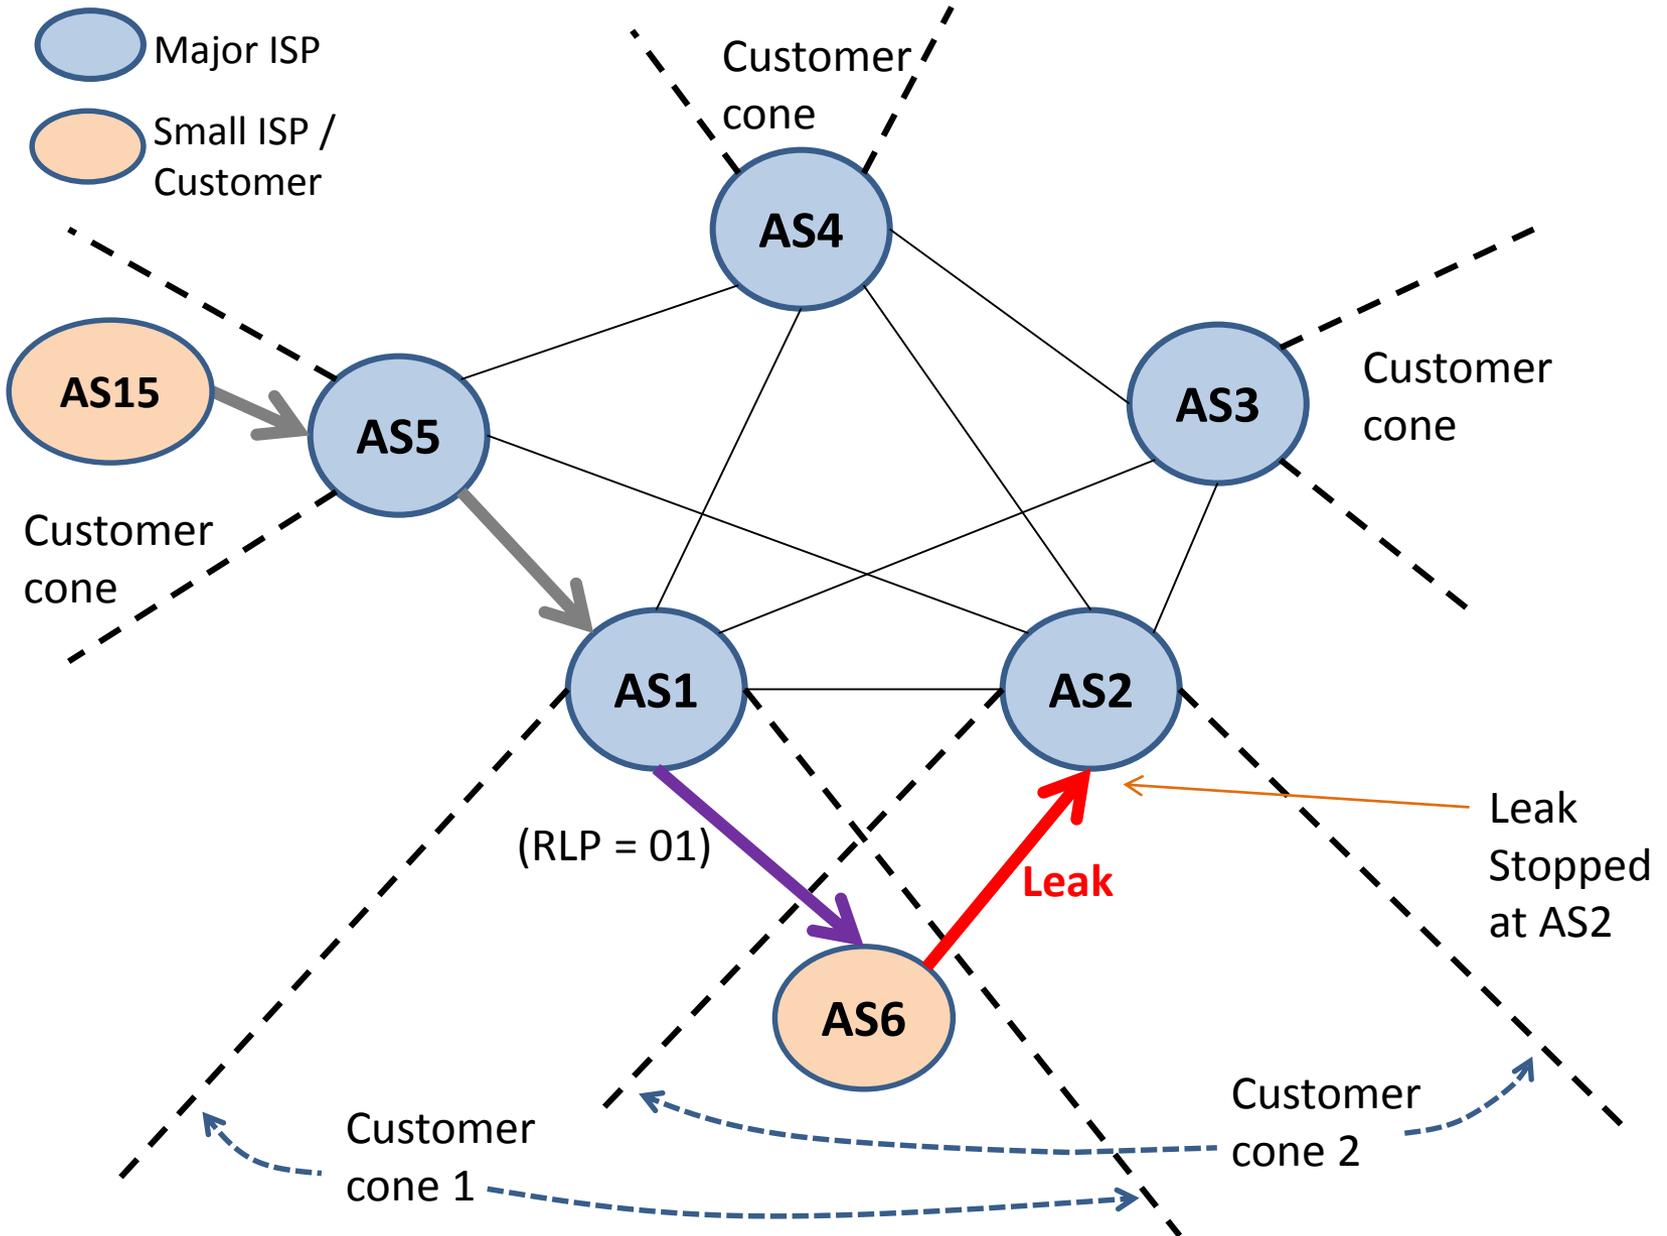
Summary and Conclusion

- Identified categories of route leaks
- Some of these are already mitigated in OV or BGPSEC
- Presented an enhancement of BGP that detects and mitigates all route leaks (when combined with Origin Validation)
- The RLP field can be carried in a Transitive Community Attribute or in BGPSEC Flags field
- The RLP field may need to be protected in order to prevent tampering and/or malicious route leaks

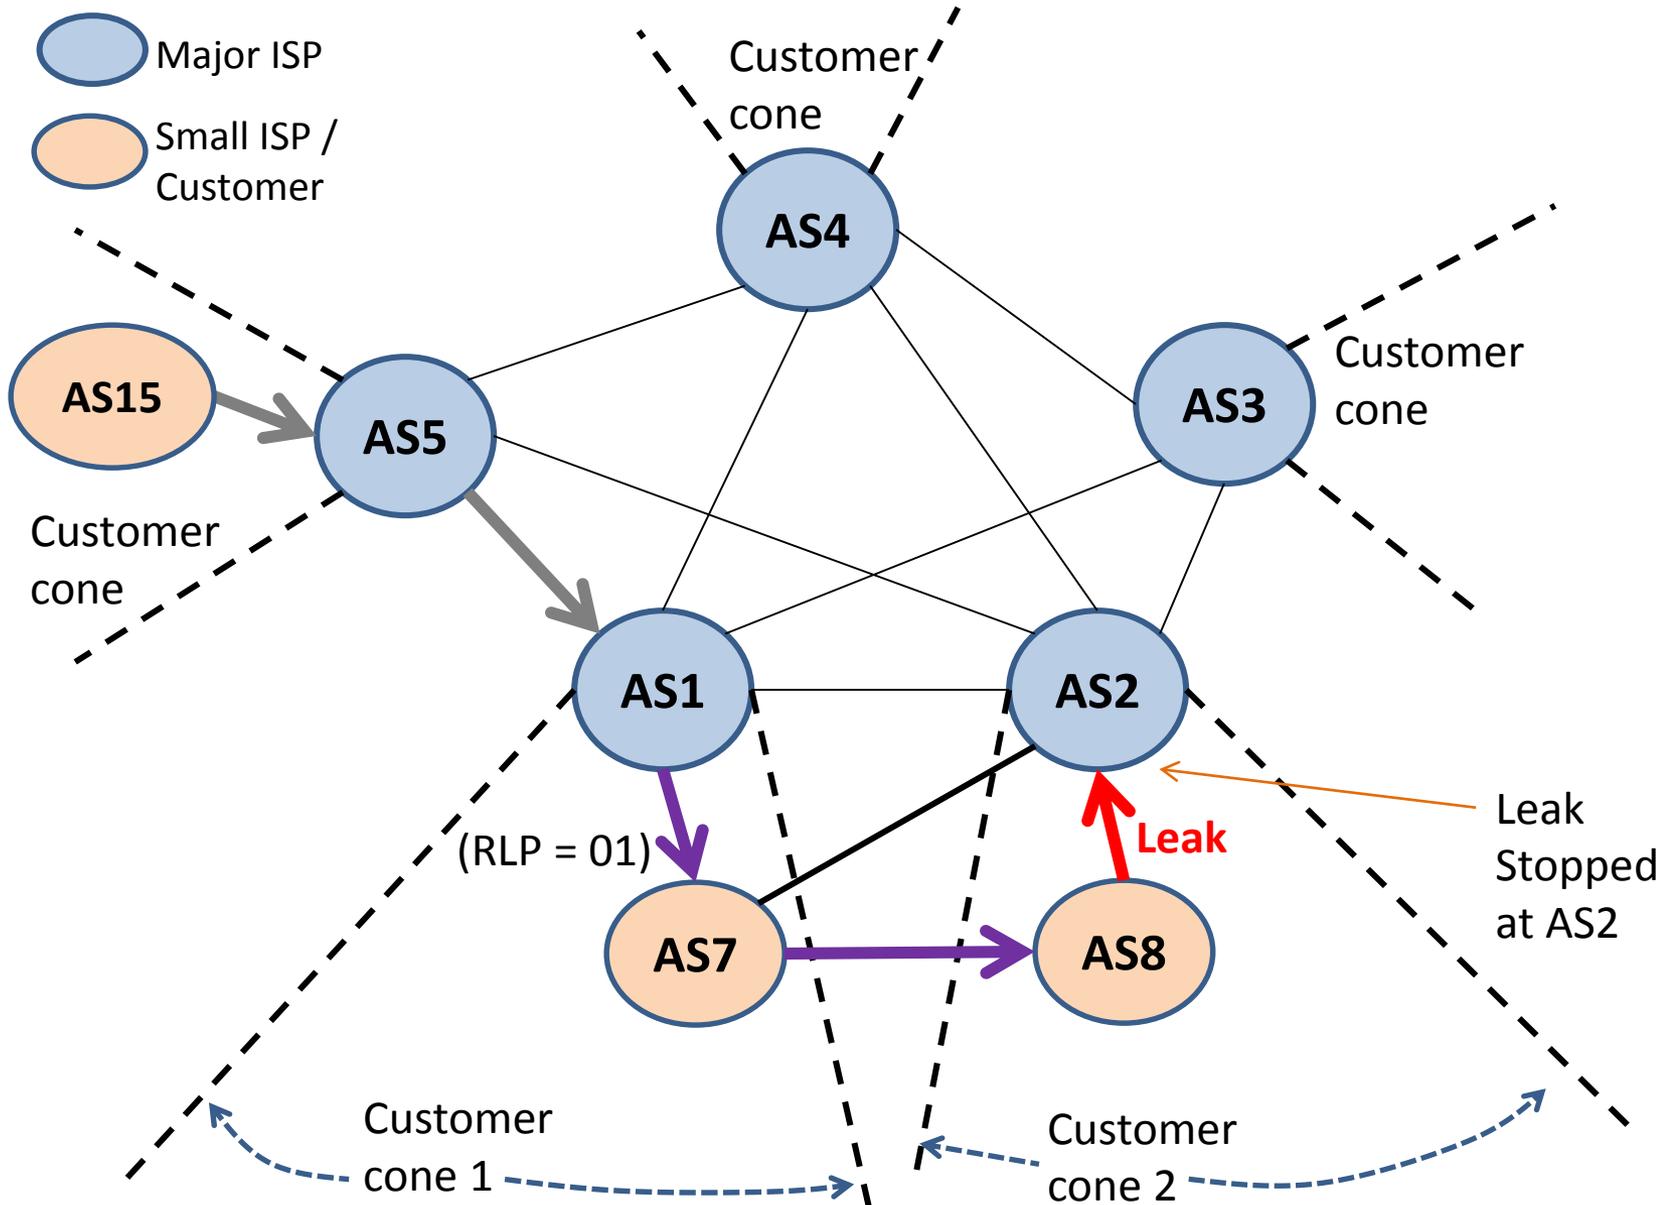
Backup Slides

Discussion & Examples – How it works!

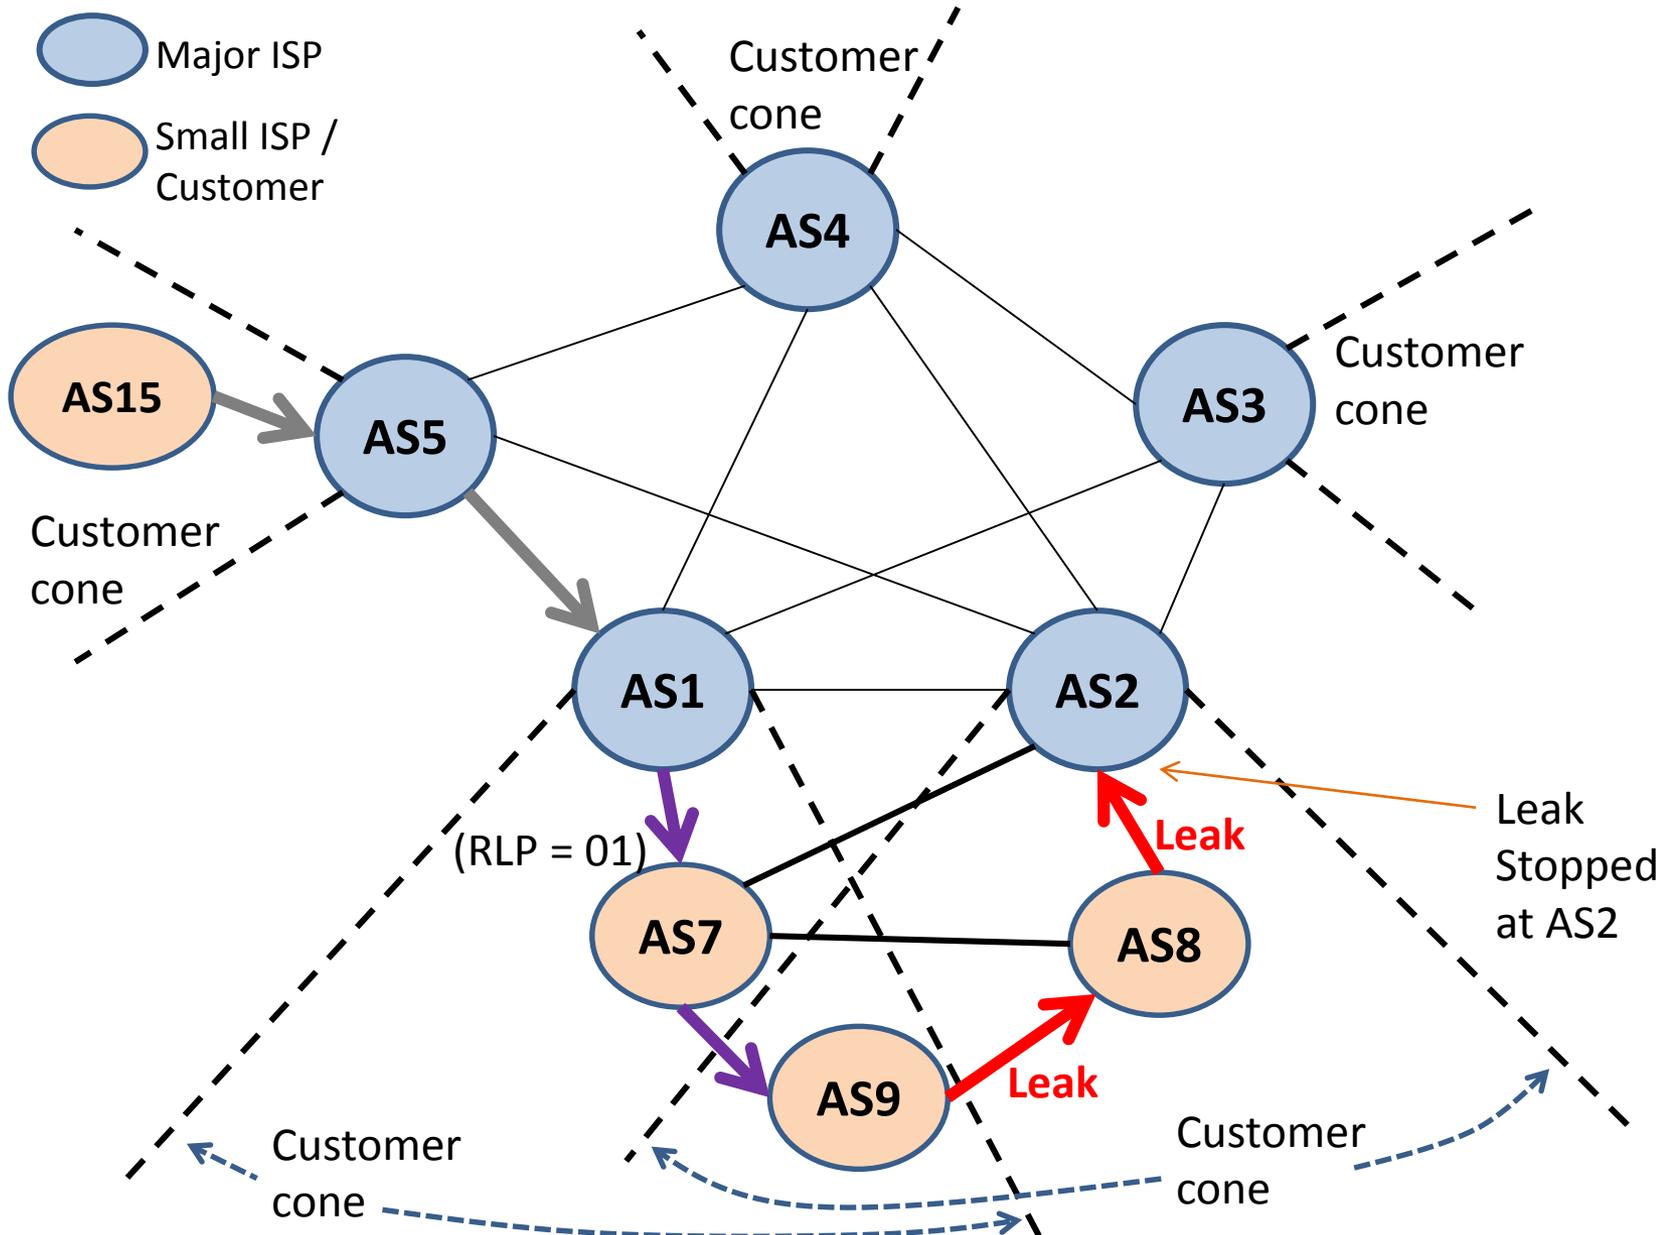
Example 1: Multi-homed Customer Leak



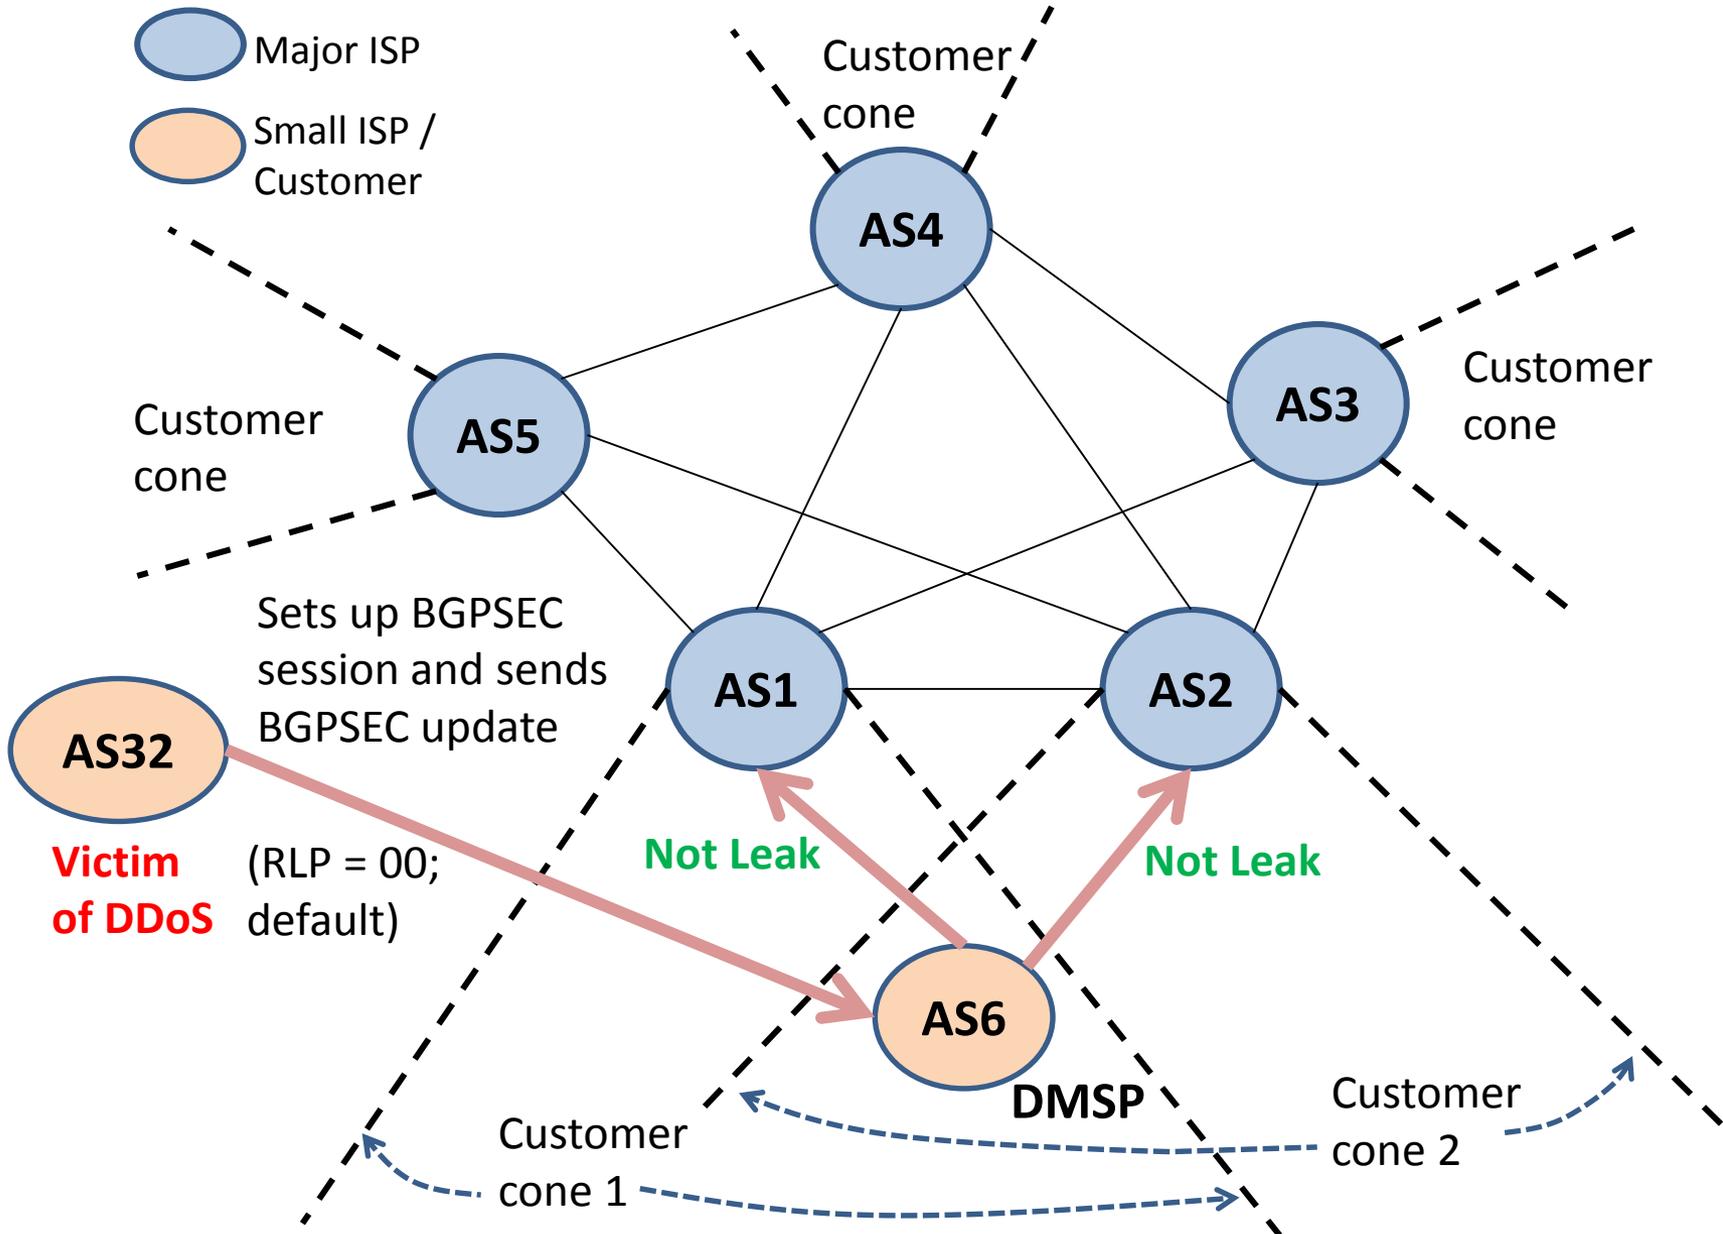
Example 2: Lateral Across Customer Cones and Then Leaked Up to Other ISP



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