Enhancement to BGPSEC for Detection/Mitigation against Route Leaks

K. Sriram and D. Montgomery
NIST

IETF 90, Toronto, Canada
July 25, 2014

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

Illustration of Basic Notion of a Route Leak

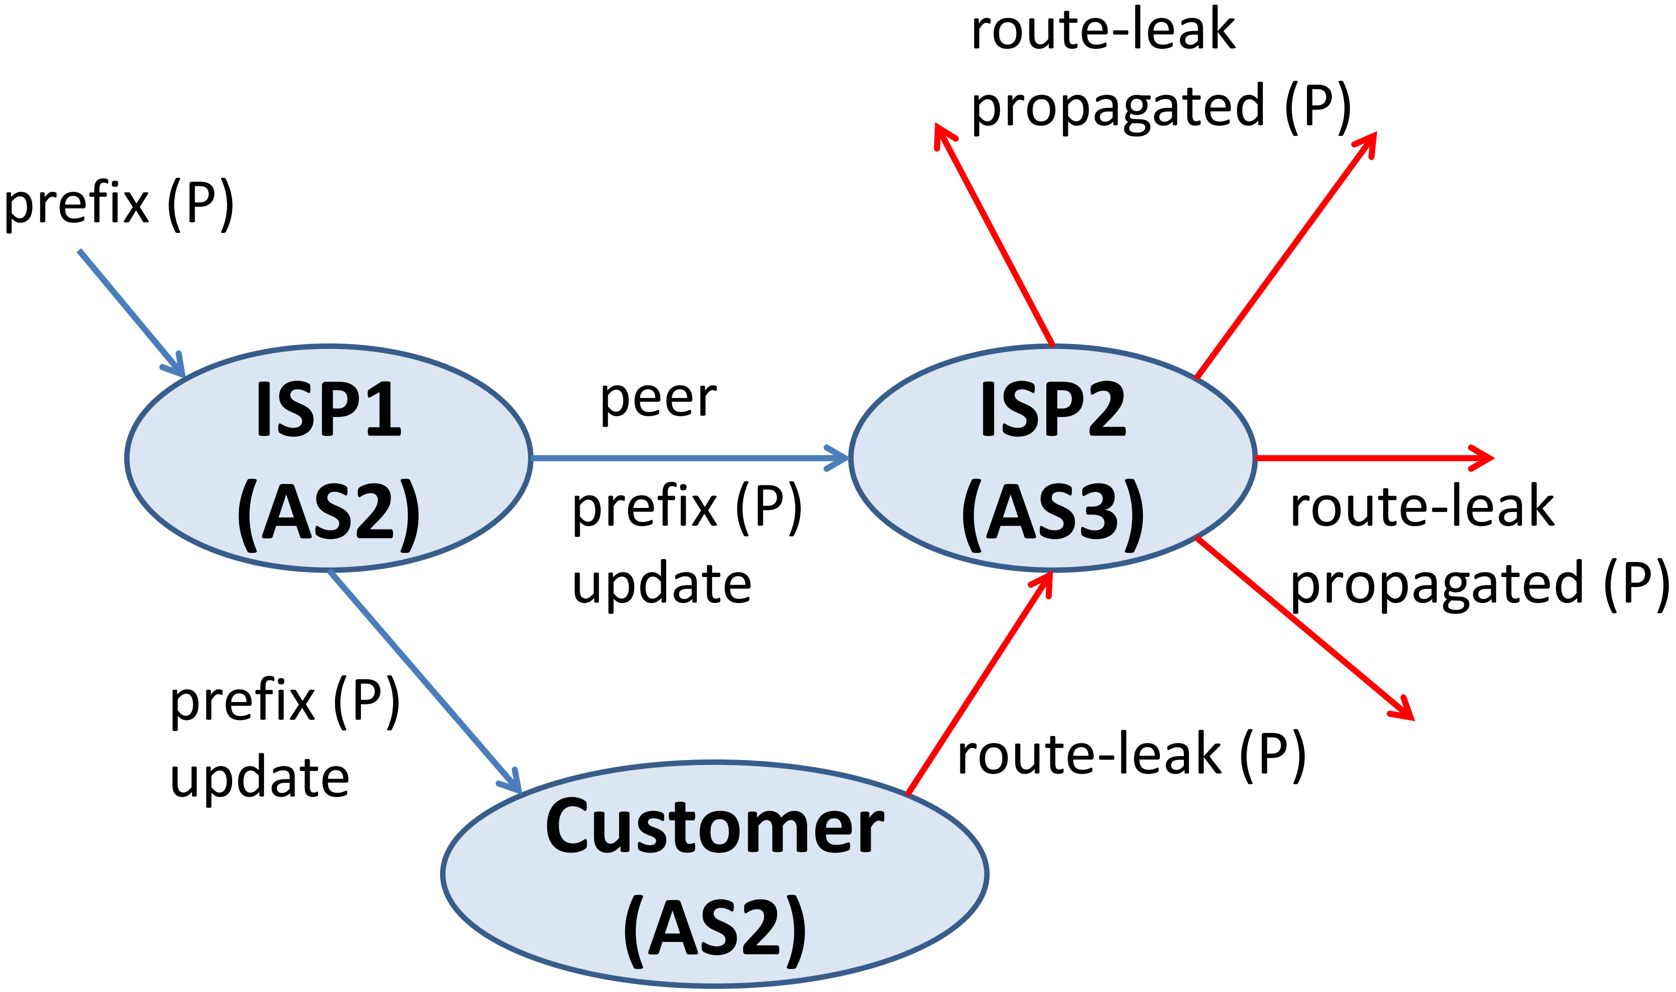

In general, ISPs prefer customer route announcements over those from others.

Anatomy of a Route Leak: Four Types

Type 1: Prefix Hijack with Data Path to Legitimate Origin

A multi-homed AS learns a prefix route from one upstream ISP and re-originates it towards another upstream ISP. This amounts to straightforward hijacking.

- ➢ Somehow (not attributable to path poisoning by the attacker) a reverse path is present, and data packets reach the legitimate destination via the offending AS (e.g. China Telecom (2008), Iceland (2013), Belarus (2013) incidents).

Type 2: U-Turn with More Specific Prefix

A multi-homed AS learns a prefix route from one upstream ISP and announces a sub-prefix (subsumed in the prefix) to another upstream ISP.

- ➢ The update basically makes a U-turn at the attacker's multi-homed AS but a subprefix is propagated. Having the subprefix maximizes the success of the attack.
- ➢ Reverse path is kept open by the path poisoning techniques as in [Kapela-Pilosov].

Anatomy of a Route Leak: Four Types

Type 3: U-Turn with Full Prefix

A multi-homed AS learns a prefix route from one upstream ISP and simply propagates the prefix to another upstream ISP.

- The update basically makes a U-turn at the attacker's multi-homed AS.
- Neither the prefix nor the AS path in the update is altered.
- This is similar to a straight forward path-poisoning attack [Kapela-Pilosov], but with full prefix. Examples: Google-Moratel (2012), Dodo’s AS38285 (2012).

Type 4: Leak of Internal Prefixes

- An offending AS simply leaks its internal prefixes to one or more of its provide ASes.

Route Leak Detection/Mitigation in BGPSEC

- BGPSEC protocol already offers detection and mitigation capability against Types 1, Type 2, and Type 4 ☺
- Can BGPSEC be enhanced to provide protection against Type 3 also?
  - ➢ The answer seems like ‘Yes’

Begin Sender Specification

(Simple Enhancement to Existing 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
- Protected in BGPSEC under path signatures
- 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)

- RLP is proposed to be a 2-bit field set by each AS along the path
- Protected in BGPSEC under path signatures
- 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. Whether the actual encoding method is RLP bits or Transitive Community, etc. – can be decided later.

End of Sender Specification.

Sending Router’s Intent

- Note: There is no disclosure about the nature of a peering relationship.
- (In Choice 1) By setting RLP indication to 01, merely asserting that this prefix-update that I’ve forwarded to my neighbor SHOULD not be propagated to a provider AS by said neighbor or any subsequent AS in the path of update propagation.

Recommendation for Receiver Action for Detection of Route Leaks of Type 3

(When Sender is using Method 1)

Receiving BGPSEC 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) It is Valid in accordance with the BGPSEC protocol.

c) 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” – if customer’s RLP field is set to 01, that would indicate an error condition.

An Example Receiver Action for Mitigation of Route Leaks of Type 3

(When Sender is using Method 1)

- If an update from a customer AS is marked as a Route-Leak, then the receiving router SHOULD prefer a Valid signed update from a peer or an upstream provider over the customer's update.

Discussion & Examples – How it works!

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.

Example 1: Multi-homed Customer Leak

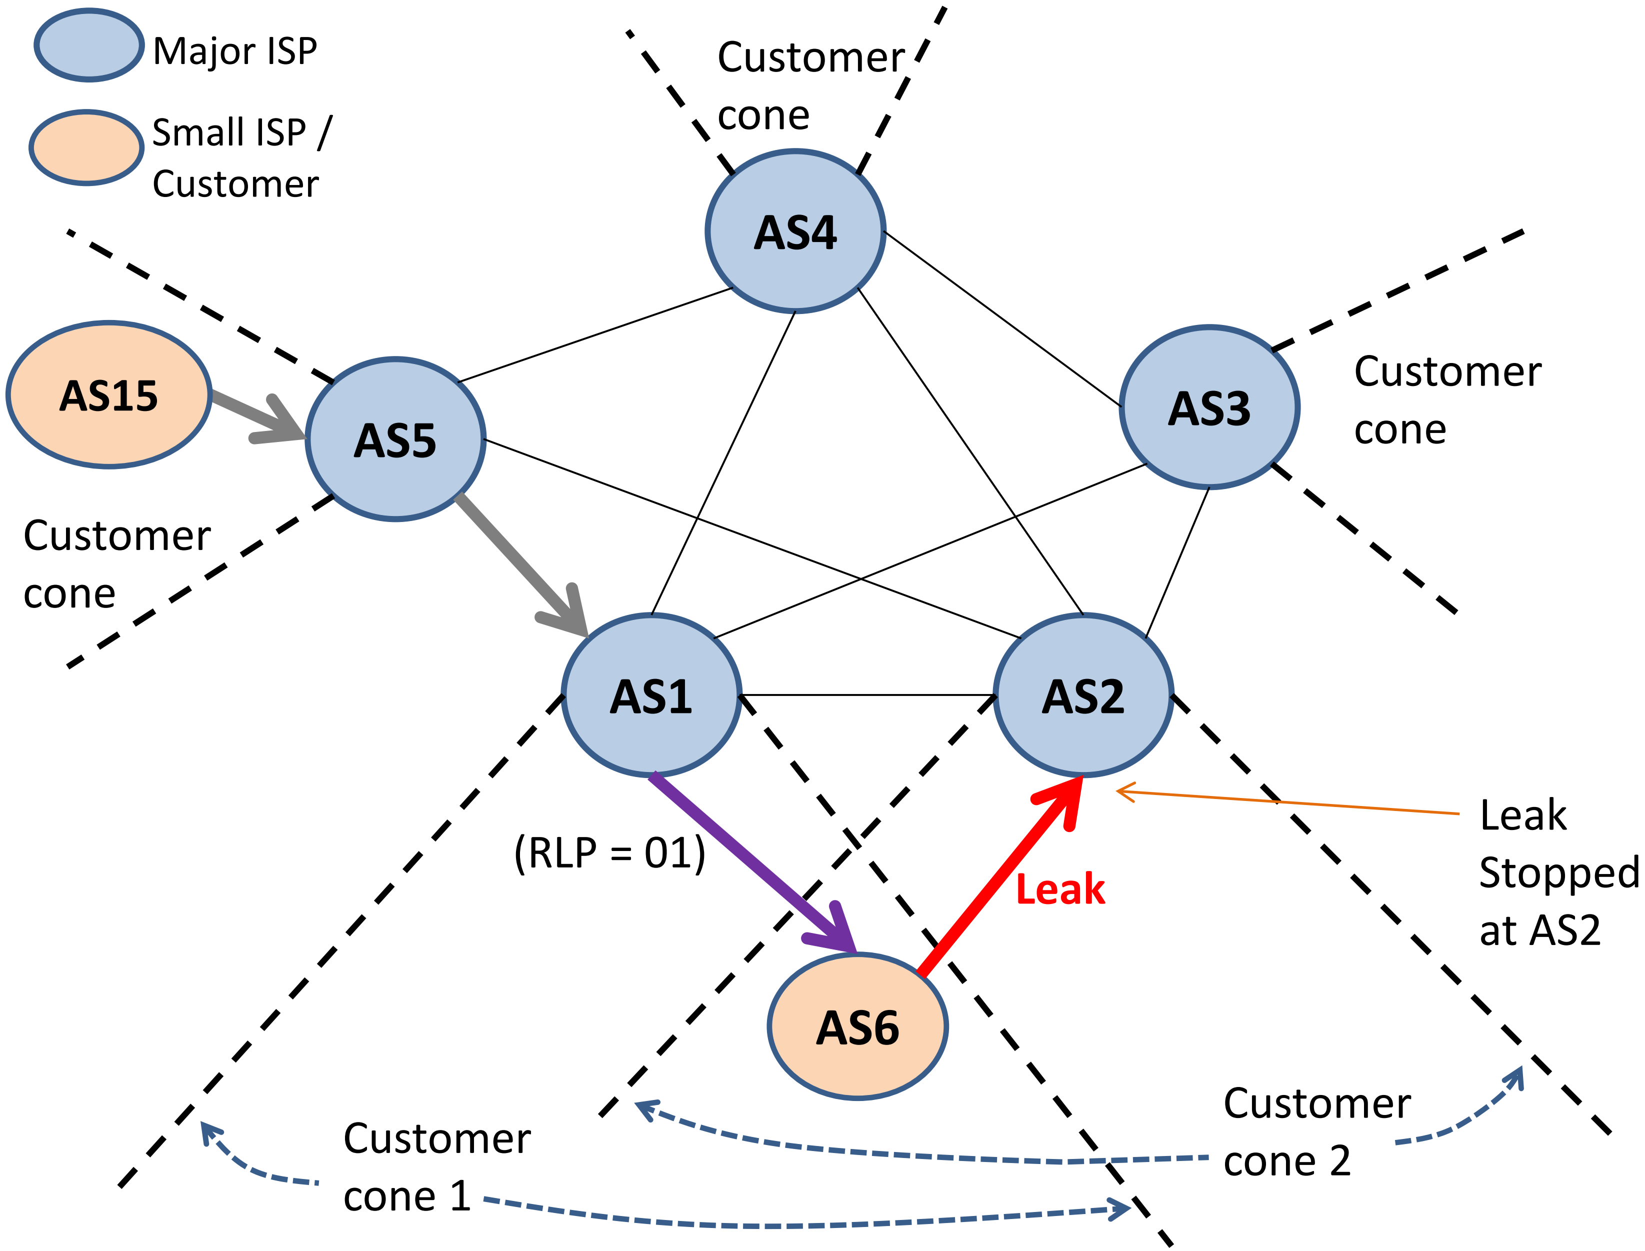

Example 2: Lateral Across Customer Cones and Then Leaked Up to Other ISP

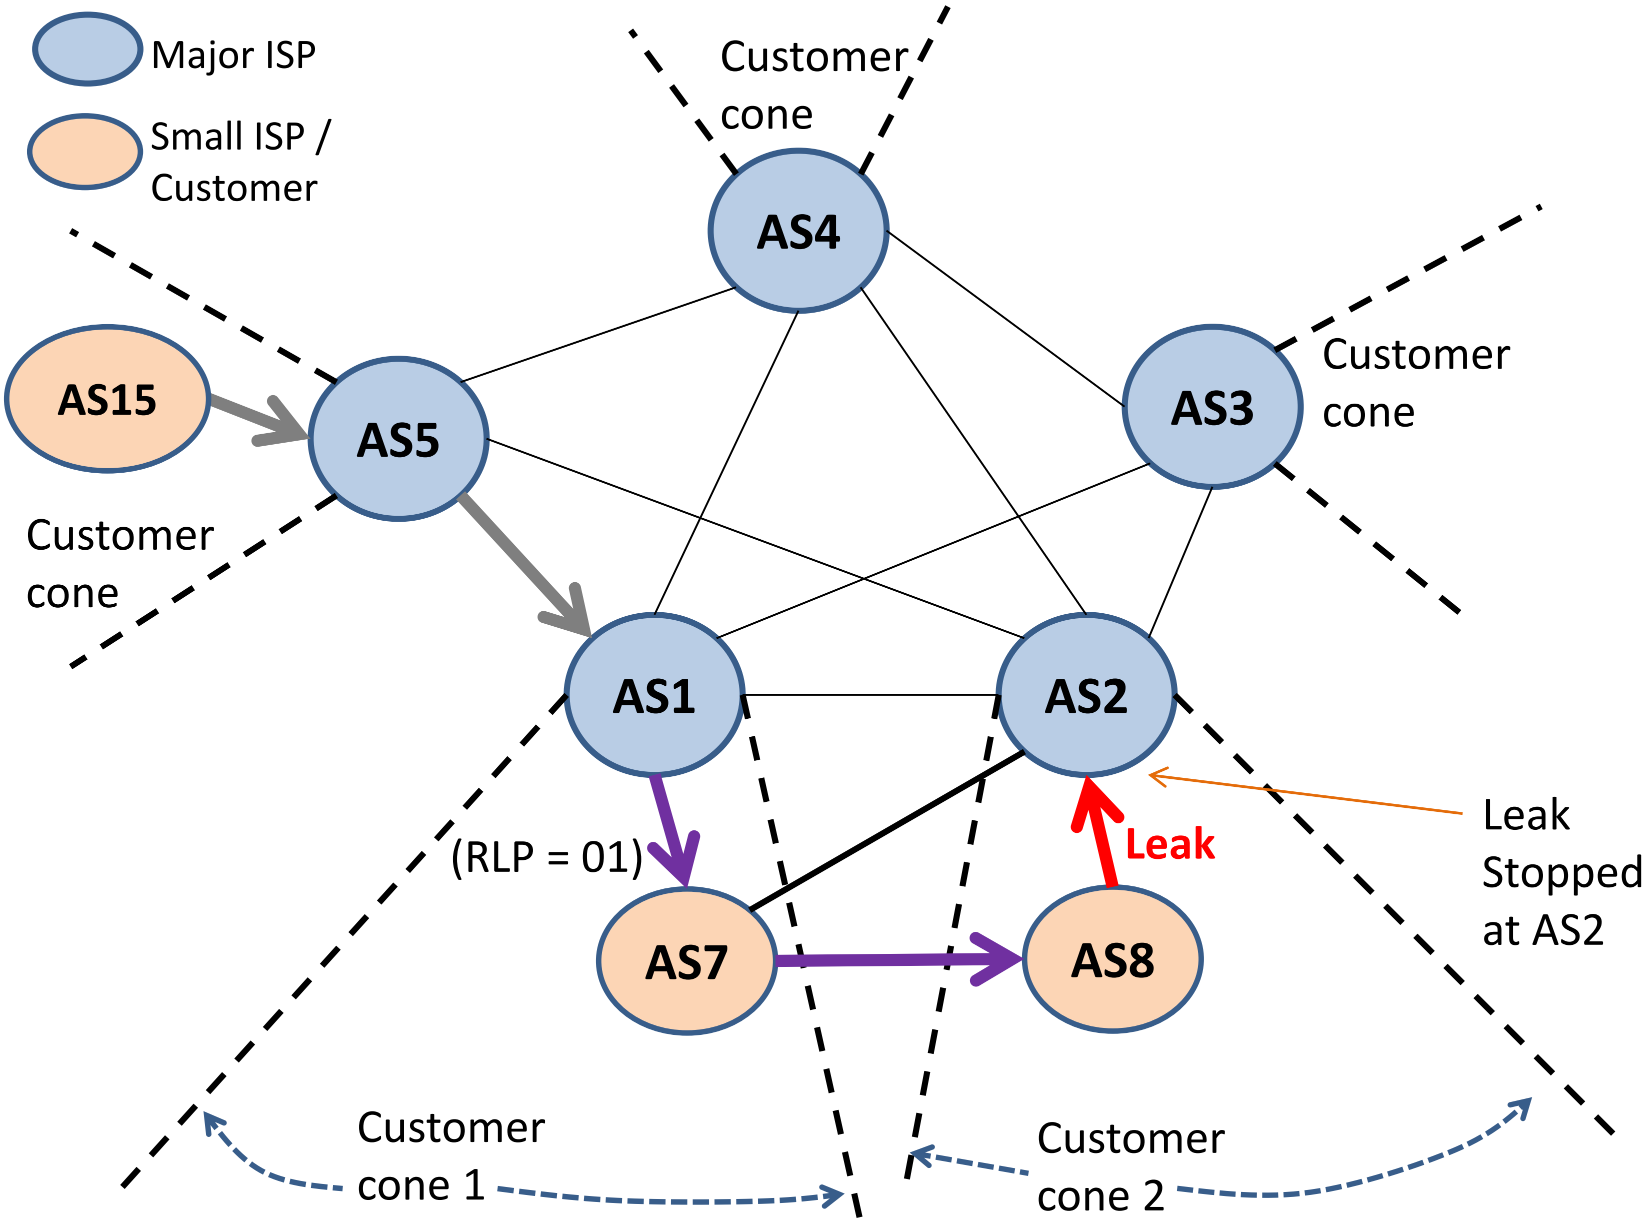

Example 3: Customer's Customer is Multi-homed and Leaks

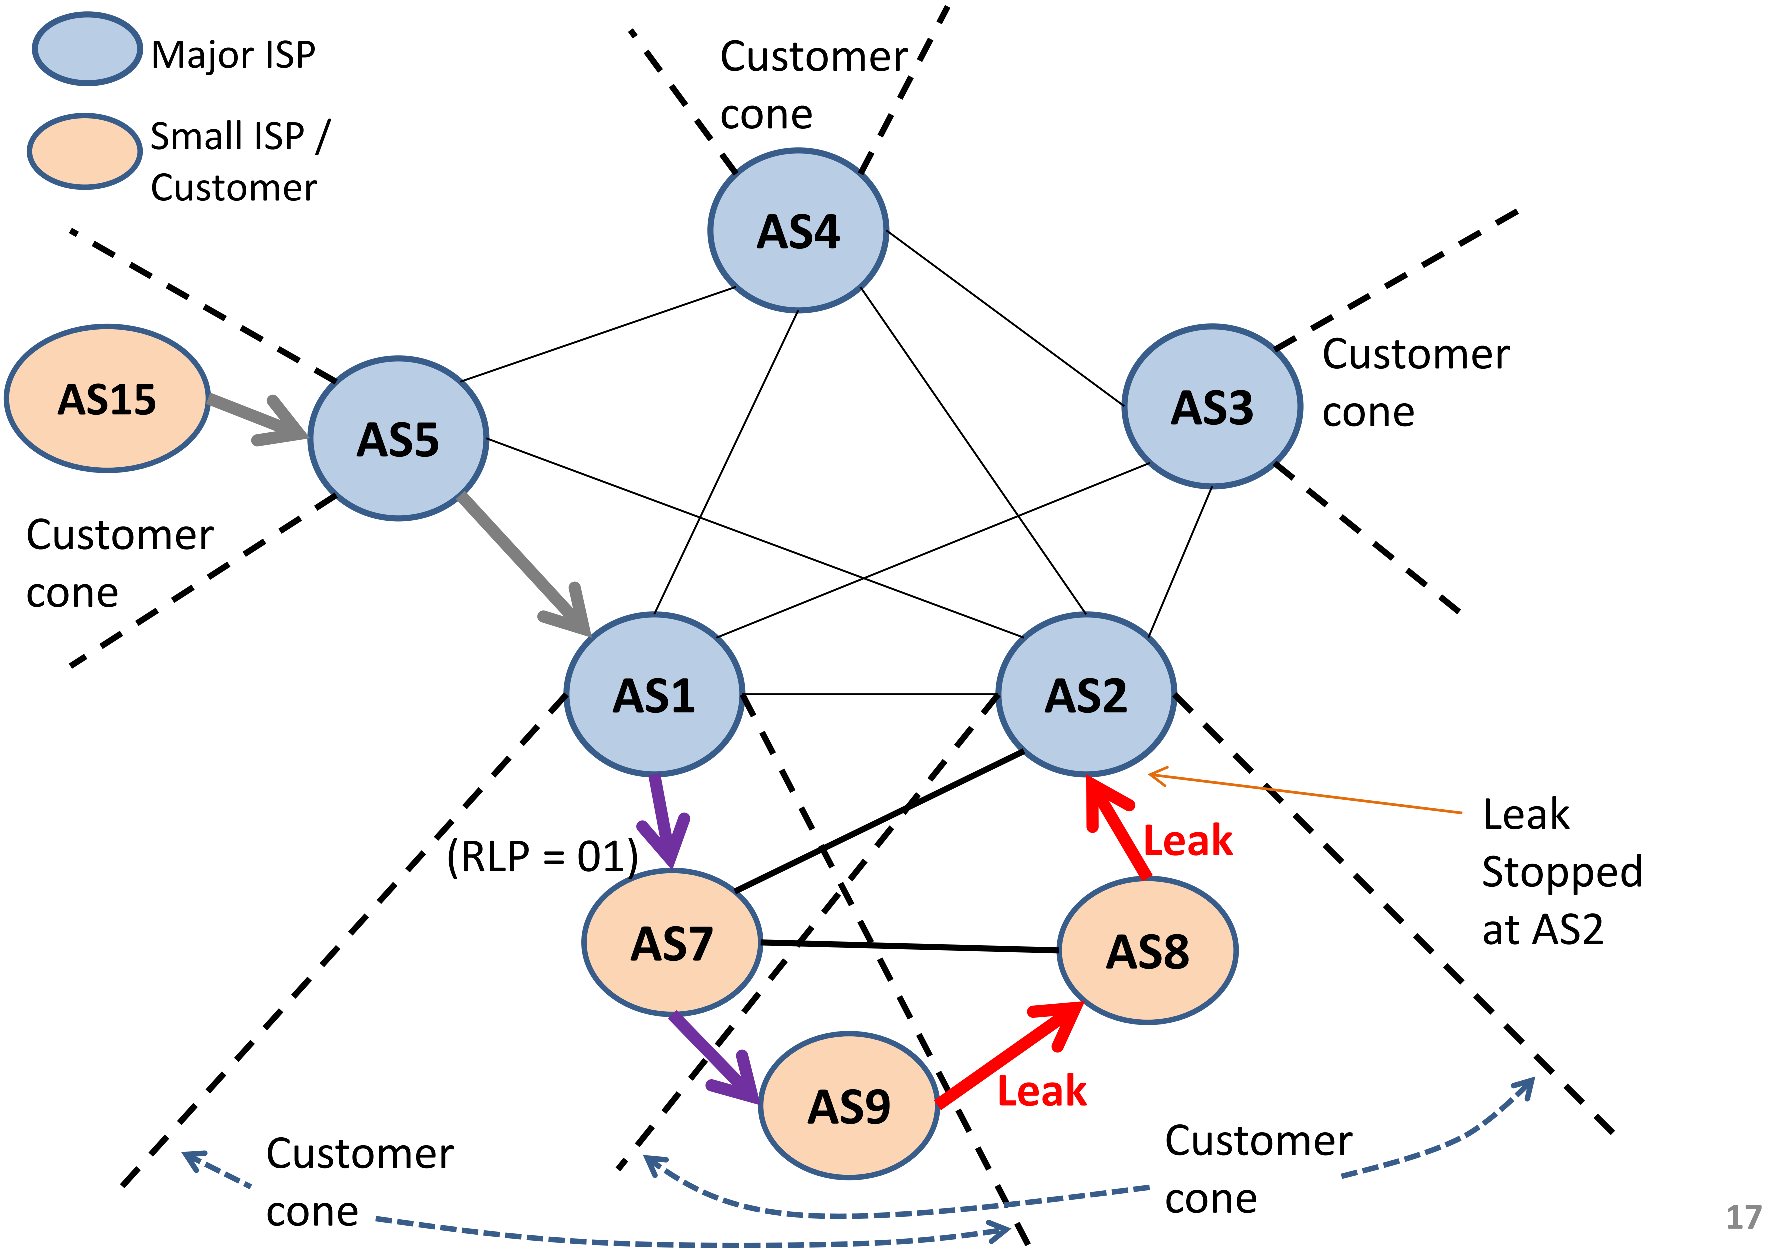

Consideration of DDoS Mitigation Service Provider

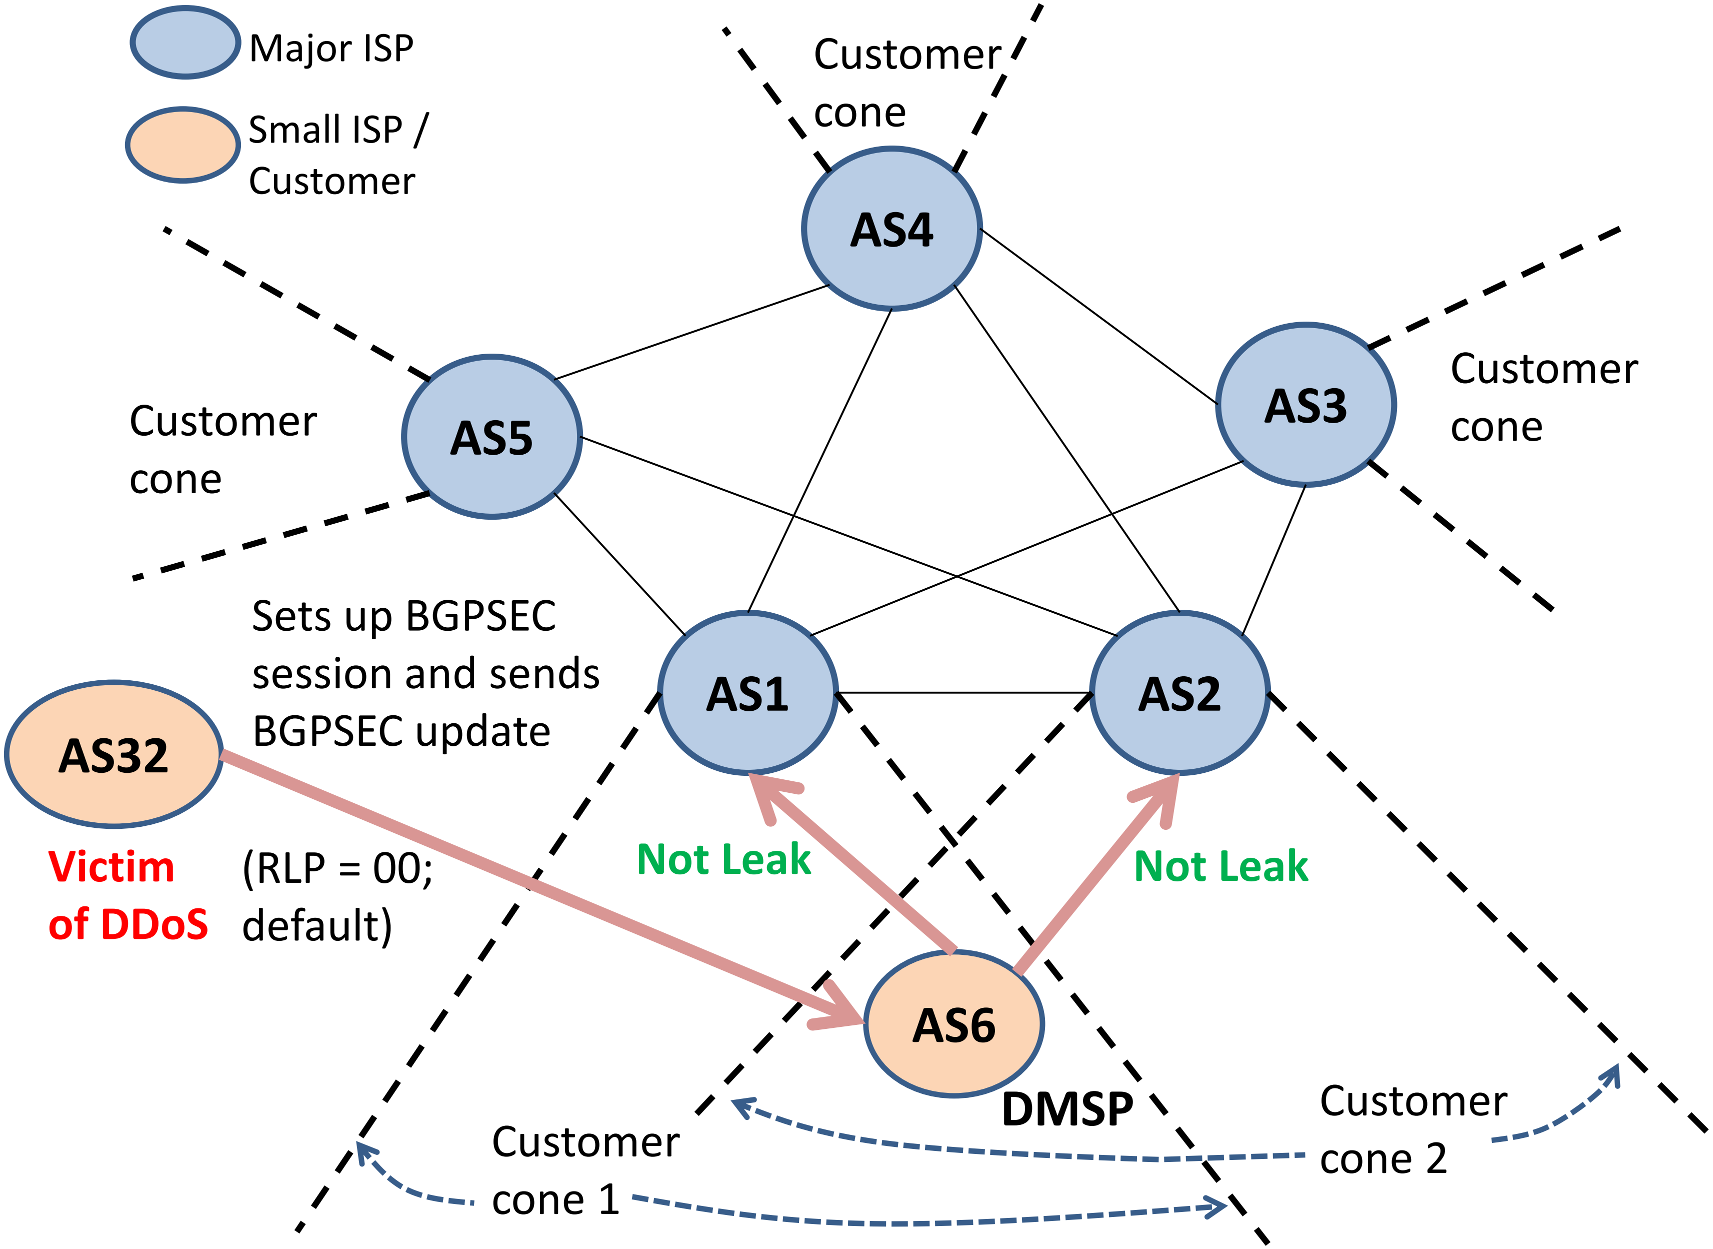

Stopgap Solution when Only Origin Validation is Deployed

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

Summary and Conclusion

- Identified four categories of route leaks
- Three of these are already mitigated in BGPSEC
- Presented an enhancement of BGPSEC that protects against the remaining type of route leaks
- When only Origin Validation is deployed, the construction of a customer prefix filter list from ROAs can help mitigate route leaks
- Offered some suggestions for special consideration for DDoS Mitigation Service Providers (DMSP)