BGP Route-Leak Protection Community

Jakob Heitz, Cisco

Nanog 71, October, 2017

Gao - Rexford

- BGP ASes have 2 types of relationships: transit-customer or peer-peer.
- A neighbor of an AS can be either transit provider, peer or customer.
- If an AS receives a route from a non-customer and sends it to a non-customer, then it is leaking that route.
- An AS can choose to transit some routes, but not others, so the relationship can be different for different routes.
- Note that an IXP route server has no relationship with its clients.
- A route server does not add its ASN to the AS-path.
- Route server clients have relationships with other route server clients.

Concept

Community to indicate the transit providers that the advertising AS uses for the indicated route and (optionally) the more specifics.

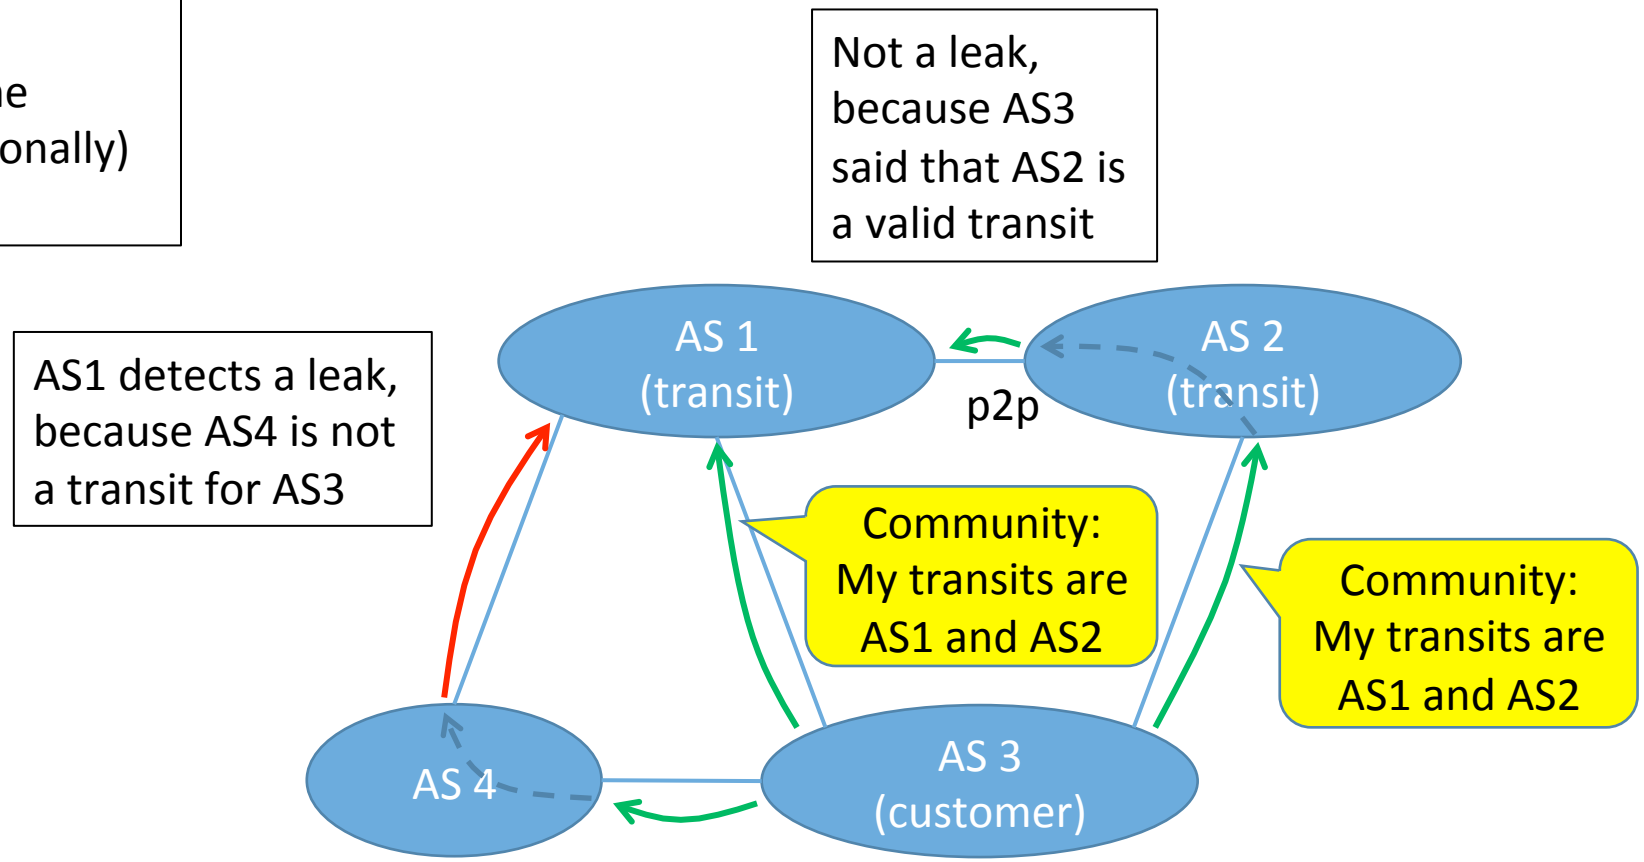

RLP Community and Large Community

- Need a new well known Large Community. Suppose it's 4100100100.
- If an AS wants route leak protection for its routes, then it will attach RLP communities to the routes to indicate what its transits are for those routes.
- The format is WKN:subject:transit, where
  - WKN is the well known number that indicates this is an RLP Large Community.
  - subject is the subject ASN.
  - transit is a transit provider of the subject ASN. It is 0 if the subject has no transits.
- For example, if AS 64500 has transit provider ASes 64501 and 64502, then it will attach the following large communities:
  - 4100100100:64500:64501
  - 4100100100:64500:64502
- AS 64501 has no transit providers, so it will attach the large community:
  - 4100100100:64501:0
- To extend the protection an AS will pass the RLP on to the next AS. An AS can choose to accept an RLP community or not.
- For those that do not have Large Communities (RFC8092) implemented, a regular community can be used in a limited way. Neighbor ASes must agree on a community value to use, say 65000:x. This way, the sender says "x is my transit".
- For the previous example, AS 64500 would attach the communities:
  - 65000:64501
  - 65000:64502

Leak of more specifics

- An AS may send more specific routes to a peer AS for traffic engineering. The more specific is not sent anywhere else.
- To detect the leak of a more-specific prefix, need an RLP community that specifies a transit for the more specific prefixes derived from the indicated route .

Finding Leaks

Consider the route 2001:DB8:1::/48

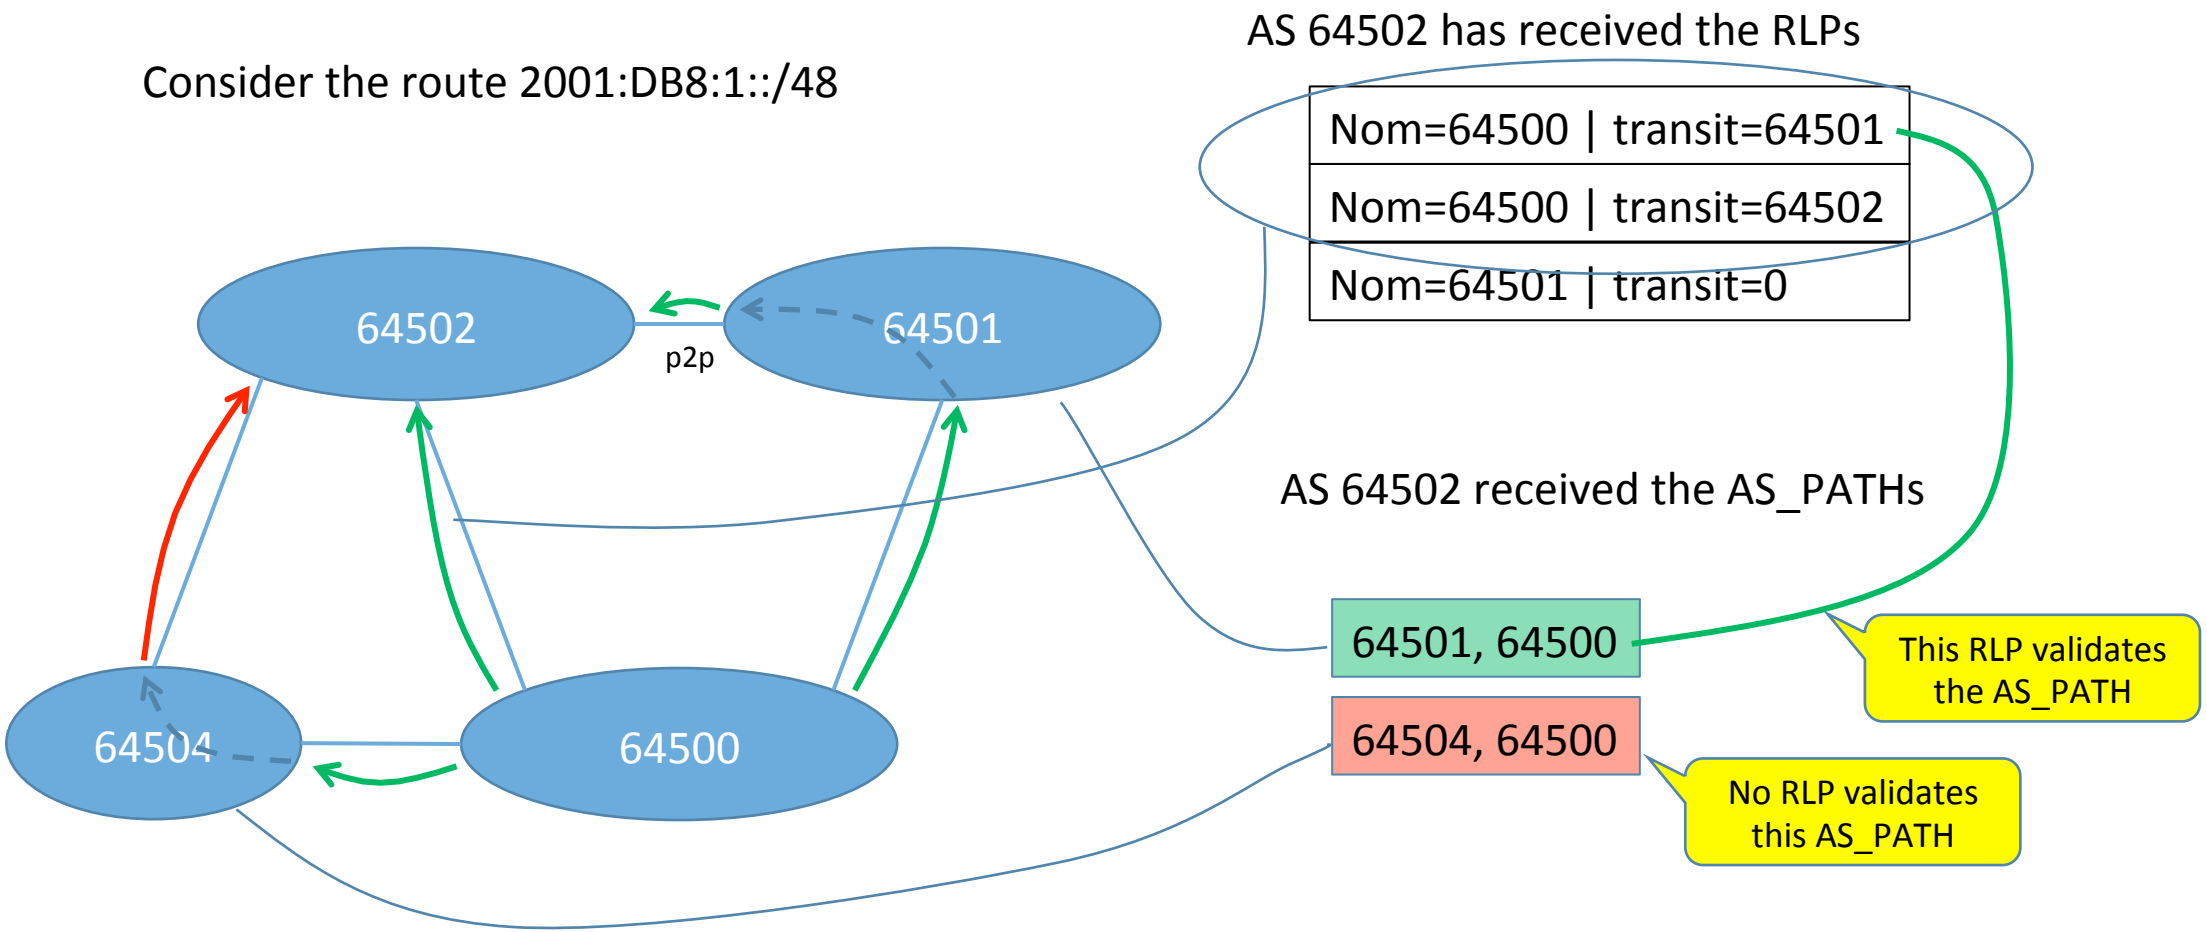

Benefits

- The leaking AS need not participate.
- The customer AS needs no new software: just to set the RLP community.
- AS relationships may be different for different routes.
- Works across IXP route servers.
- An ISP can sell Route Leak Protection to customers without having to rely on other ASes to cooperate. If the ISP can cooperate with its peers, then protection improves.

Examples

- The blue ovals are BGP autonomous systems (AS).
- Each arrow is a route announcement of the same prefix. One prefix is announced at the first arrow. Follow the arrows for the propagation of the announcement. The same prefix is announced to multiple ASes. A green arrow is a correct announcement. A red arrow is a leak.
- The yellow callouts indicate the relevant attributes in the announcement: the as-path and the RLP large communities.

Leak of transit to transit

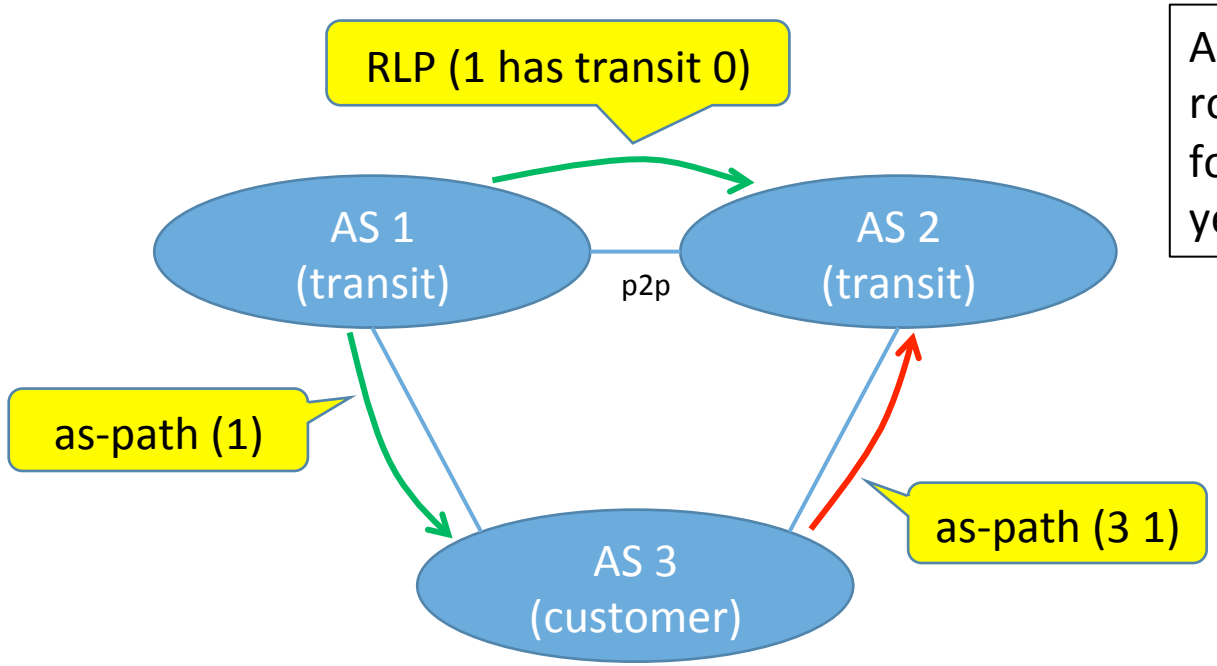

AS 2 detects a leak, because route is received with 3 following 1 in the as-path, yet RLP says 1 has no transits

Route from transit is not a leak

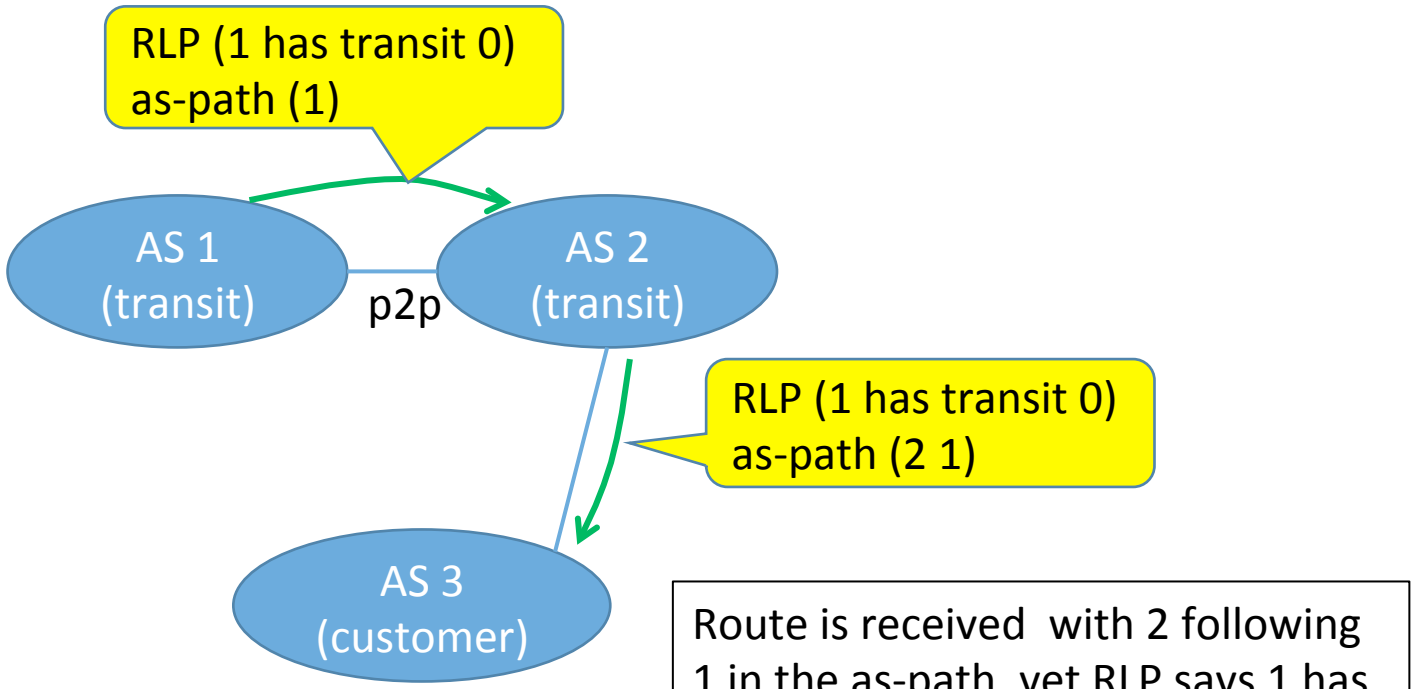

Route is received with 2 following 1 in the as-path, yet RLP says 1 has no transits. However, this is not a leak, because route is received from a transit AS.

Leak of transit to peer

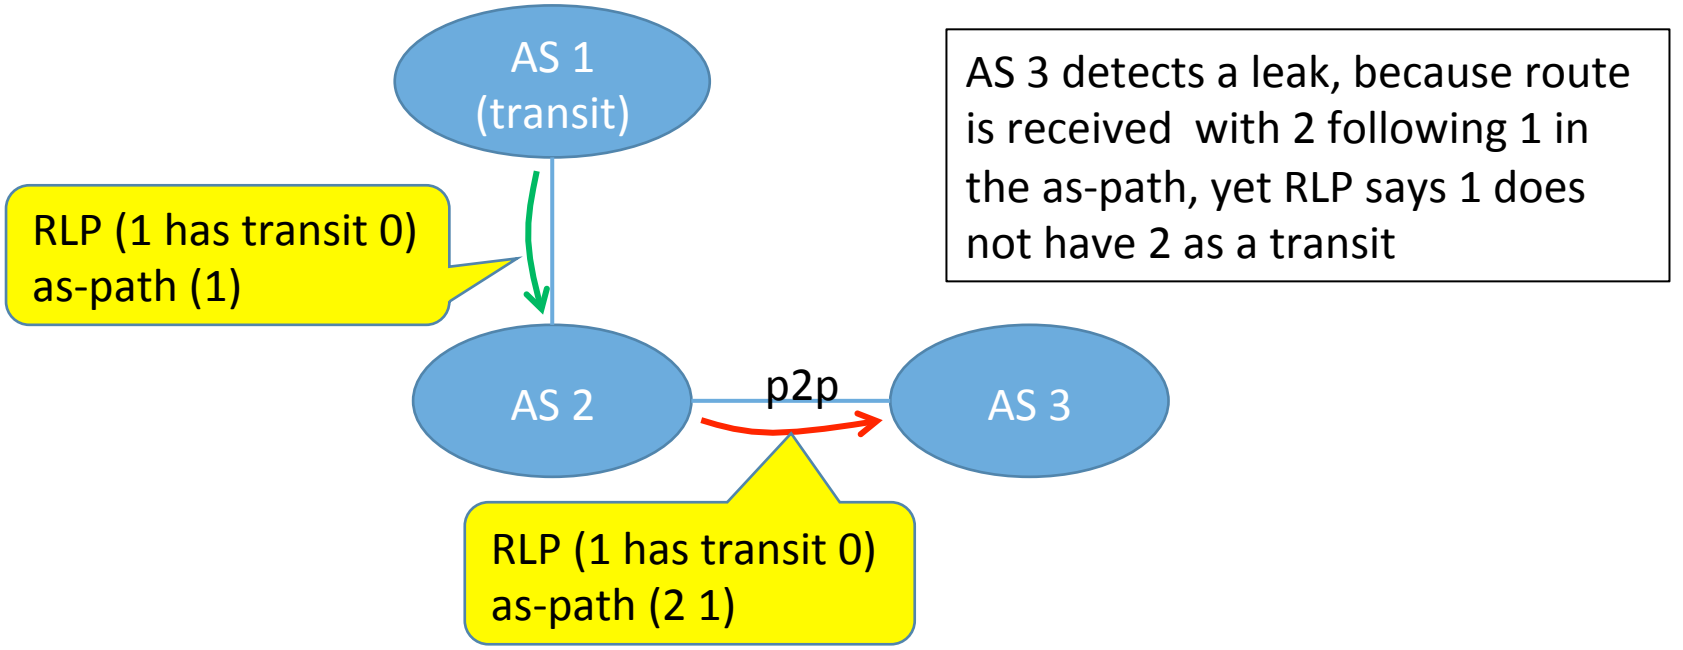

Leak of peer to transit

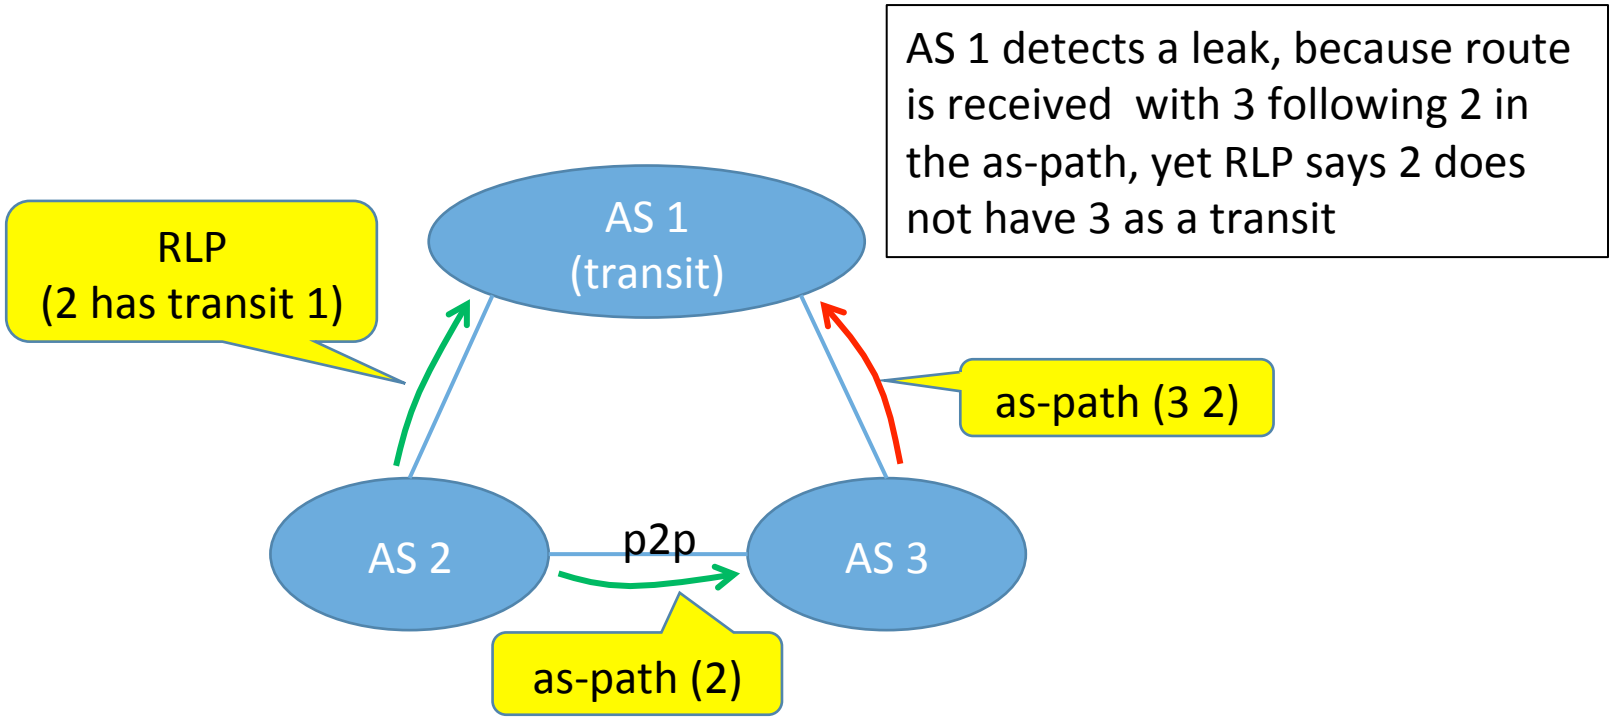

Leak of more-specific

AS 1 detects a leak, because 10::/13 is received with 3 following 2 in the as-path, yet CRLP says 2 has no transits for more-specifics of 10::/12

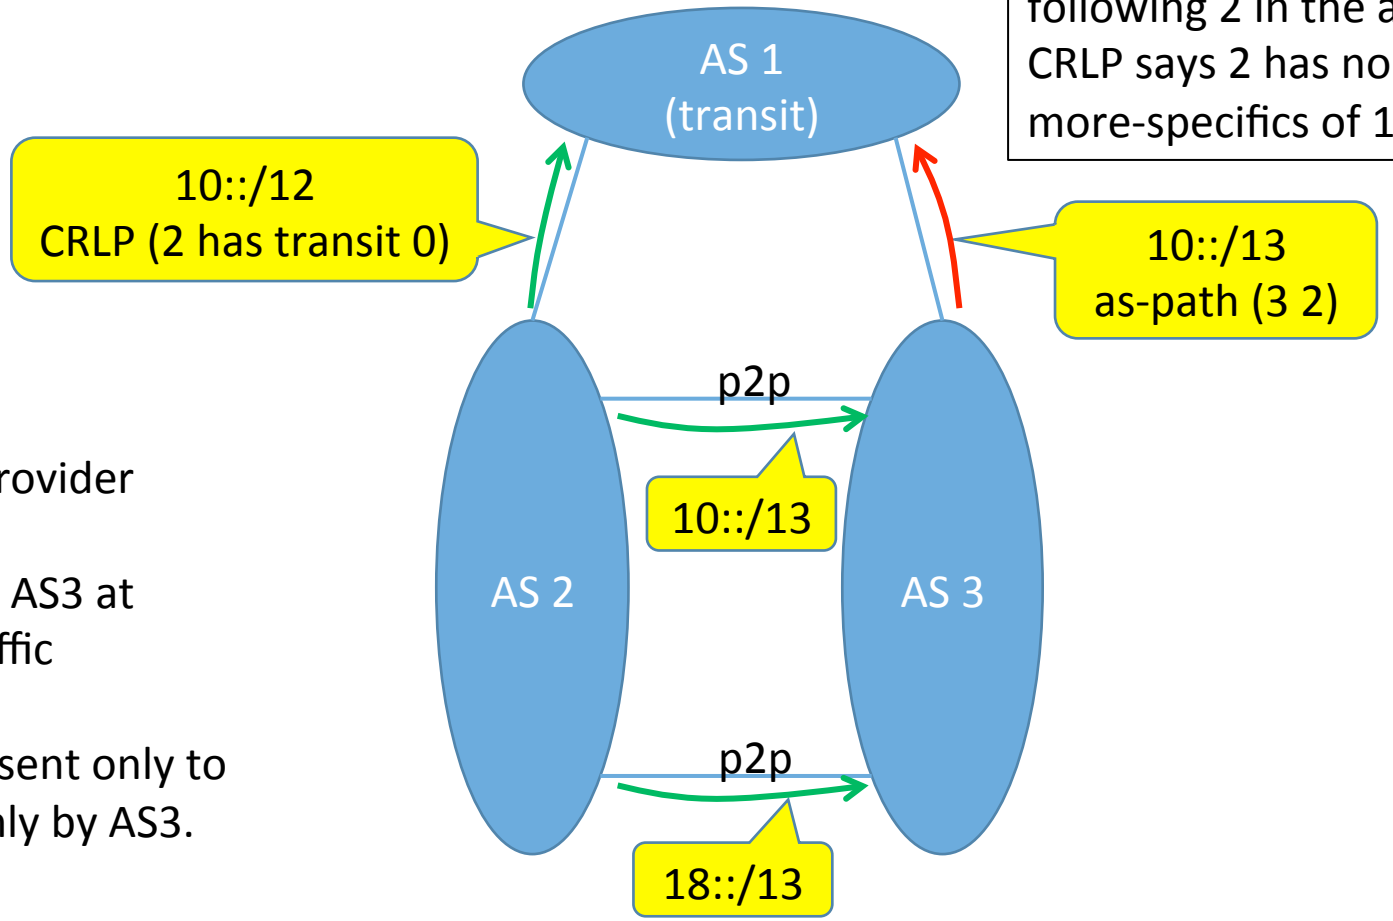

AS2 peers with content provider AS3.

It sends specific routes to AS3 at different locations for traffic engineering.

These specific routes are sent only to AS3 and meant for use only by AS3.

Leak of peer to peer

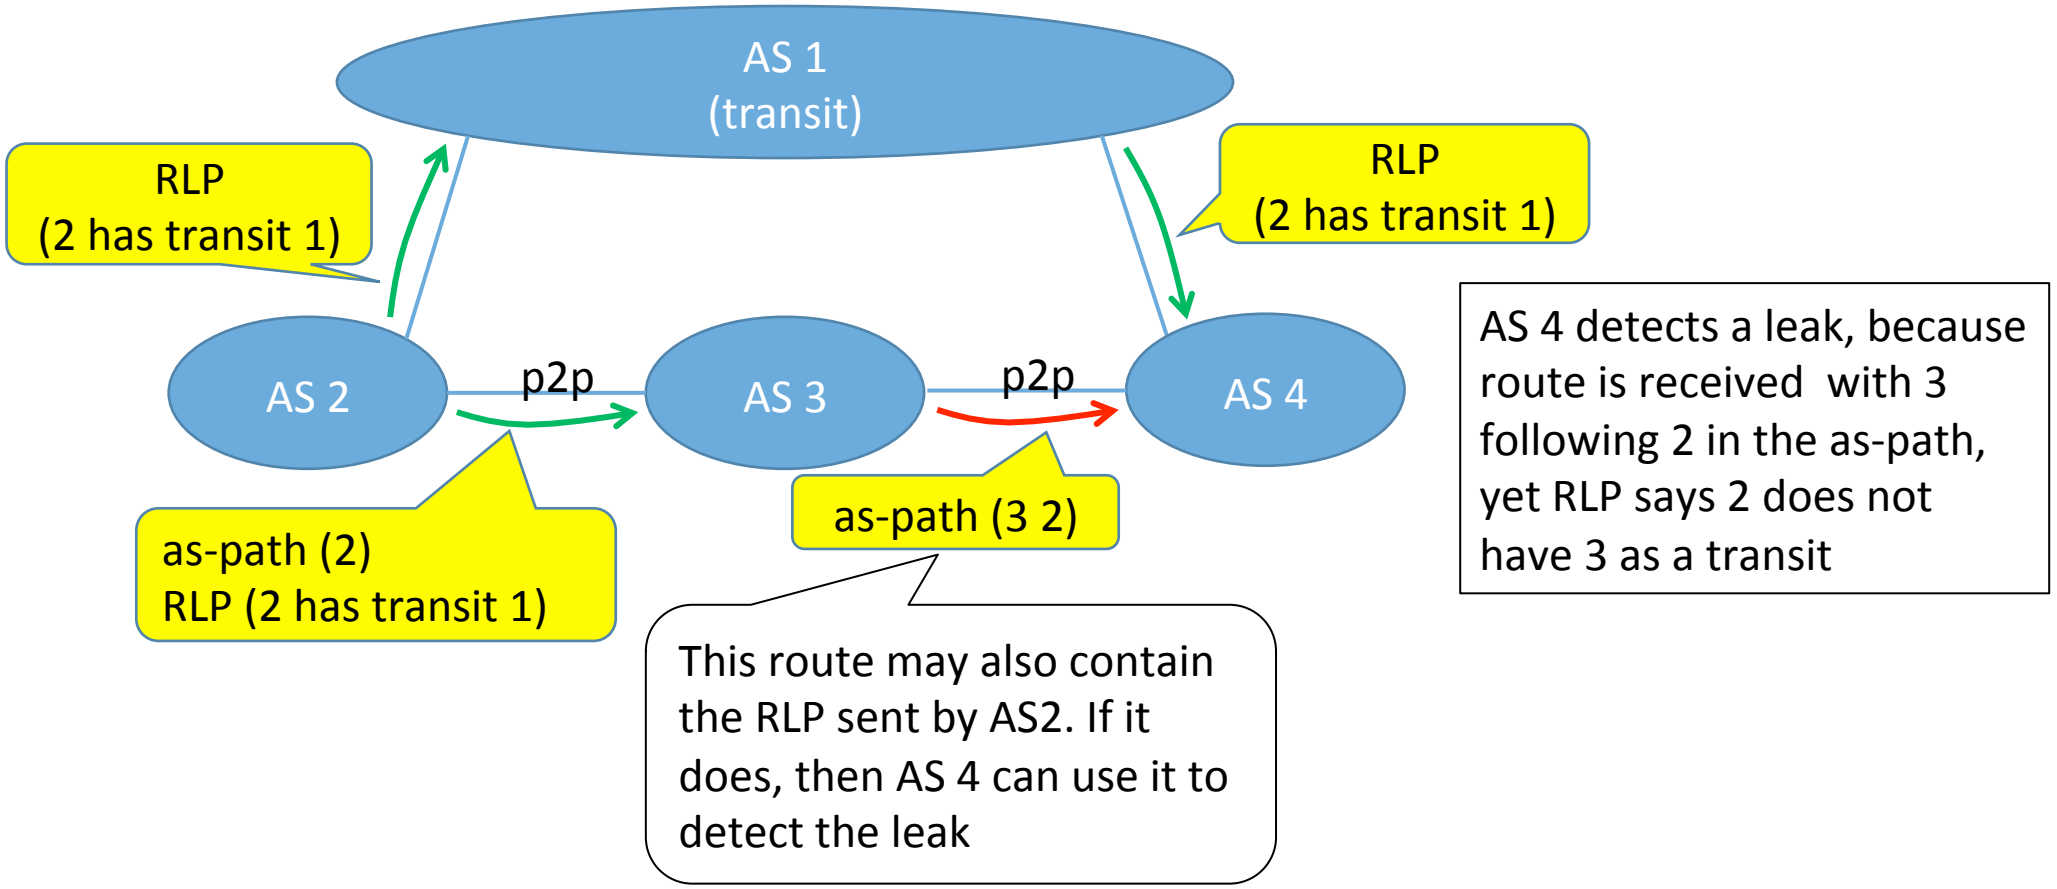

Transits share RLP

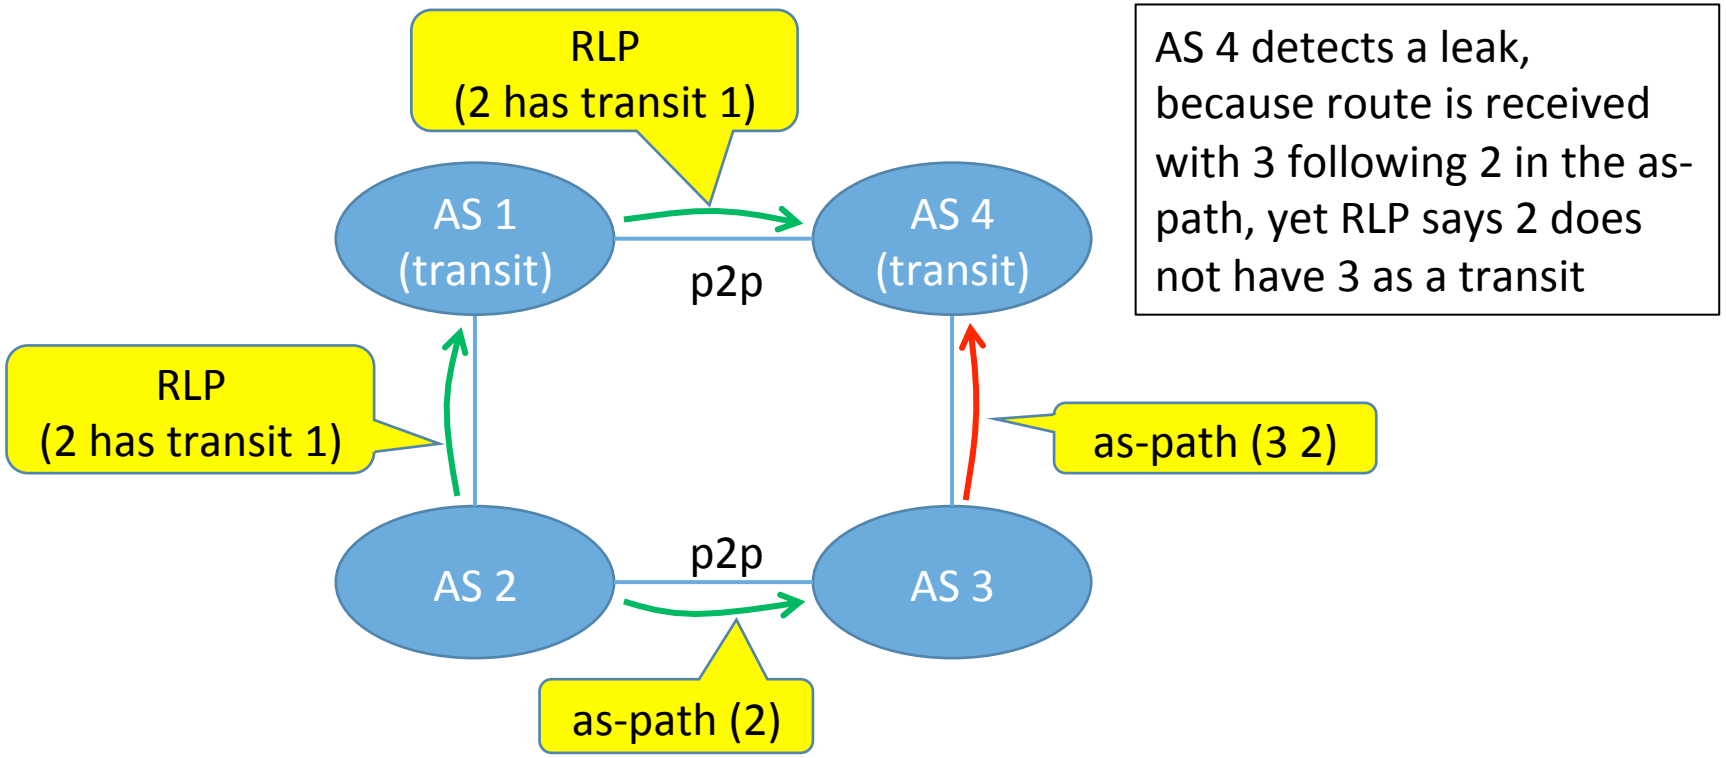

Route Server has no Relationship

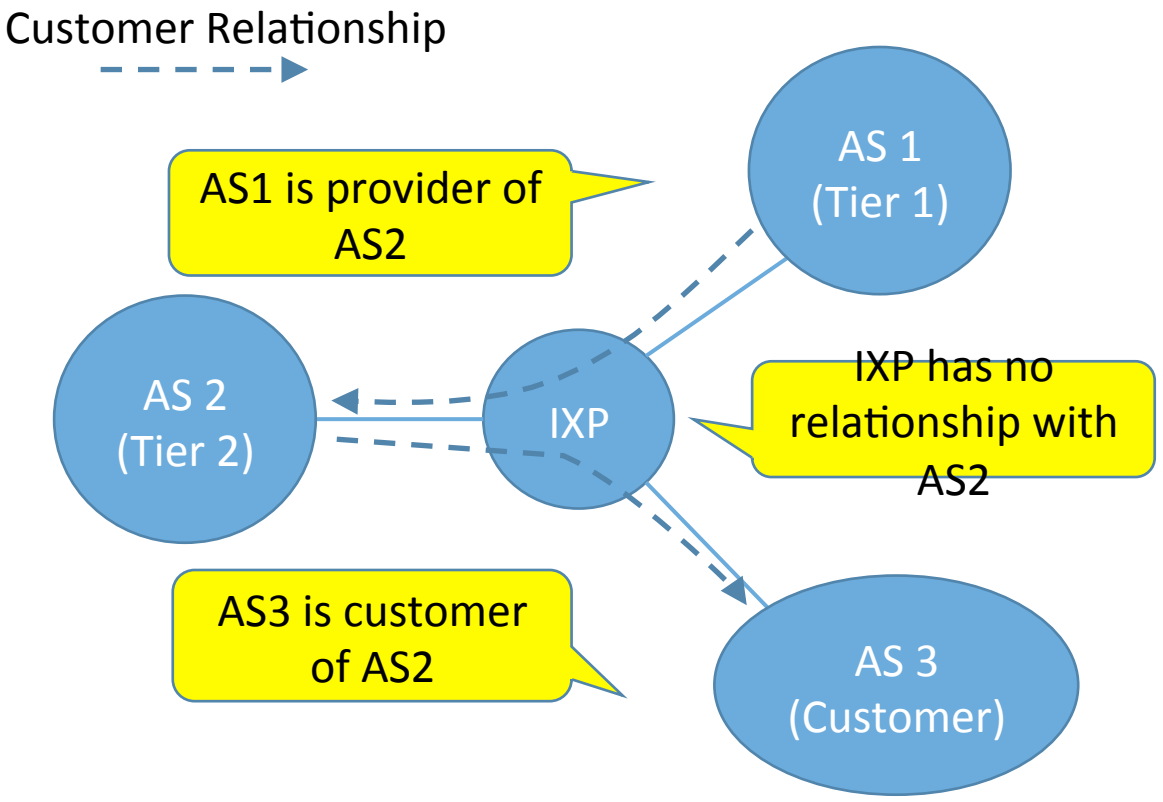

- A route server does not add its ASN to the as-path.
- A route server does not need a transit/customer/ peer relationship with its clients
- Route server clients have relationships with each other.

Instructions for an AS with no transit providers (Tier-1 ISP)

- Assume the large community for RLP is 4100100100:*:*.
- When the policy tests `as-path-validity`, the router will test the as-path against all RLPs available to determine the leak state.
- This example is for ASN 64501.
- Apply this policy at neighbor in:

```
route-policy nbr_in
  # if not accepting forwarded RPL
  delete large-community in (4100100100:not-peeras:*)
  # low local pref for leaked routes
  if as-path-validity is leaked then
    set local-preference 1
  endif
end-policy
```

- Apply this policy at neighbor out:

```
route-policy nbr_out
  # if not forwarding RPLs
  delete large-community in (4100100100:*:*)
  # set my own RLP
  set large-community (4100100100:64501:0) additive
end-policy
```

Instructions for an AS with transit providers

- This example is for ASN 64500. Assume it has transit providers AS64501 and AS64502.
- Apply this policy at neighbor in:

```
route-policy nbr_in
  # if not accepting forwarded RPL
  delete large-community in (4100100100:not-peeras:*)
  # low local pref for leaked routes
  # do not check as-path-validity at transit neighbors
  if as-path-validity is leaked then
    set local-preference 1
  endif
end-policy
```

- Apply this policy at neighbor out:

```
route-policy nbr_out
  # if not forwarding RPLs
  delete large-community in (4100100100:*:*)
  # set my own RLP
  set large-community (4100100100:64500:64501, 4100100100:64500:64502) additive
end-policy
```

Instructions for an AS wanting protection, but does not support large communities or RLP

- Assume the community for RLP is 64519:*
- The peer router will automatically convert the received RLP community into an RLP large community.
- This example is for ASN 64500. Assume it has transit providers AS64501 and AS64502.
- Apply this policy at neighbor out:

```
route-policy nbr_out
  # set my own RLP
  set community (64519:64501, 64519:64502) additive
end-policy
```

Reference

- This automates the concept of Peer Locking described in https://www.nanog.org/sites/default/files/Snijders_Everyday_Practical_Bgp.pdf
- It is submitted as a draft to IETF IDR working group
- https://tools.ietf.org/html/draft-heitz-idr-route-leak-community-00