



Efficient Technique for Enforcing Internet Peering Policies

PLNOG#4 - Warsaw March 5th, 2010



2010-02-10


Klaudiusz Staniek

Network Consulting Engineer, Cisco

kstanie@cisco.com

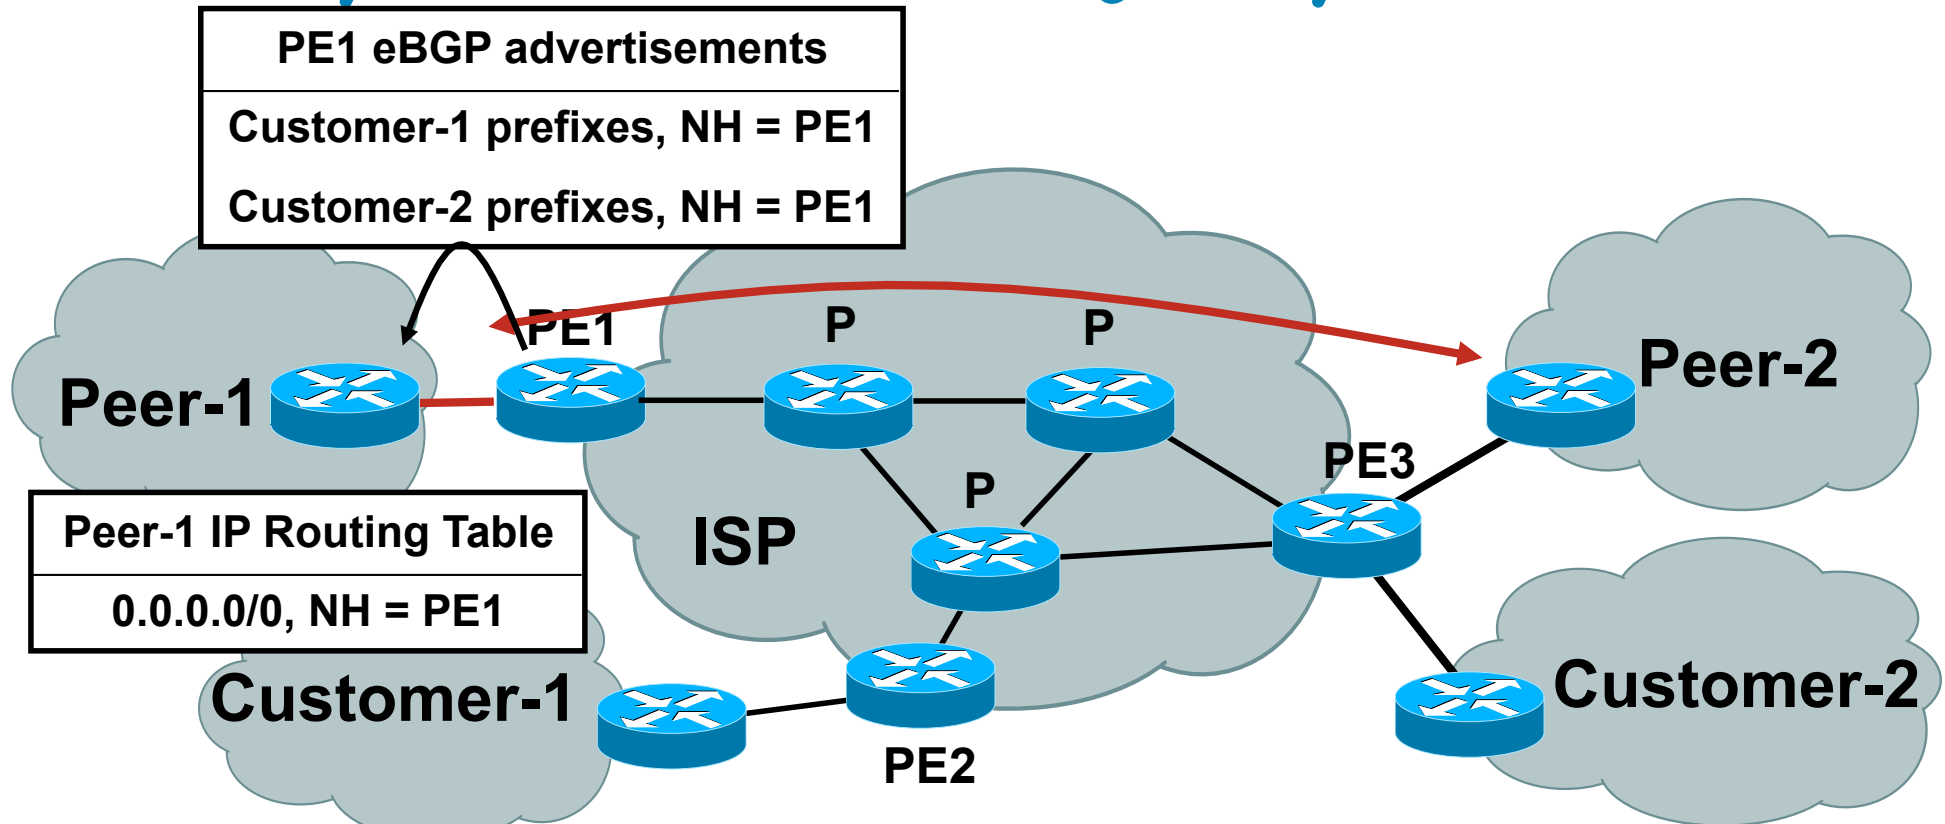


Internet Peering Policy Overview



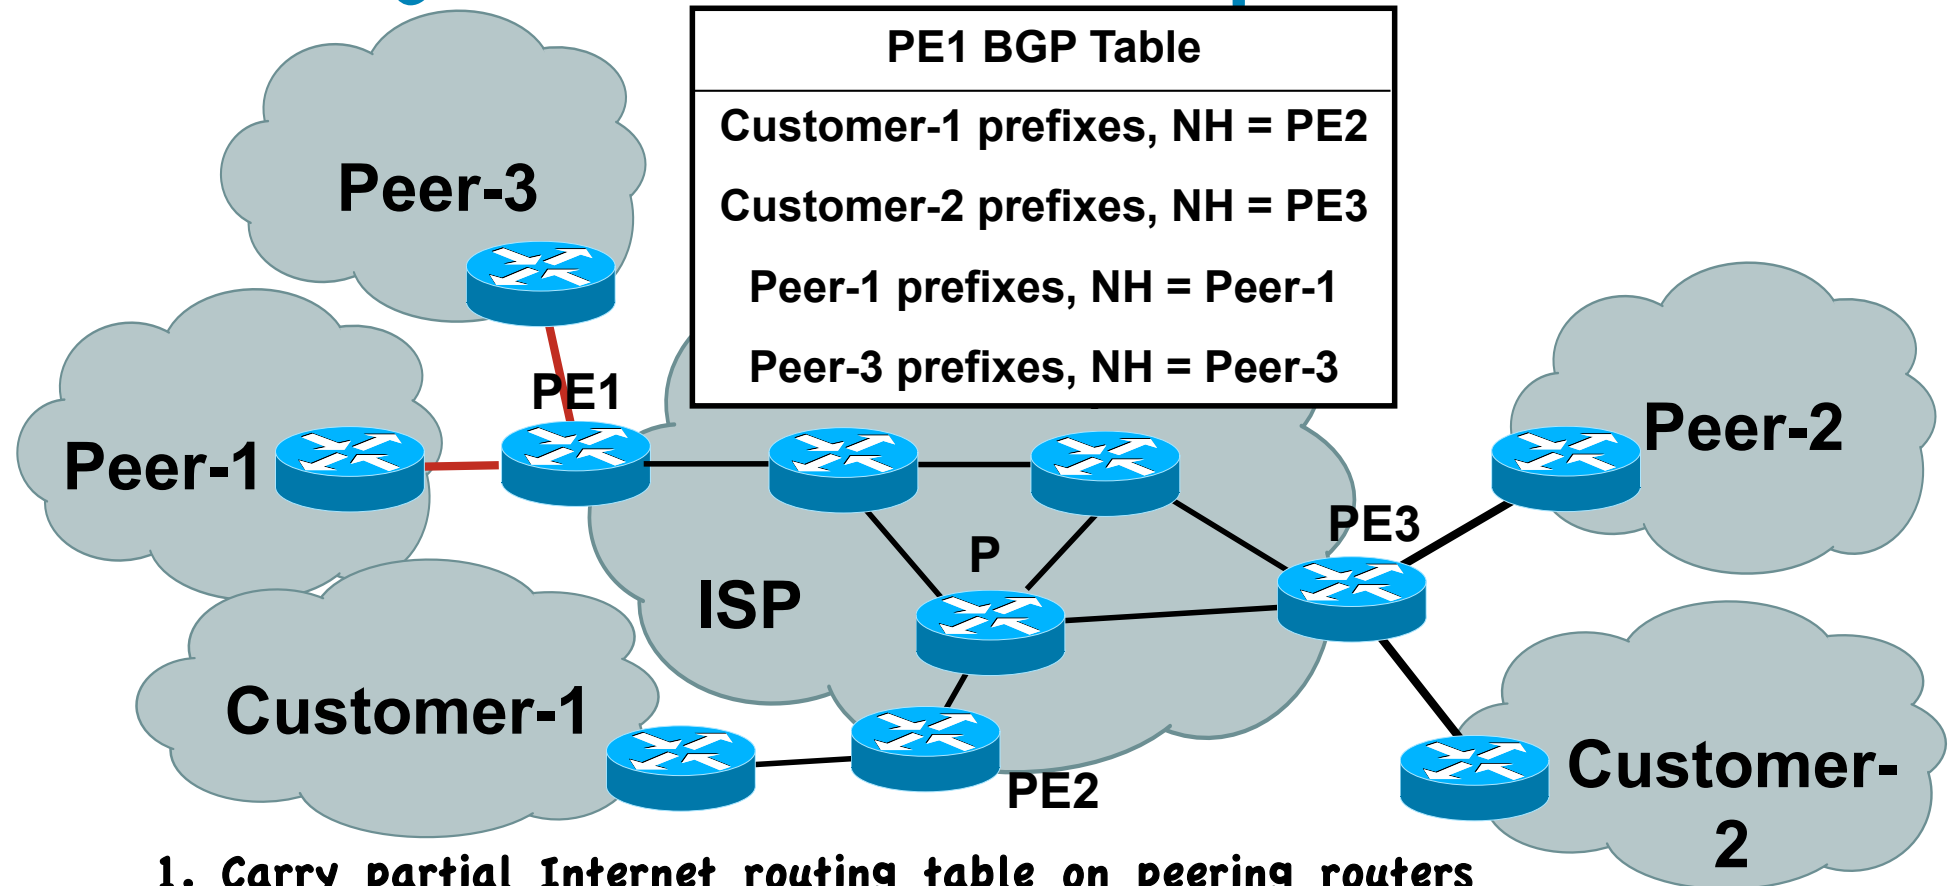
- **Peers should only have IP reachability to & from ISP's customer prefixes**
 - For example, traffic between Peer-1 and Customer-1 is permitted within the ISP and Peer-1 peering policy
- **Peers should not use the ISP as transit to one another**
 - For example, traffic between Peer-1 and Peer-2 is in violation of the ISP and Peer-1 peering policy (as well as the ISP and Peer-2 peering policy)

Policy Enforcement Using Only BGP



- BGP control plane techniques only filter prefix advertisements
- If a peer uses IP routing tricks (e.g., default routing), it may bypass BGP policies and steal bandwidth from the ISP peer
 - for example, using the peer as transit to another peer
- This is possible because BGP policies are only enforced within the IP control plane and *not* within the IP data forwarding plane

Challenges with Alternate Options



1. Carry partial Internet routing table on peering routers

- For example, filter Peer-2 prefixes from being carried on PE1
- Does not prevent IP reachability between peers connected to the same local peering router (e.g., Peer-1 and Peer-3)

2. Interface ACLs - not scalable or operationally efficient

- Adds, moves or changes to ISP customer and downstream provider address ranges force updates to static ACL policies

Proposed Technique

1. ISP tags peer prefixes uniquely within its BGP and FIB tables

- Peer prefixes set with community attribute (X) and tag (X') in BGP and FIB tables, respectively
- Customer prefixes set with community attribute (Y) and tag (Y') in BGP and FIB tables, respectively

2. ISP tags external packets that ingress peering interconnects based upon longest prefix match within FIB

- Tag (X') for packets received from peer and destined to a prefix in the FIB with tag (X)
- Tag (Y') for packets received from peer and destined to a prefix in the FIB with tag (Y)

3. ISP forwards or discards packets that ingress peering interconnects based upon associated packet tag value

- Packets with tag (X') are discarded since destined to peer prefix
- Packets with tag (Y') are forwarded since destined to customer prefix

Not A Futurist Talk

- **Proposed technique available today**
 - 12000 E3/E5 using IOS 12.0S
 - XR 12000 using IOS XR 3.6
 - CRS-1 using IOS XR 3.6
 - Other IOS routers also
- **Router Configuration**
 1. FIB prefix tagging via BGP (i.e., IOS table-map CLI)
 2. Packet tagging via QPPB (i.e., IOS bgp-policy CLI)
 3. Packet classification via MQC (i.e., IOS service-policy CLI)
- **QPPB glues the IP control plane policy (i.e., BGP) with the IP data plane policy (i.e., MQC)**
 - Prefix-based QoS provided by QPPB (QoS Policy Propagation for BGP) includes packet filtering

IOS Config Illustration of Proposed Technique


```
class-map peer-prefix
  match qos-group 66
!
policy peer-in
  class peer-prefix
    police 8000 conform-action drop exceed-action drop
!
interface pos3/1
  description peering interconnect to Peer-1
  bgp-policy destination ip-qos-map
  service-policy input peer-in
!
router bgp 1
!
  table-map set-prefix-type
!
ip bgp-community new-format
!
ip community-list 1 permit 1:66
!
route-map peer permit 10
  set community 1:66
!
route-map set-prefix-type permit 10
  match community 1
  set ip qos-group 66
```

(2) Enable destination-based QPPB which glues BGP control plane with data plane QoS policy

(3) Traffic received from Peer-1 and destined to any peer prefix is discarded

(1) Set prefix-type within FIB based on BGP community attribute (e.g. 66 for peer prefixes)

Example - Demo



Policy

- i.** Any Peer to Peer communication (transit) blocked at the edge
- ii.** Any Customer to Any Customer communication allowed
- iii.** Any Customer to Any Peer communication allowed

Benefits of Proposed Technique

- **Enforcement of Internet peering policies within the IP data forwarding plane protects against an Internet peer using routing tricks to bypass BGP control plane policies**
 - Traffic received from a peer and destined to a peer (local or remote) is dropped
 - Traffic received from a peer and destined to a customer prefix is forwarded normally
- **Proposed technique glues the IP control plane policy (e.g., BGP) with the IP data plane policy (e.g., MQC)**
 - No ACLs required;
 - Prefix tagging within the FIB (e.g., peer versus customer prefixes) possible through standard BGP policies
 - BGP topology and policy changes automatically reflected within the IP data forwarding plane
- **Complements other BGP control plane applications commonly used today including RTBH, sinkholes, etc.**



CISCO