Efficient Technique for Enforcing Internet Peering Policies

David J. Smith – dasmith@cisco.com

Internet Peering Policy Overview

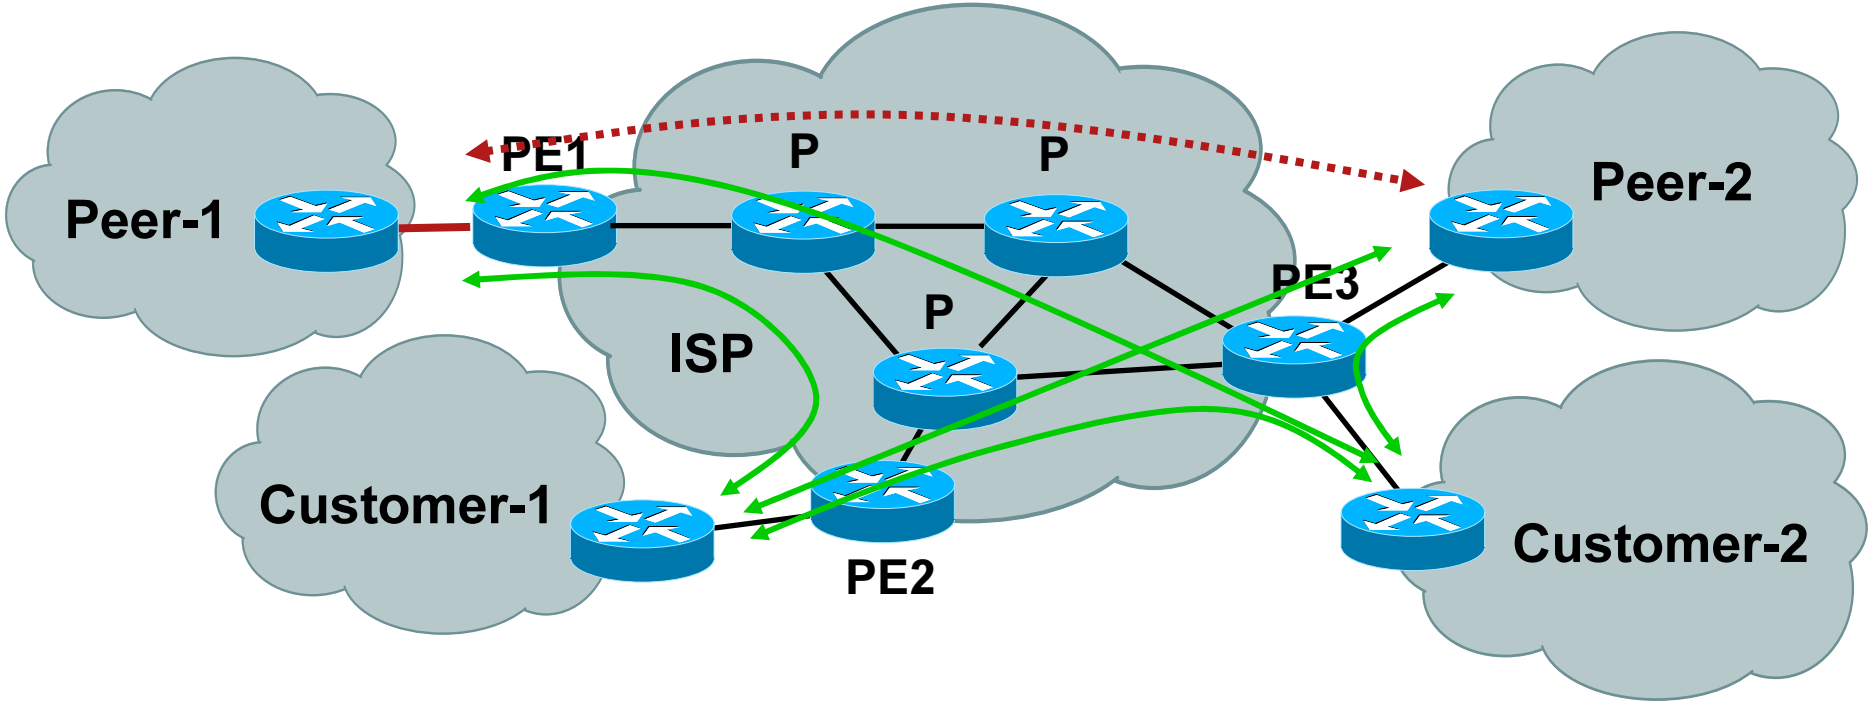

- **Peers should only have IP reachability to & from ISP's customer prefixes**
  - For example, traffic between Peer-1 and Customer-1 is permitted within the ISP and Peer-1 peering policy
- **Peers should not use the ISP as transit to one another**
  - For example, traffic between Peer-1 and Peer-2 is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

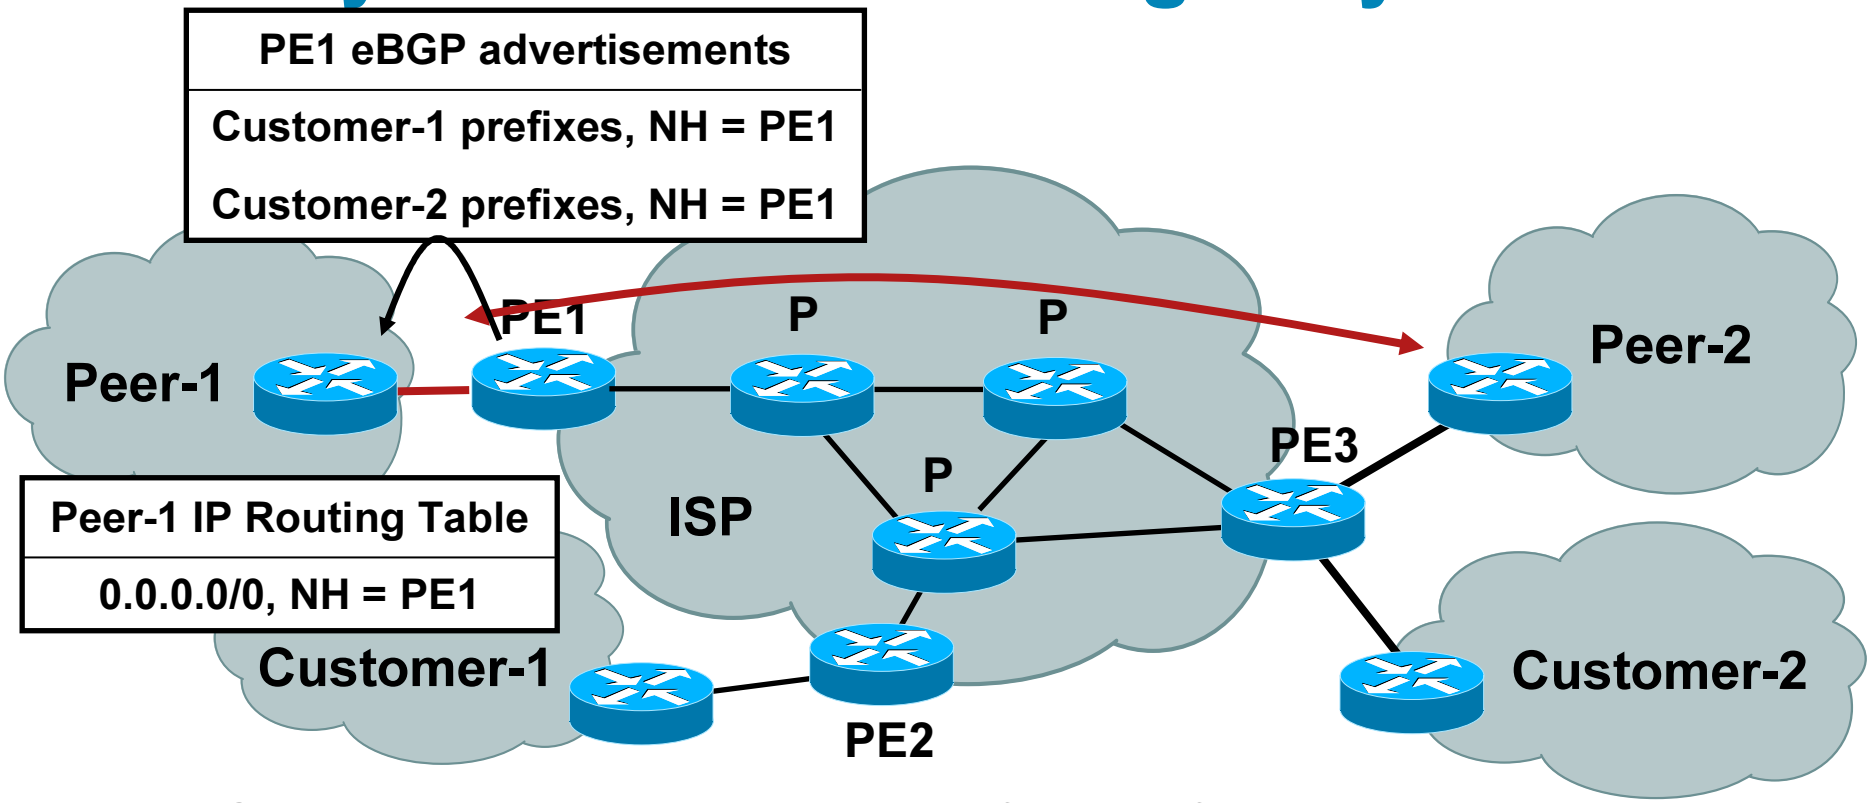

- **BGP control plane techniques only filter prefix advertisements**
- **If a peer uses IP routing tricks (e.g., default routing), it may bypass BGP policies and steal bandwidth from the ISP peer**
  - For example, using the peer as transit to another peer
- **This is possible because BGP policies are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

Challenges with Alternate Options

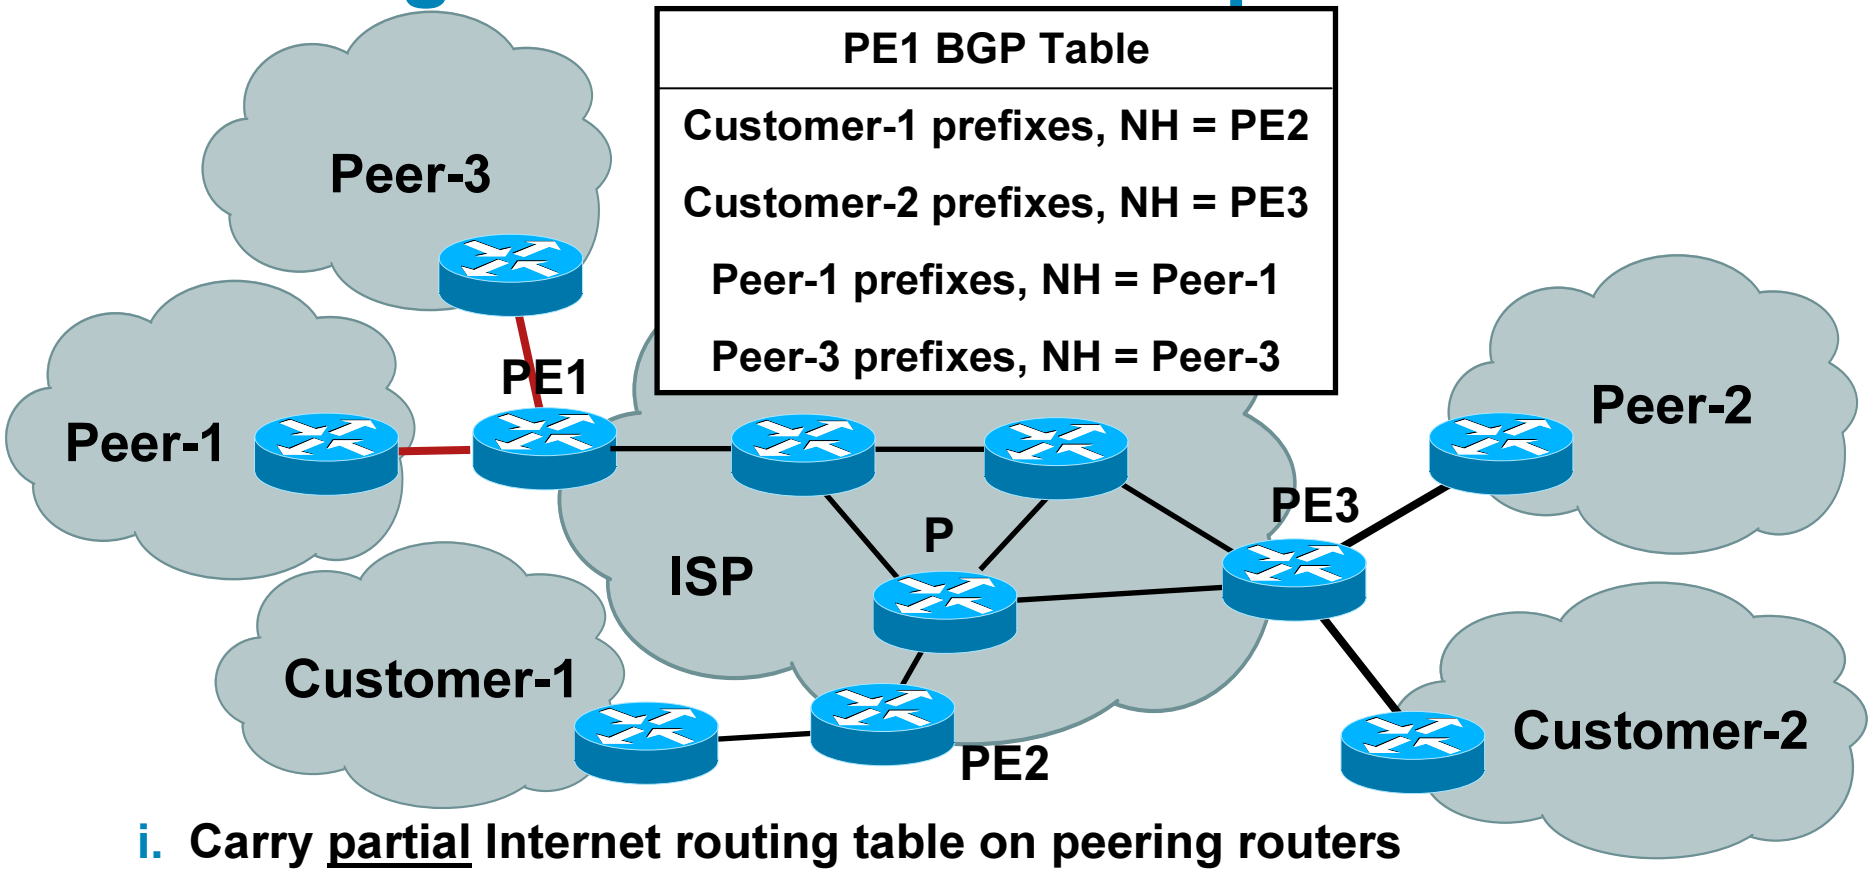

i. **Carry <u>partial</u> Internet routing table on peering routers**
   - For example, filter Peer-2 prefixes from being carried on PE1
   - Does not prevent IP reachability between peers connected to the same local peering router (e.g., Peer-1 and Peer-3)

ii. **Interface ACLs – not scaleable 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

Benefits of Proposed Technique

- **Traffic received from a peer and destined to a peer is dropped in accordance with Internet peering policies**
  - Traffic received from a peer and destined to a customer prefix is forwarded normally
- **Traffic between peers attached to the same ISP peering router is also filtered in accordance with Internet peering policies**
  - Peering routers can carry the full Internet routing table
- **Prefix tagging within the FIB (e.g., peer versus customer prefix) possible through standard BGP policies**
- **Proposed technique glues the IP control plane policy (e.g., BGP) with the IP data plane policy (e.g., MQC)**
  - No ACLs required
  - BGP topology and policy changes automatically reflected within the IP data forwarding plane
- **Enforcement of Internet peering policies within the IP data forwarding plane protects against an Internet peer using routing tricks to bypass control plane policies**

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*

Illustration of Proposed Technique

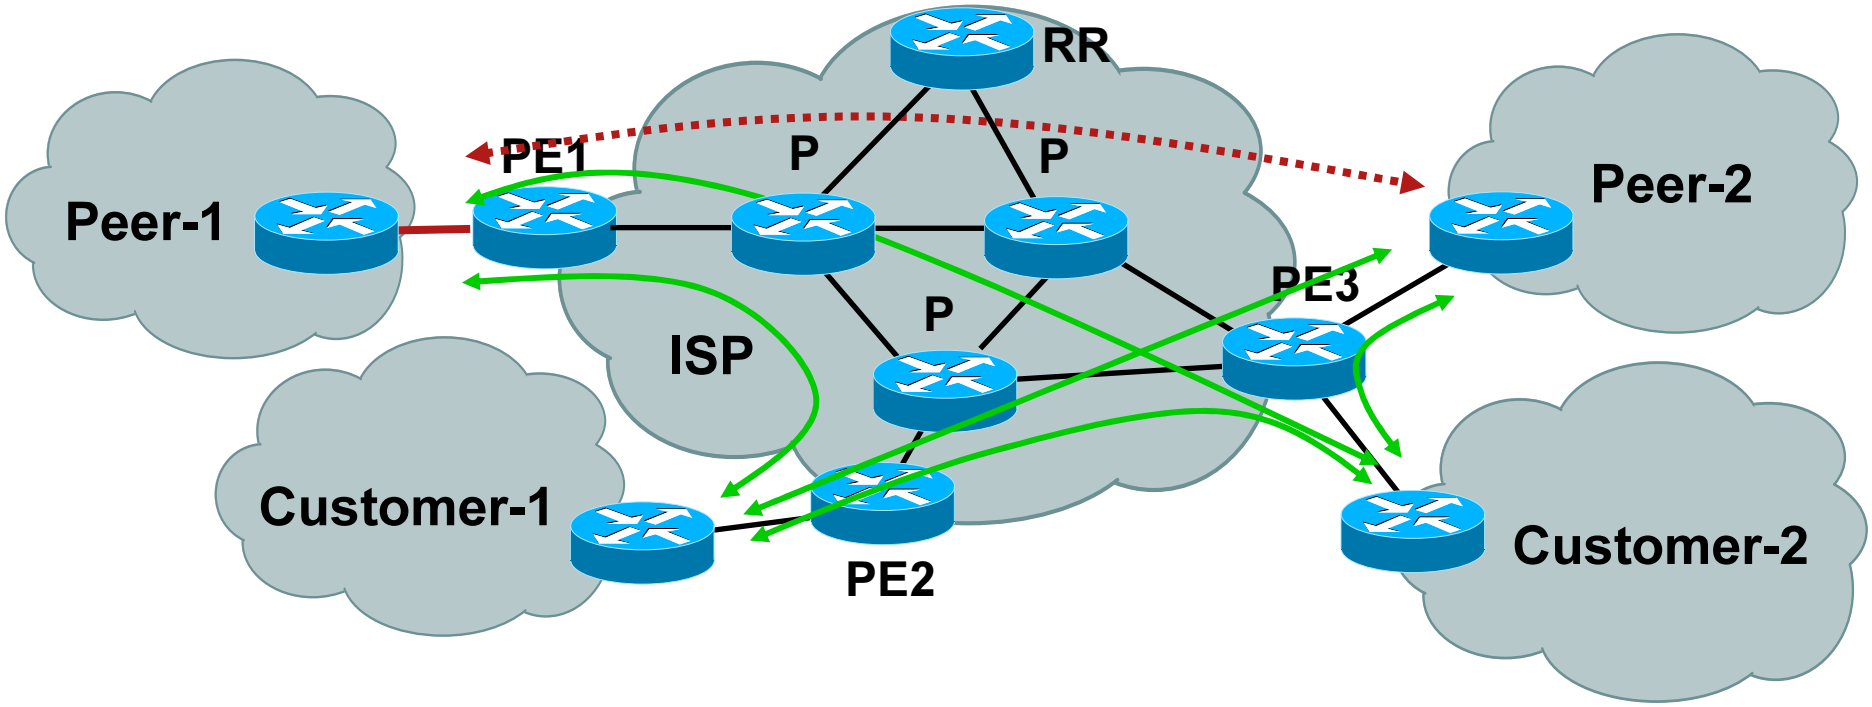

- **ISP tags peer prefixes uniquely within its BGP and FIB tables**
  - Both peer-1 and peer-2 prefixes are tagged with BGP community attribute {ISP-ASN}:66 and FIB tag 66
  - Both customer-1 and customer-2 prefixes are tagged with BGP community attribute {ISP-ASN}:77 and FIB tag 77

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 {isp-asn}
!
  table-map set-prefix-type
!
  ip bgp-community new-format
!
route-map set-prefix-type permit 10
  match community 1
  set ip qos-group 66
!
ip community-list 1 permit {isp-asn}: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

Conclusion

- **QPPB glues the IP control plane policy (i.e., BGP) with the IP data forwarding plane policy (i.e., MQC)**
  - Complements other BGP control plane signaling applications (e.g., RTBH, sinkholes, etc.) commonly used today
- **QPPB provides a scaleable and operationally efficient technique to enforce Internet peering policies**
  - Packet level policy enforcement using BGP/FIB prefix classification
  - No ACLs required
  - BGP topology and policy changes automatically reflected within the IP data forwarding plane
  - Traffic between peers attached to the same ISP peering router is also filtered in accordance with Internet peering policies
  - Available today

References

- Scholl, T., N. Patrick, A. Shaikh, and R. Steenbergen. *Peering Dragnet: Anti-Social Behavior Amongst Peers, and What You Can Do About It*. NANOG 38, 2006.
- Schudel, G., and D.J. Smith. *Router Security Strategies: Securing IP Network Traffic Planes*. Cisco Press, 2008. ISBN-10: 1-58705-336-5.

