



# PEERING BOF XVII

NANOG 42



# What is this peering BOF thing?

- \* Explain what the peering BOE<sub>is</sub><sub>Text</sub>
- \* Who attends these things?
- \* What should we expect?
- \* What do Peering People do anyway?

Anonymous Video contribution

# The Premise

ValueBOF~Who there

<show Peering Community Movie>

\* 8-9 yrs. -->Community Seating - fosters discussion <sup>Interesting topics</sup>

\* Discussions preferred over presentation

\* several microphones, please stand up

\* historically debates, voting, beer, drunks, opinions, bad ideas

\* Most Important

\* Get the right people into the room, face to face afterwards.

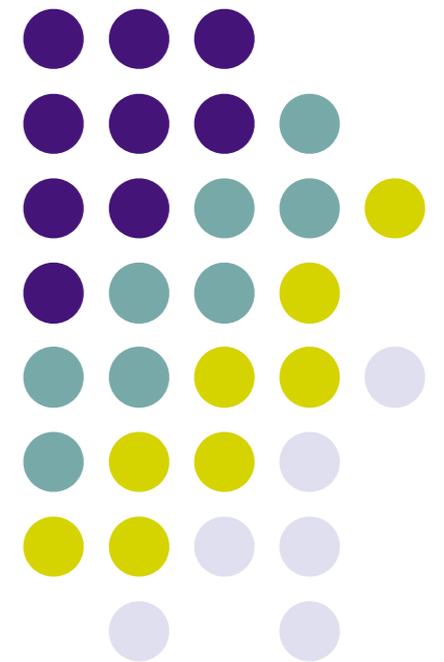
Beat the bushes

# Peering BOF XVII

- \* 4:00 -4:05 What is this Peering BOF thing? - Bill Norton
- \* 4:05-4:10 Peering Survey - Greg Hankins and Ren Provo
- \* 4:10-4:25 Efficient Technique for Enforcing Internet Peering Policies, David J. Smith, Cisco Systems
- \* 4:25-4:35 Summary proxy sociological analysis of Peering Coordinator Interaction - Chris Malayter
- \* 4:35-4:40 asr Transit Playbook
- \* 4:40-4:50 Peering Personals - batch #1
- \* 4:50-4:55 Ken Florance -Video & Peering
- \* 4:55-5:10 Jeffrey Payne - p2p, topology awareness and peering
- \* 5:10-5:25 The Great Debate "Does Peering Make Sense Anymore?" - shockg/ianai
- \* 5:25-5:30 - Peering Personals People come to front of room - meet peers

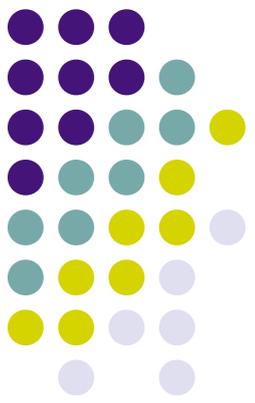
# Peering Survey 2008

Greg Hankins  
Ren Provo  
Tom Scholl



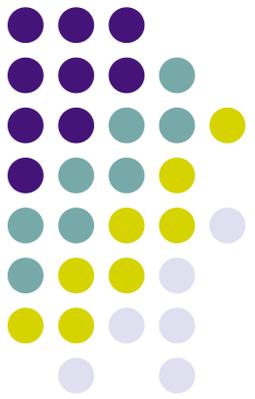
Man **peering** through **survey** equipment

# What is it?

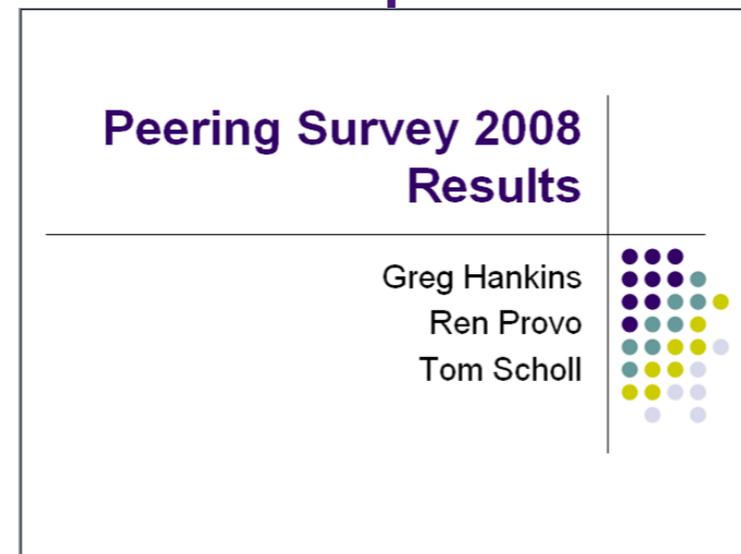
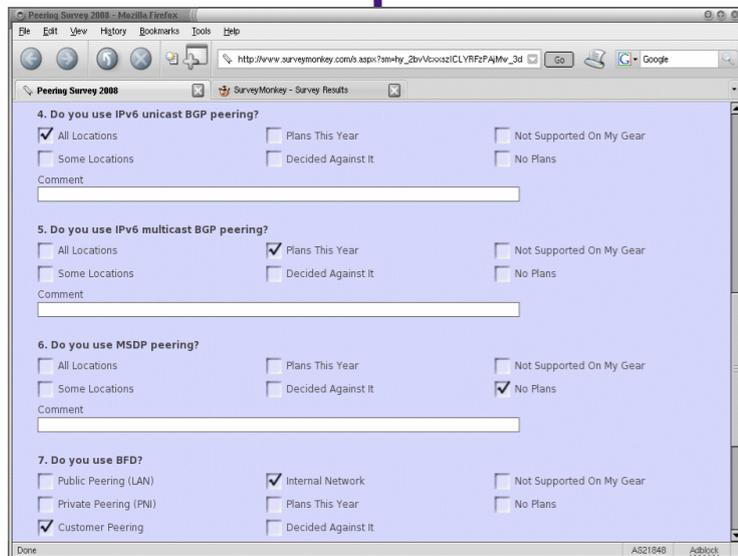


- New survey on how people configure peering!
- Featuring technical questions on what protocols and features are used
- No questions on facilities, IXs, Peter Cohen or who you'd like to depeer most

# How does it work?



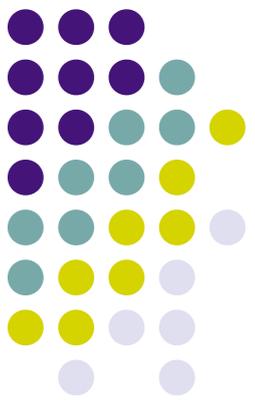
**Now!**



**You take the survey.**

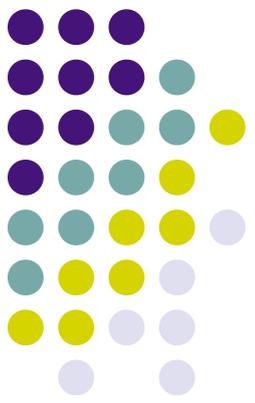
**We present the results at  
GPF 3.0.**

# We need you to take the survey!



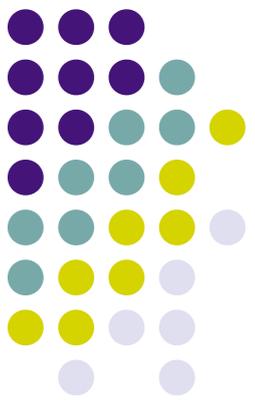
- It will take you under 10 minutes to answer 24 easy questions
- All you do is go to SurveyMonkey here:
  - <http://tinyurl.com/3xoa6g>  
([http://www.surveymonkey.com/s.aspx?sm=hy\\_2bvVcxxszICLYRFzPAjMw\\_3d\\_3d](http://www.surveymonkey.com/s.aspx?sm=hy_2bvVcxxszICLYRFzPAjMw_3d_3d))
- We need lots of people to participate, please!
  - No personal or identifying information will be used

# Sample Questions



- Do you use IPv6 unicast BGP peering?
- Do you use BFD?
- Do you use four byte ASNs (RFC 4893)?
- What is the largest frame size you use on peering links?
- What is your biggest concern when deploying a new feature?

# Questions?

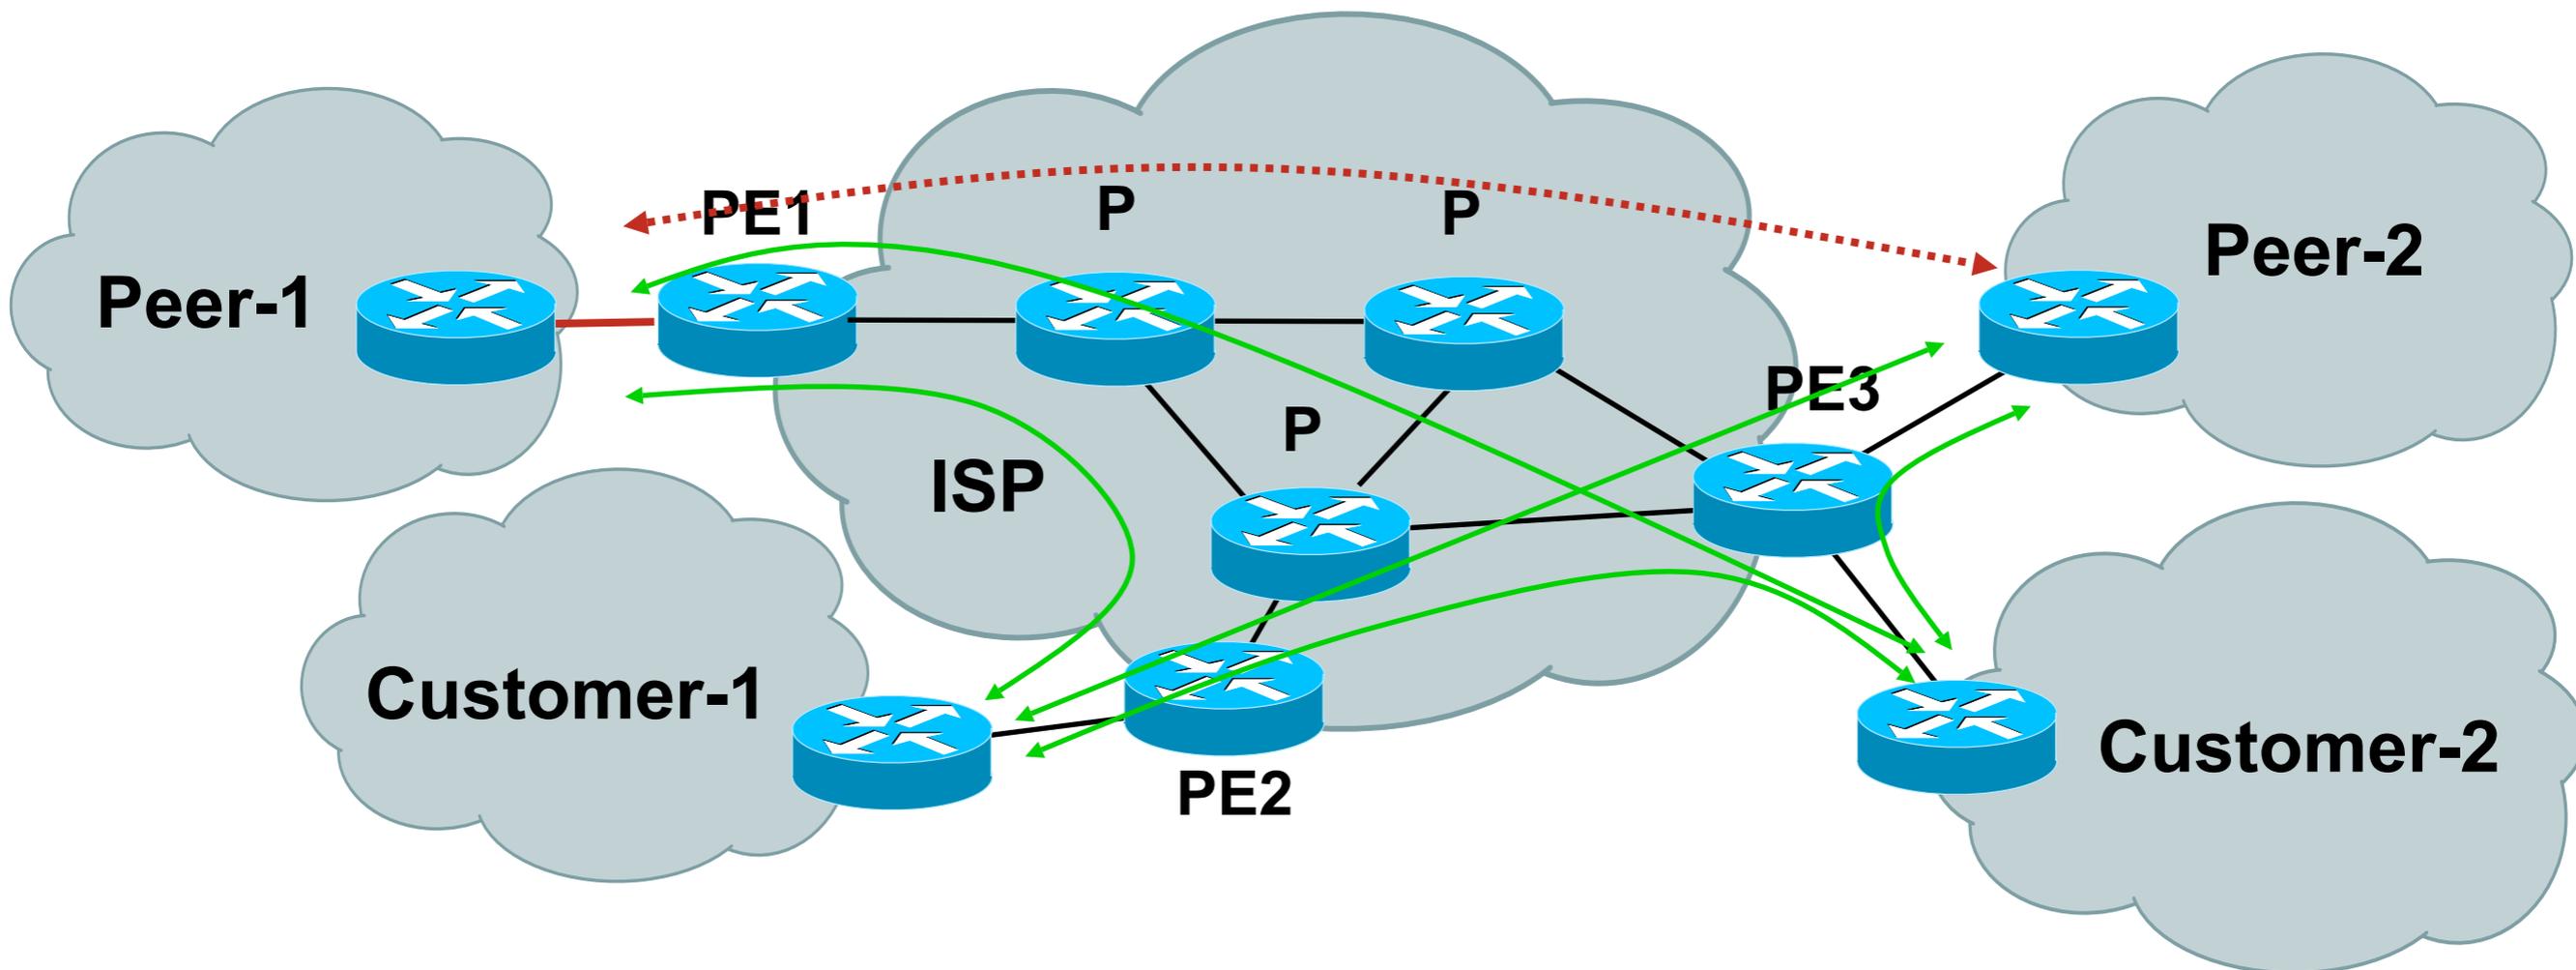


- Thanks for participating
- Get in touch with questions or comments
  - Greg <[ghankins@mindspring.com](mailto:ghankins@mindspring.com)>
  - Ren <[ren.provo@gmail.com](mailto:ren.provo@gmail.com)>
  - Tom <[tscholl@gmail.com](mailto:tscholl@gmail.com)>

# Efficient Technique for Enforcing Internet Peering Policies

**David J. Smith** – [dasmith@cisco.com](mailto:dasmith@cisco.com)

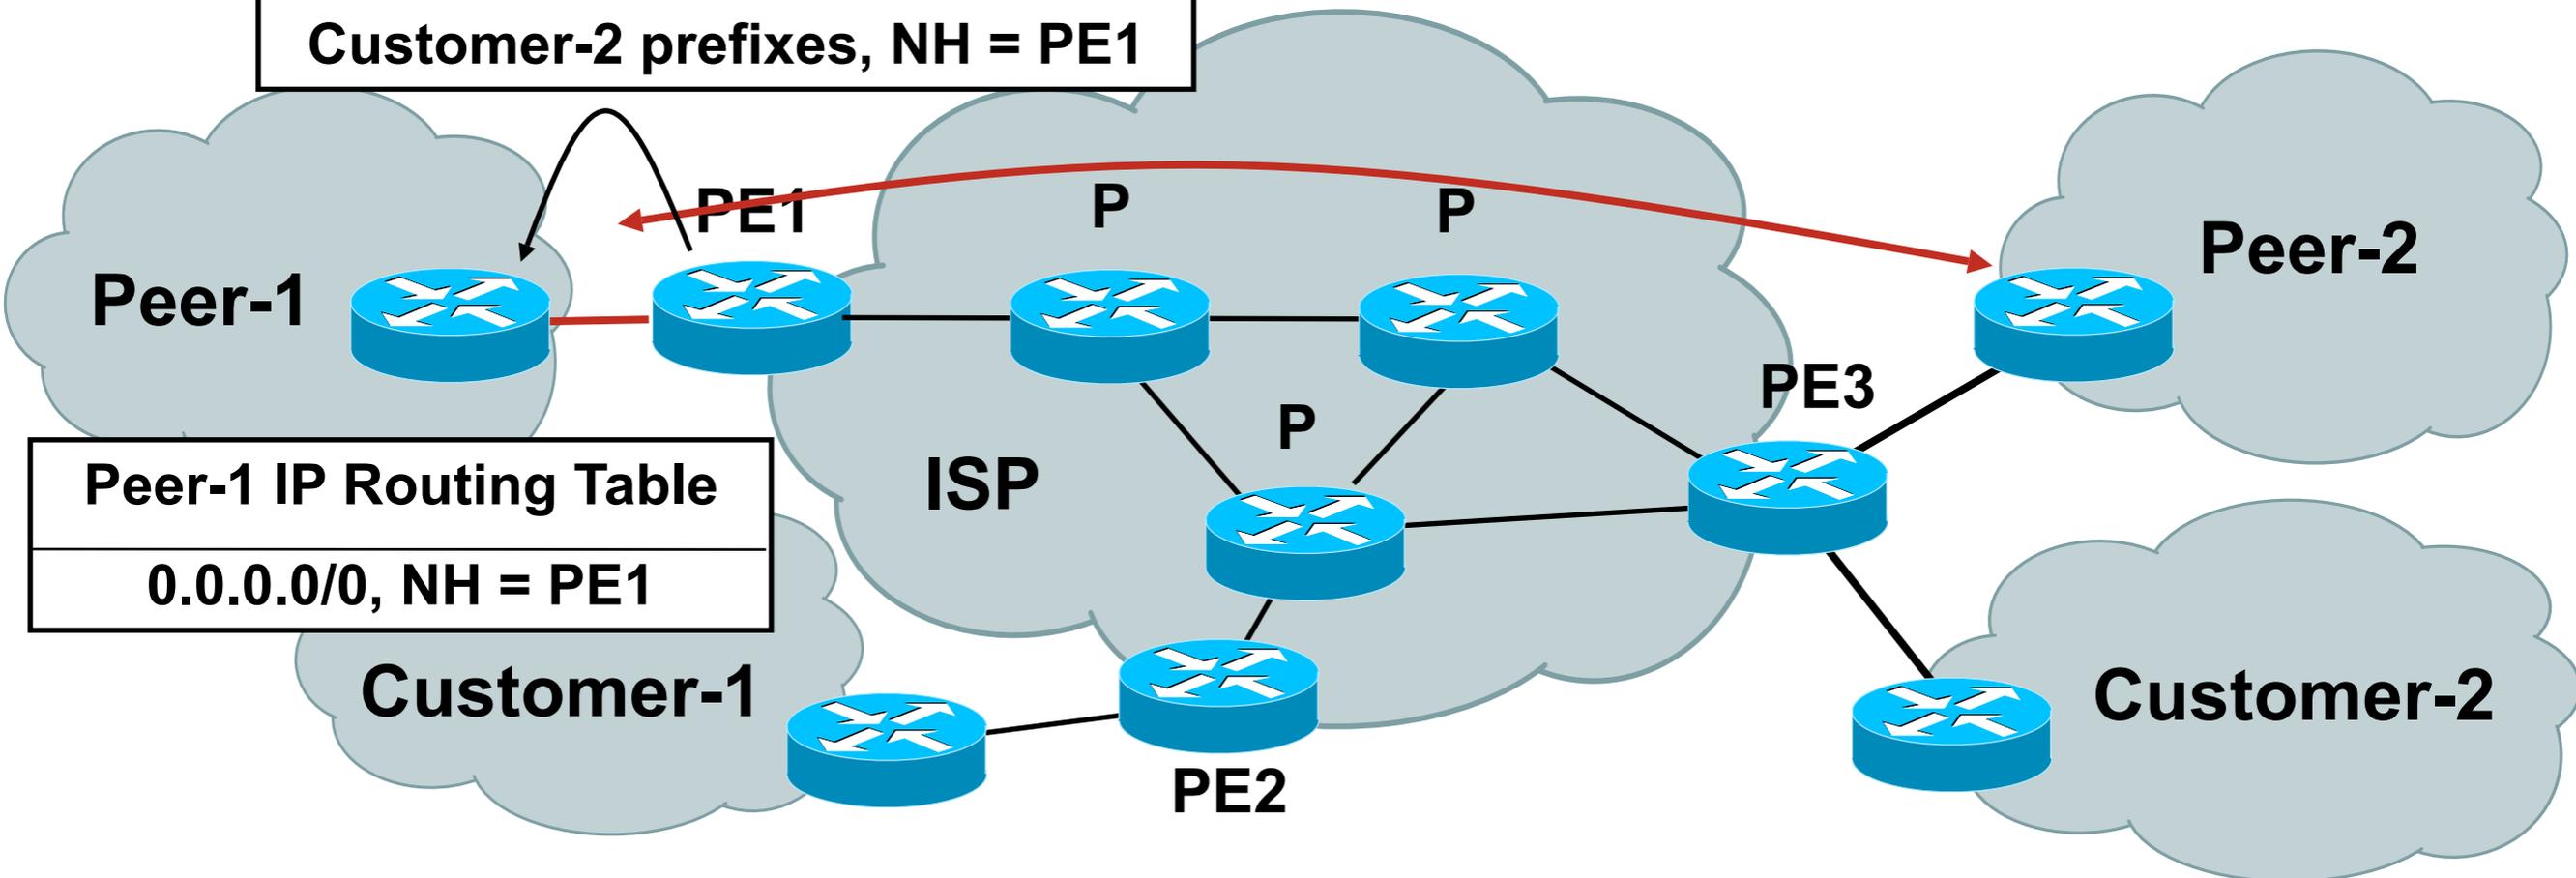
# Internet Peering Policy Overview



- **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)

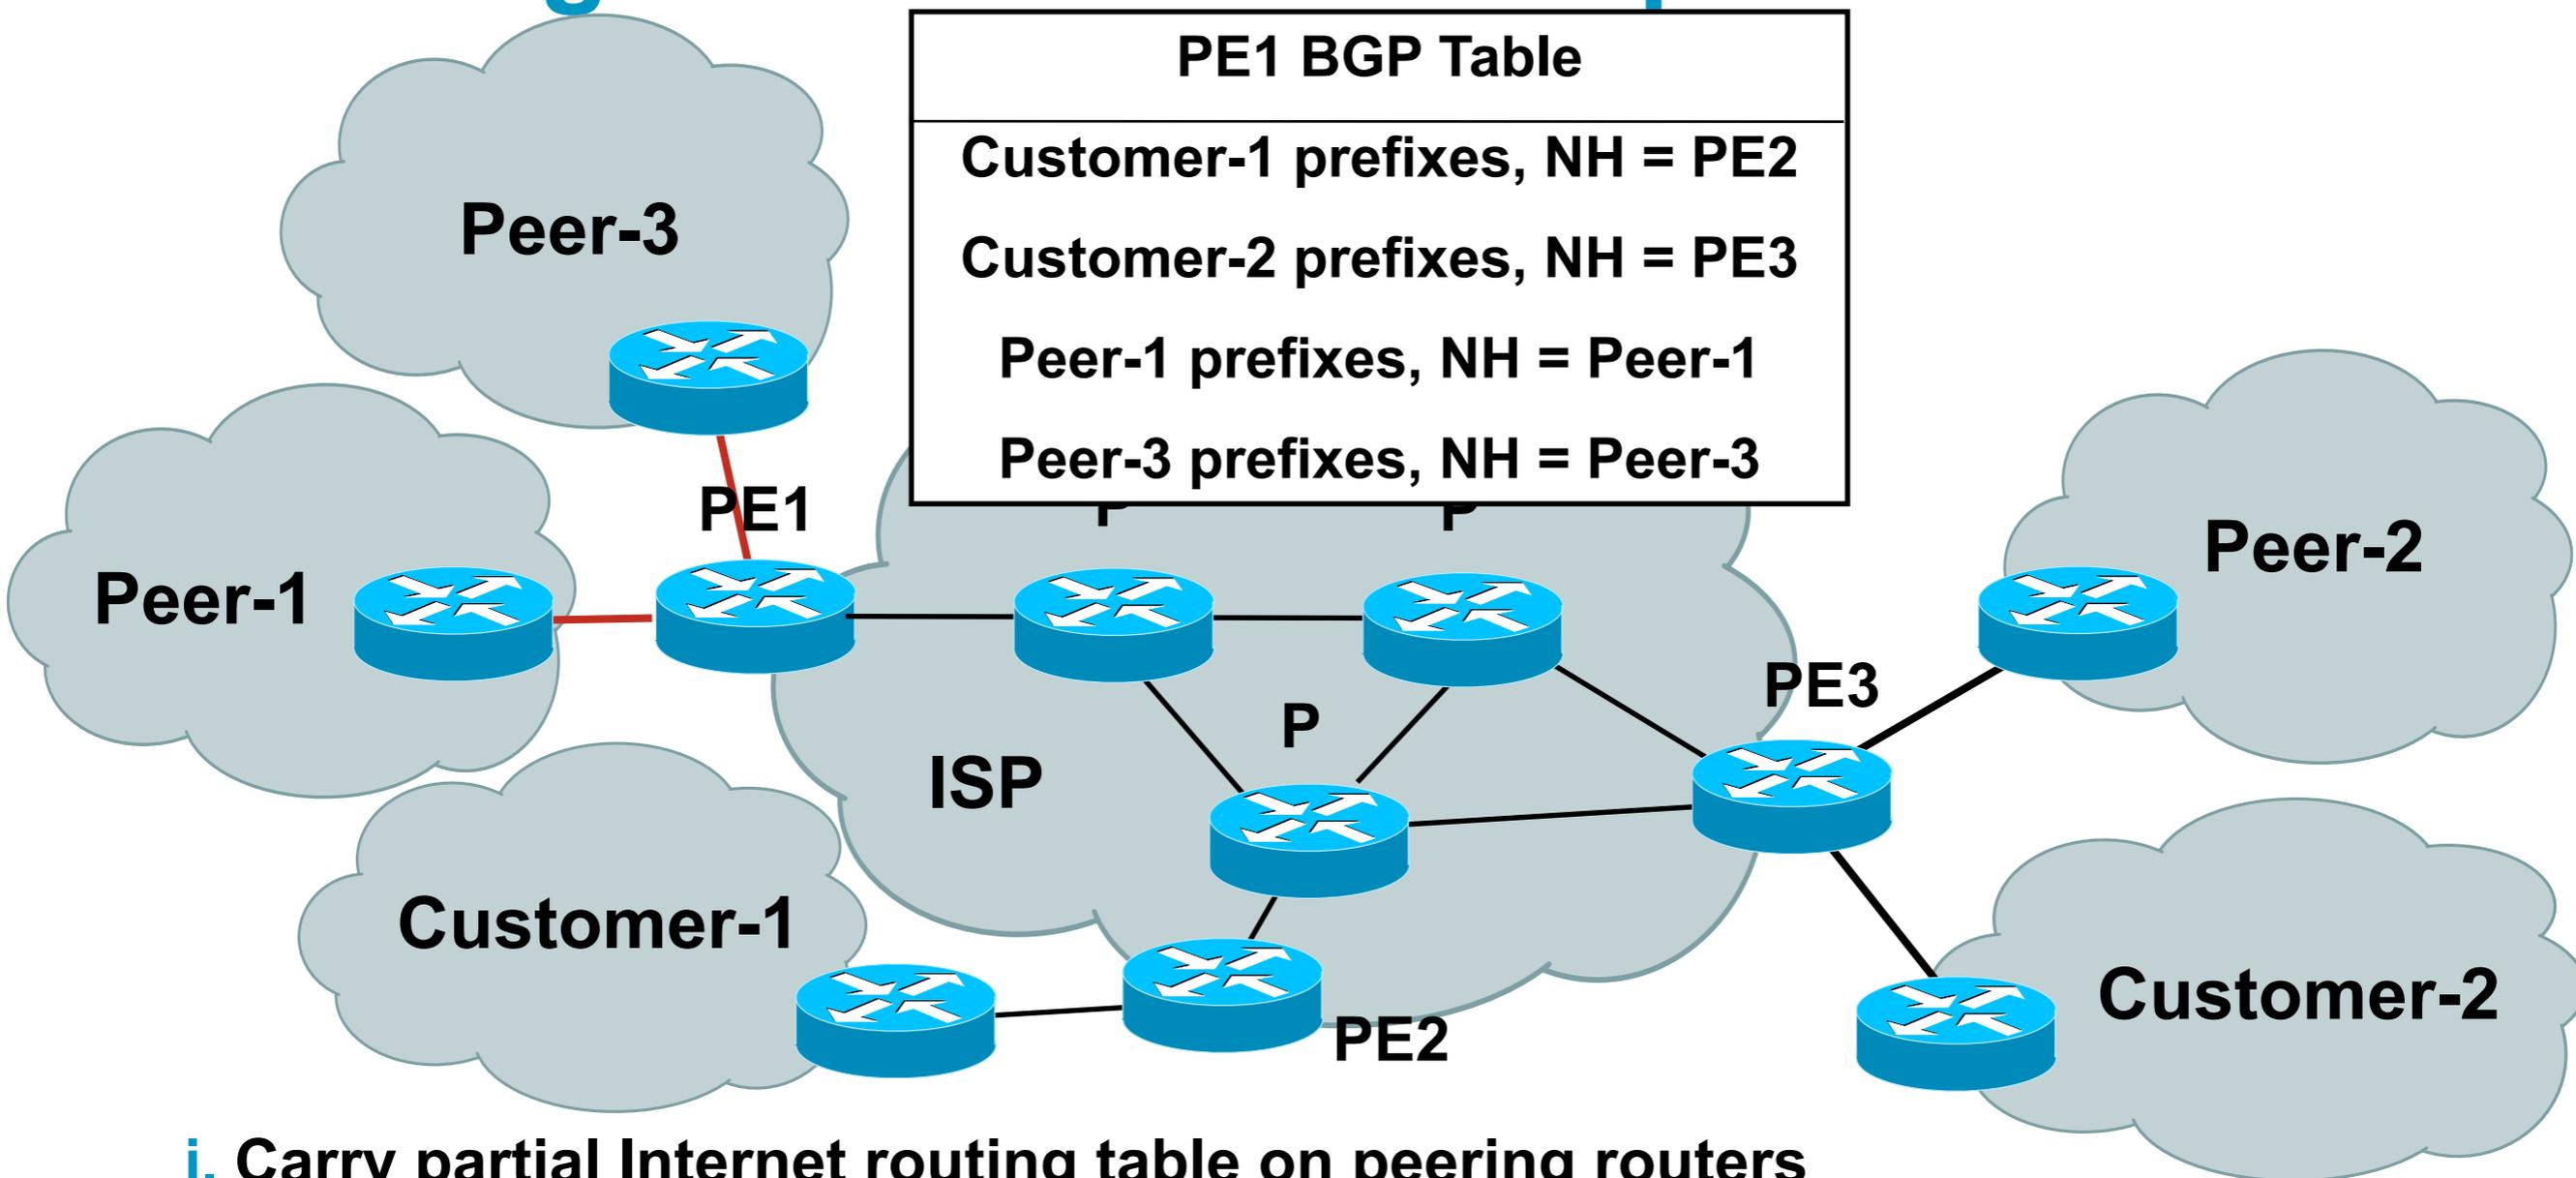
# Policy Enforcement Using Only BGP

PE1 eBGP advertisements
Customer-1 prefixes, NH = PE1
Customer-2 prefixes, NH = PE1



- 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 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)

## 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

# 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 {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 (e.g., 66 for peer prefixes)**

# 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.**

\* SUMMARY PROXY SOCIOLOGICAL ANALYSIS OF PEERING  
COORDINATOR INTERACTION - CHRIS MALAYTER



---

---

# PEERING PERSONALS

---

---



Batch #1 - Ethern Lin - AS9264  
Matt Peterson - VideoBox - 36472  
Simon Ferrett - Veoh - AS 40415  
Bryan Berg - Imeem - 36119

# Ethern M., Lin AS9264

◆ Peering Locations Now:  
HKIX, Hong Kong

JPIX, Japan

PAIX, Palo Alto

Peering Locations Planned:  
None

Why should  
people want to  
peer with you?

◆ Transit Load Today: 700Mbps

Traffic Balanced

Will do publics or privates

---

We provide IP transit service for academic and research networks.

Ethern M., Lin, <[ethern@ascc.net](mailto:ethern@ascc.net)>

Network Division

Peer with us you can get those routes in Taiwan, and

Computing Centre, Academia Sinica

you could reduce the cost of IP transit.

Phone: +886-2-2789-9953

# Matt Peterson - VideoBox - AS36472

Peering Locations Now:  
365 Main, PAIX, San Jose  
EQIX, SF-MIX

Transit Load Today:

Peering Locations Planned:  
None

Traffic Outbound

Why should  
people want to  
peer with you?

Free membership with each peering session.

# Simon Ferrett - Veoh - AS40415

Simon Ferrett <[sferrett@veoh.com](mailto:sferrett@veoh.com)>

Peering Locations Now:  
EQIX LA3

Transit Load Today: 75Gbps

Peering Locations Planned:  
TBD

Traffic Outbound

Why should  
people want to  
peer with you?

What do you  
look for in a  
peer?

# Bryan Berg - Imeem - AS36119

Transit Load Today: 15G

[Bryan@imeem.com](mailto:Bryan@imeem.com)

Peering Locations Now: Public: EQIX ASH, EQIX San Jose

Private: EQIX CHI, EQIX LA, PAIX Palo Alto

Peering Locations Planned: Any2

Traffic Outbound

Why should  
people want to  
peer with you?

Looking for eyeballs in Philippines, Malaysia, Singapore,

South America, Portugal

---

---

# \* JEFFREY PAYNE - PEER-TO-PEER Q&A

---



Top 5 Q's

---



# The Great Debate

\* Does Peering Make Sense Anymore?

\* Pro: Patrick Gilmore - Presenting the view that "Peering still makes sense"

Con: Guy Tal - Presenting the view that "Peering doesn't make sense anymore"

These two debaters have graciously volunteered to share the strongest sides of the arguments regardless of how they or their employers view the issue. This is an educational exercise to explore and then discuss the arguments on both side of the issue.

---

---

# PEERING PERSONALS

Finishing up



# Brandon Ross - AS14302, 32459, 23293

[<brross@xiocom.com>](mailto:brross@xiocom.com)

Transit Load Today: 10Mbps

Peering Locations Now: 56 Marietta St. Atlanta

Peering Locations Planned: TBD

Traffic Balanced

Why should  
people want to  
peer with you?

We'll be growing rapidly

# Cheol-Hee Yun-AS3786-LG Datacom

[yich3425@chollian.net](mailto:yich3425@chollian.net)

Transit Load Today: 4Gbps

Peering Locations Now: PAIX Palo Alto, EQIX San Jose

Peering Locations Planned: TBD

Traffic Outbound

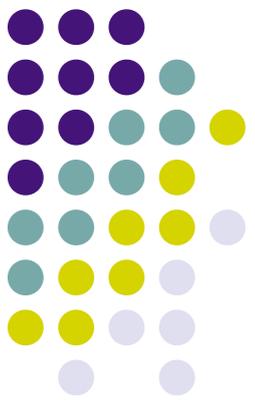
Why should  
people want to  
peer with you?

We have #1 IDC at KOREA. High-Speed  
Internet user reach 1.5 million also.

What are you  
looking for in a  
peer?

T-Systems

# We need you to take the survey!



- It will take you under 10 minutes to answer 24 easy questions
- All you do is go to SurveyMonkey here:
  - <http://tinyurl.com/3xoa6g>  
([http://www.surveymonkey.com/s.aspx?sm=hy\\_2bvVcxxszICLYRFzPAjMw\\_3d\\_3d](http://www.surveymonkey.com/s.aspx?sm=hy_2bvVcxxszICLYRFzPAjMw_3d_3d))
- We need lots of people to participate, please!
  - No personal or identifying information will be used