Peering at Peerings:
On the Role of IXP Route Servers

Contact: Philipp Richter (prichter@inet.tu-berlin.de)
Paper: net.t-labs.tu-berlin.de/~prichter/imc238-richterA.pdf

Philipp Richter
TU Berlin

Georgios Smaragdakis
MIT / TU Berlin

Anja Feldmann
TU Berlin

Nikolaos Chatzis
TU Berlin

Jan Boettger
TU Berlin

Walter Willinger
NIKSUN

Agenda

- **Introduction: IXPs and Route Servers**
- **IXP Route Server architecture**

- Empirical study
  - Peering offerings
  - Connectivity & traffic
  - Usage patterns

- Route Server Peering Strategies

IXPs are…

Physical locations that offer a shared (often distributed) layer-2 switching fabric for members (networks) to exchange traffic with one another.

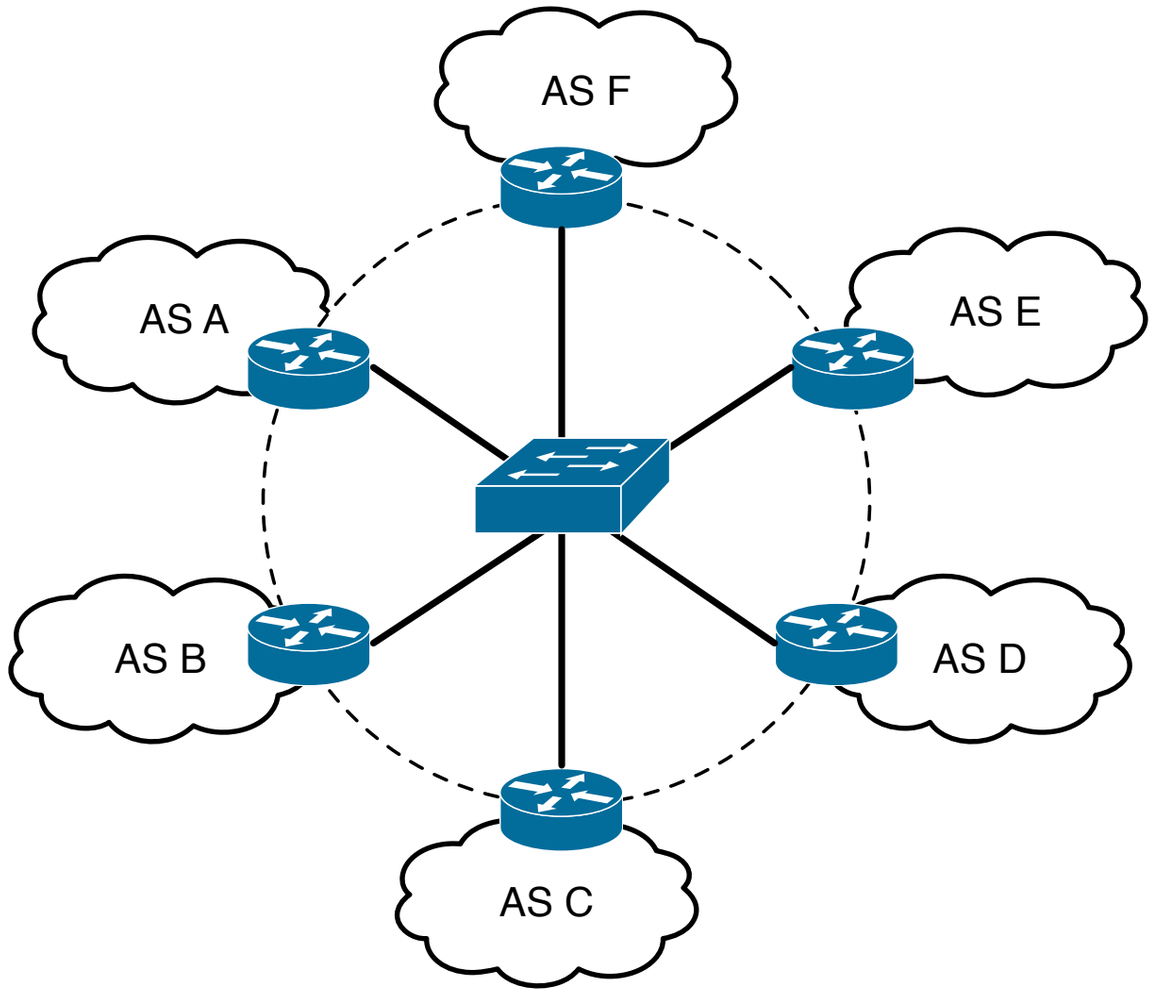

IXPs on the Increase

- Members benefit from peering opportunities
  - Reduced transit costs
  - Increased performance
  - Increased redundancy
- 350+ IXPs in the world
- Largest IXPs: 600+ members, 3 Tbps peak traffic

IXPs…

- Emerged as critical components in today's Internet
  - Establish large number of the Internet's peering links
    [Ager at al., SIGCOMM '12, Giotsas et al., ConEXT '13]
  - Key entities to bring content closer to the user
    [Labovitz et al., SIGCOMM '10, Chatzis at al., IMC '13]

- Fuel a more diverse peering ecosystem
  [Lodhi et al., CCR '14, Giotsas et al., IMC '14]

- Are eager to innovate
  - Resellers, Remote Peering [Castro et al., CoNEXT '15]
  - **Free use of Route Server**

IXP Route Servers

- What are IXP RSes?
- How do RSes work?
- What peering opportunities do RSes offer?
- How much connectivity do they set up?
- How do networks make use of them and why?

Peering at IXPs (bi-lateral)

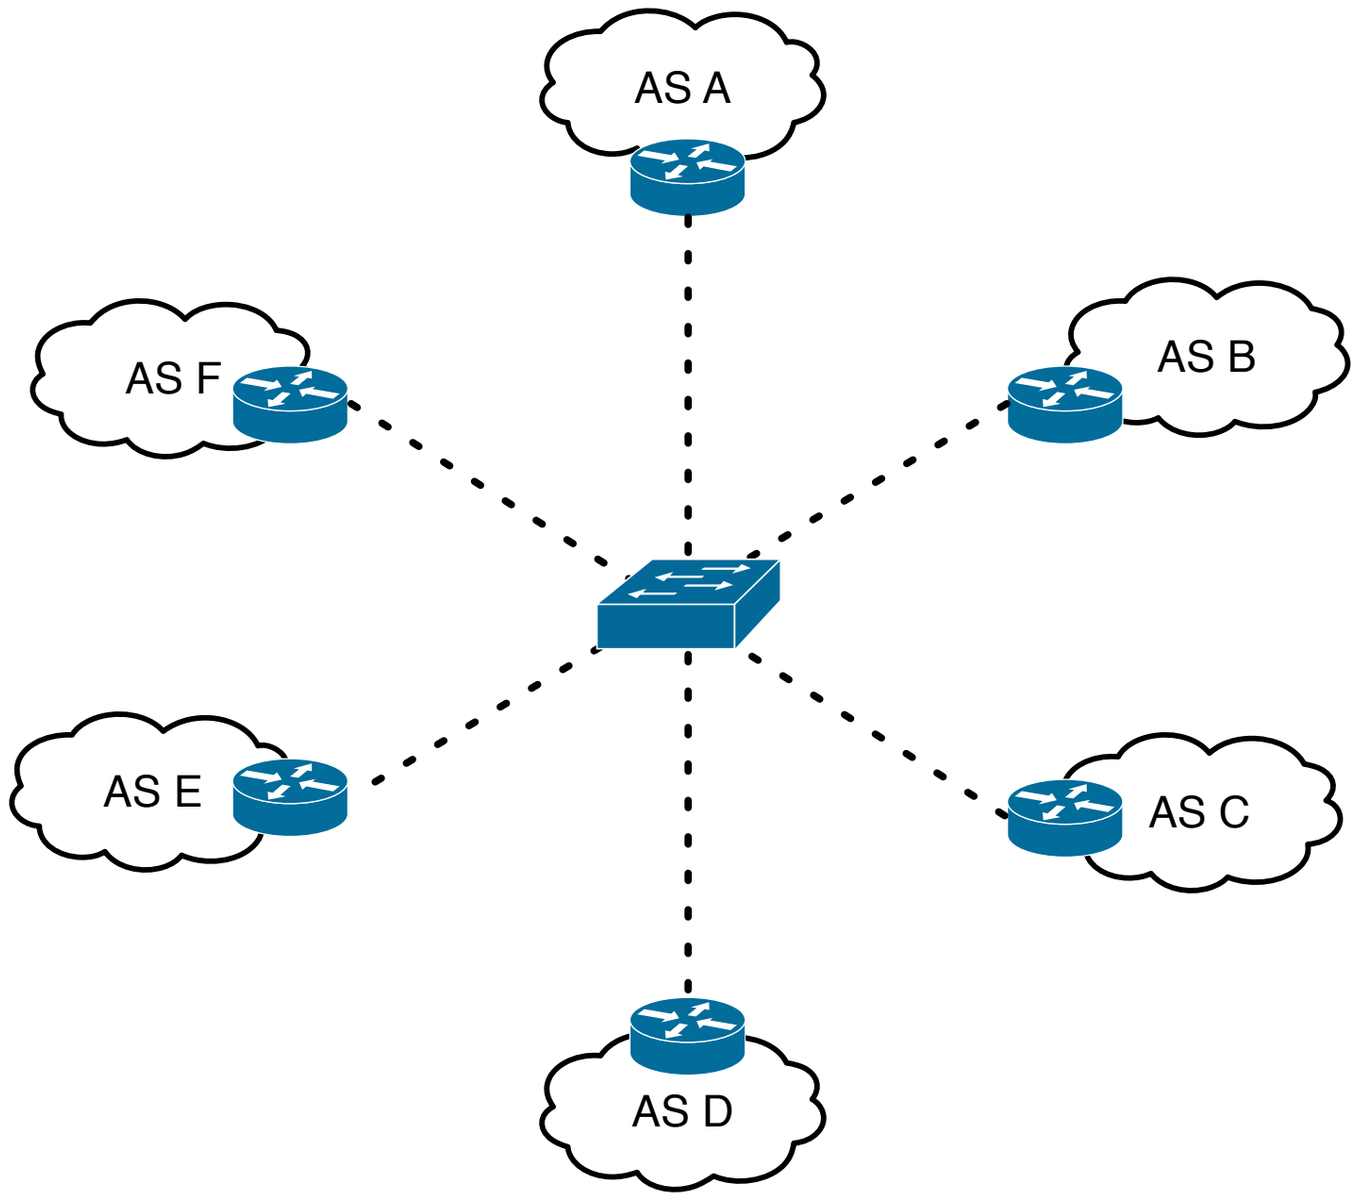

Peering at IXPs (bi-lateral)

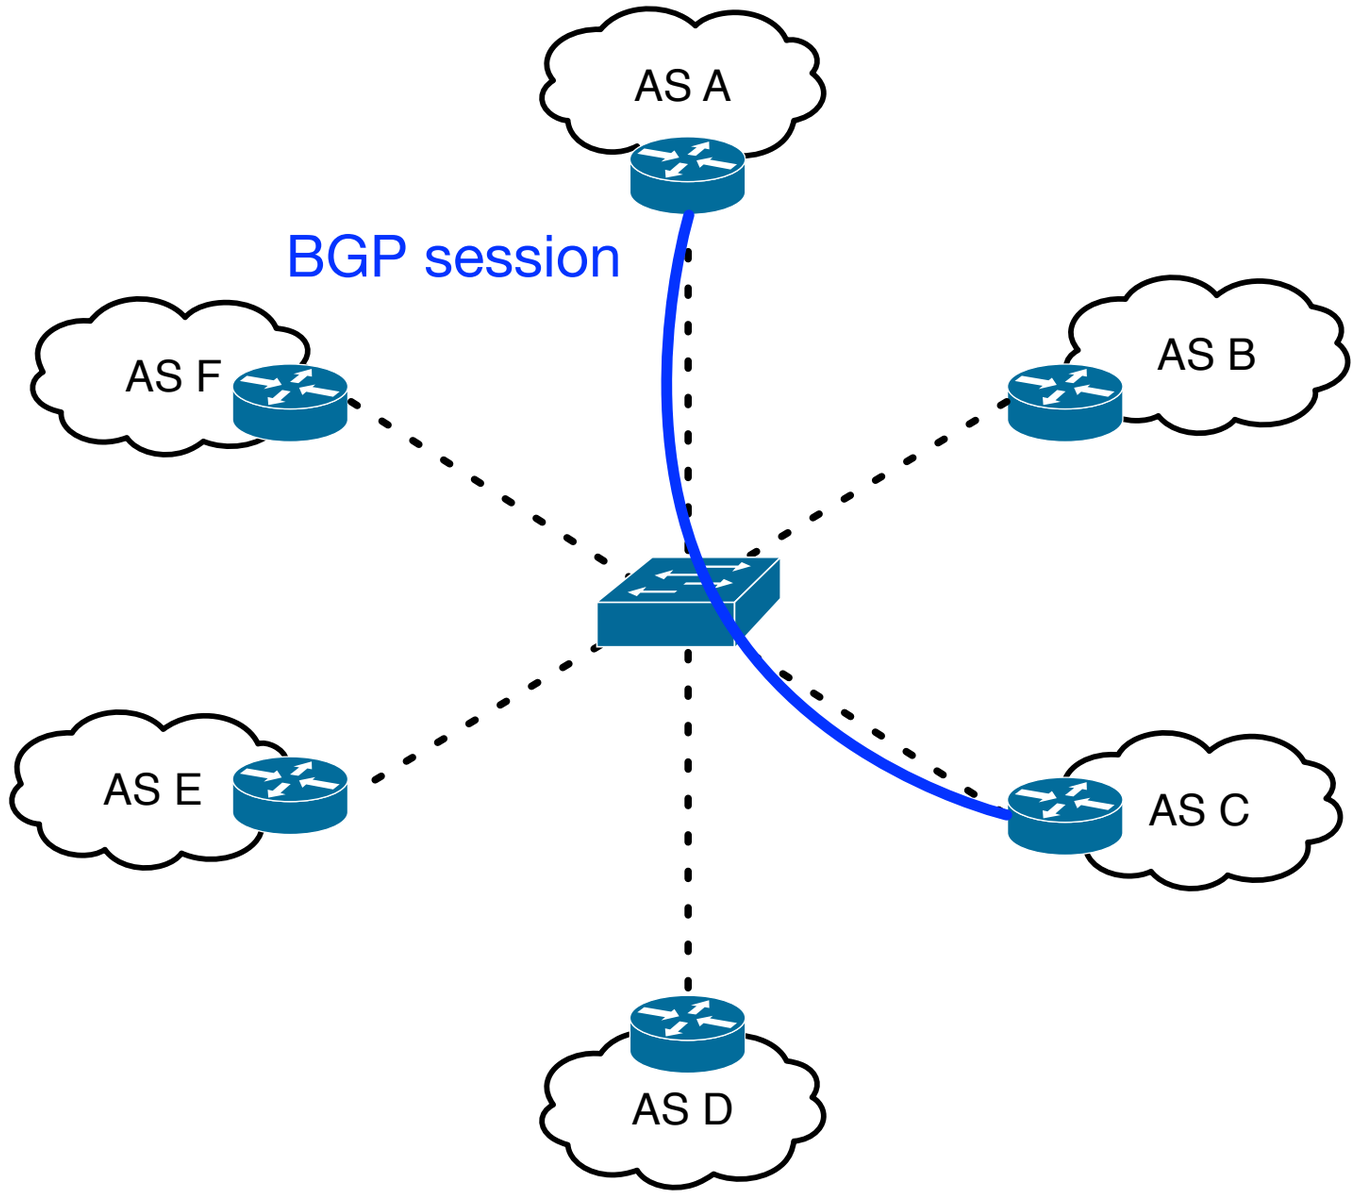

(1) Establish BGP session

Peering at IXPs (bi-lateral)

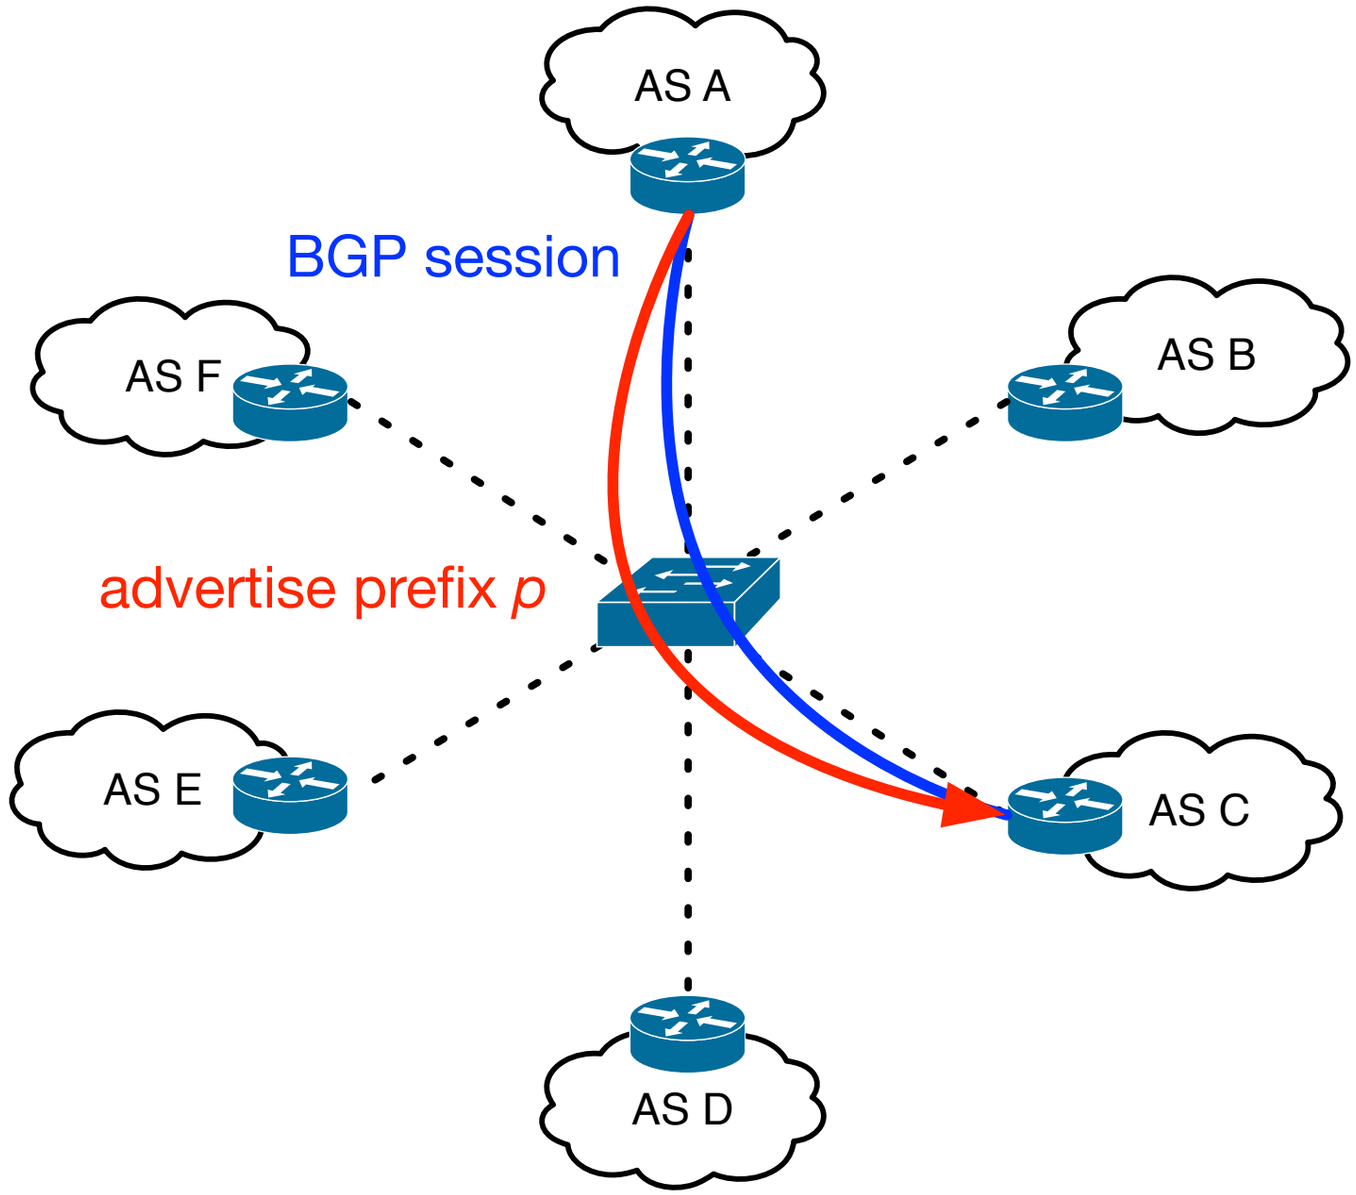

(2) Advertise prefix(es)

Peering at IXPs (bi-lateral)

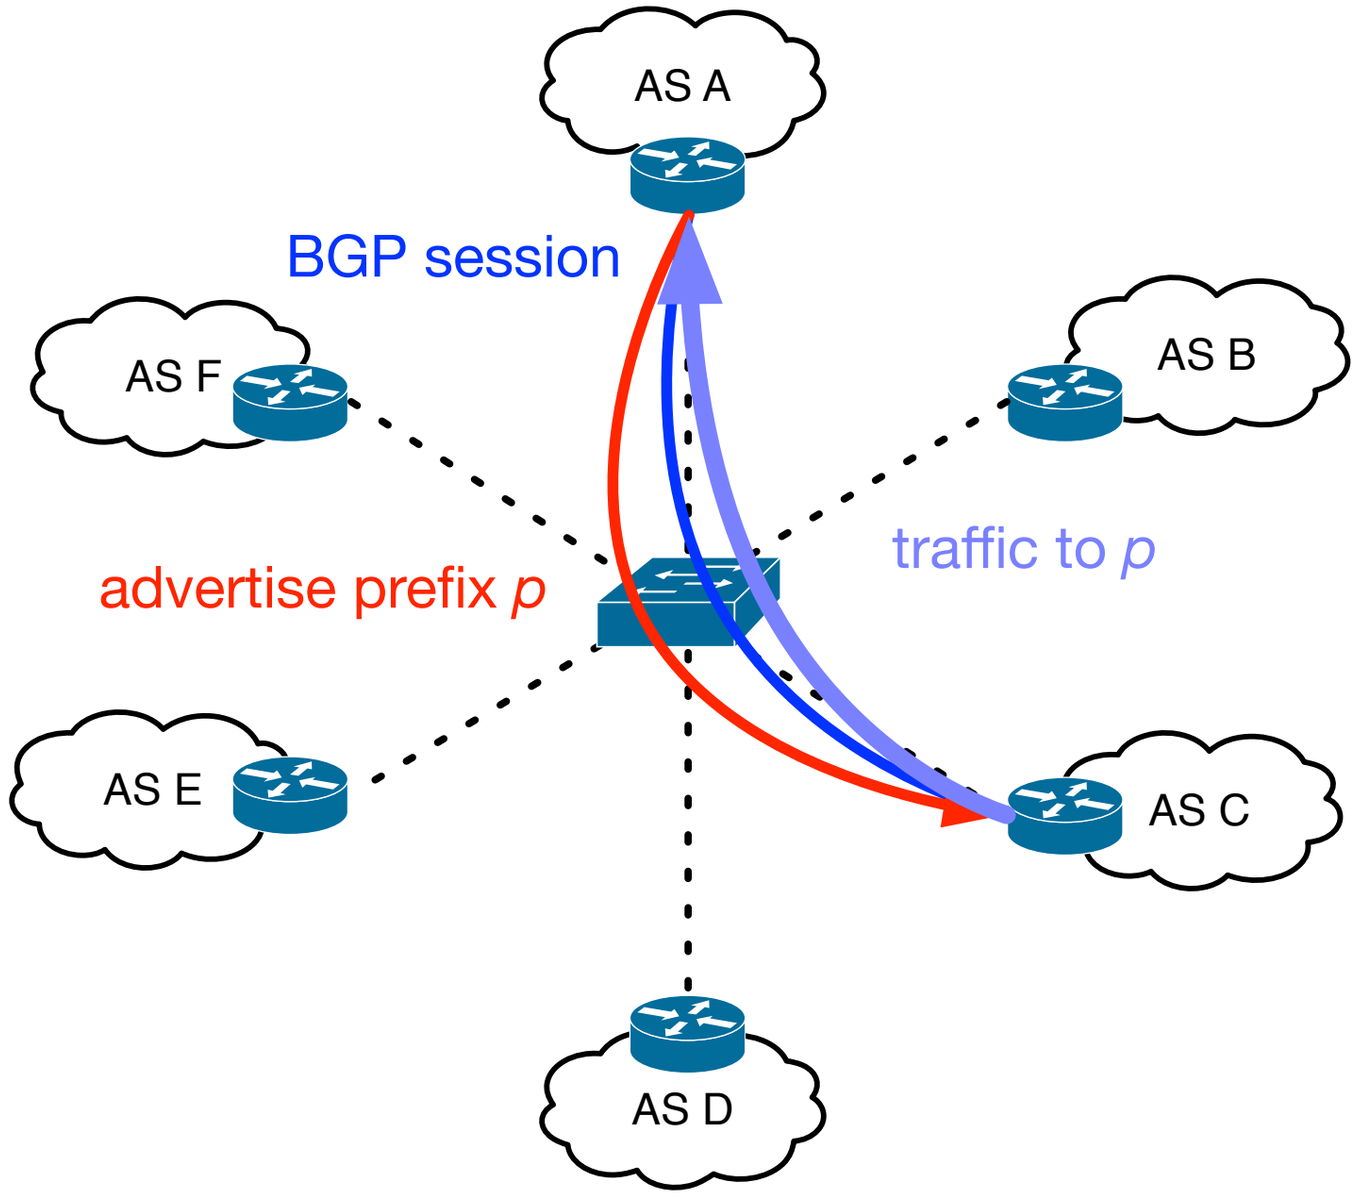

(3) Exchange Traffic

Peering at IXPs (bi-lateral)

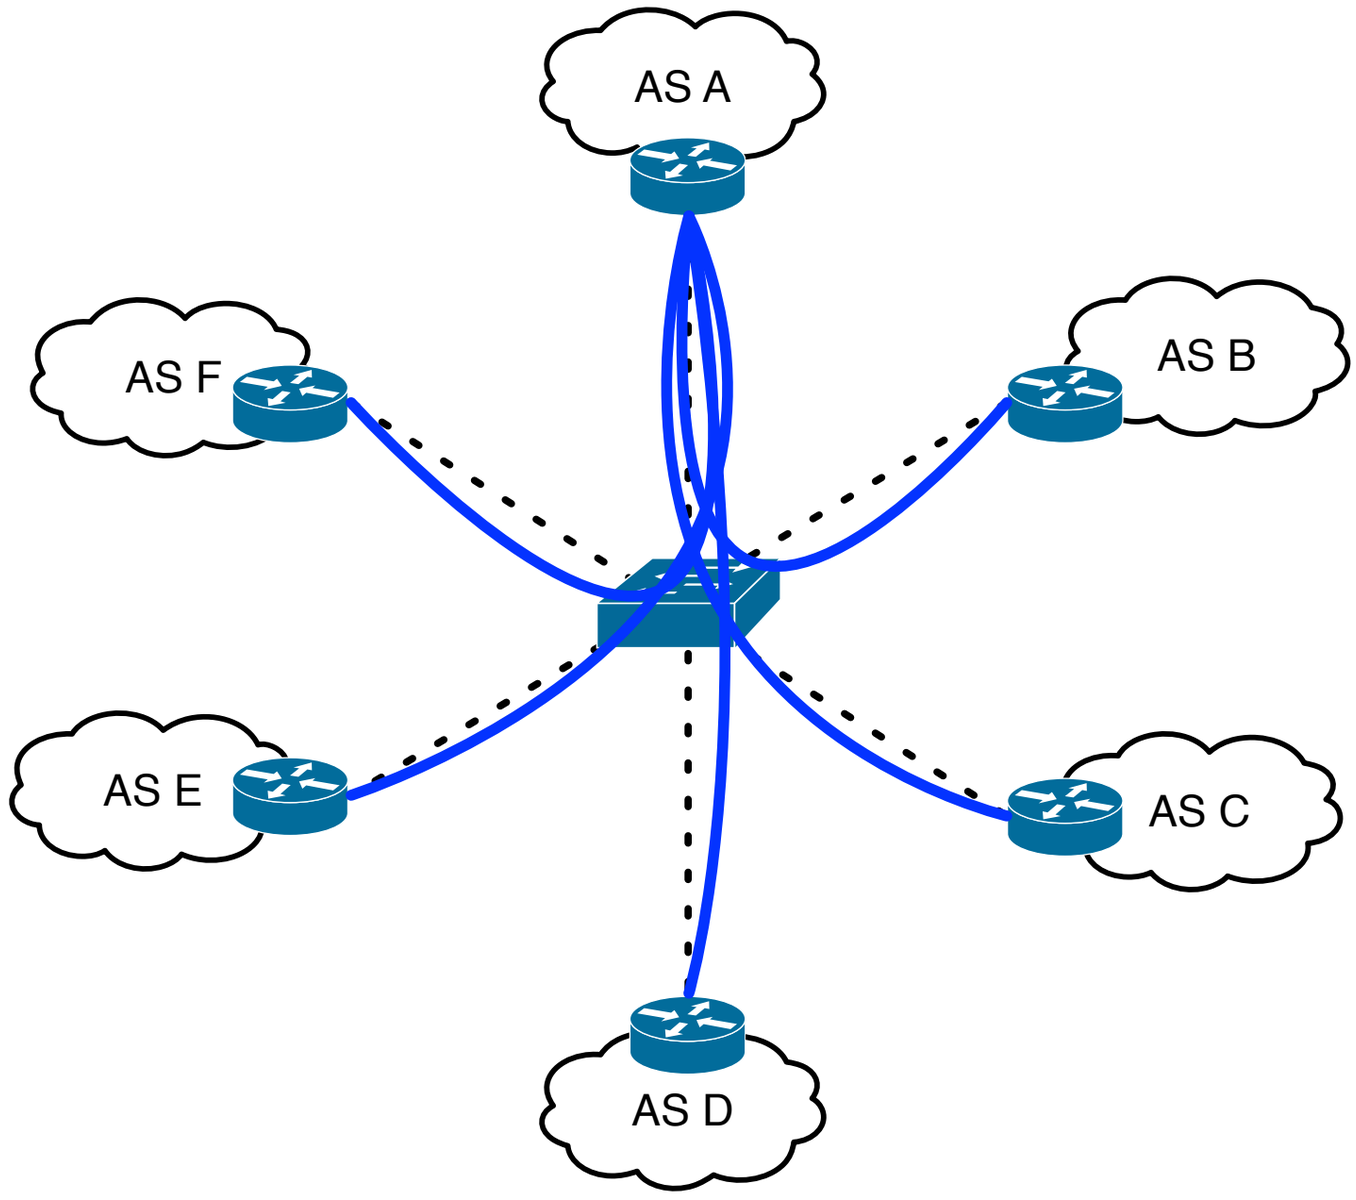

AS A needs 5 BGP sessions to peer with all other members.

Peering at IXPs (bi-lateral)

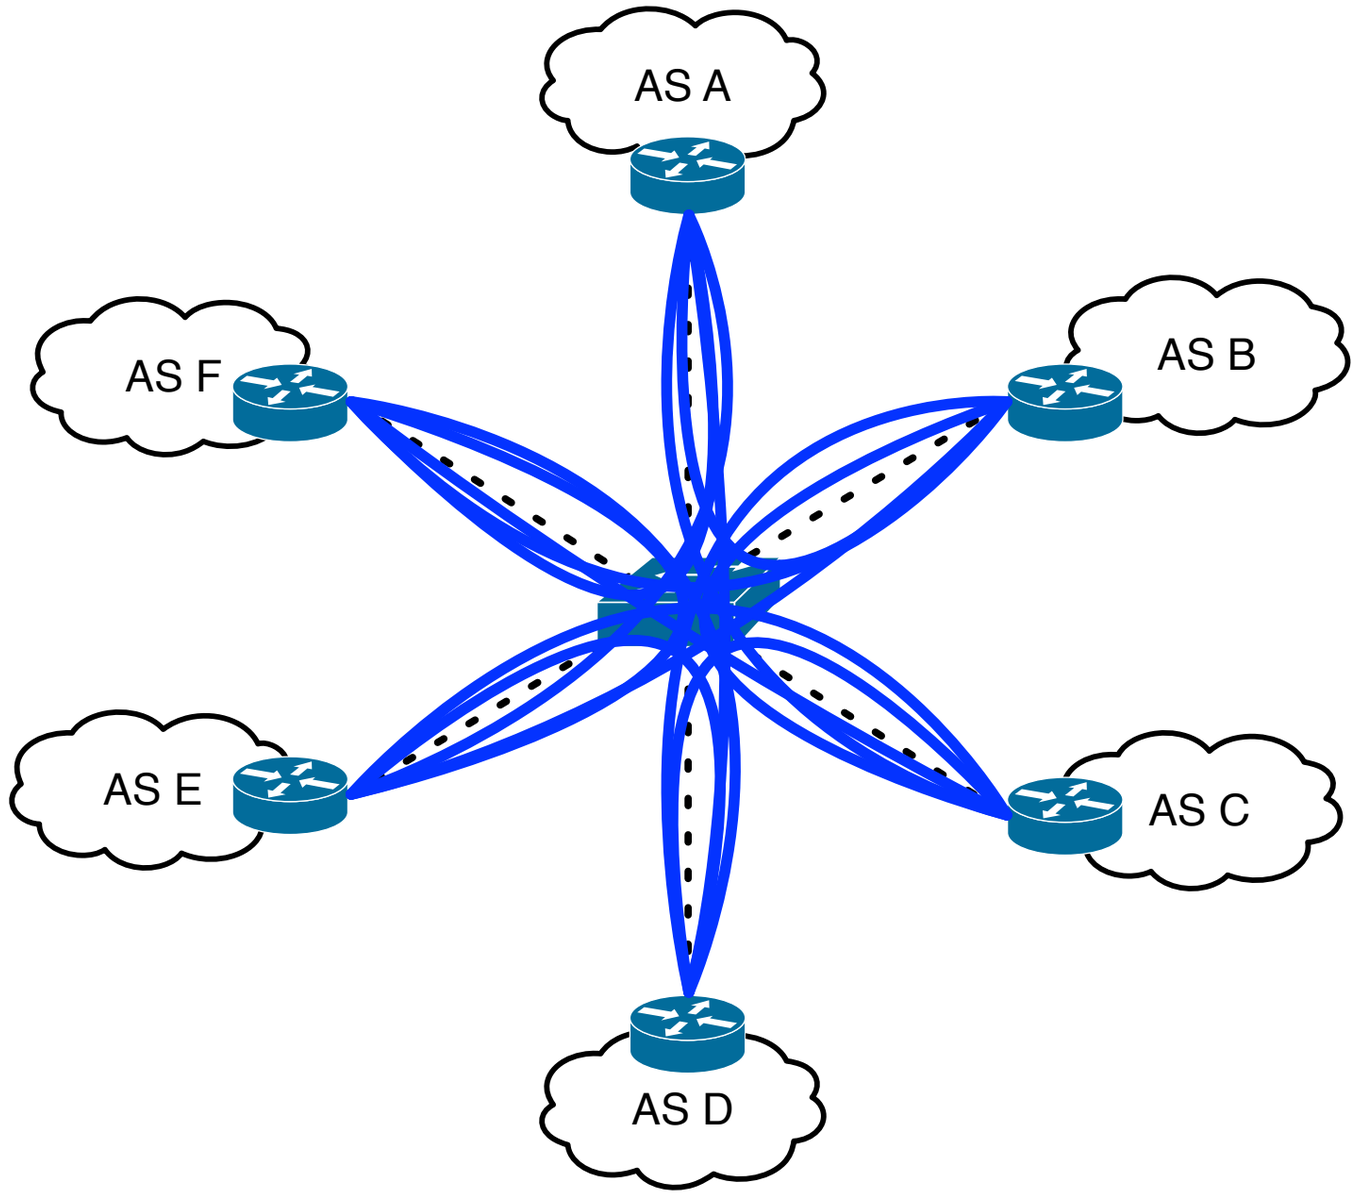

6 members: 15 sessions — 600 members: 180K sessions.

Peering at IXPs

- More peerings -> more benefit for each member

- Setting up peerings requires effort
  - Coordination between operators
  - Hardware limitations (early routers)

- Solution offered by IXPs: Route Servers
  - Instant peering with hundreds of networks

Peering at IXPs (multi-lateral)

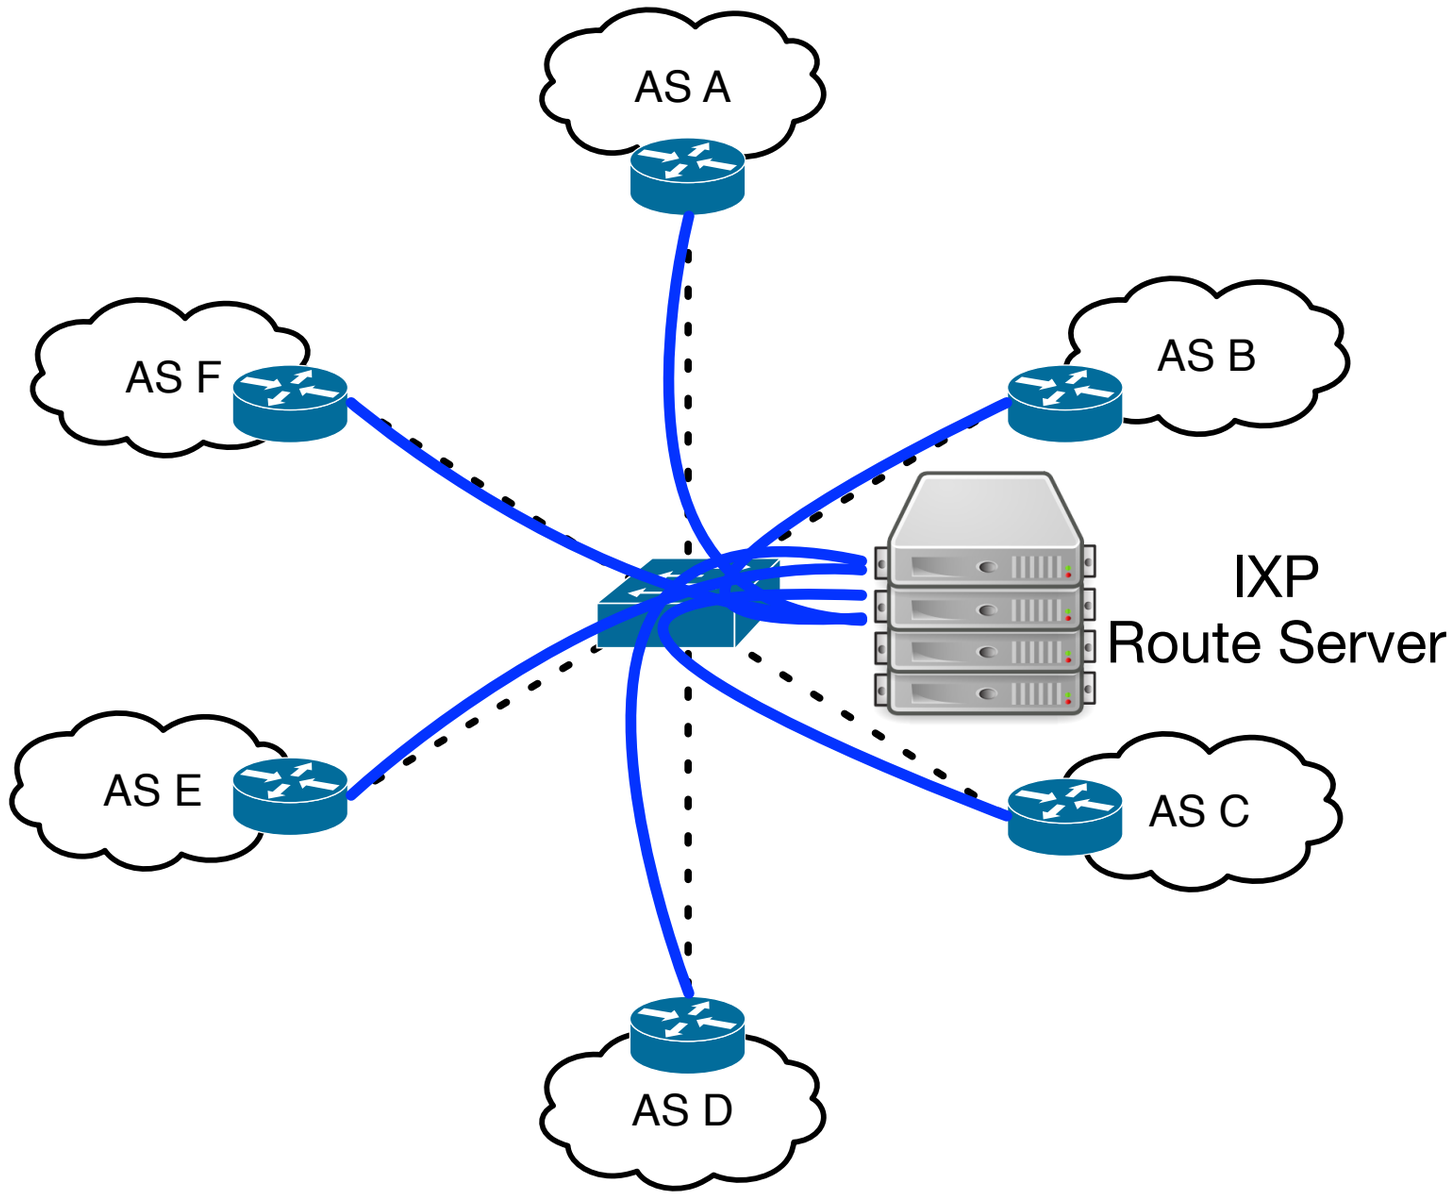

Route Servers make peering easy.

A Modern RS Architecture

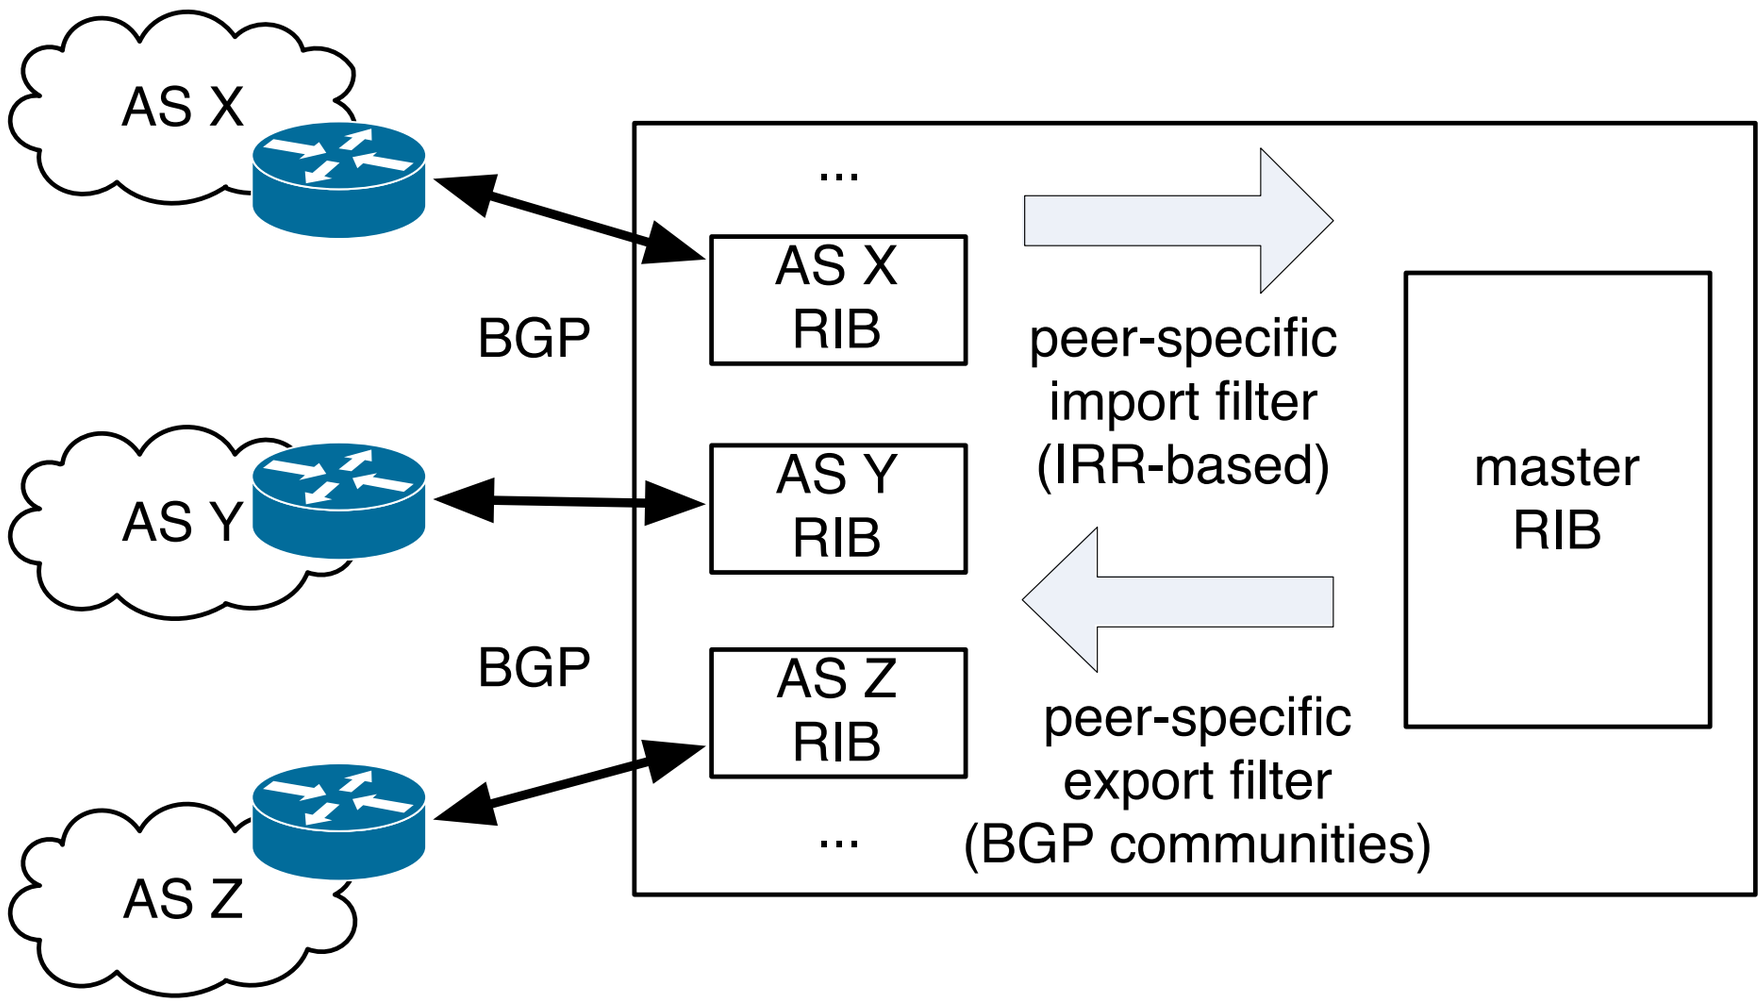

A Modern RS Architecture

AS X advertises prefix *p* (standard BGP)

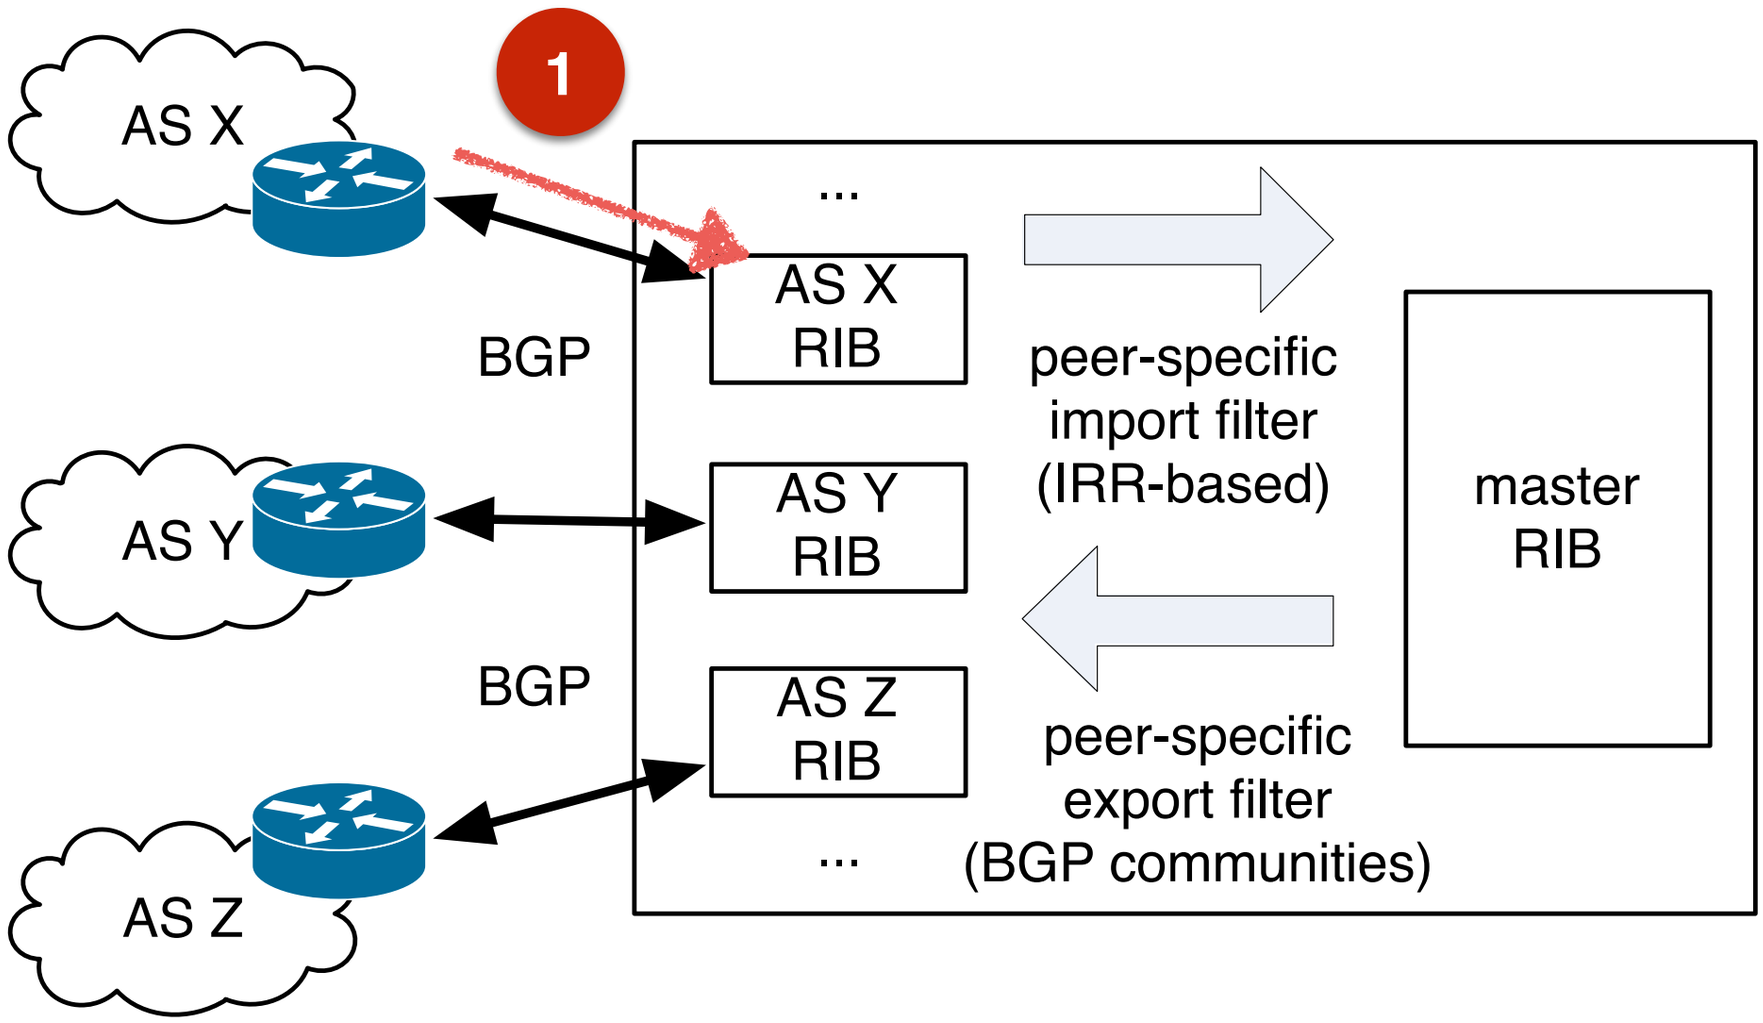

A Modern RS Architecture

AS X advertises prefix *p* (standard BGP)

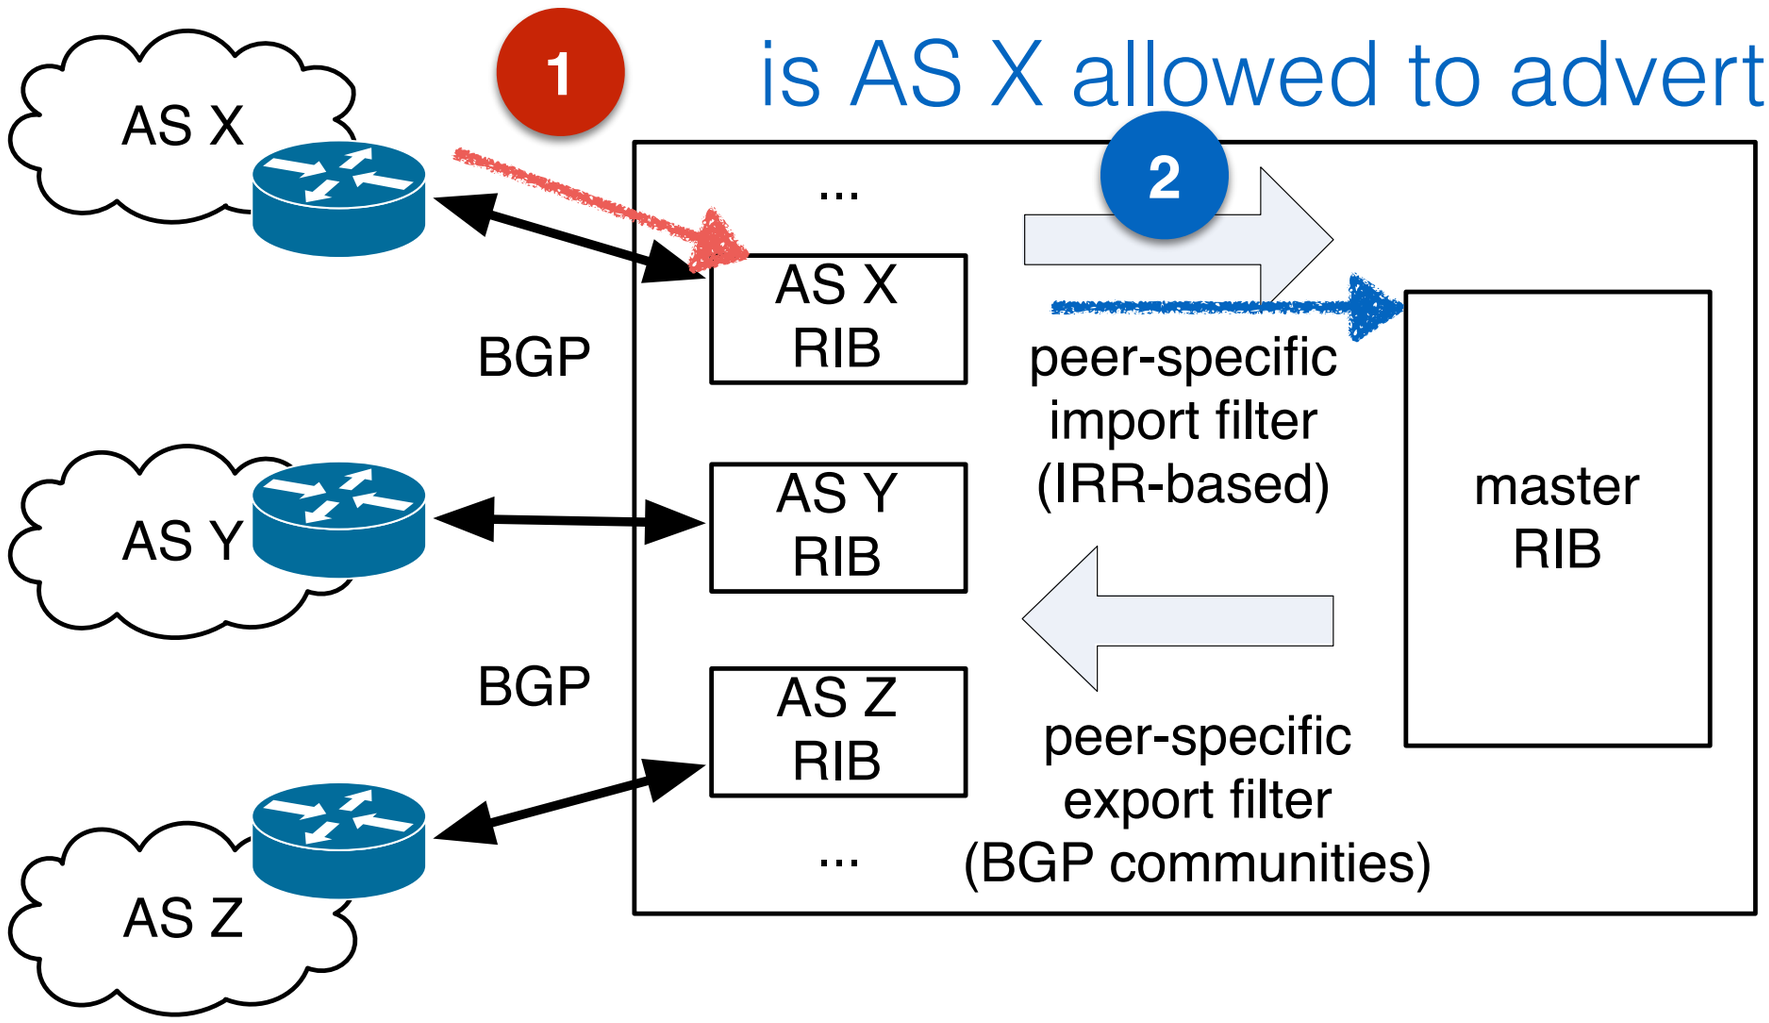

A Modern RS Architecture

AS X advertises prefix *p* (standard BGP)

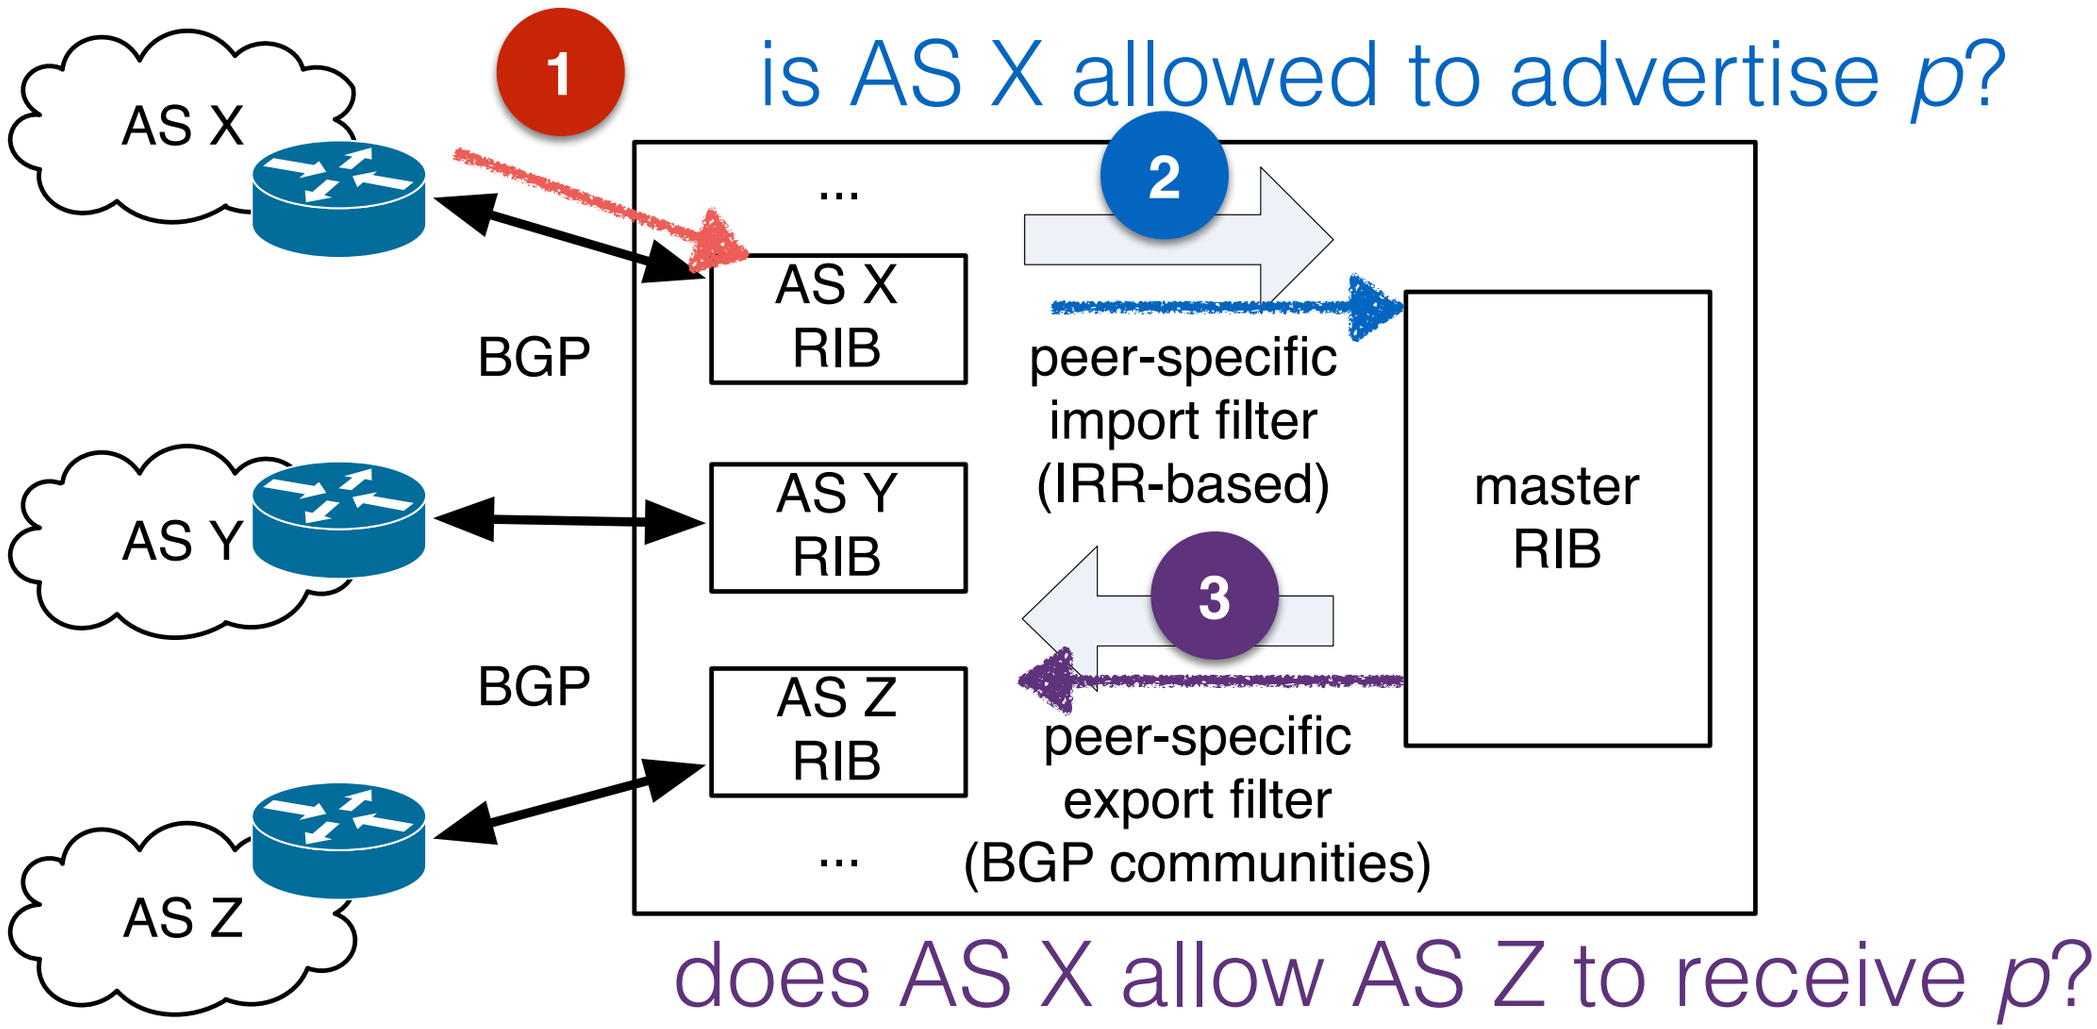

does AS X allow AS Z to receive *p*?

A Modern RS Architecture

AS X advertises prefix *p* (standard BGP)

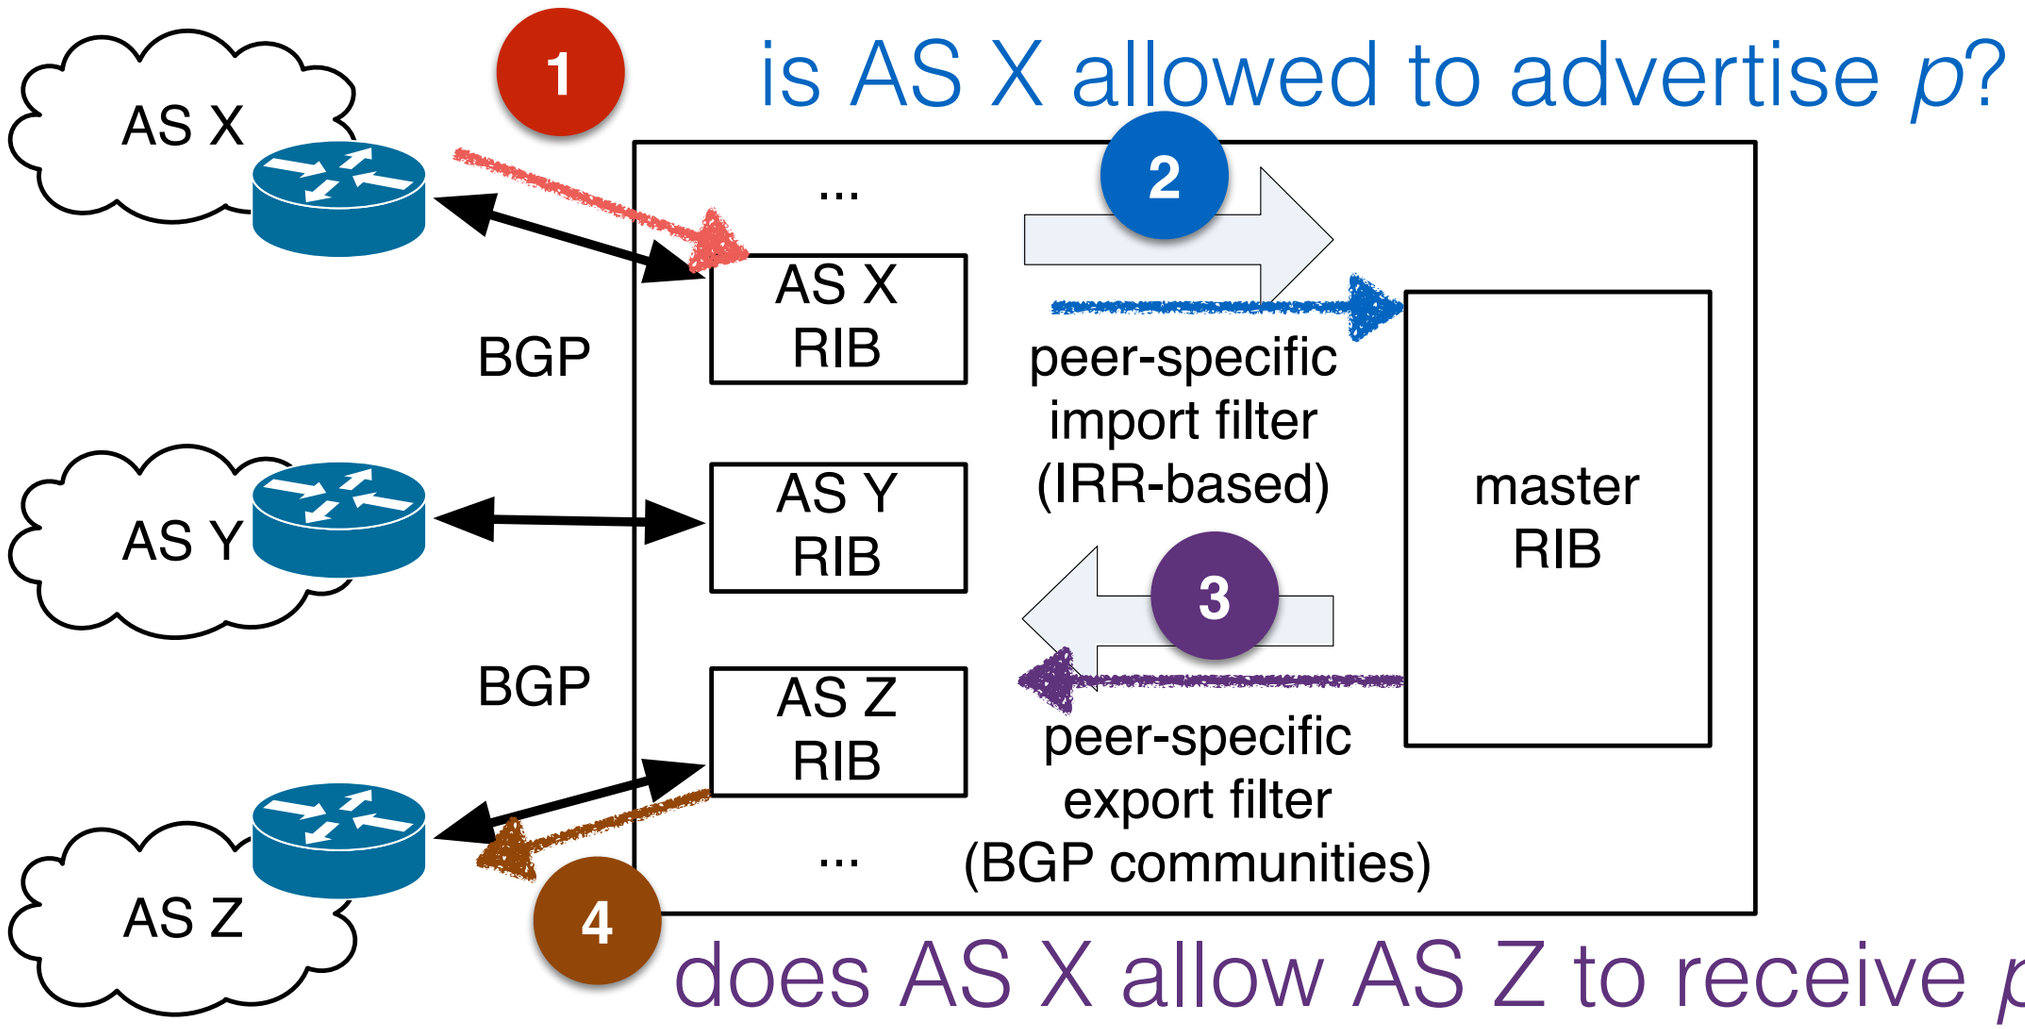

RS advertises *p* to AS Z with AS X as next hop.

5 multi-lateral peering between AS X and AS Z

Agenda

- Introduction: IXPs and Route Servers
- IXP Route Server architecture
- **Empirical study**
  - Peering offerings
  - Connectivity & traffic
  - Usage patterns
- Route Server Peering Strategies

IXPs and Datasets

	L-IXP	M-IXP
Member ASes	496	101
Peak Traffic	3 Tbps	250 Gbps
Route Server Usage	410 members (83%)	96 members (95%)
Data: Route Server	RS dumps	RS dumps
Data: Traffic	sFlow records 4 weeks 2013-09	sFlow records 4 weeks 2013-12

Most IXP members connect with the RS.

Route Server: Prefixes

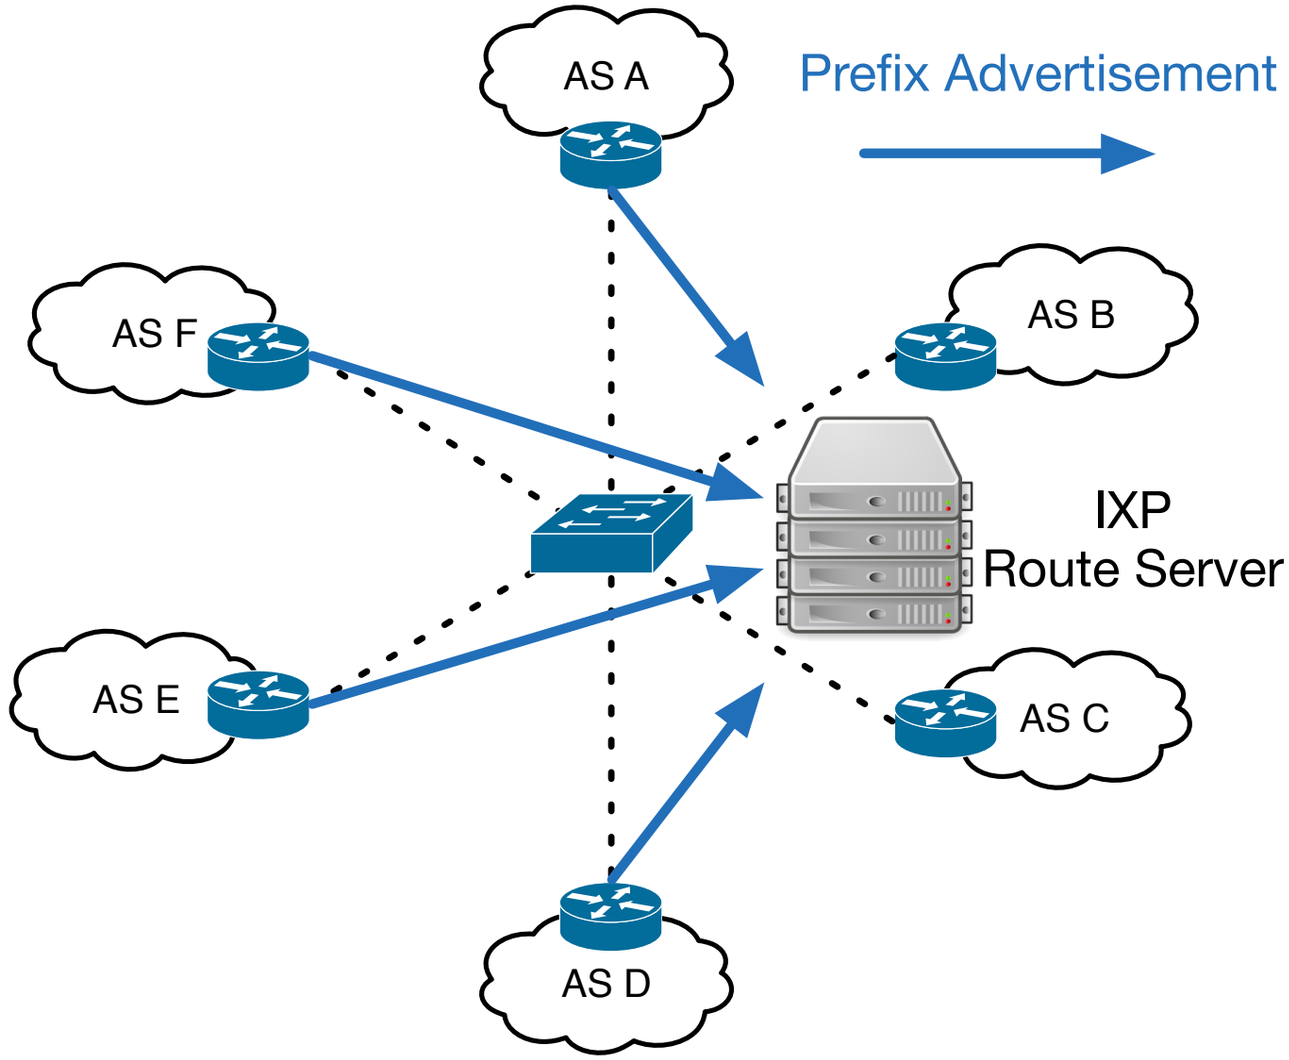

(1) Members advertise their prefixes to the RS.

Route Server: Prefixes

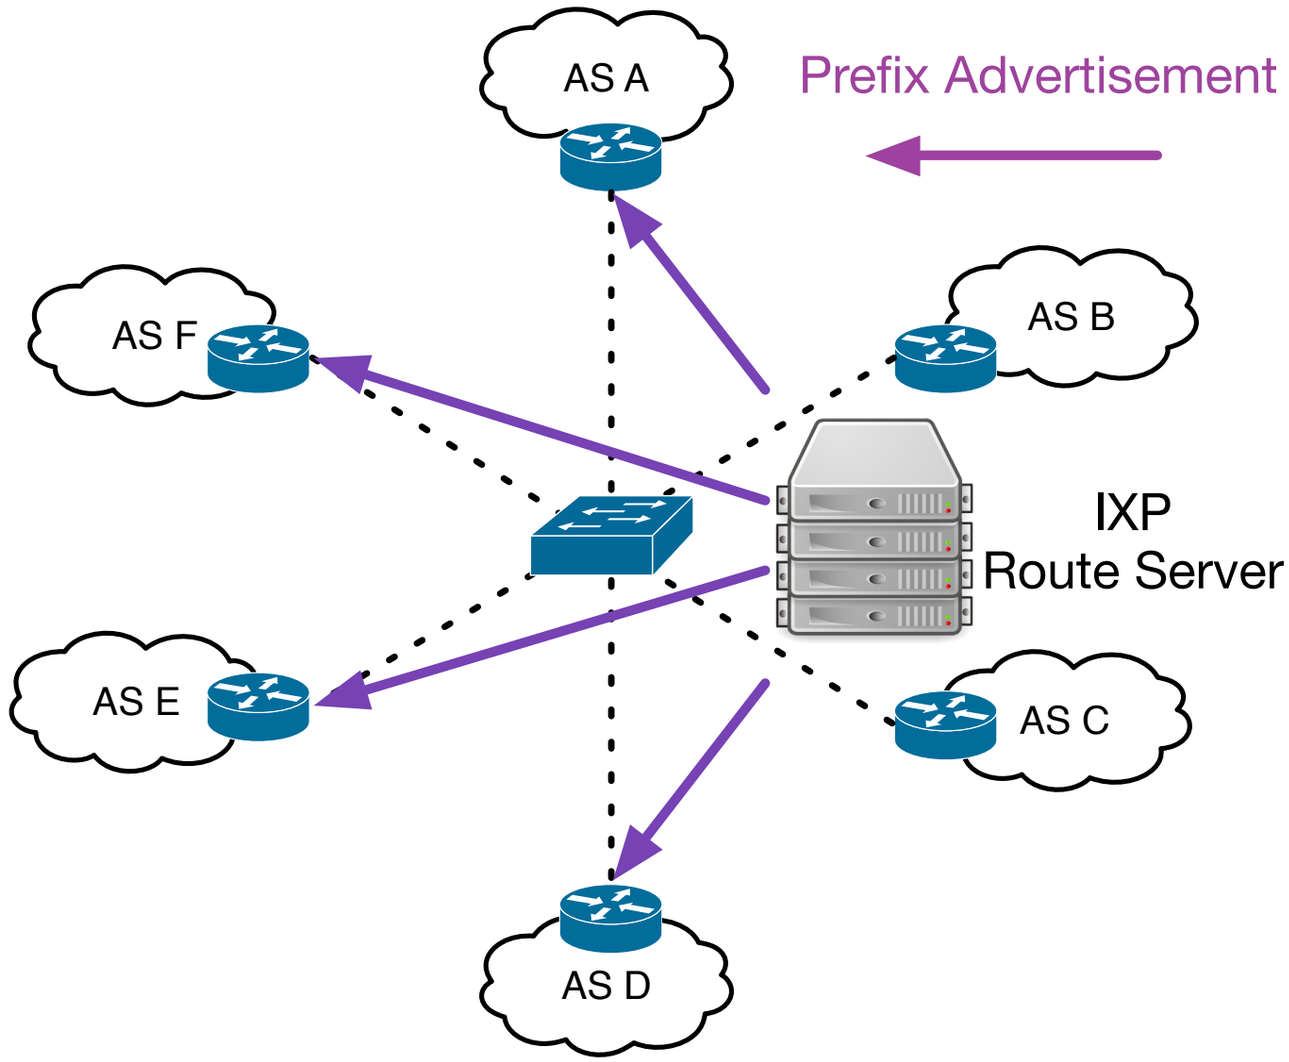

(2) RS re-advertises prefixes.

What do networks advertise? What do they receive?

Route Server: Prefixes

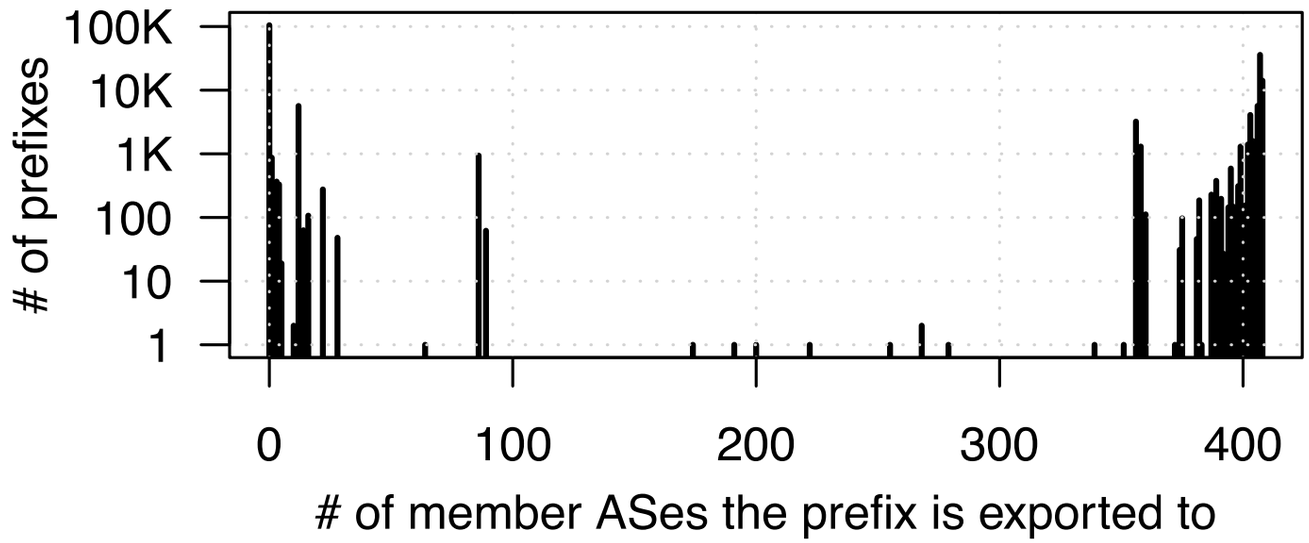

	L-IXP		**M-IXP**	
Export to % of peers	$< 10\%$	$> 90\%$	$< 10\%$	$> 90\%$
Prefixes	112.5K	68.0K	171	12.6K
/24 Equivalent	1.97M	819K	7.4K	337K
Origin ASes	13.06K	11.1K	44	3.0K

Route Server: Prefixes

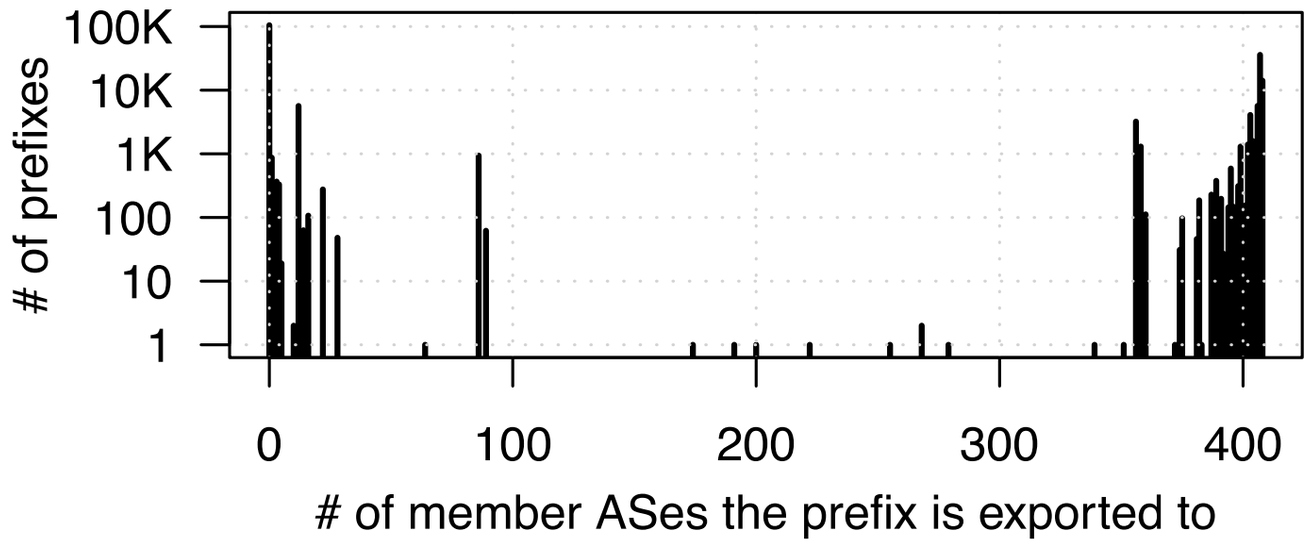

	L-IXP		M-IXP	
Export to % of peers	< 10%	> 90%	< 10%	> 90%
Prefixes	112.5K	68.0K	171	12.6K
/24 Equivalent	1.97M	819K	7.4K	337K
Origin ASes	13.06K	11.1K	44	3.0K

this is what a member instantly gets

Open access to a substantial fraction of routes.

Route Server: Traffic

10% of the L-IXPs traffic (<5% M-IXP)

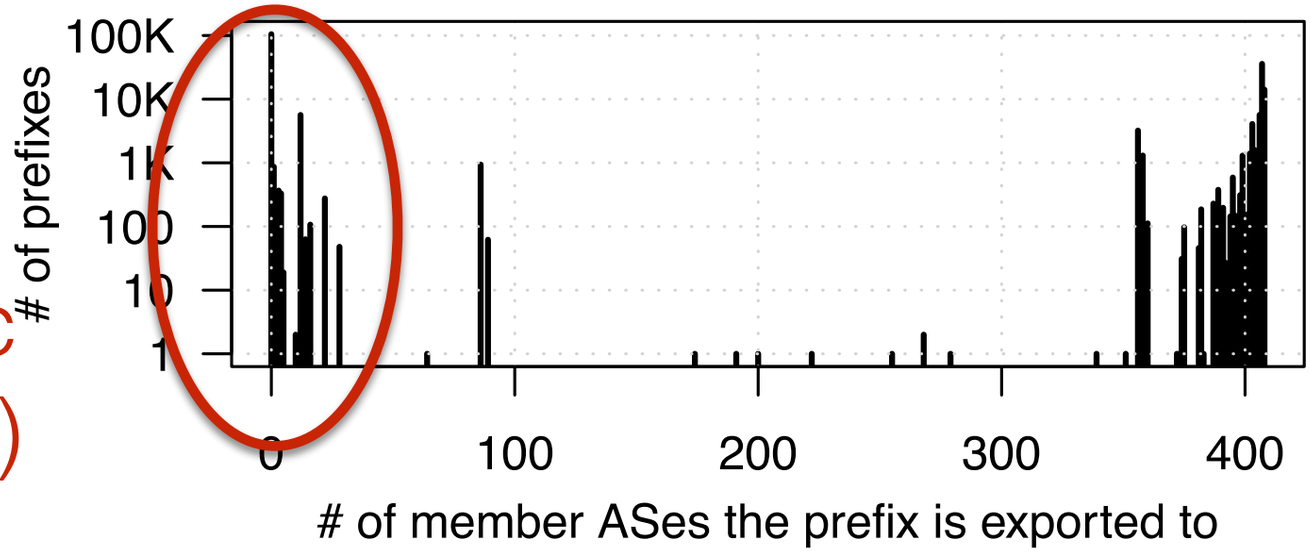

	L-IXP		M-IXP	
Export to % of peers	< 10%	> 90%	< 10%	> 90%
Prefixes	112.5K	68.0K	171	12.6K
/24 Equivalent	1.97M	819K	7.4K	337K
Origin ASes	13.06K	11.1K	44	3.0K

this is what a member instantly gets

Route Server: Traffic

10% of the L-IXPs traffic (<5% M-IXP)

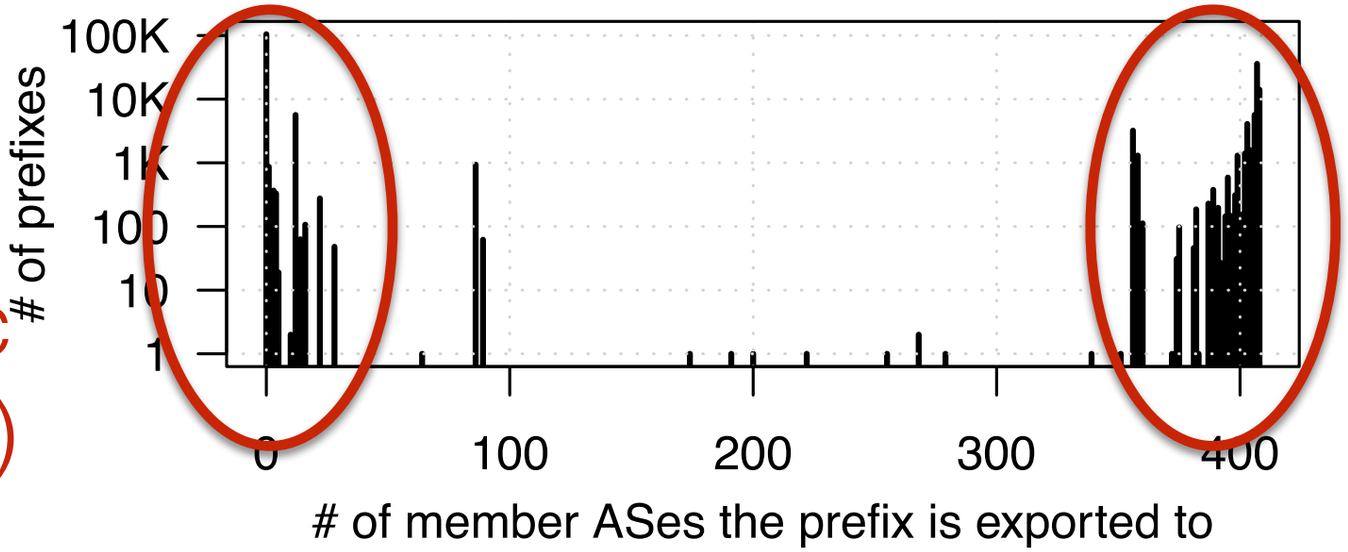

70% of the L-IXP's traffic (>90% M-IXP)

	L-IXP		M-IXP	
Export to % of peers	< 10%	> 90%	< 10%	> 90%
Prefixes	112.5K	68.0K	171	12.6K
/24 Equivalent	1.97M	819K	7.4K	337K
Origin ASes	13.06K	11.1K	44	3.0K

this is what a member instantly gets

Openly peered prefixes receive largest share of traffic.

Detecting Peerings

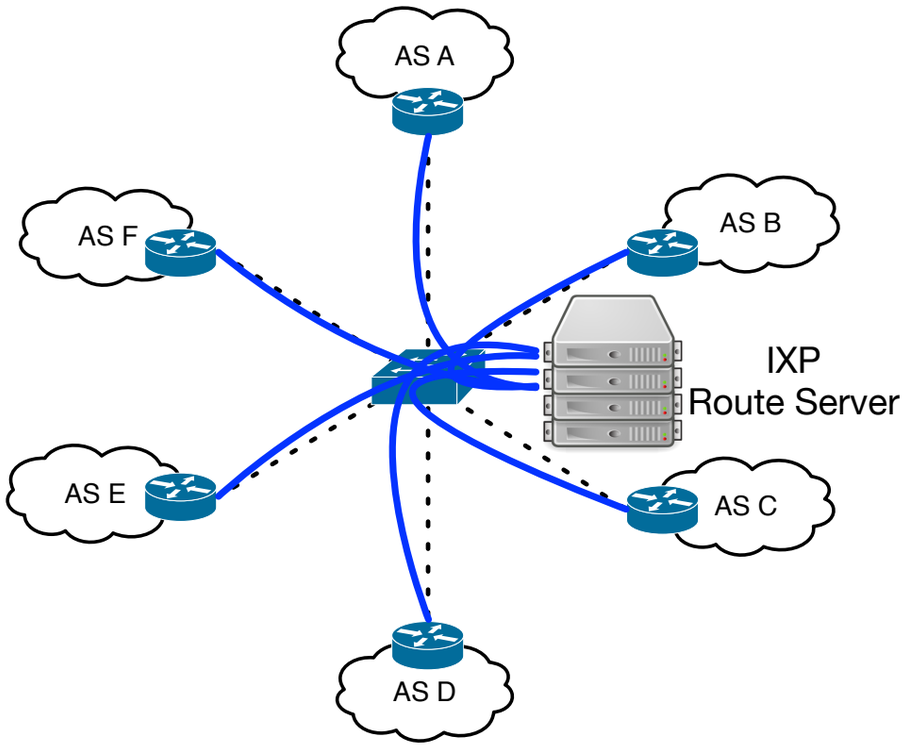

Multi-lateral

Access to RS RIBs
(* publicly available
using looking glasses)

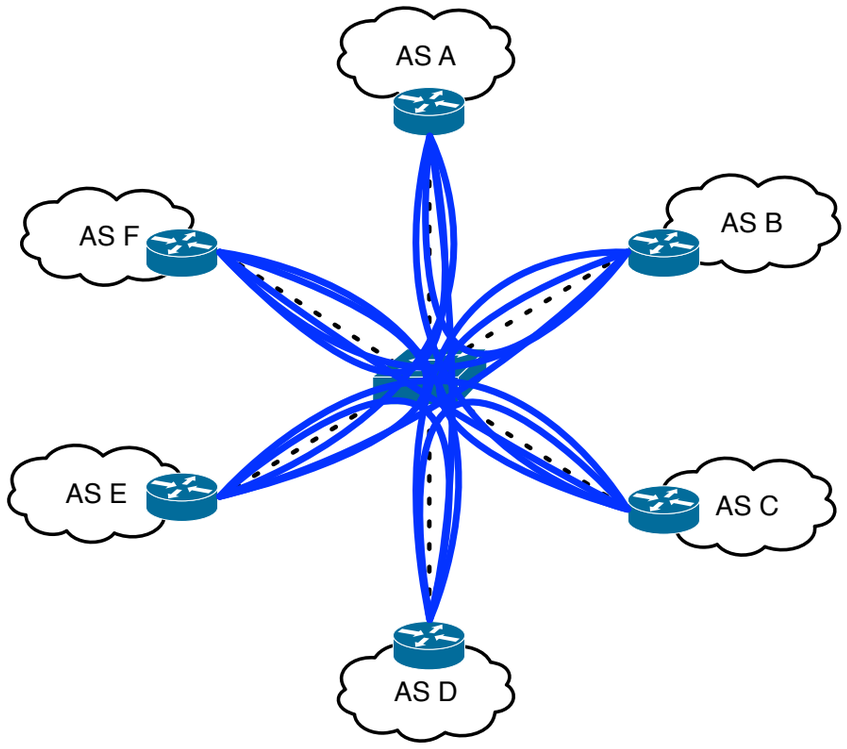

Bi-lateral

Sampling BGP packets
between border routers.

Peerings: ML vs. BL

	L-IXP	M-IXP
Bi-Lateral	20K	450
Multi-Lateral	80K	3.7K
Total	85K	3.8K

Table: Peering Links.

Ratio ML-to-BL peerings:

4:1 (L-IXP)
8:1 (M-IXP)

>95% of new peerings in last 2 years are ML!

IXP connectivity is clearly dominated by multi-lateral peering.

Traffic: ML vs. BL

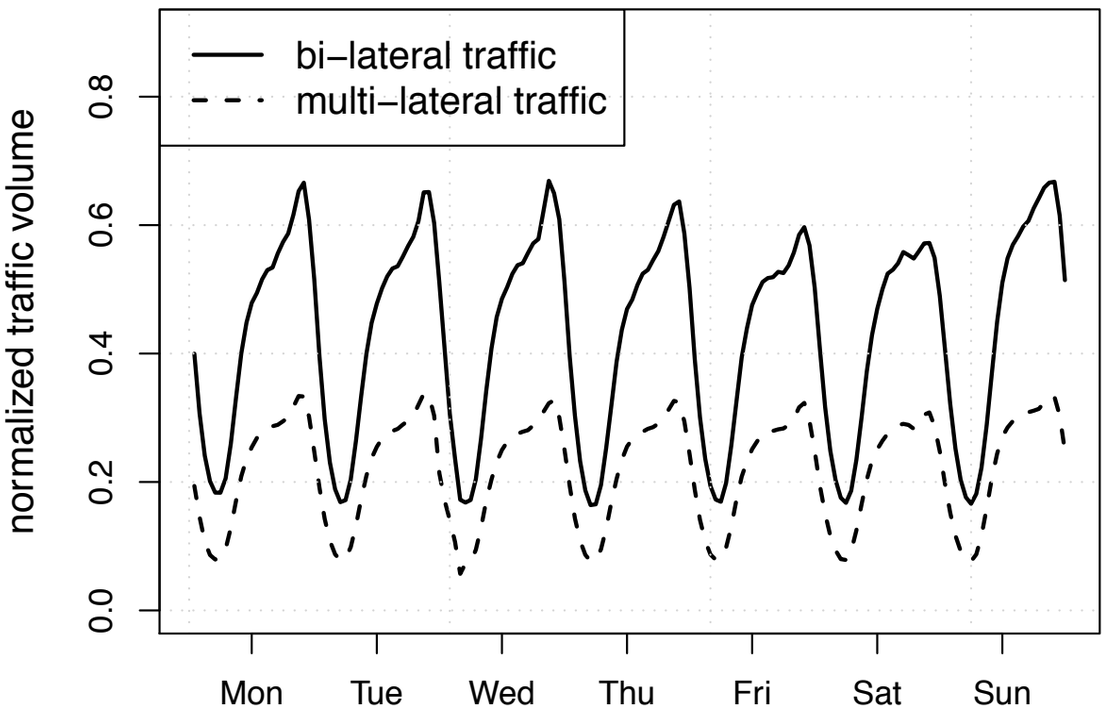

Ratio ML-to-BL traffic:

1:2 (L-IXP)
1:1 (M-IXP)

- BL more likely to carry traffic and carry more traffic
- Some heavy-hitters are ML!

IXP traffic is dominated by fewer bi-lateral peerings.
But RS-prefixes receive most traffic. How come?

Understanding RS Usage

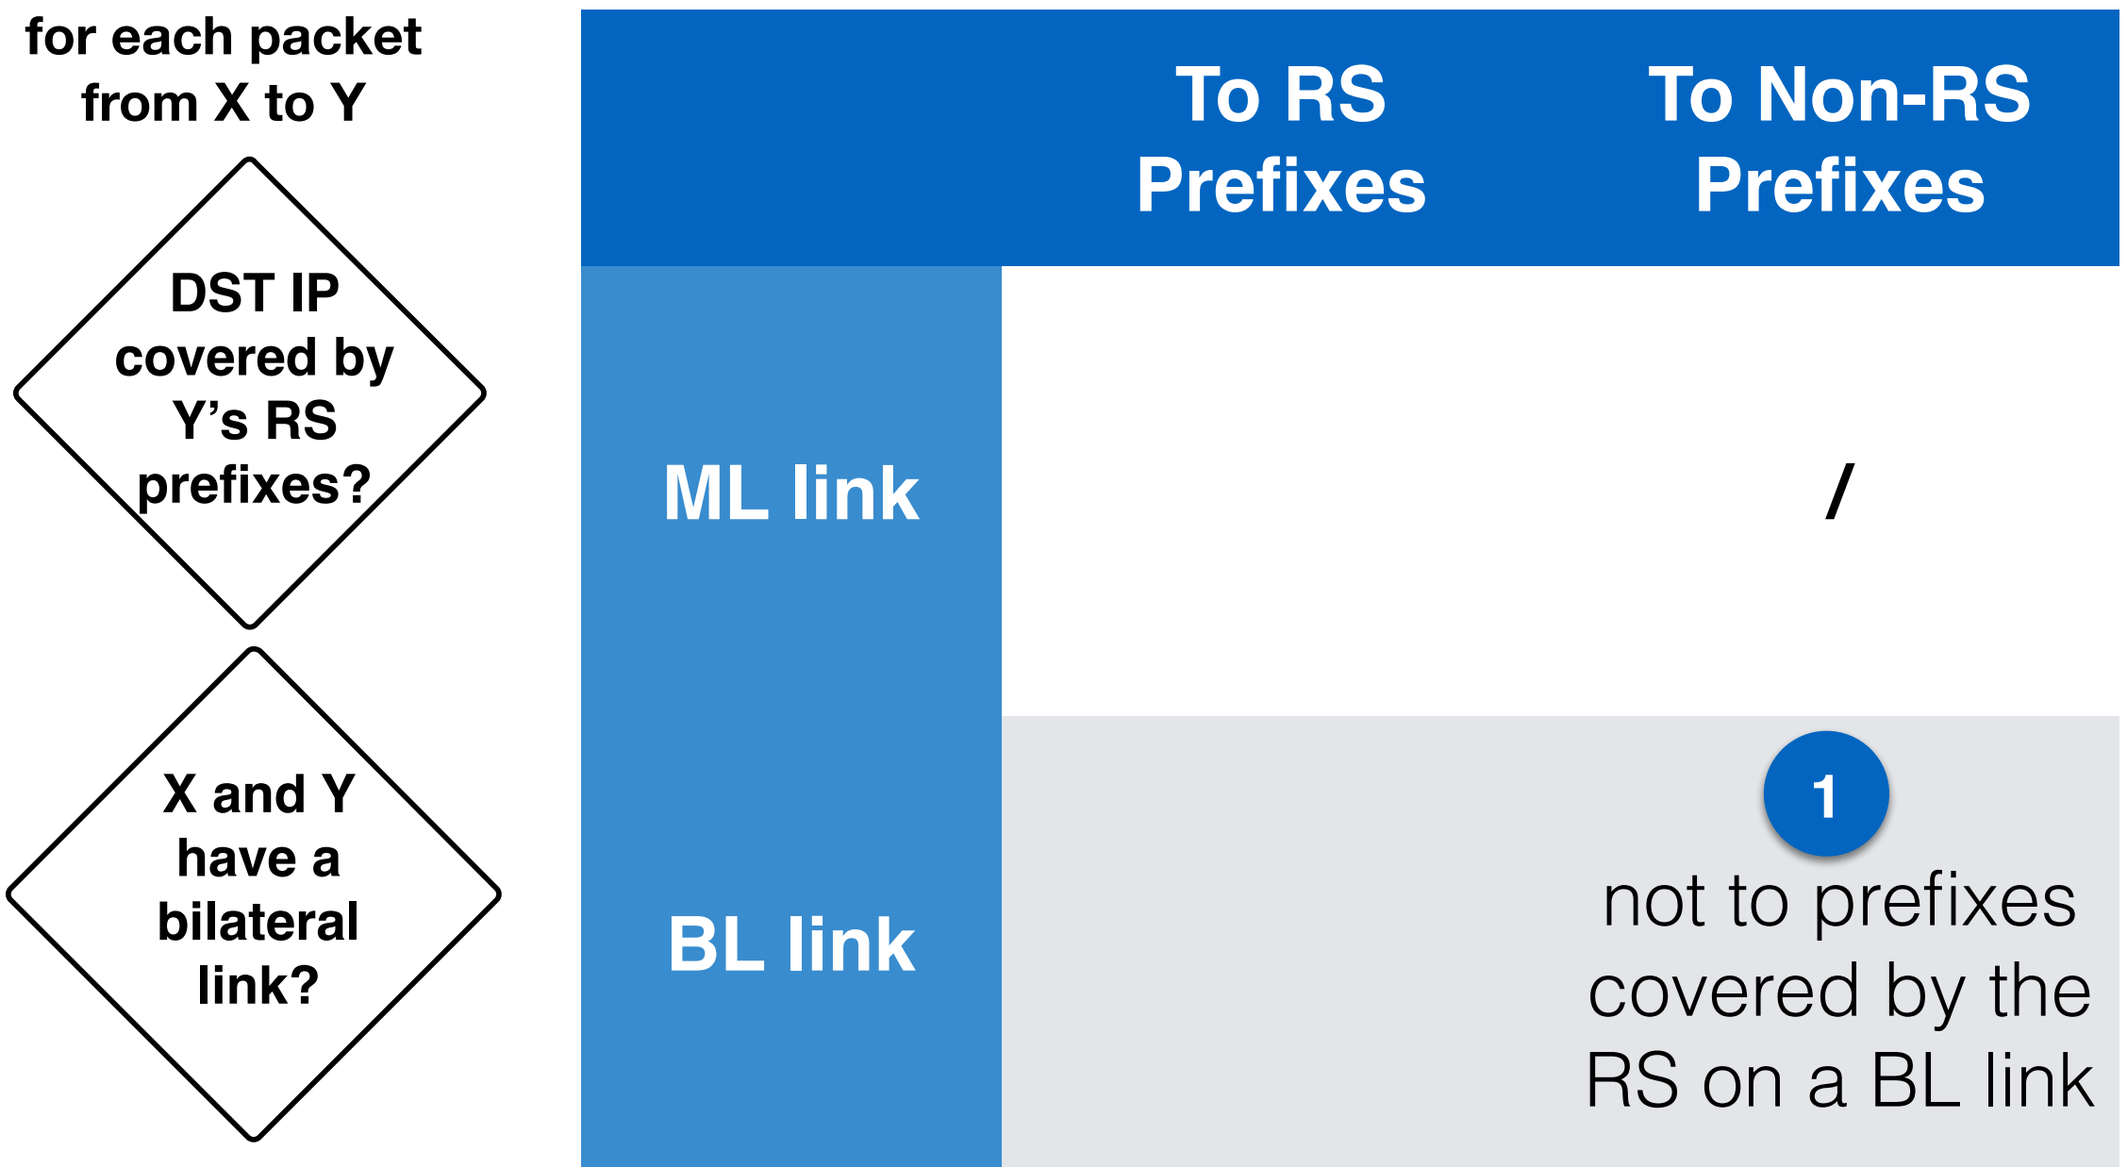

1 Vanilla bi-lateral peering

Understanding RS Usage

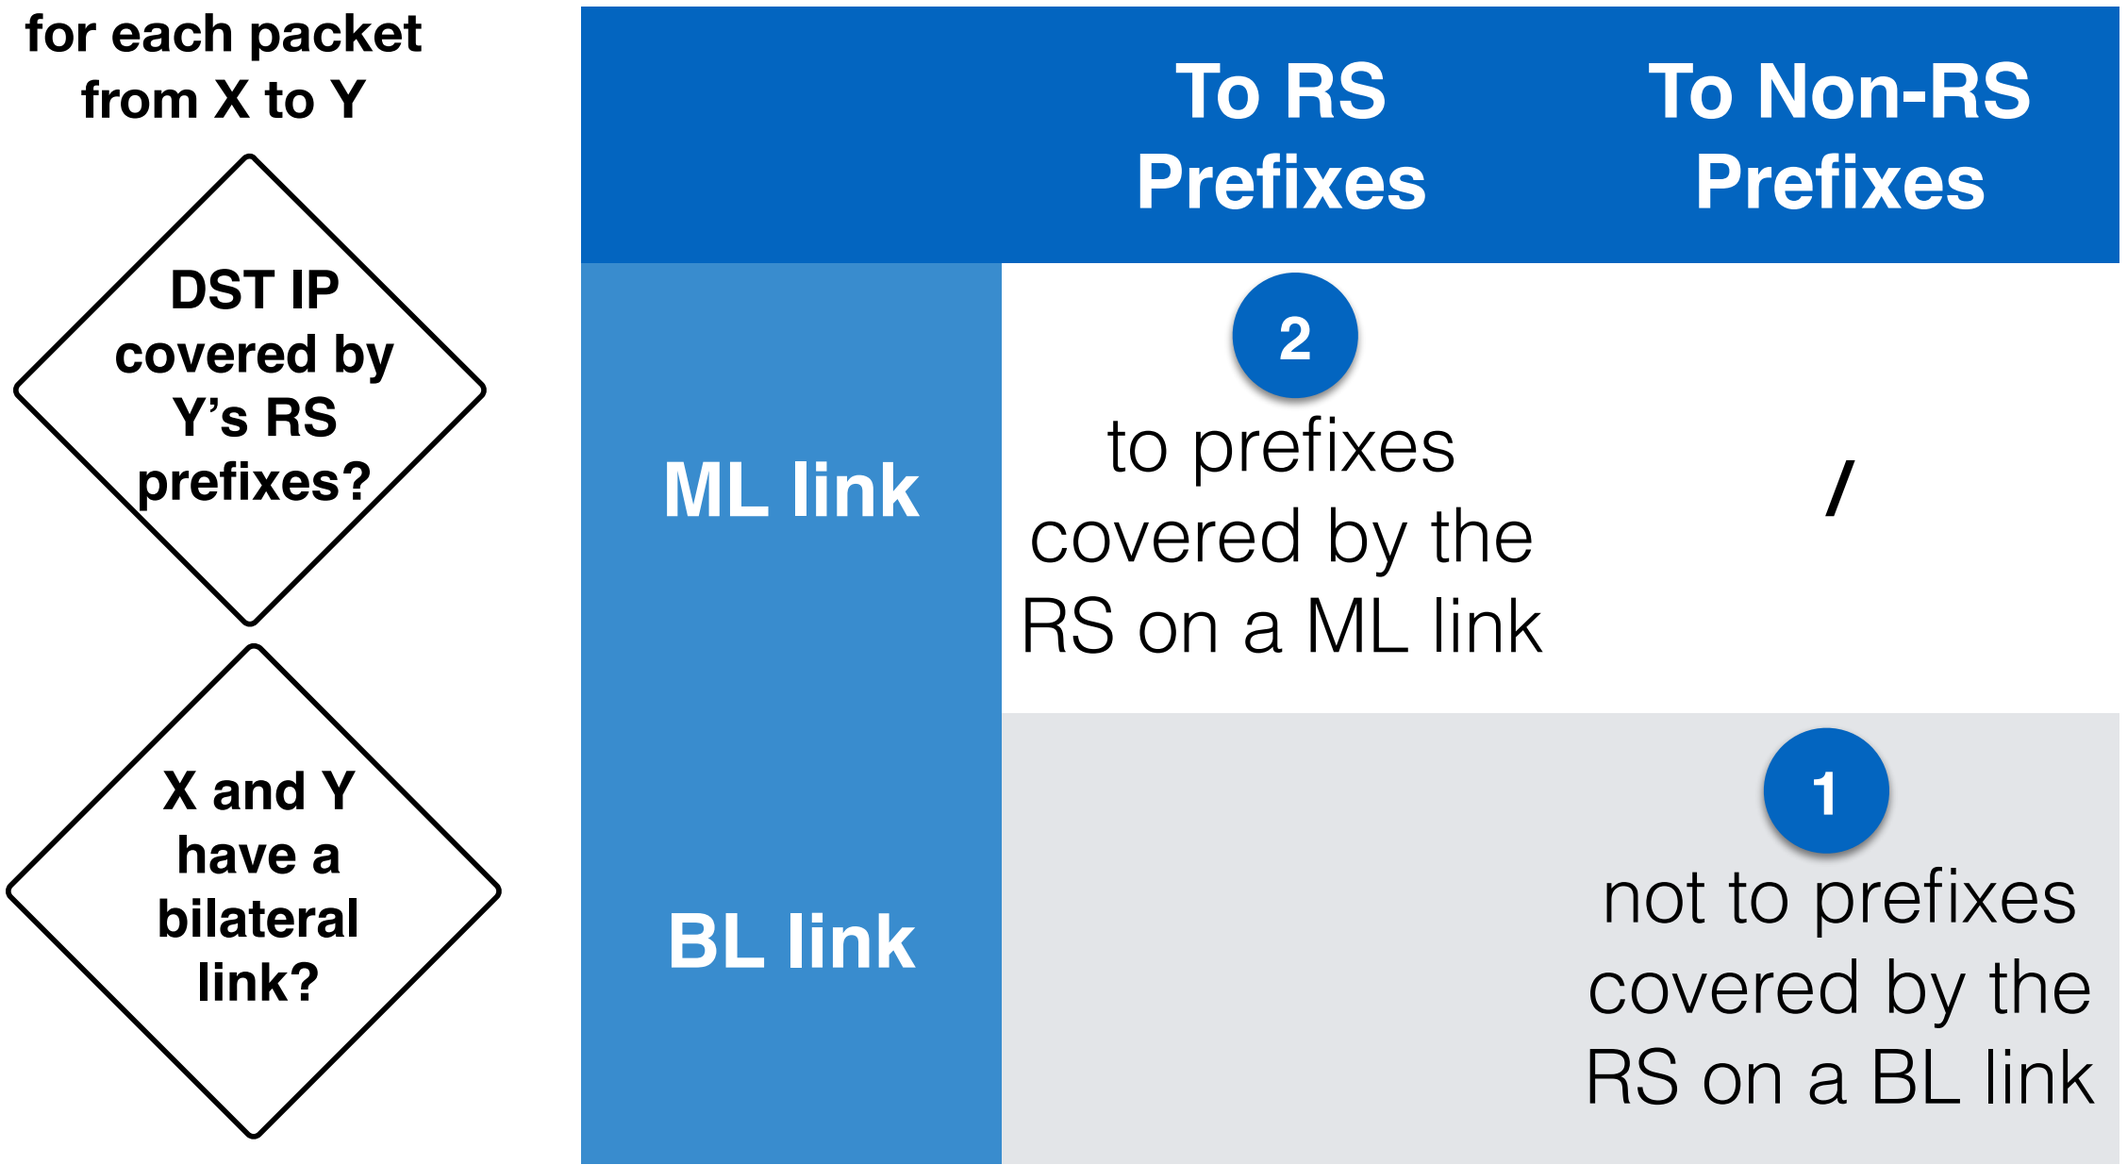

2 Vanilla multi-lateral peering

Understanding RS Usage

for each packet from X to Y

DST IP covered by Y's RS prefixes?

X and Y have a bilateral link?

	To RS Prefixes	To Non-RS Prefixes
ML link	2 to prefixes covered by the RS on a ML link	/
BL link	3 to prefixes covered by the RS on a BL link	1 not to prefixes covered by the RS on a BL link

3 Possible multi-lateral peering, yet bi-lateral links

Understanding RS Usage

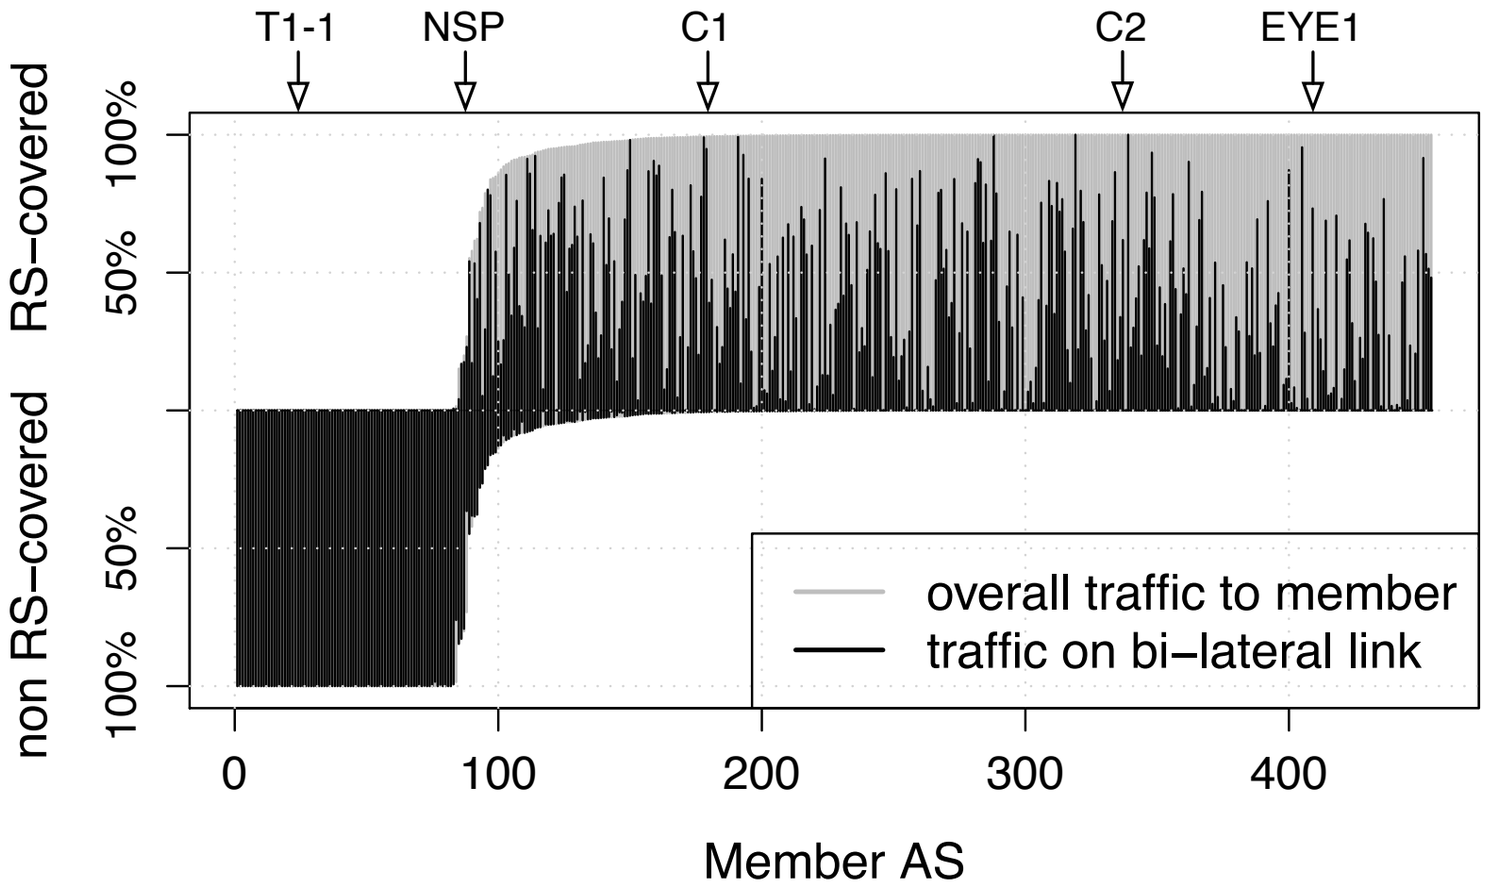

Understanding RS Usage

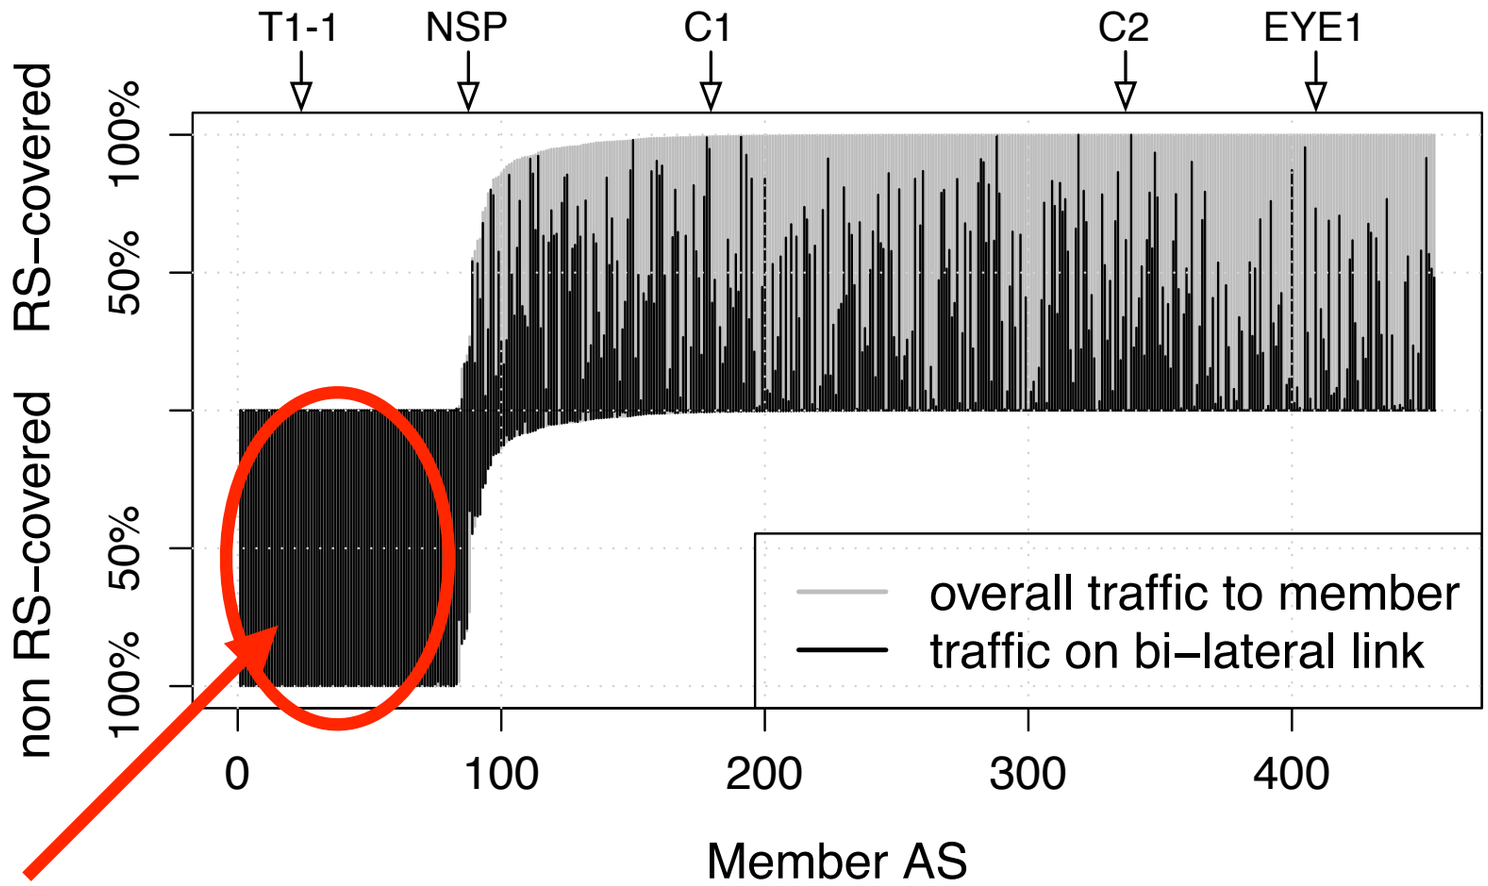

No RS usage.
only bilateral peering

Understanding RS Usage

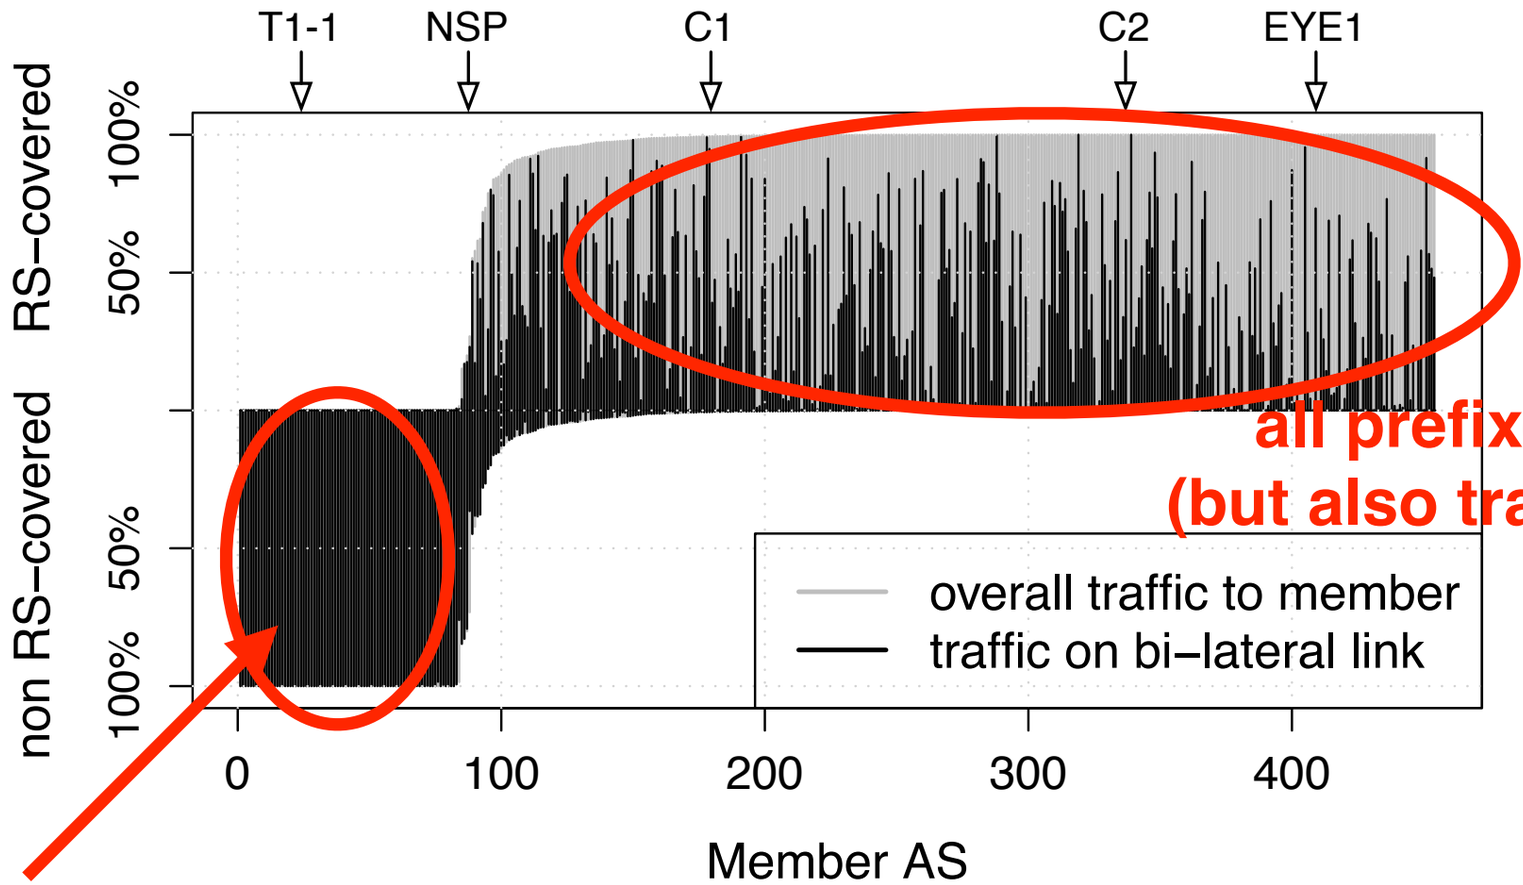

all prefixes advertised
(but also traffic on BL links)

No RS usage.
only bilateral peering

Most members advertise all prefixes to the RS.

Understanding RS Usage

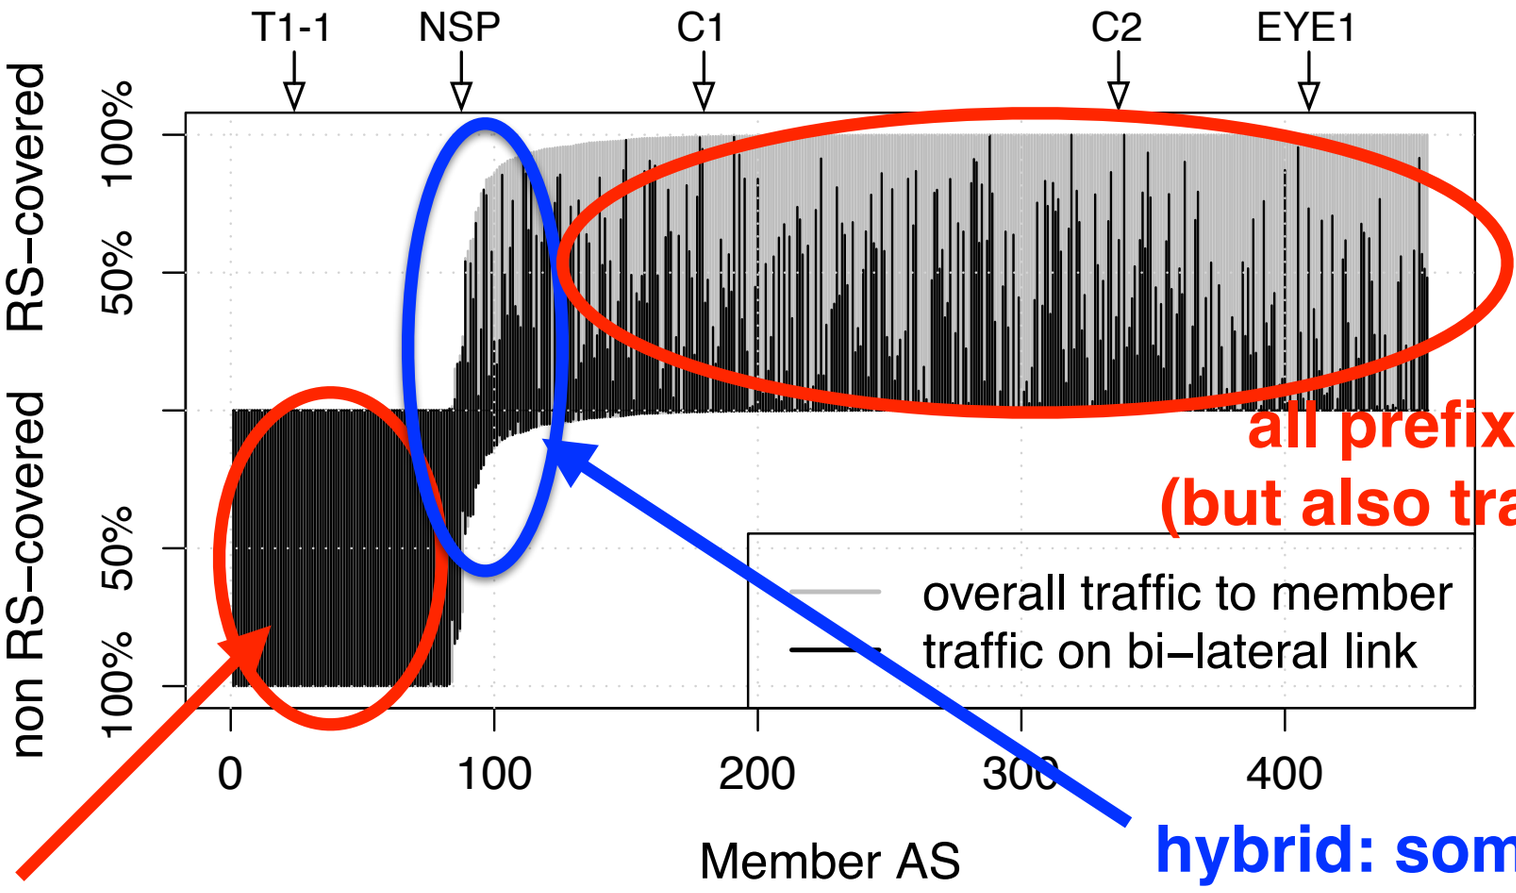

all prefixes advertised
(but also traffic on BL links)

No RS usage.
only bilateral peering

hybrid: some prefixes
advertised to the RS
but other routes via BL peerings

Agenda

- Introduction: IXPs and Route Servers
- IXP Route Server architecture
- Empirical study
  - Peering offerings
  - Connectivity & traffic
  - Usage patterns
- **Route Server Peering Strategies**

Case Studies: Big Players

C1, C2: Major Content Providers

Open peering via the RS at both IXPs
C1's traffic mainly on BL peerings, C2 promotes ML peering

EYE1, EYE2: National Eyeball Providers

Open peering via the RS at both IXPs,
yet mainly bi-lateral peerings

OSN1, OSN2: Two Popular Online Social Networks

OSN1 peers only bi-laterally, OSN2 only using the RS

T1-1,T1-2: Large Transit Providers

T1-1 doesn't peer with the RS, T1-2 does, but doesn't export prefixes

RSes are used by (almost) all types of networks.

Case Studies: Hybrid Peering

NSP: A Large Transit Provider

- Open peering with everyone at the IXP for some prefixes
- Large superset advertised via BL peerings (likely customers)
  - Open peering for some prefixes
  - Restricted peering for others

significant traffic contribution

CDN: Mid-sized CDN Provider

- Some prefixes openly advertised via RS
- Different prefixes on BL sessions with path prepending
  - Complex traffic engineering of CDNs

significant traffic contribution

Networks already implement advanced RS peering strategies.

Peering: RS or Non-RS

- Peering policies of content providers (e.g., Google)
  - ML peering with small networks
  - Subsequent BL peering if traffic significant
- Reasons for Non-RS peering:
  - Session monitoring
  - Traffic engineering
    - Inbound: Prefix deaggregation, MEDs, etc.
    - Outbound: Best path selection by RS

Bi-lateral still preferred for traffic-intensive peerings.

RSes, Peering, and Innovation

- Innovation in inter-domain routing
- Make peering easy and scalable
- Heavily used by all different types of networks

- Central components with large impact
- Make deployment of new technologies possible
- Better traffic engineering capabilities needed
  - e.g., by leveraging SDN (SDX) [Gupta et al., SIGCOMM ‘14]

Route Servers key components in the peering ecosystem.

Conclusion

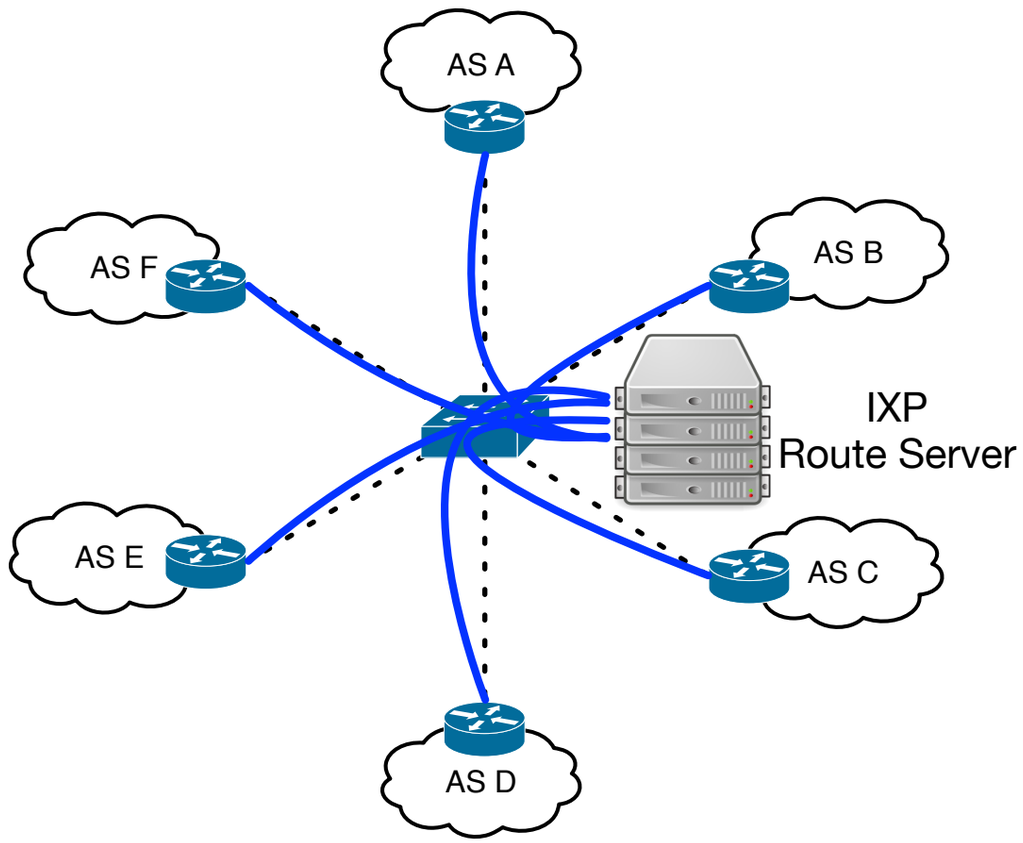

Route Servers

- Make peering easy
- Heavily used
- Great places for innovation

Contact: Philipp Richter (prichter@inet.tu-berlin.de)

Paper: net.t-labs.tu-berlin.de/~prichter/imc238-richterA.pdf