



RIPE NCC

RIPE NETWORK COORDINATION CENTRE

BGP Security Webinars

BGP Filtering

July 2025

RIPE NCC Learning & Development



**This session is
being recorded**

Agenda



- Introduction to BGP Route Filtering
- BGP Filters (BGP Policies)
- Prefix Filtering Recommendations
- **Demo:** Filtering too specific prefixes



Introduction to BGP Filtering

Section 1



What is BGP route filtering?

- The most basic **protection** mechanism against malicious or accidental BGP incidents
- Technique used to control prefixes exchanged between BGP peers
 - Which prefixes will you **accept** into your network?
 - Which prefixes will you **advertise** to your peers?



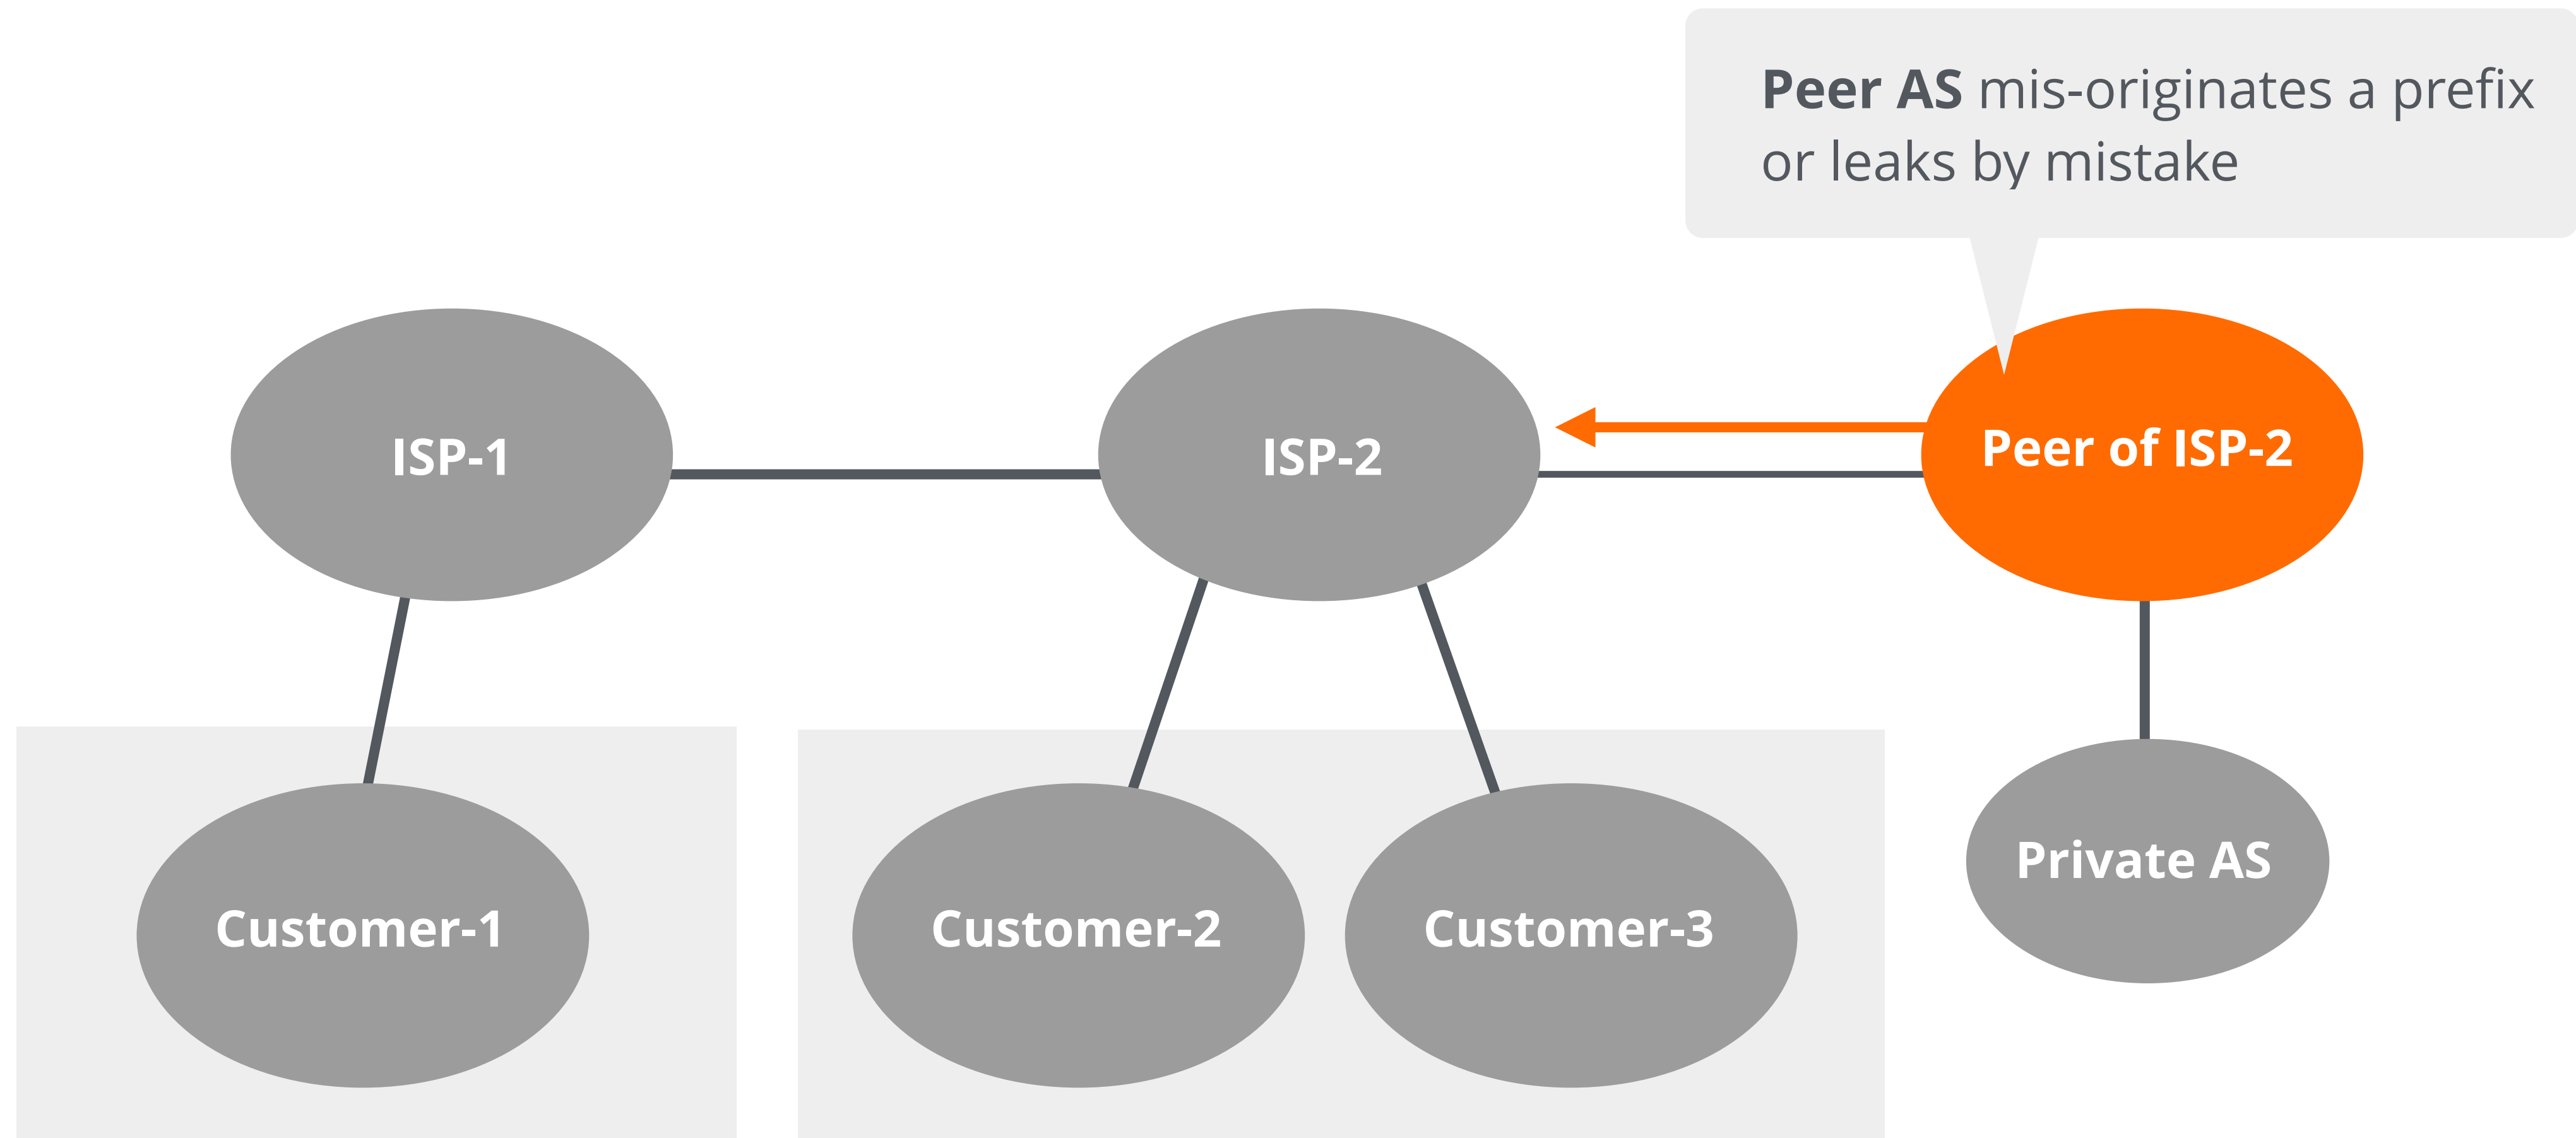


Why filter BGP prefixes?

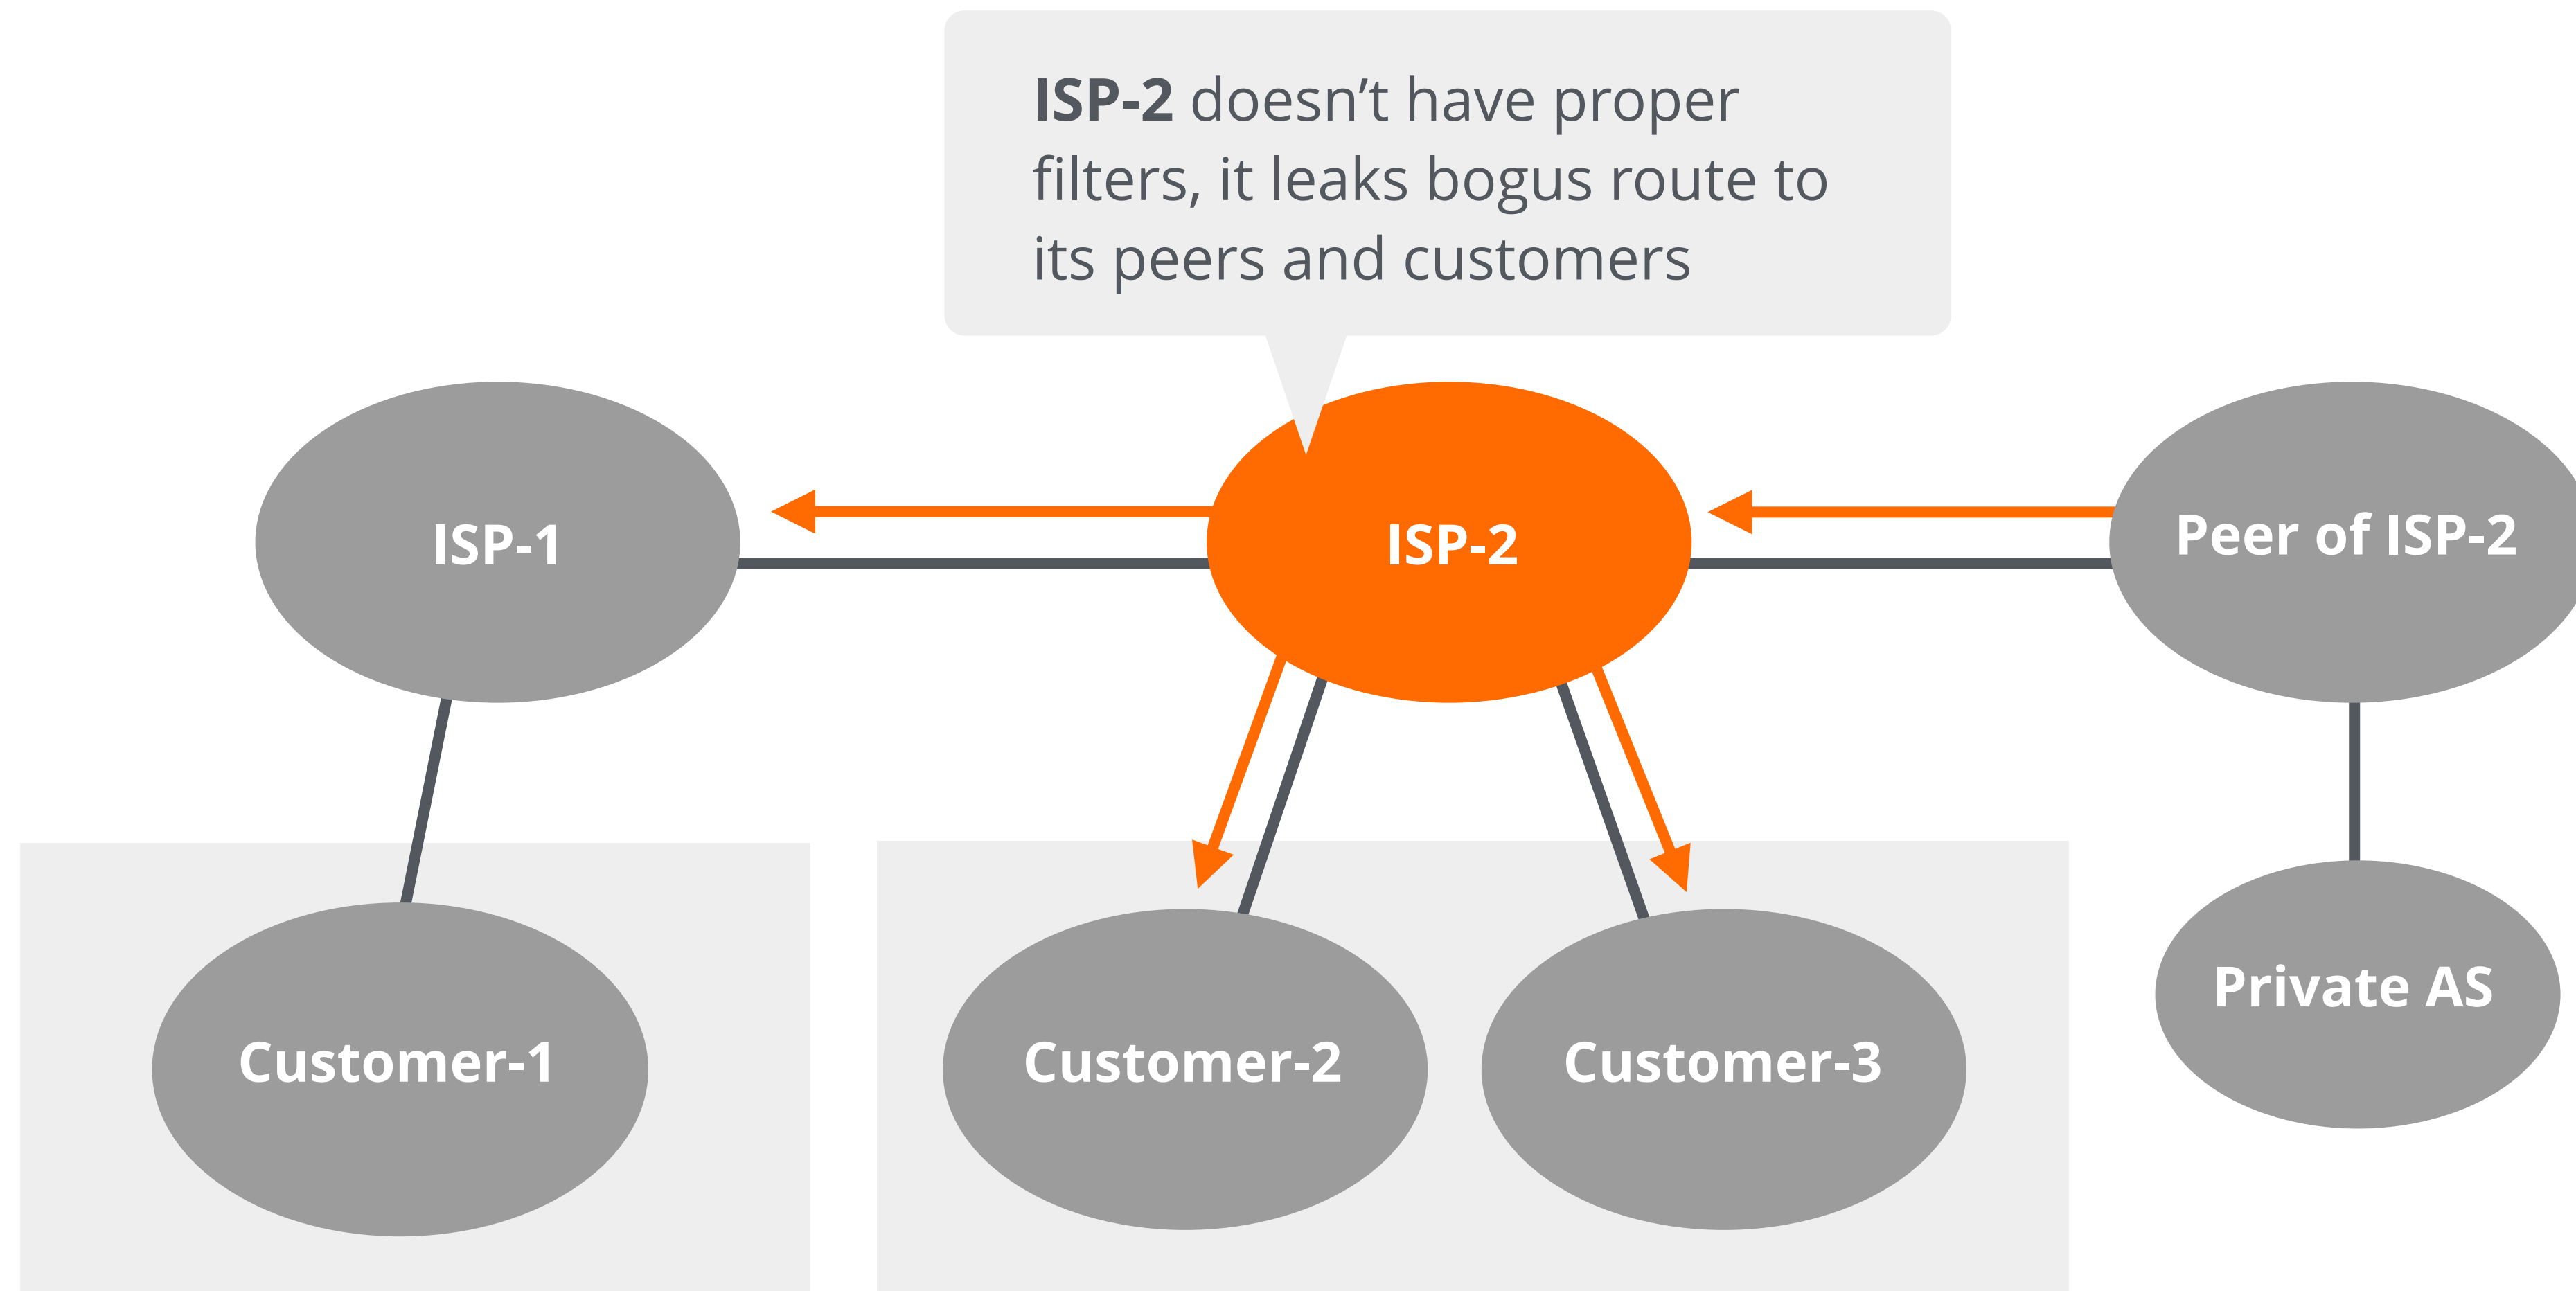
- **Essential for routing security!**
 - Your first line of defence!
- Because you can only control what you're announcing
- Increases the security and stability of Internet routing
 - Prevents **route leaks**
 - Mitigates the impact of **BGP hijacks**



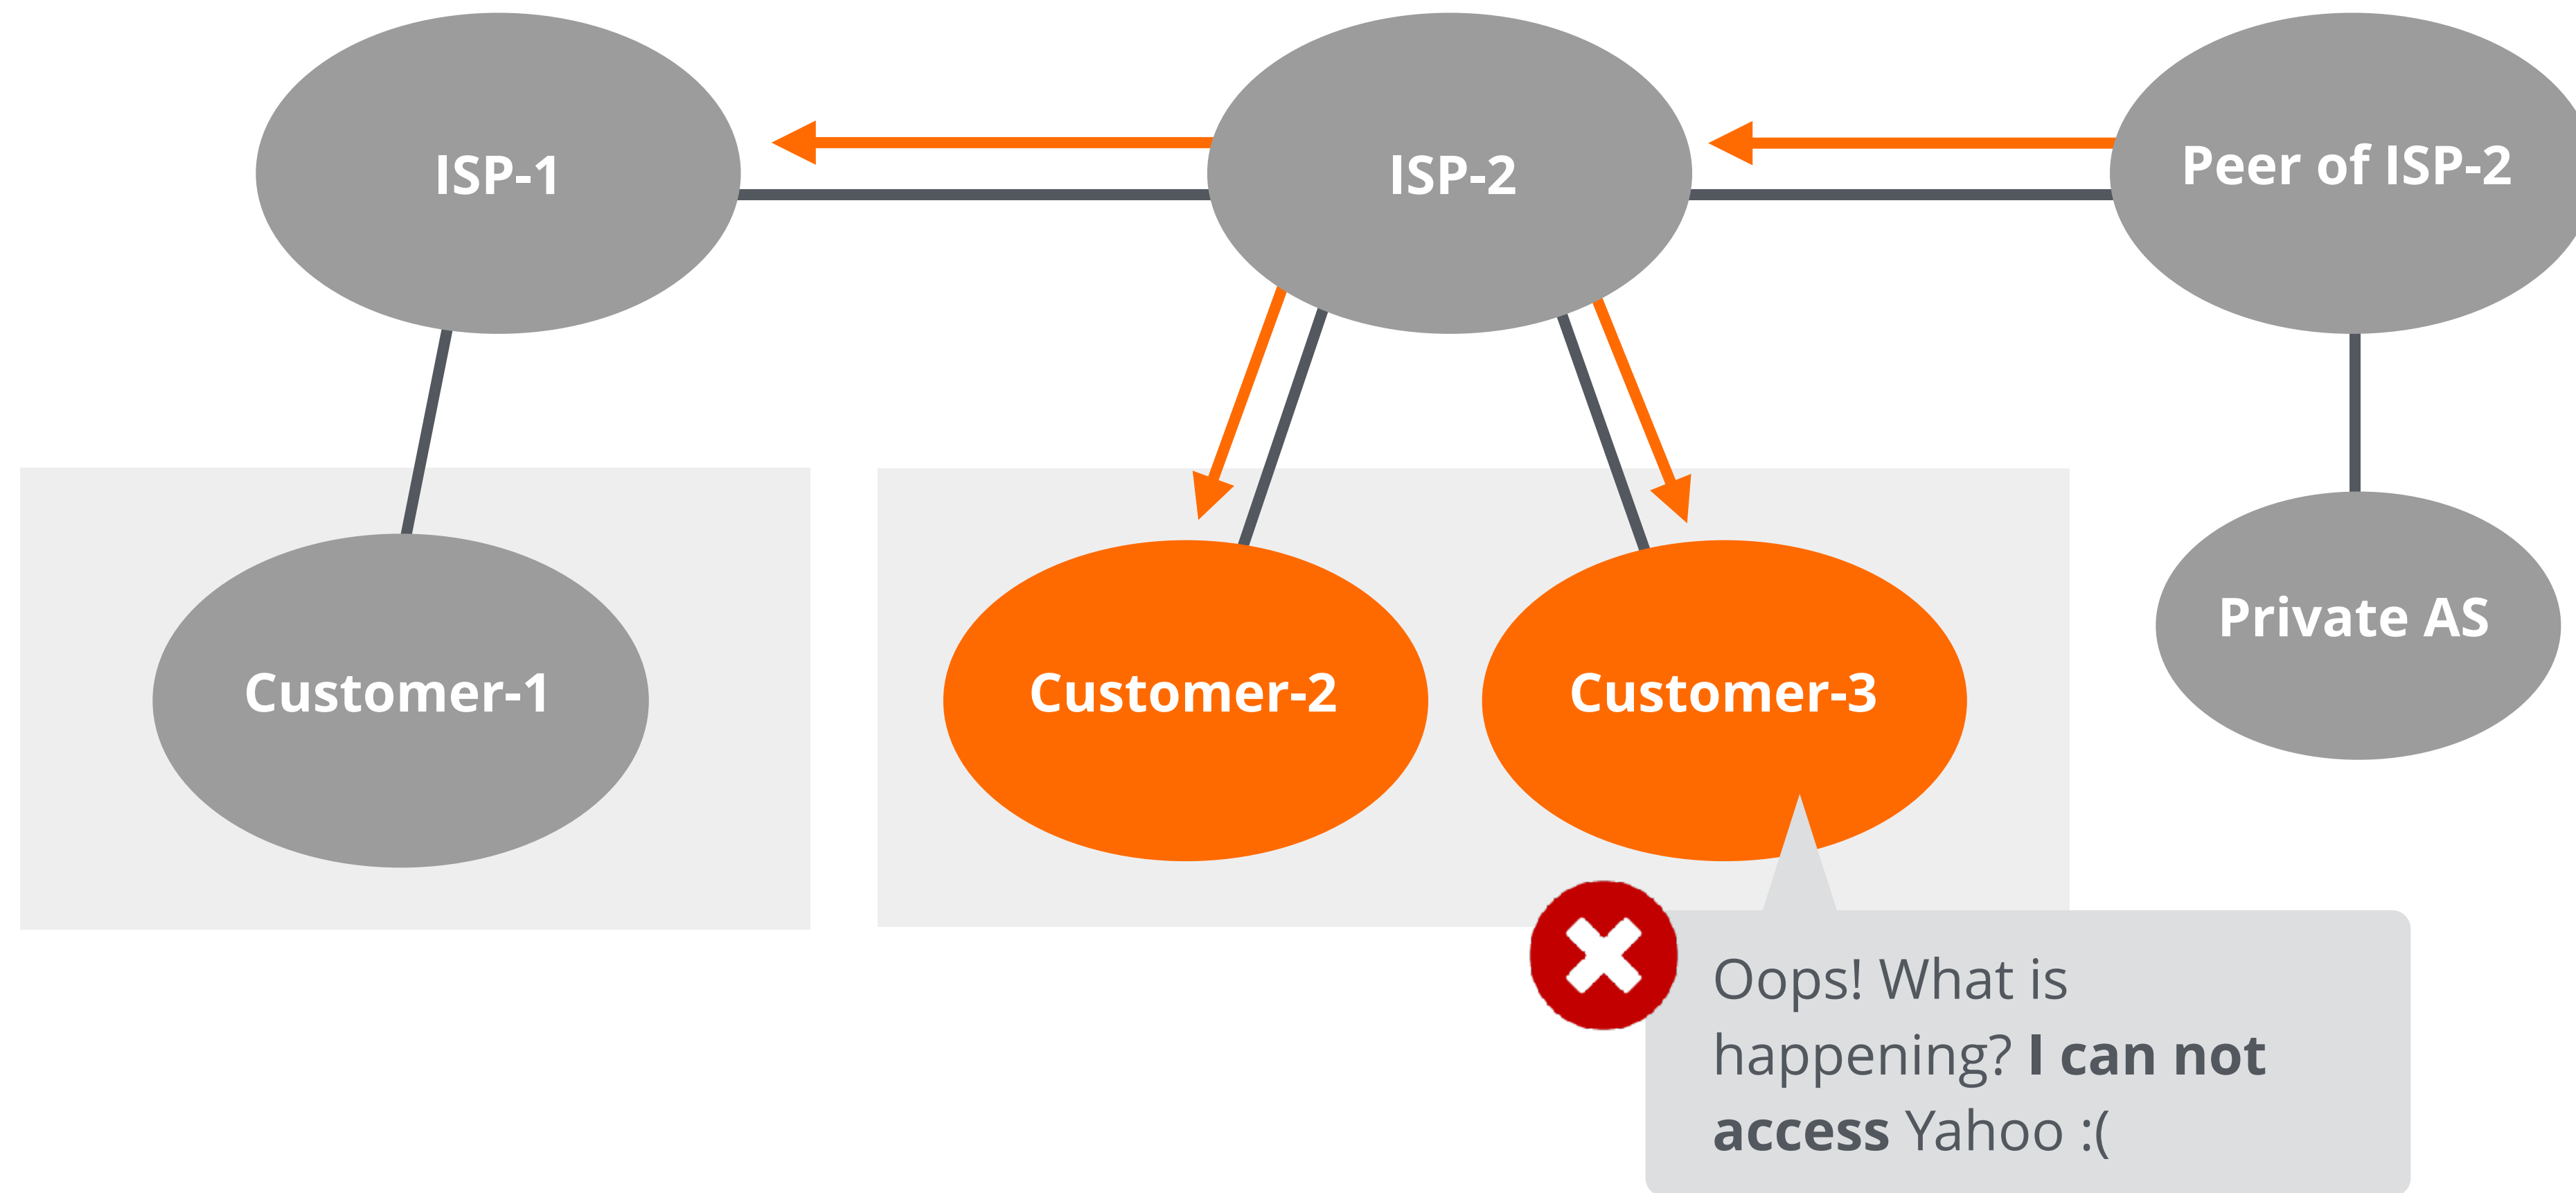
Why filter BGP prefixes?



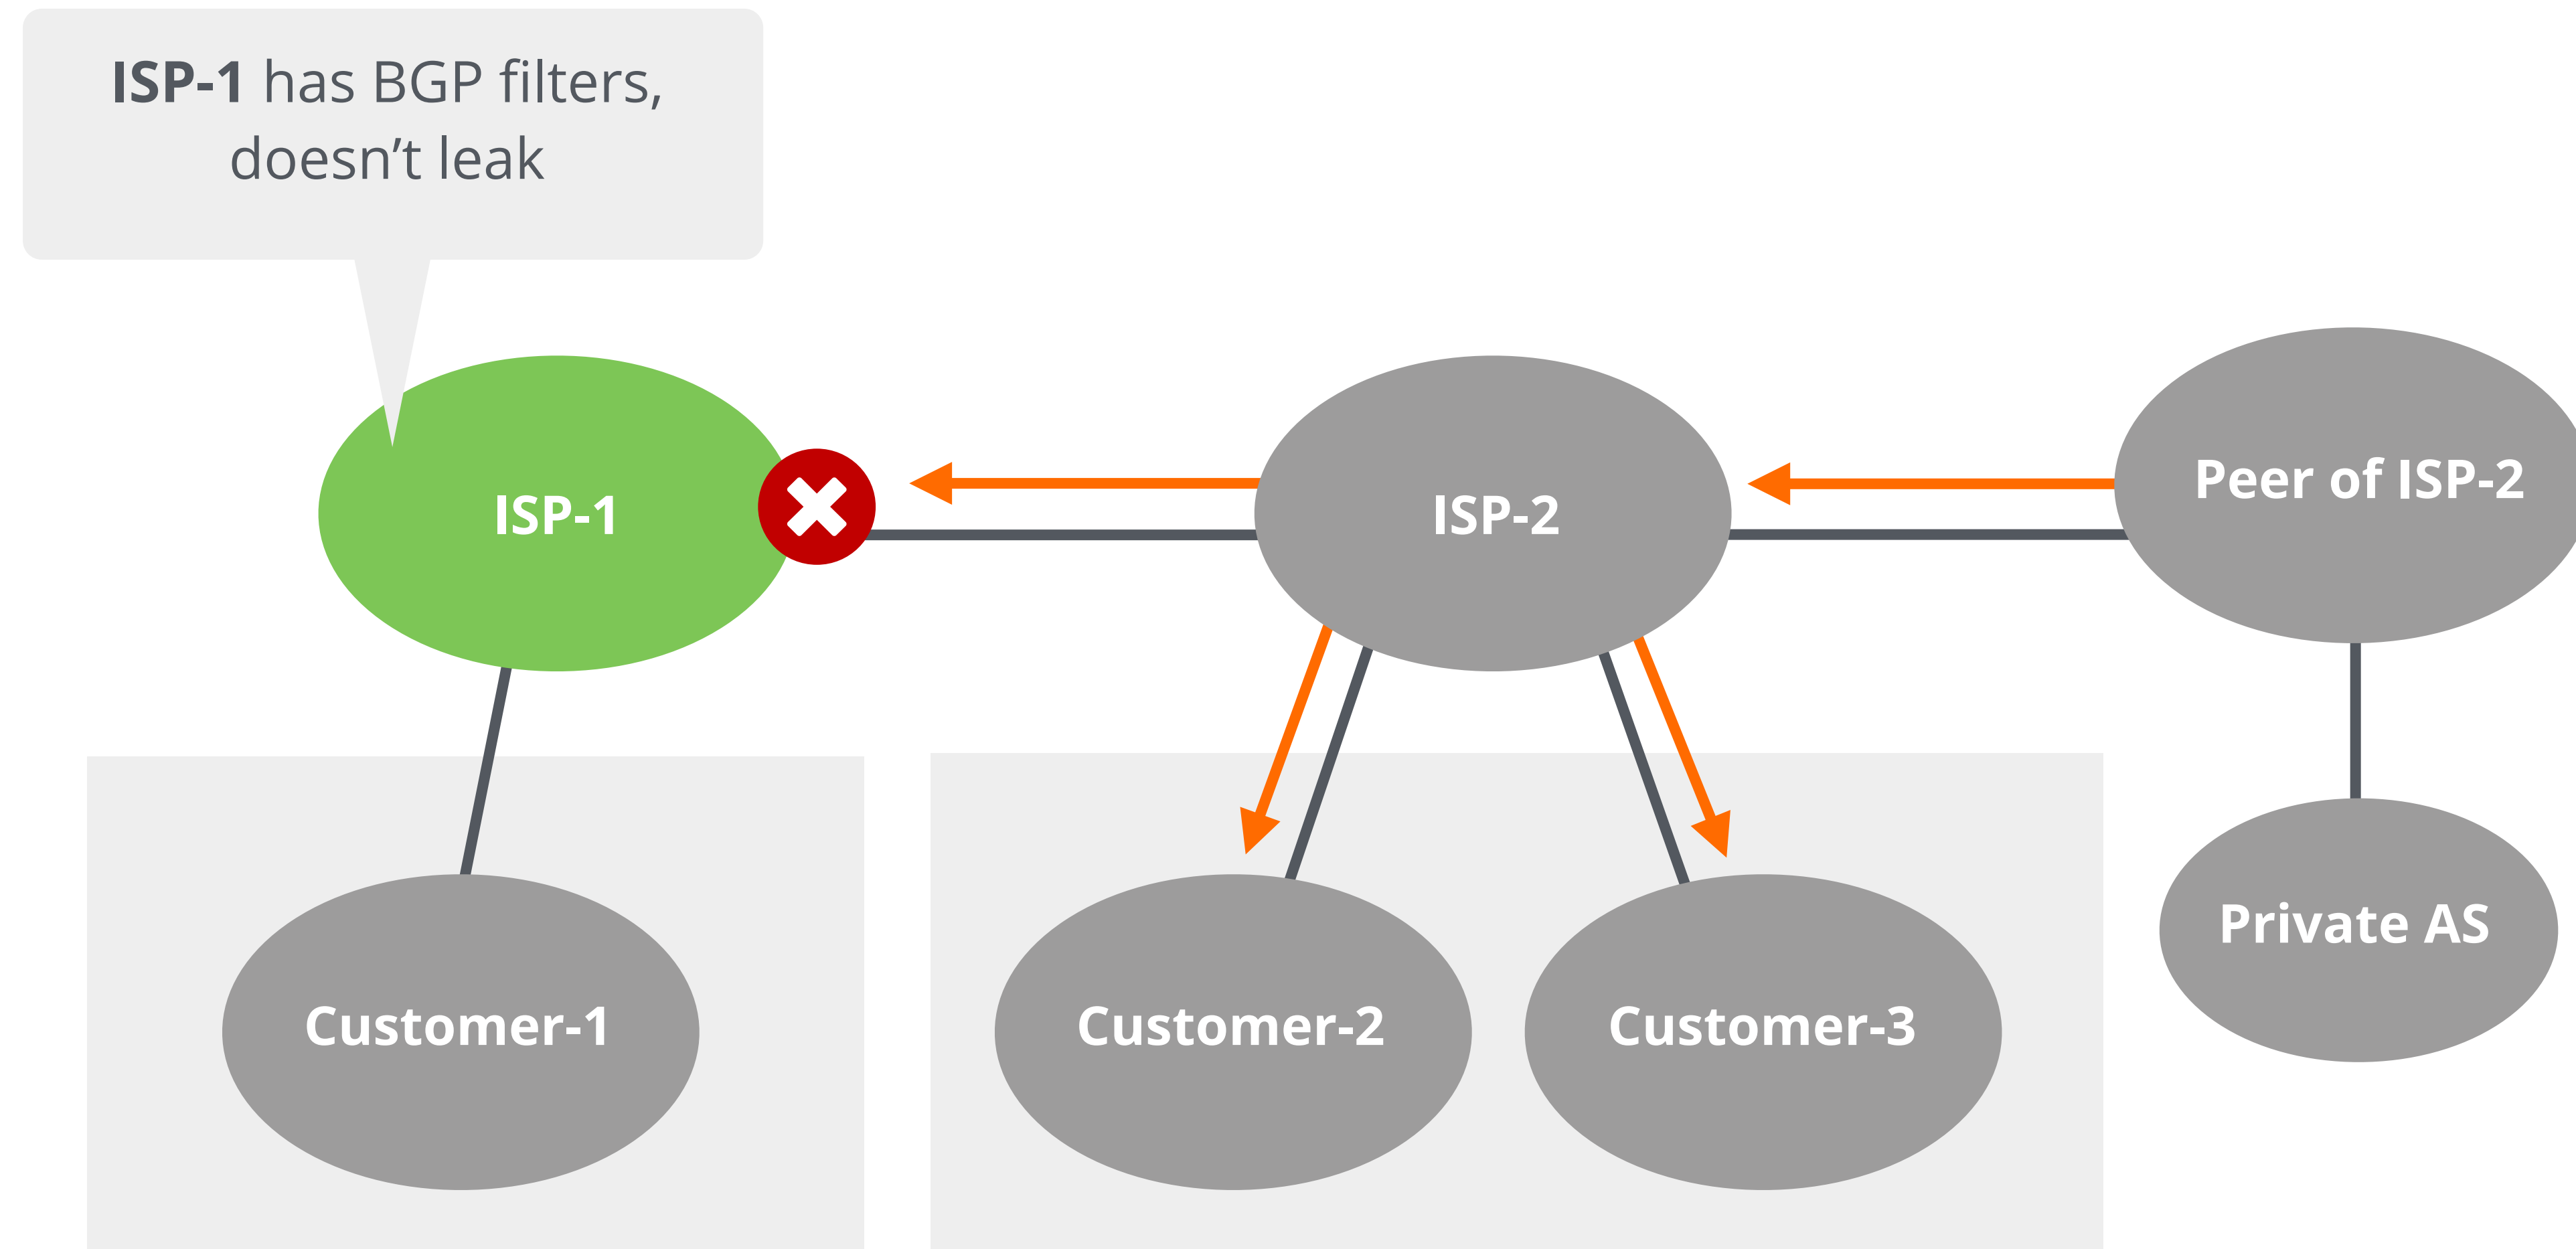
Why filter BGP prefixes?



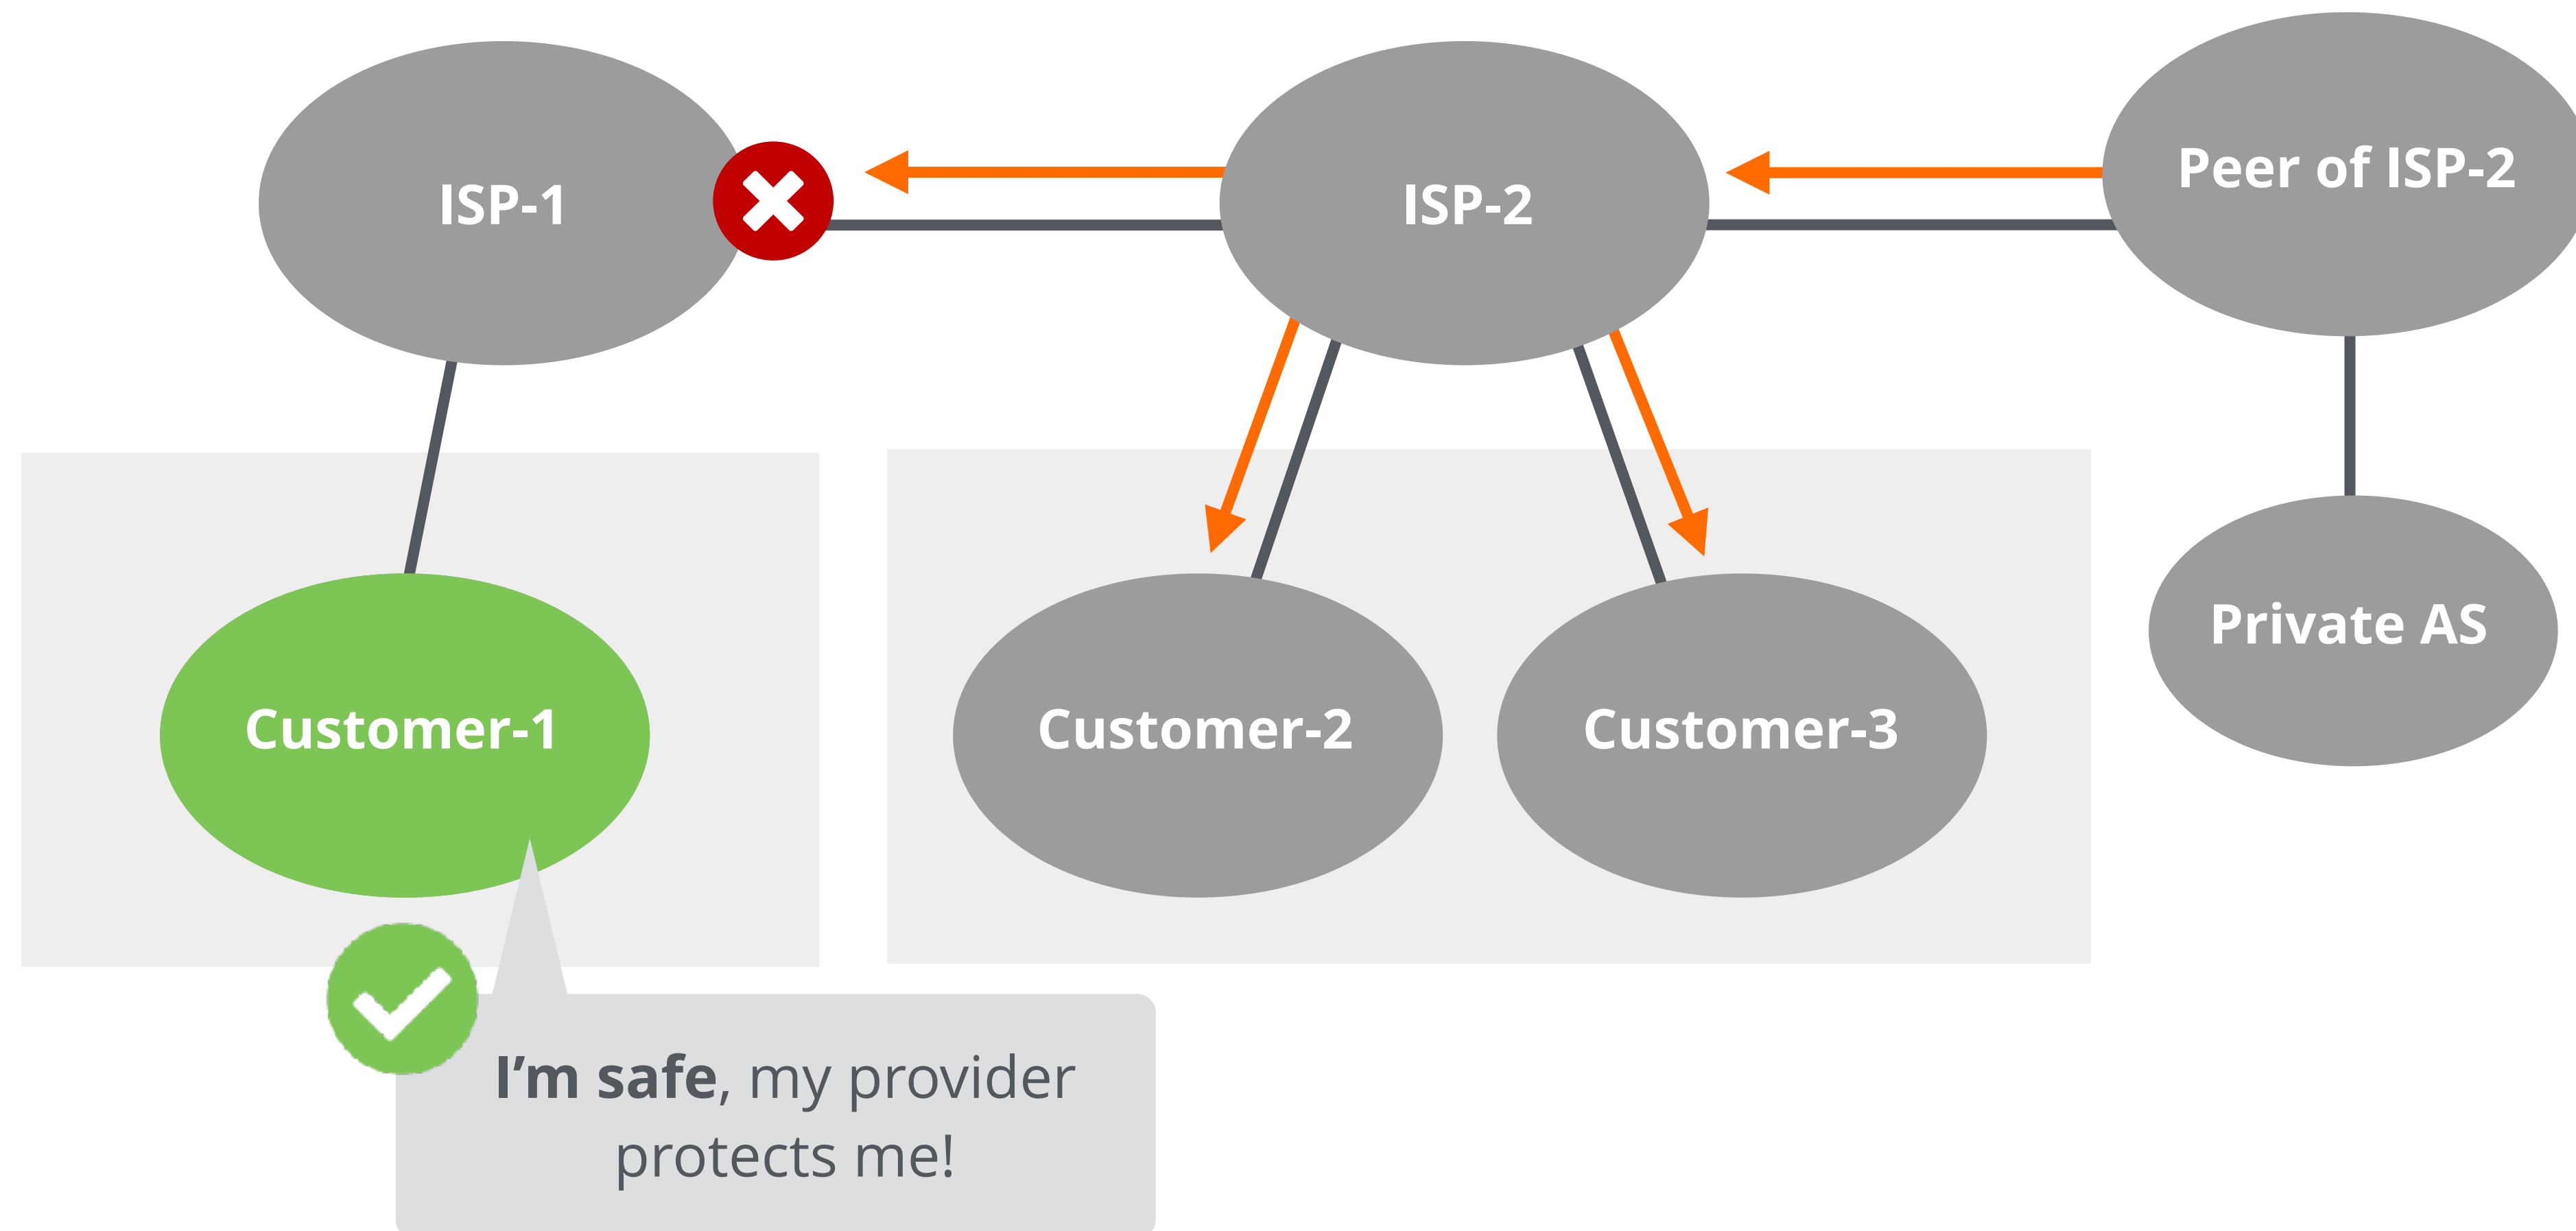
Why filter BGP prefixes?



Why filter BGP prefixes?



Why filter BGP prefixes?





Recent BGP Incidents

- YouTube (2008)
- AWS route leak (2016)
- Google prefix leak (2018)
- Akamai, Amazon, Alibaba (2020) ...



Having BGP filters could have mitigated the impact of these incidents!



Other reasons to use filtering ...

- **Business relationships**
 - Customer-provider, peer-to-peer
- **Technical reasons**
 - Reduce memory utilisation, scalability
- **Traffic engineering**
 - Manipulate traffic flows and influence best path selection



BGP Filters


(BGP Policies)

Section 2



BGP Filters (BGP Policies)

- Used to filter prefixes exchanged between BGP peers
- Describe BGP peers and routing relationships with them
- Filters can match on
 - IP prefixes
 - AS paths
 - Or any other BGP attributes (e.g. MED, BGP communities, ...)





BGP Filters (BGP Policies)

- Inbound policy

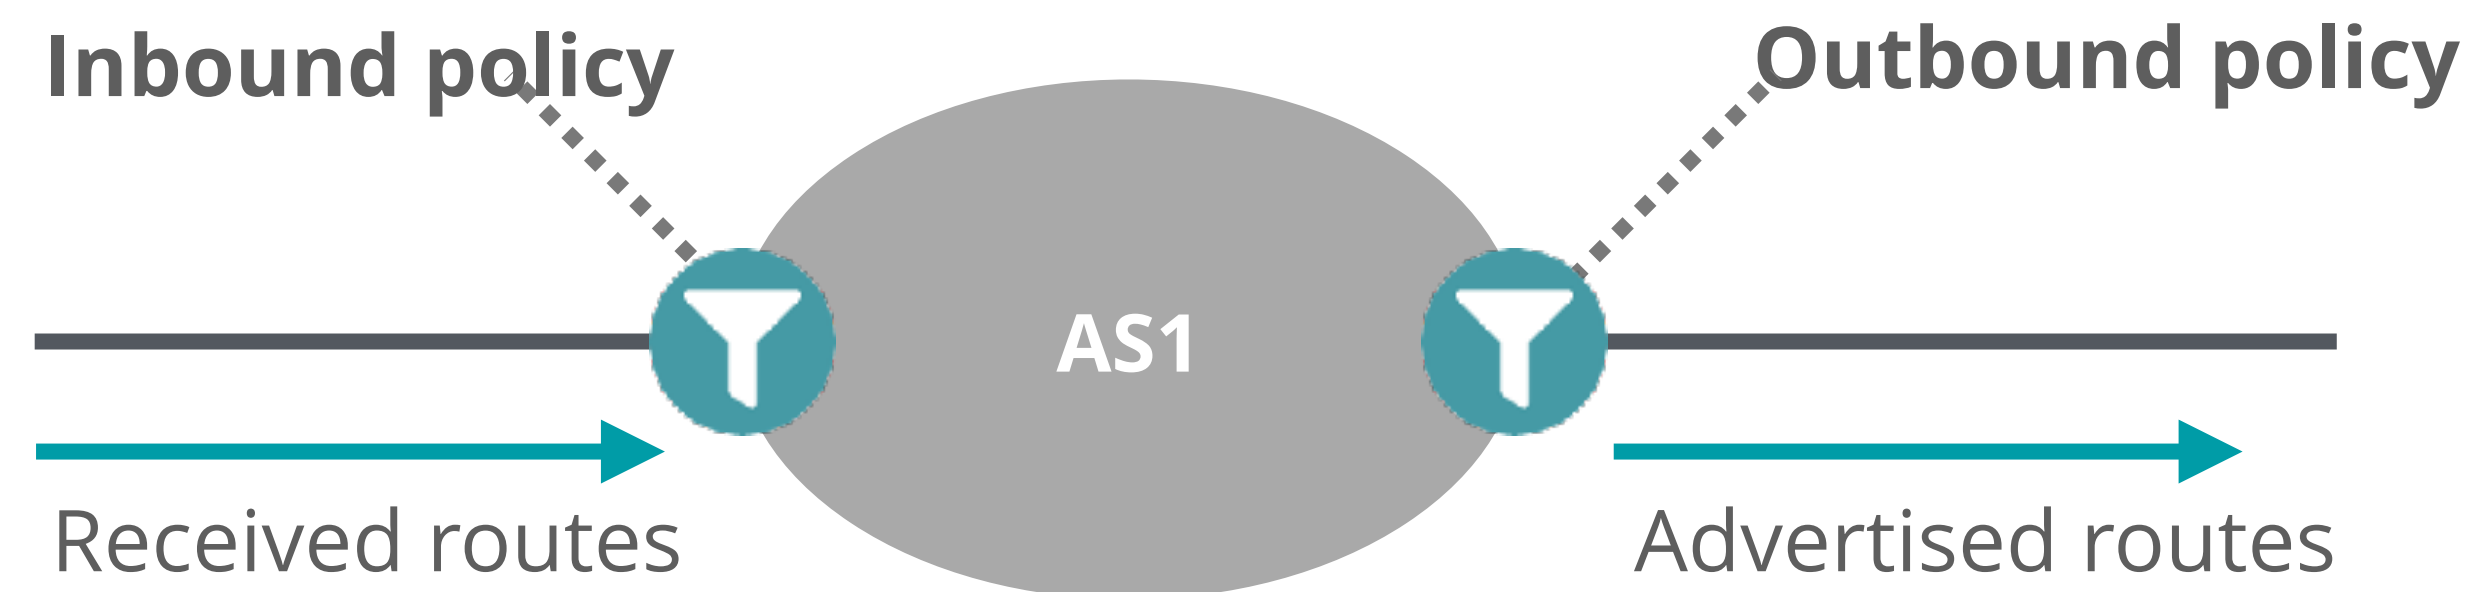
- For **incoming** (received) routes
- Detects configuration mistakes and attacks

- Should be applied on each eBGP peer

- Both on ingress and egress

- Outbound policy

- For **outgoing** (advertised) routes
- Limits propagation of routing information





Filtering Principles

- Filter **as close to the edge** as possible
- Filter **as precisely** as possible
- Two filtering approaches:
 - Explicit Permit (permit then deny any)
 - Explicit Deny (deny then permit any)

How to filter BGP routes





Prefix List

- List of routes you want to **accept** or **announce**
- You can create a list **manually** or **automatically** with data from IRRs
- You can use scripts or tools
 - Filtergen (Level3)
 - bgpq4
 - IRRToolSet
 - IRR Power Tools

Easy to use, but not highly scalable



AS Path Filtering

- Filters routes **based on AS path**
 - Permit or deny prefixes from **certain ASes**

```
router bgp 65564
  network 10.0.0.0 mask 255.255.255.0
  neighbor 172.16.1.1 remote-as 65563
  neighbor 172.16.1.1 filter-list 1 out
  neighbor 172.16.1.1 filter-list 2 in

ip as-path access-list 1 permit ^65564$
ip as-path access-list 2 permit ^65563$
```

Widely used and highly scalable

Take the poll!

Which routes should you **filter** in your BGP filter configuration?





Which routes should be filtered?


- Special-purpose prefixes (IPv4/IPv6) (Martians)
- Unallocated prefixes
- Routes that are too specific
- Prefixes belonging to the local AS
- IXP LAN prefixes
- The default route (0.0.0.0/0, ::/0)

RFC 7454 - "BGP Operations and Security"
lists the prefixes to be filtered.



Special-purpose Prefixes

- Also known as **Martians**
 - RFC 1918 Private addresses
 - Reserved space (documentation, multicast, etc.)
- Not globally routable
 - Should be **discarded** on Internet BGP peering

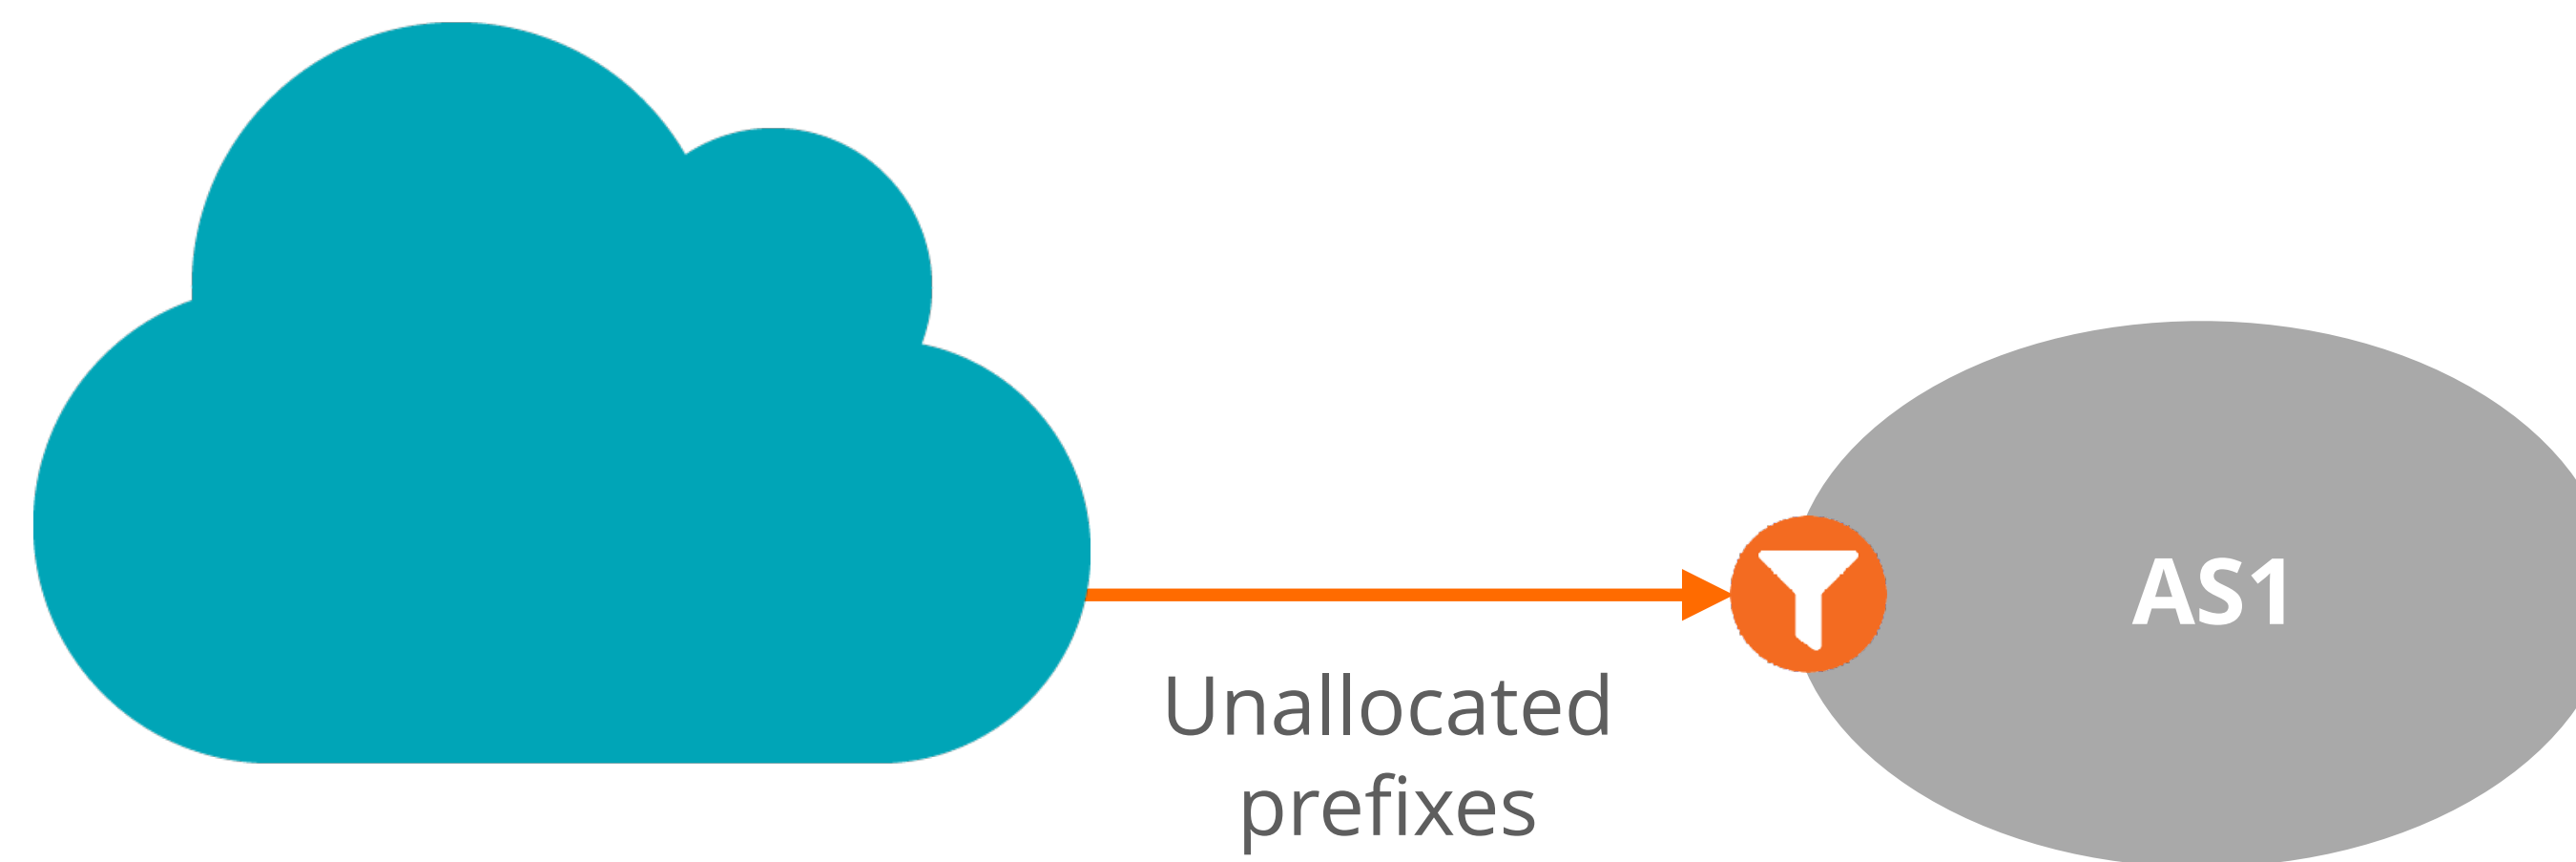


<http://www.iana.org/assignments/iana-ipv4-special-registry>
<http://www.iana.org/assignments/iana-ipv6-special-registry>



Unallocated Prefixes

- **All unallocated prefixes should be filtered**
 - Prefixes not yet allocated by IANA to RIRs (only for IPv6)
 - Prefixes allocated to an RIR but have not yet been distributed by an RIR to LIRs/End-users
- Filtering unallocated prefixes requires regular update





Longest Accepted Prefixes

- **Smaller prefixes should not be a part of global routing!**
 - /24 for IPv4 (RIPE-399)
 - /48 for IPv6 (RIPE-532)
- Those prefixes are generally neither announced nor accepted on the Internet

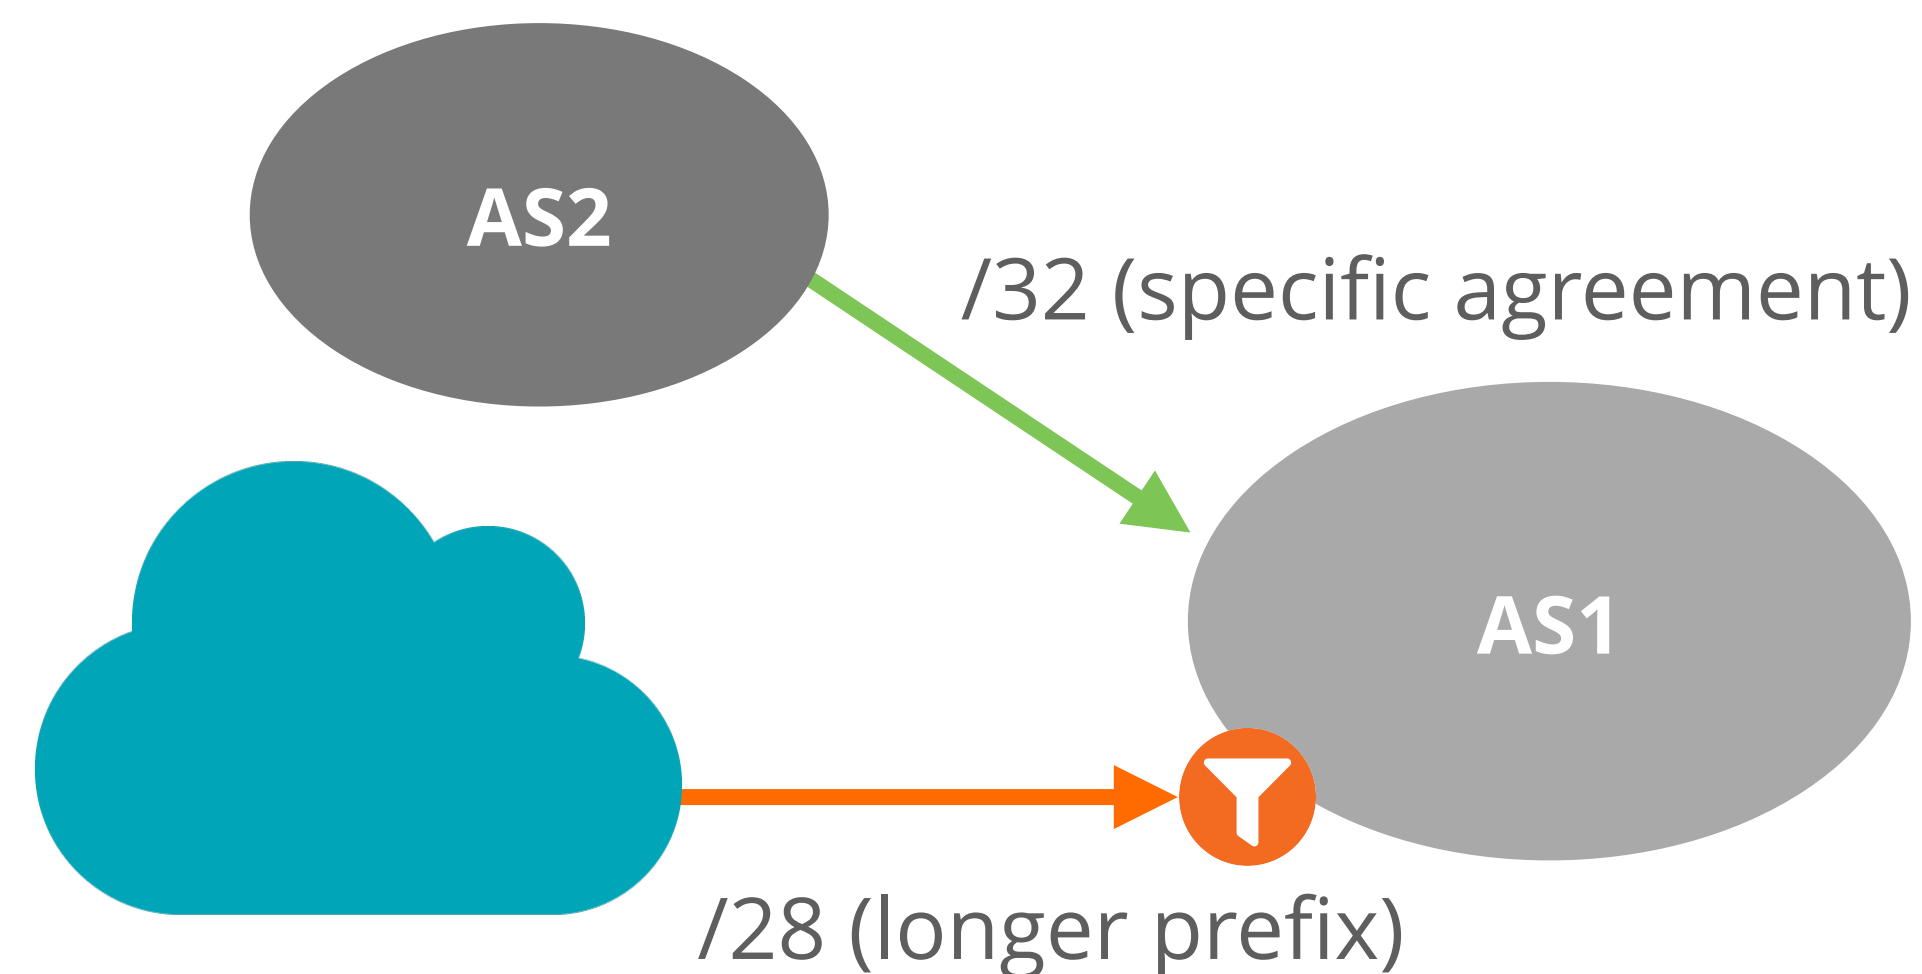
```
ip prefix-list SMALL-V4 permit 0.0.0.0/0 le 24
```

```
ipv6 prefix-list SMALL-V6 permit 2000::/3 le 48
```




Longest Accepted Prefixes

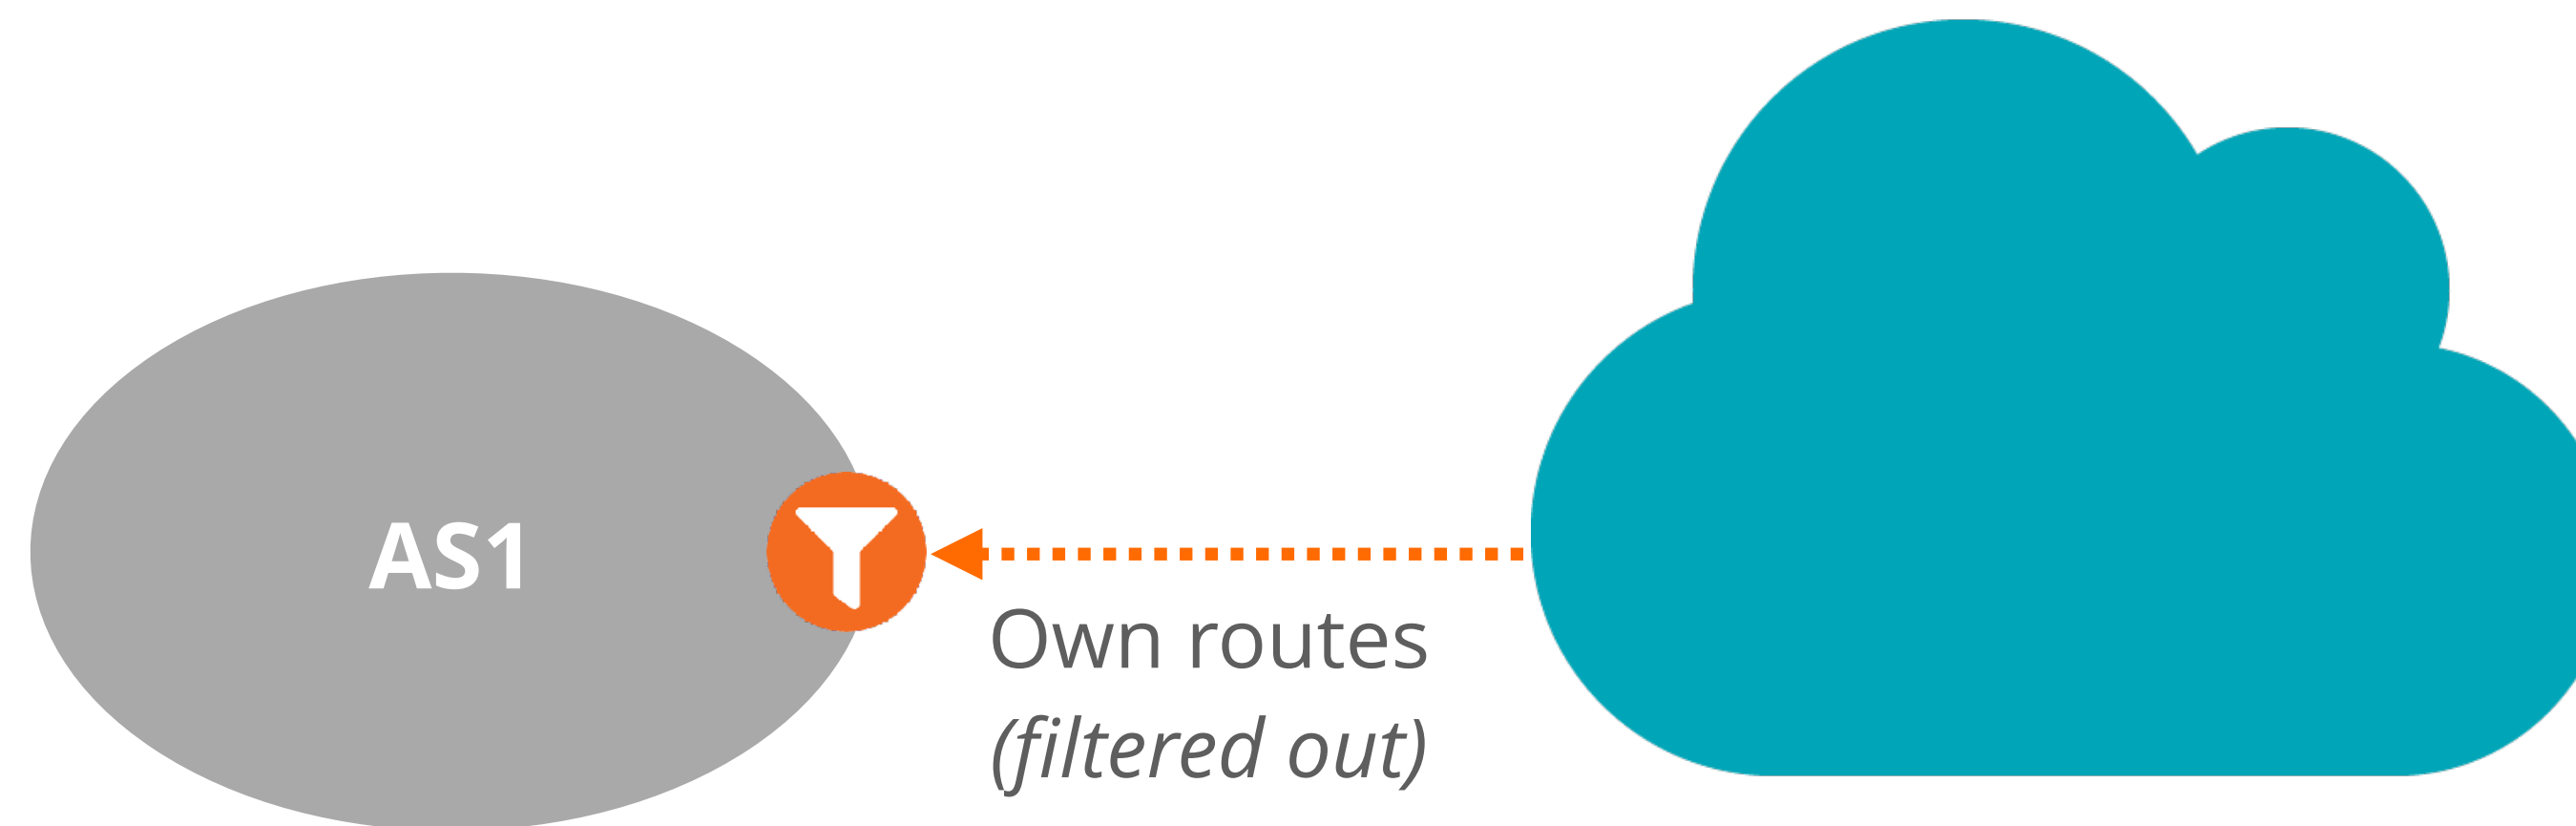
- In some cases ASes mutually agree to accept longer prefixes
 - Only for certain pre-agreed prefixes
 - e.g. flowspec is used between adjacent ASes for DDOS mitigation
- In this case, accepted prefix size should be defined for **that eBGP peer**
 - Reject prefixes exceeding the longest prefix size limit **per peer**





Prefixes Belonging to the Local AS

- You should **filter your own prefixes** on all BGP peering
 - Prevents local traffic from leaking over an external peering
- Such filters can also be configured for downstream customers' prefixes
- In case of multi-homed customer, be careful not to break redundancy mechanism





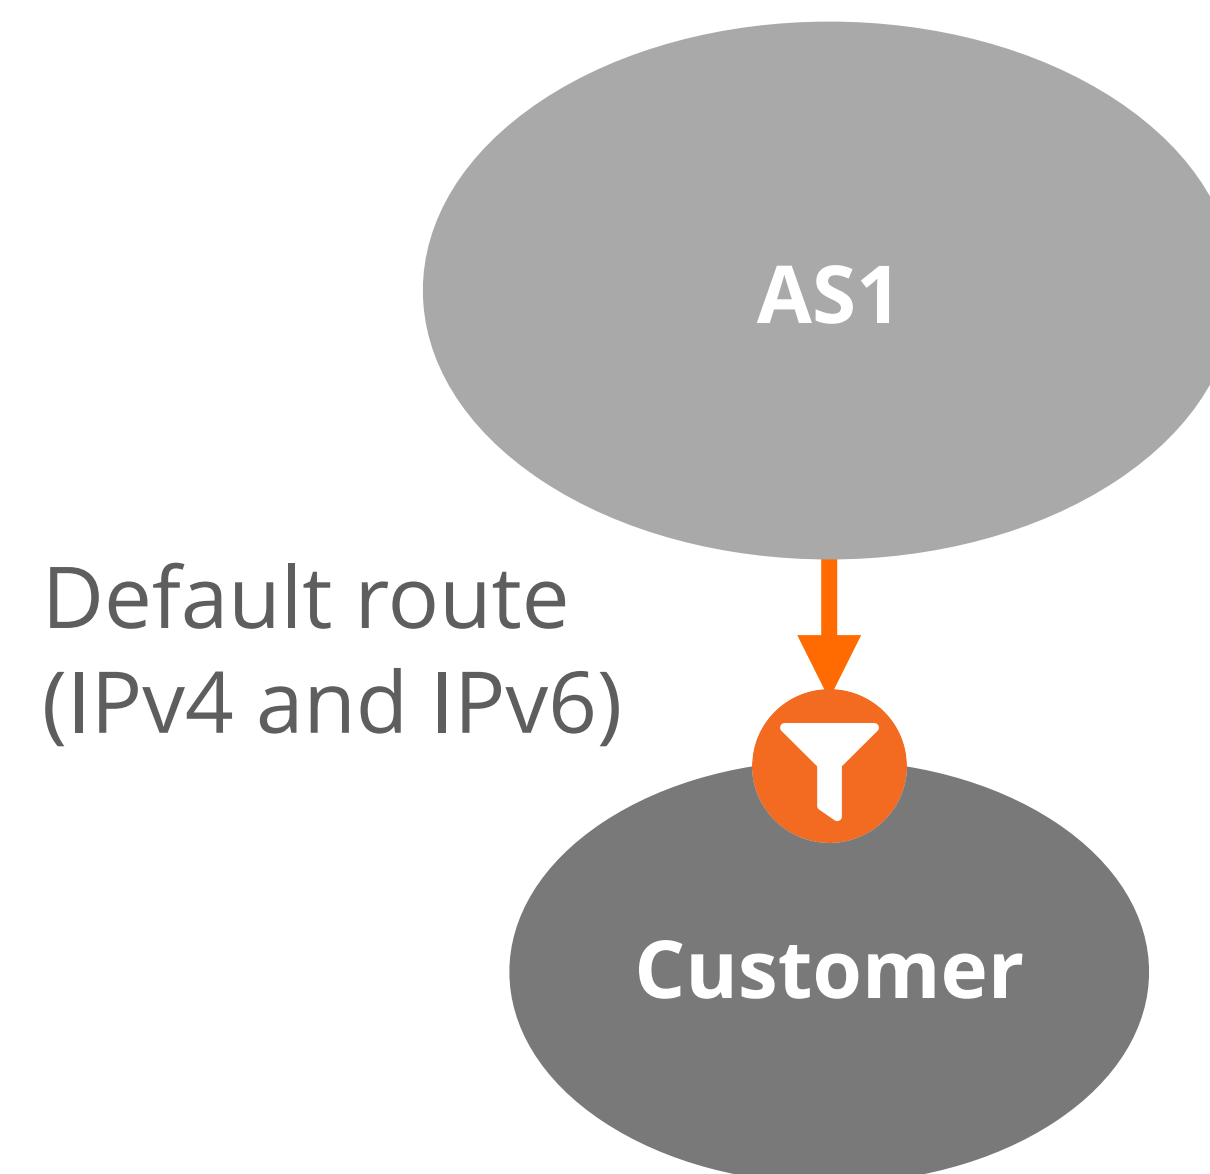
IXP LAN Prefixes

- An IXP should originate its LAN prefix
 - Advertise it from its route server to all IXP members
- **Do not accept an IXP LAN prefix from any of your eBGP peers!**
 - It may create a blackhole for connectivity to the IXP LAN
- IXP prefix announcements should pass IRR-generated filters



Default Route

- **0.0.0.0/0** (IPv4) and **::/0** (IPv6)
- Advertised or accepted only in specific customer-provider peering relationships
 - E.g. A customer with a stub network
- Should be rejected unless a special peering agreement is in place

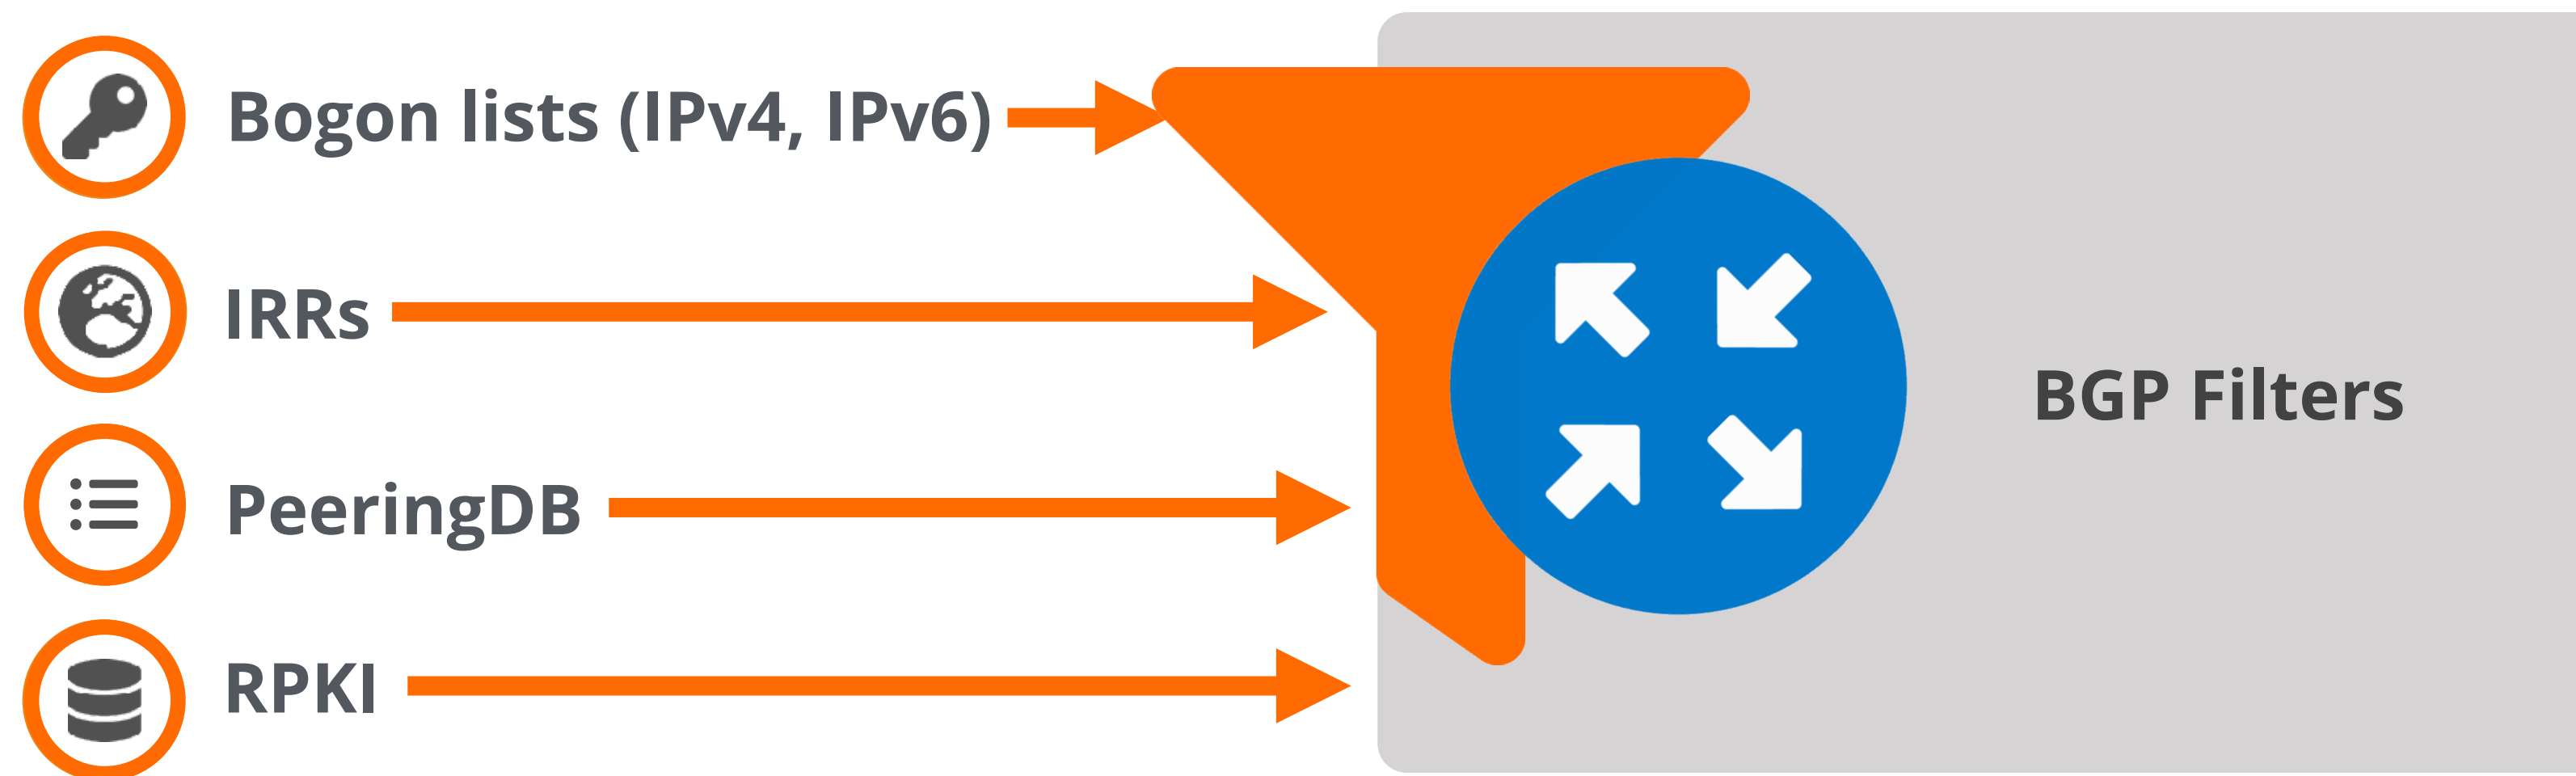


Take the poll!

Which **data sources** could be used for creating BGP filters?



Data Sources for BGP Filters





Bogon Lists

- **Bogons** are prefixes that should never appear in the Internet routing table!
 - Martians (RFC#1918 Private addresses + Reserved space)
 - IANA unallocated space
- **Full Bogons** should be filtered as well
 - Bogons + RIR unassigned prefixes
- The bogon and full bogon lists are not static!

Bogon ASN Filtering



ASNs	Status	RFC
0	Reserved	RFC7607
23456	AS_TRANS	RFC6793
64496 - 64511	Reserved for use in docs and code	RFC5398
64512 - 65534	Reserved for Private Use	RFC6996
65535	Reserved	RFC 7300
65536 - 65551	Reserved for use in docs and code	RFC5398
65552 - 131071	Reserved	IANA
4200000000 - 4294967294	Reserved for Private Use	RFC6996
4294967295	Reserved	RFC 7300



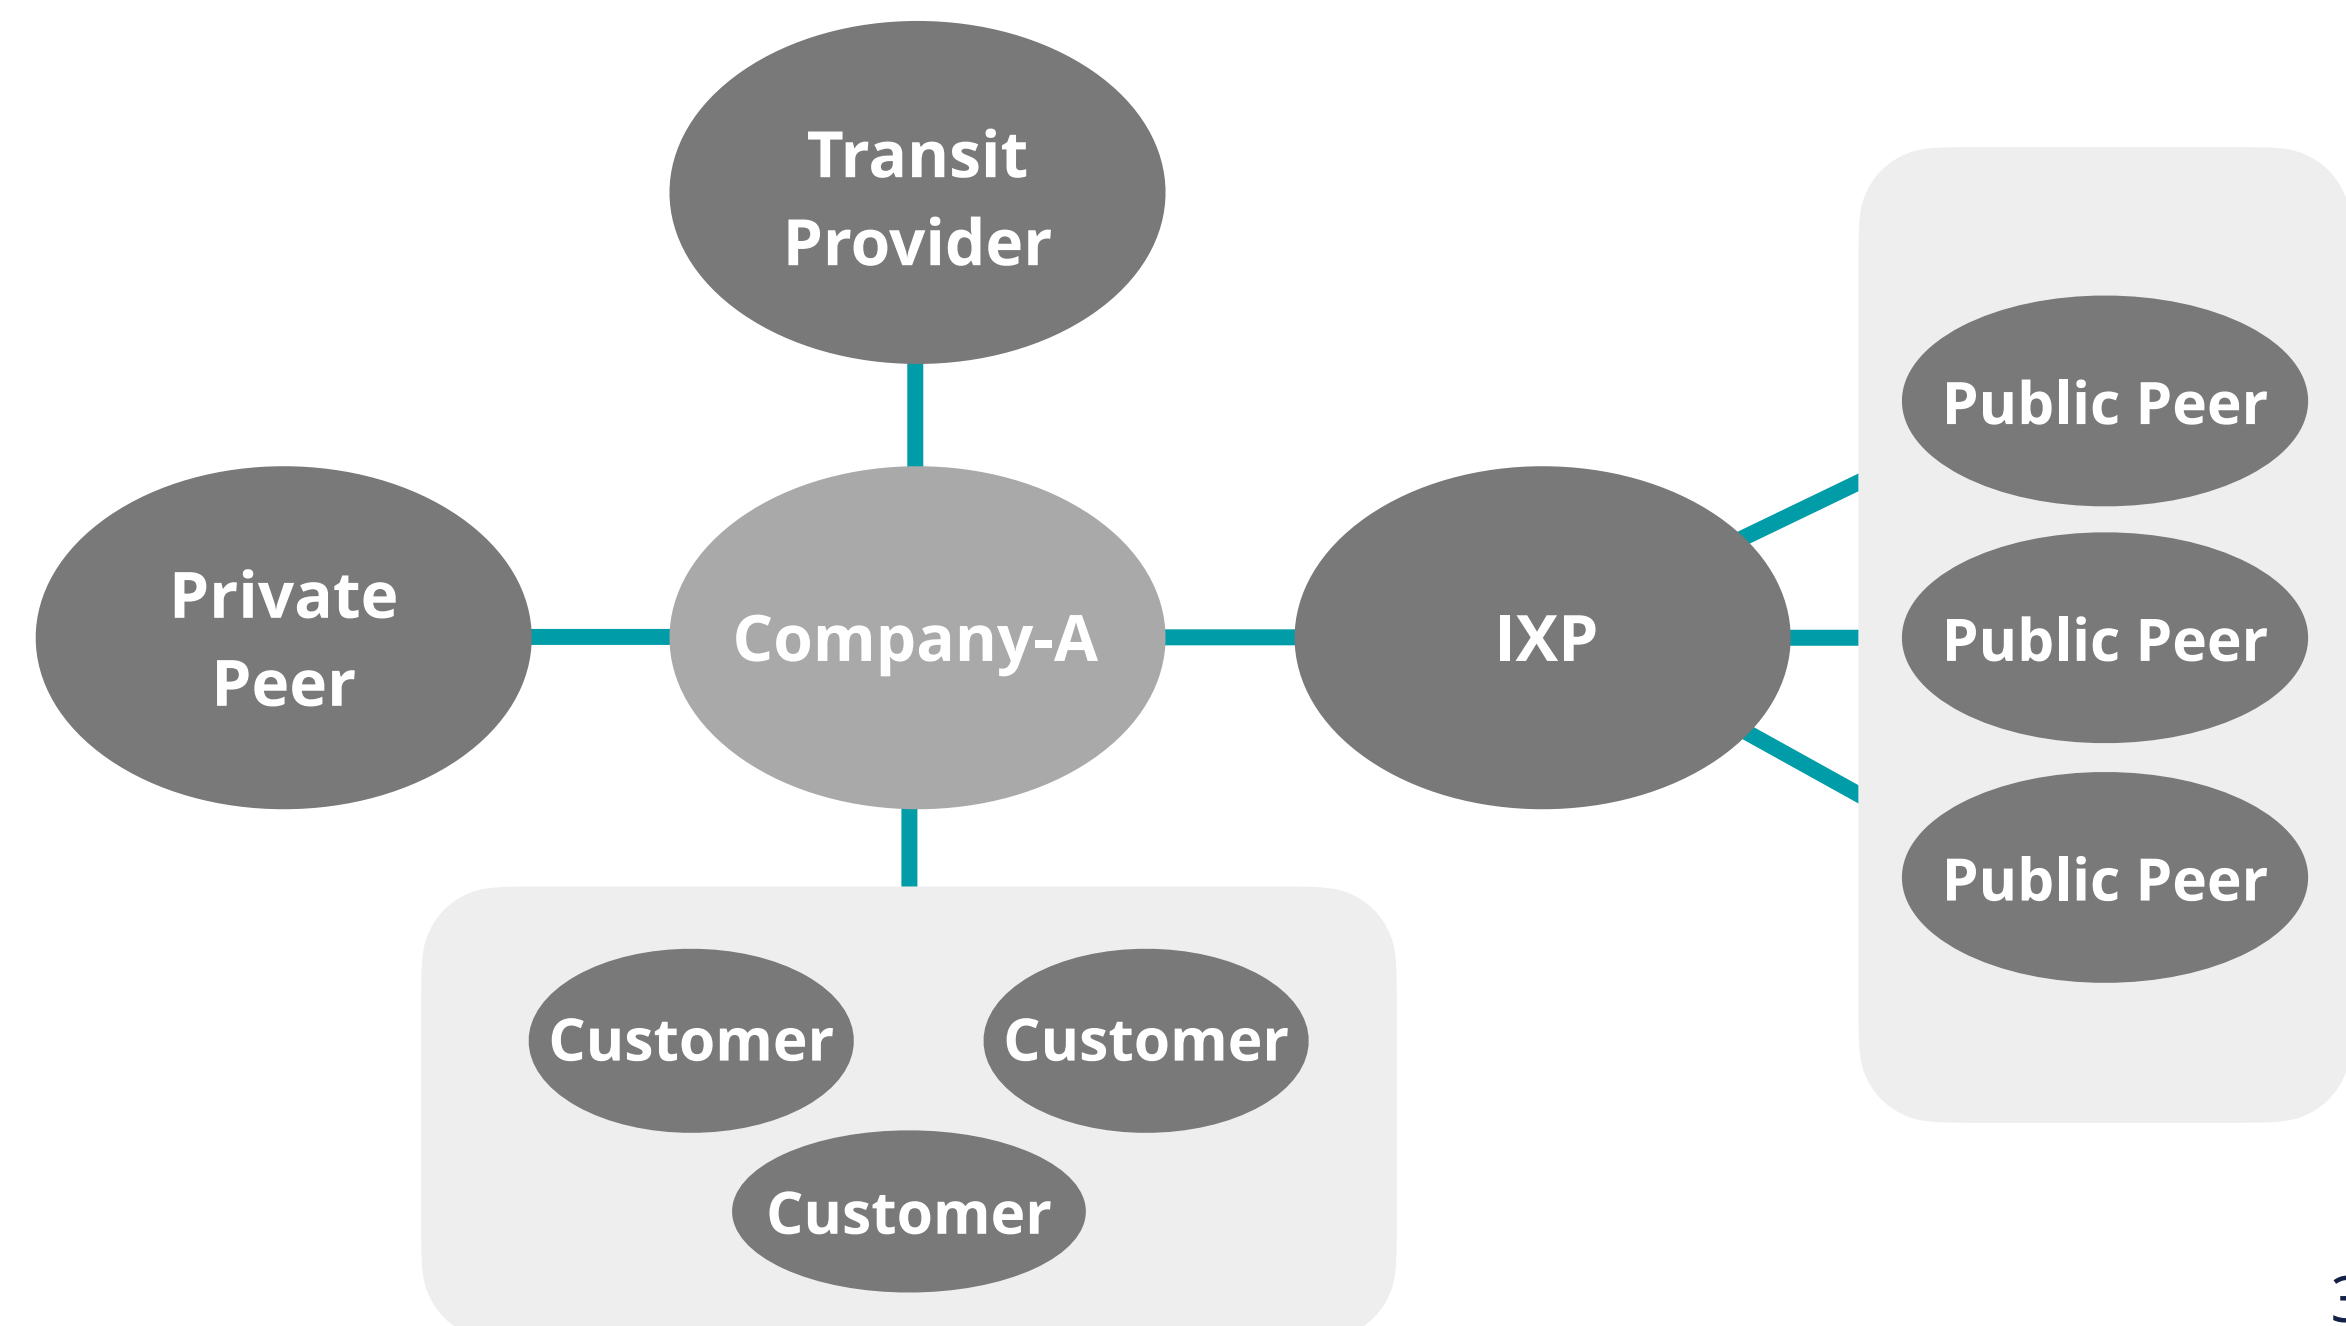
Prefix Filtering Recommendations

Section 3



Prefix Filtering Recommendations

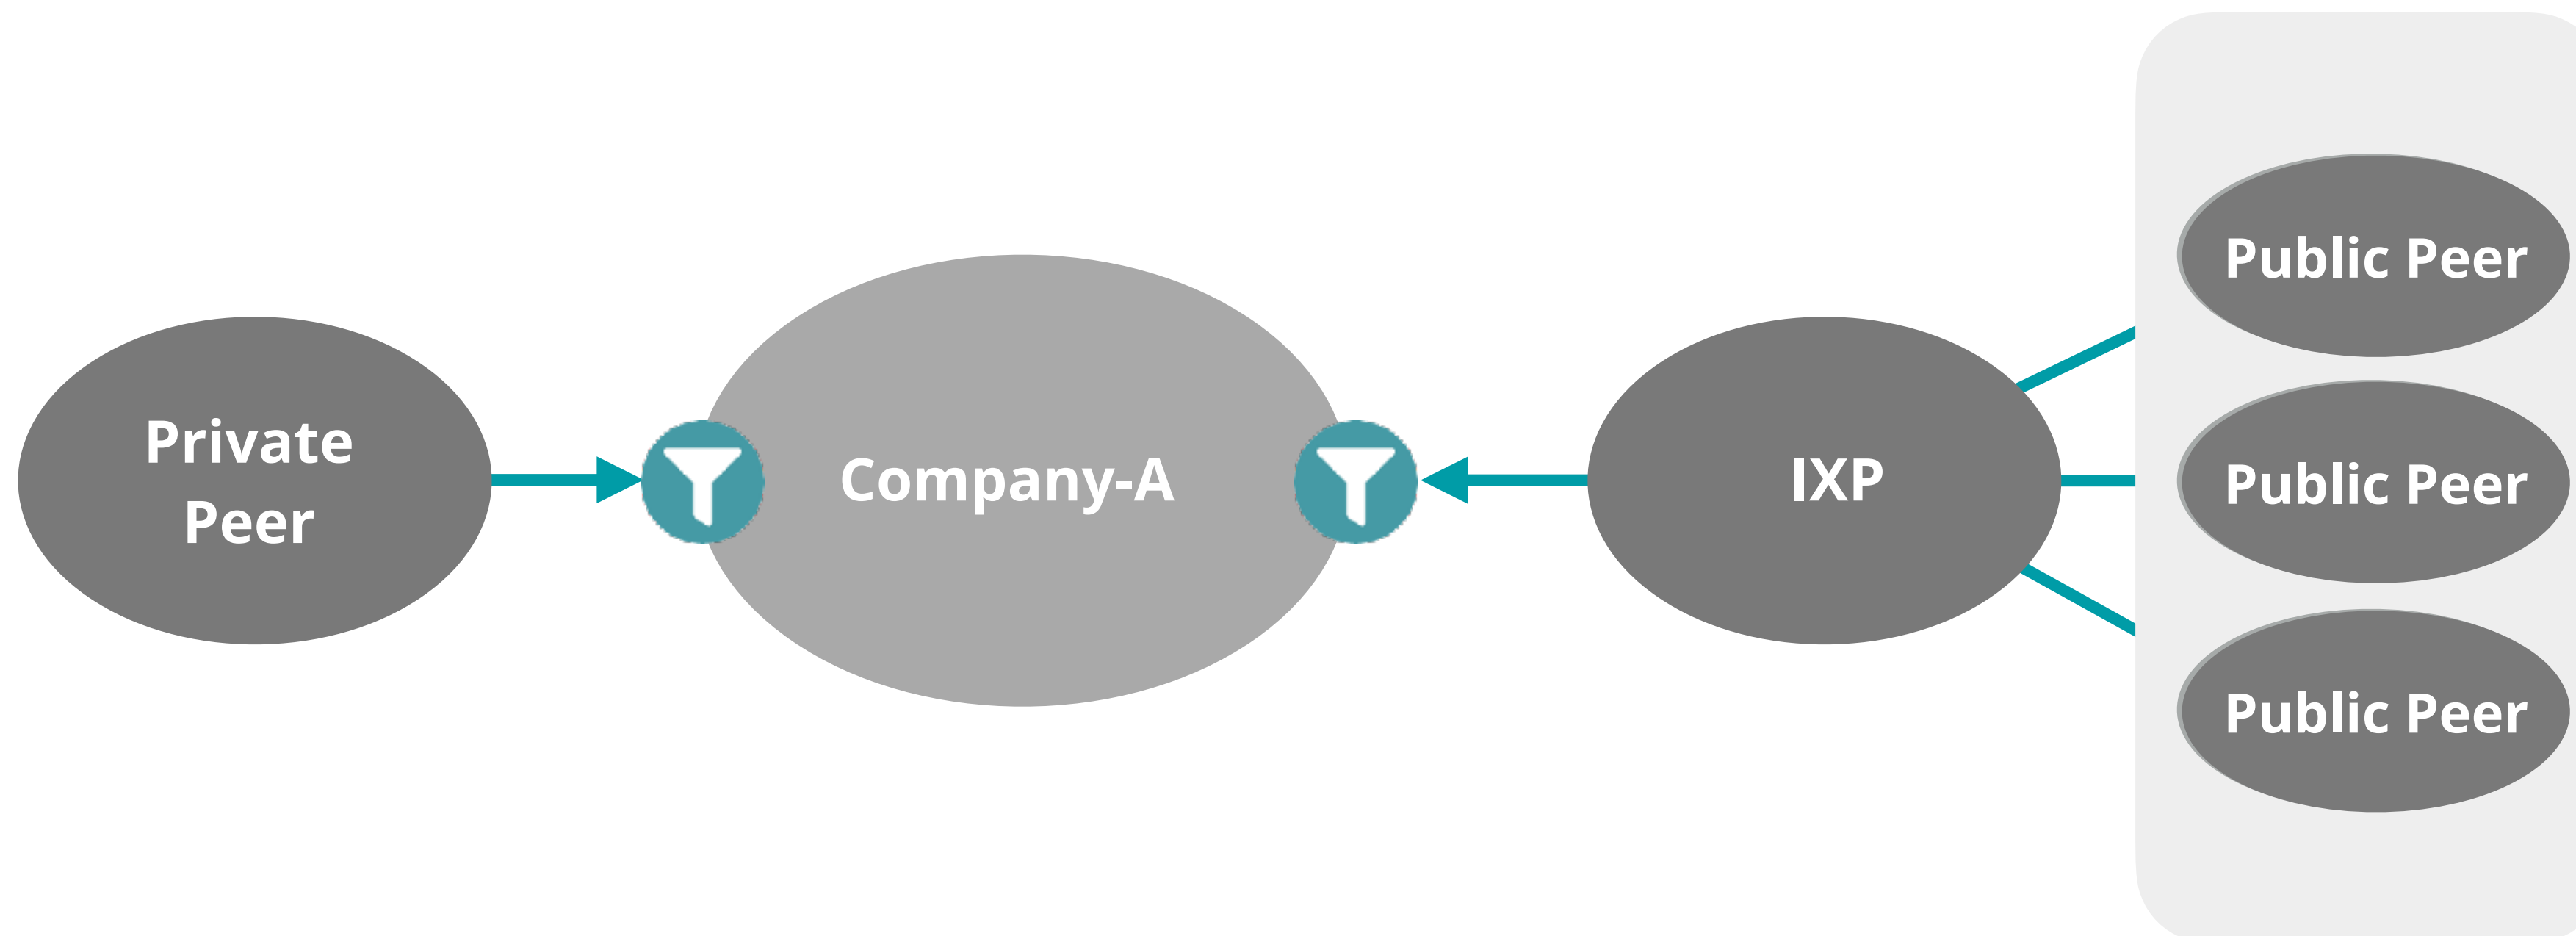
- In full routing networks, some policies should be applied:
 - On each BGP peer
 - For both **received** and **advertised** routes (inbound and outbound)
- Recommendations vary based on type of BGP peering relationships:
 - Public and Private Peering
 - Transit Provider (Upstream)
 - Customer





Filters With Peers (Inbound)

- Filters with public and private peers
- On **inbound**, strict or loose filtering could be implemented





Filters with Peers (Inbound)

- **Strict filtering:**
 - Makes sure advertisements conform to what is declared in IRRs
 - Impact should be checked before applying the policy
- **Loose filtering:**
 - Filters the routes based on RFC 7454 recommendations

Prefixes that are not globally routable

Prefixes belonging to the local AS

Prefixes not allocated by IANA (IPv6 only)

IXP LAN prefixes

Routes that are too specific

The default route



Filters With Peers (Outbound)

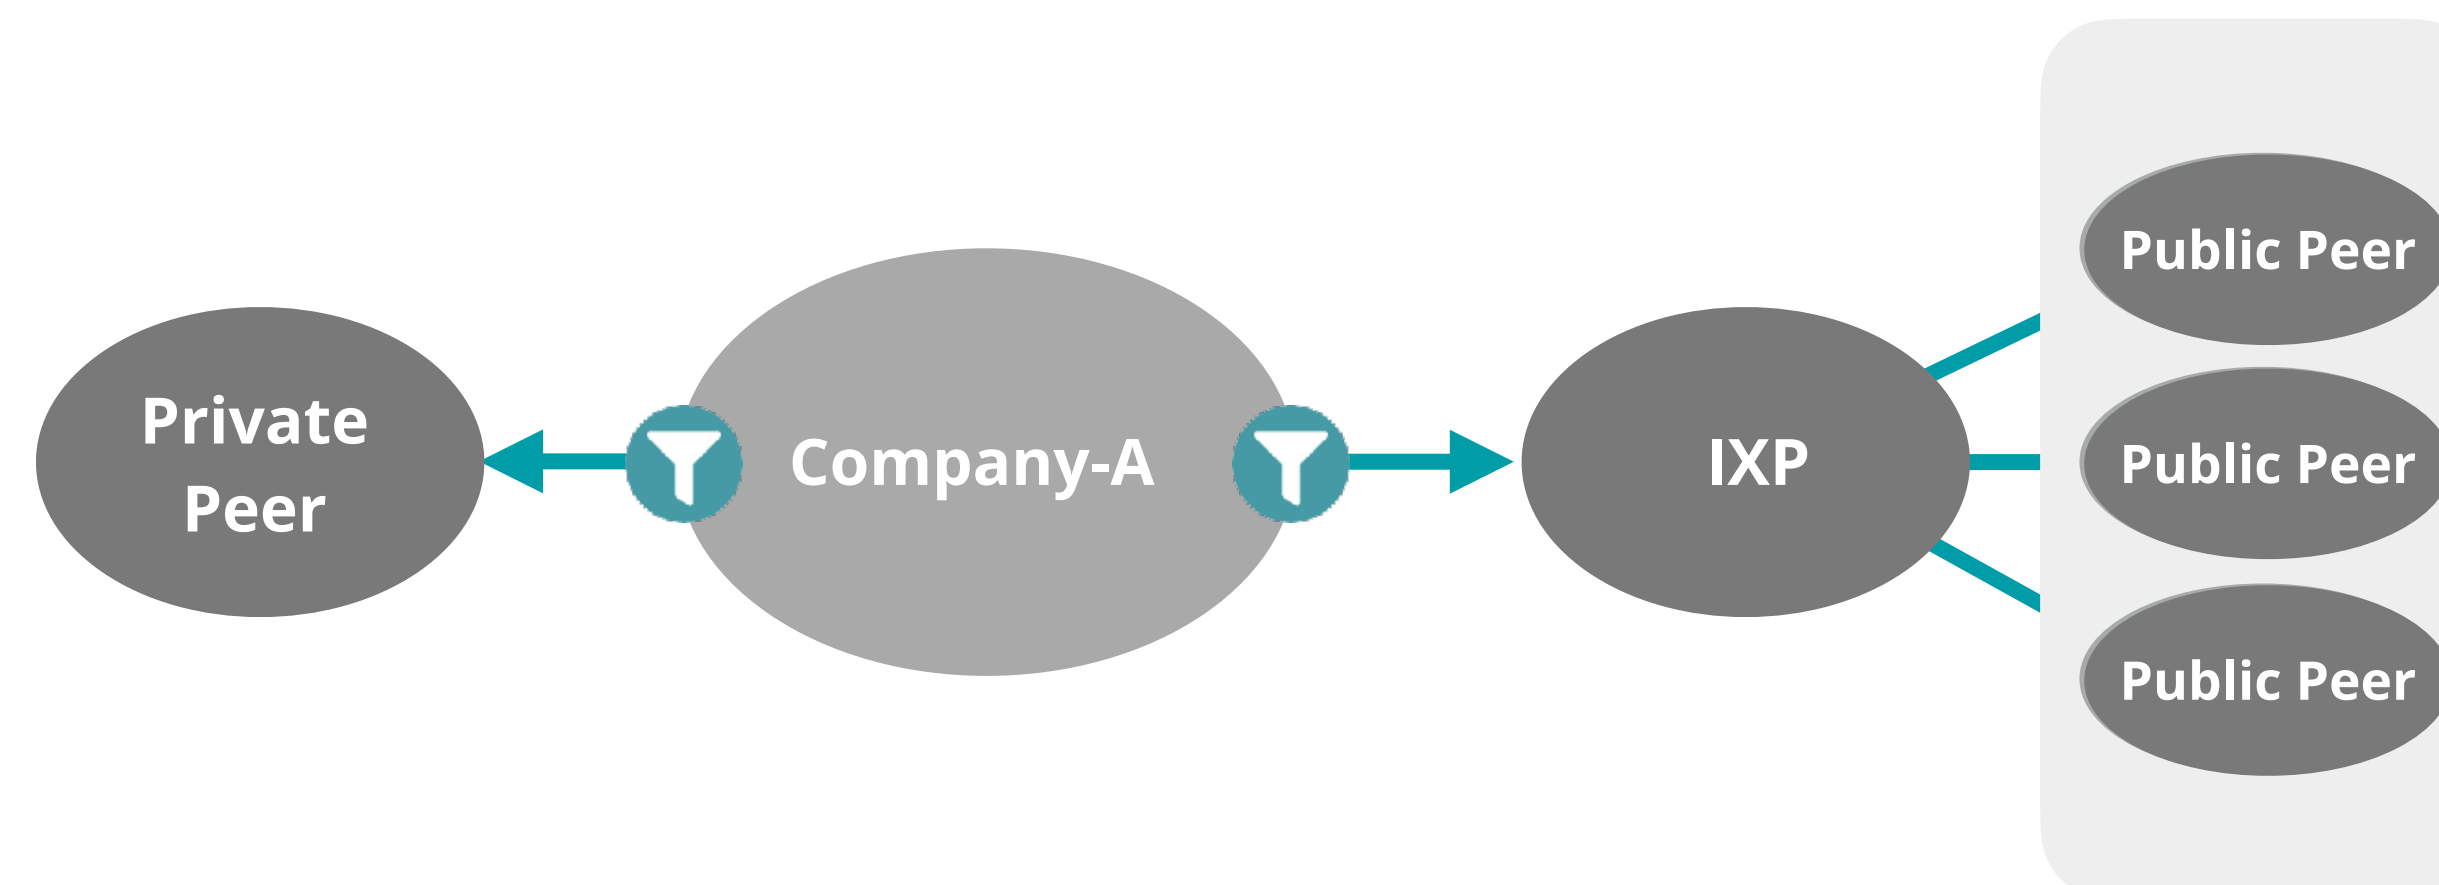
- Only locally originated and customers' prefixes should be sent
 - If possible, list the prefixes to be advertised, and deny the rest!
- Additional filters could be added to filter the following:

Prefixes that are not globally routable

IXP LAN prefixes

Routes that are too specific

The default route



Take the poll!

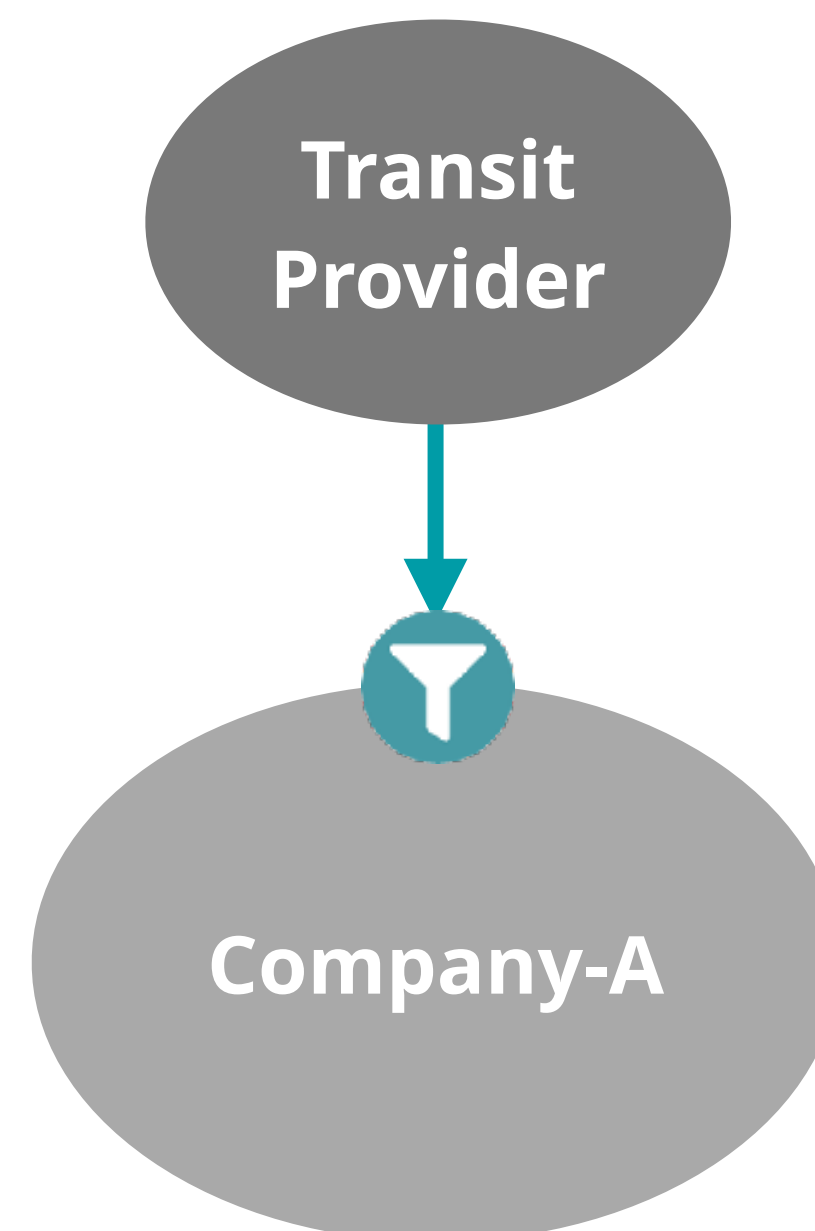
Which prefixes should be **filtered** from a **transit provider**?





Filters With Transit (Inbound)

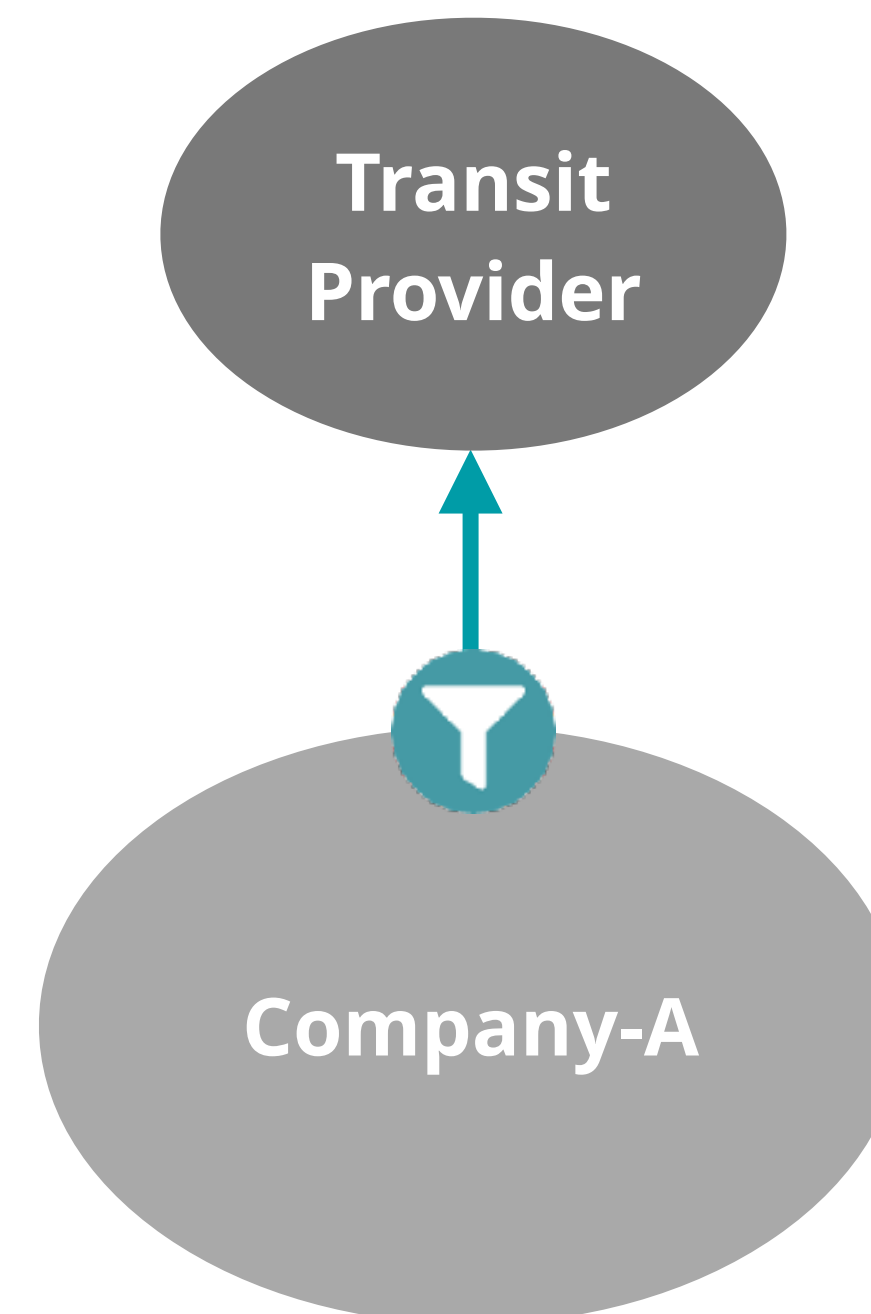
- If the full route table is desired,
 - RFC 7454 recommendations are the same with public and private peers
 - Except the default route
- If the upstream provider is supposed to announce the default route only
 - Accept only the default route





Filters With Transit (Outbound)

- The same outbound filters should be applied as those for public and private peers
- Make sure that **only authorised prefixes are sent**
 - Locally originated and customers' prefixes





Filters With Customers (Inbound)

- If all customer prefixes are known,
 - **Accept customer prefixes only** and discard the rest!
- What if you do not have this information? Filter the following:

Prefixes that are not globally routable

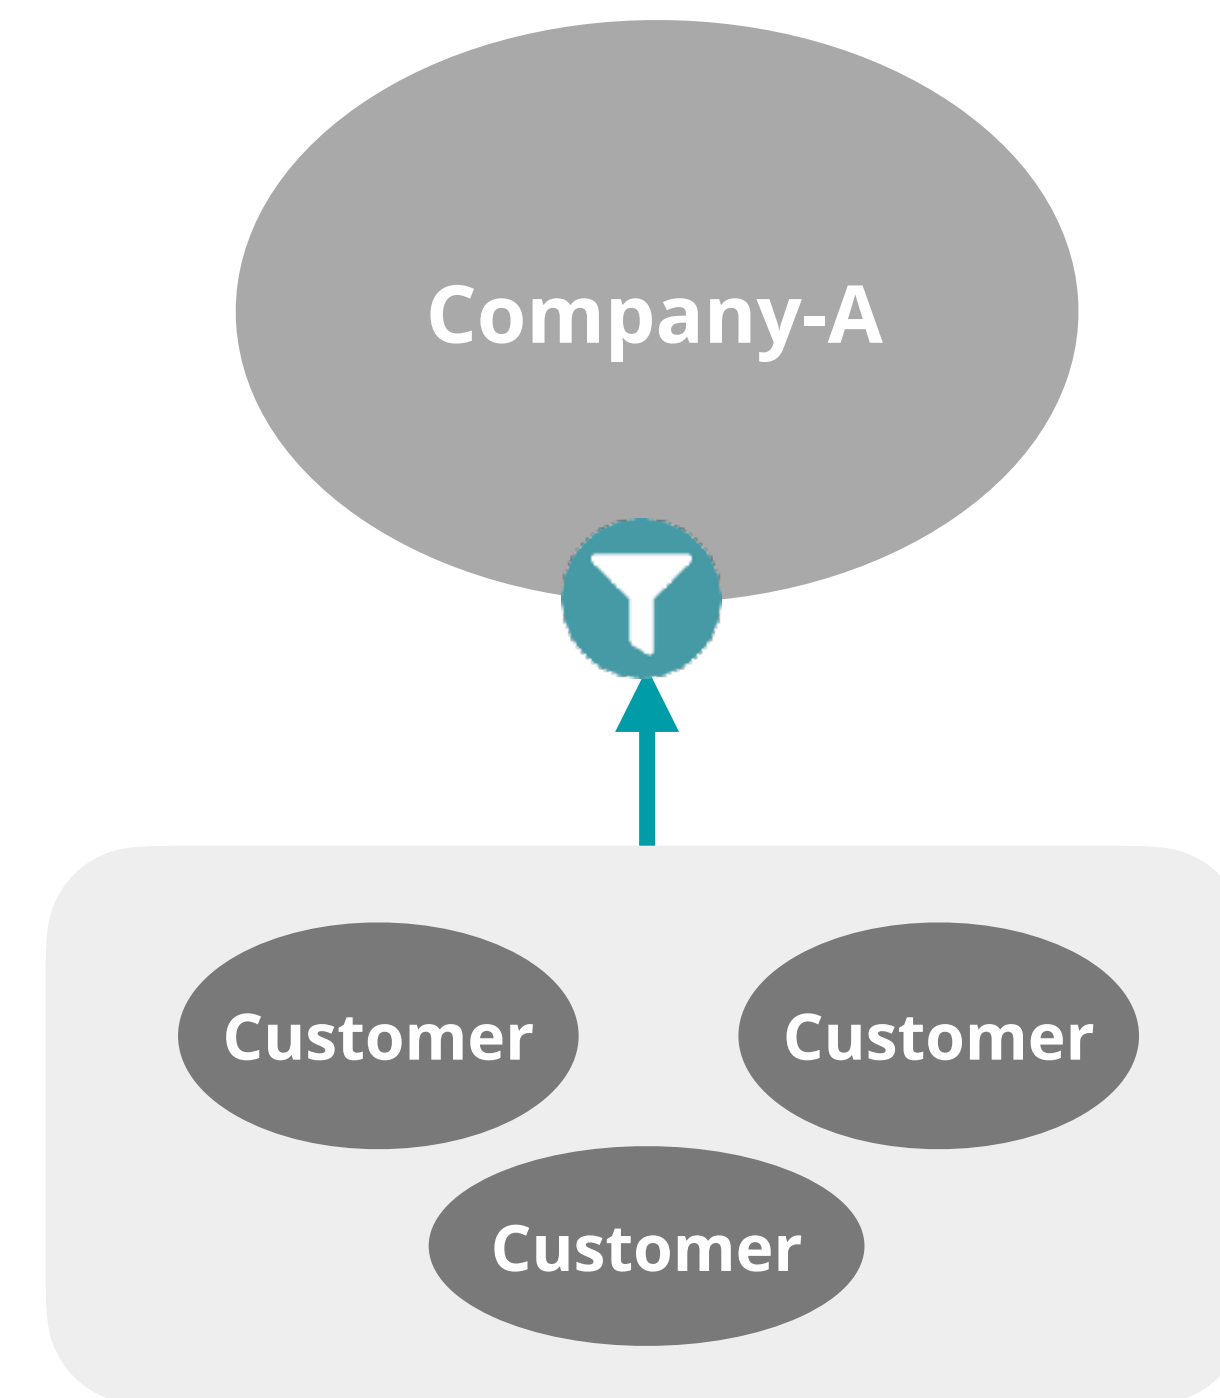
Prefixes not allocated by IANA (IPv6 only)

Routes that are too specific

Prefixes belonging to the local AS

IXP LAN prefixes

The default route





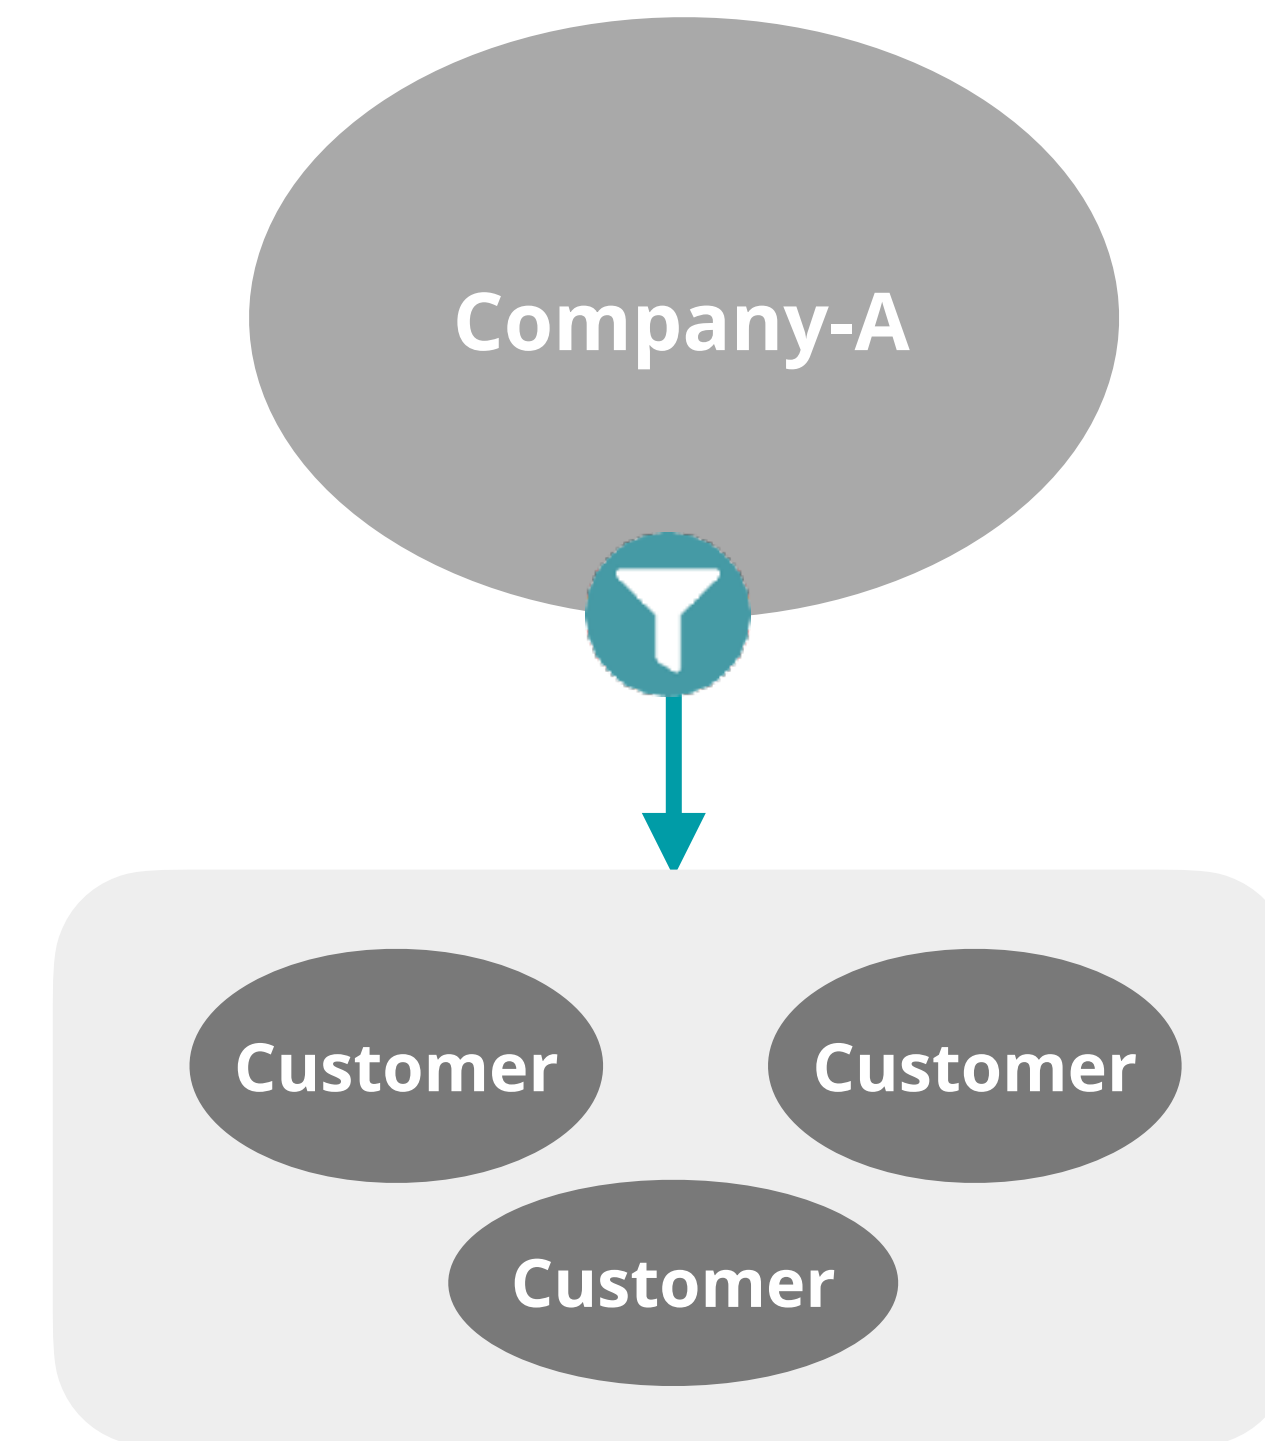
Filters With Customers (Outbound)

- According to RFC 7454, it may vary depending on customers preferences
- If a customer requests default route only, send only the default
- For other cases, filter the following prefixes:

Prefixes that are not globally routable

Routes that are too specific

The default route (?)





Leaf Customer Network (Inbound)

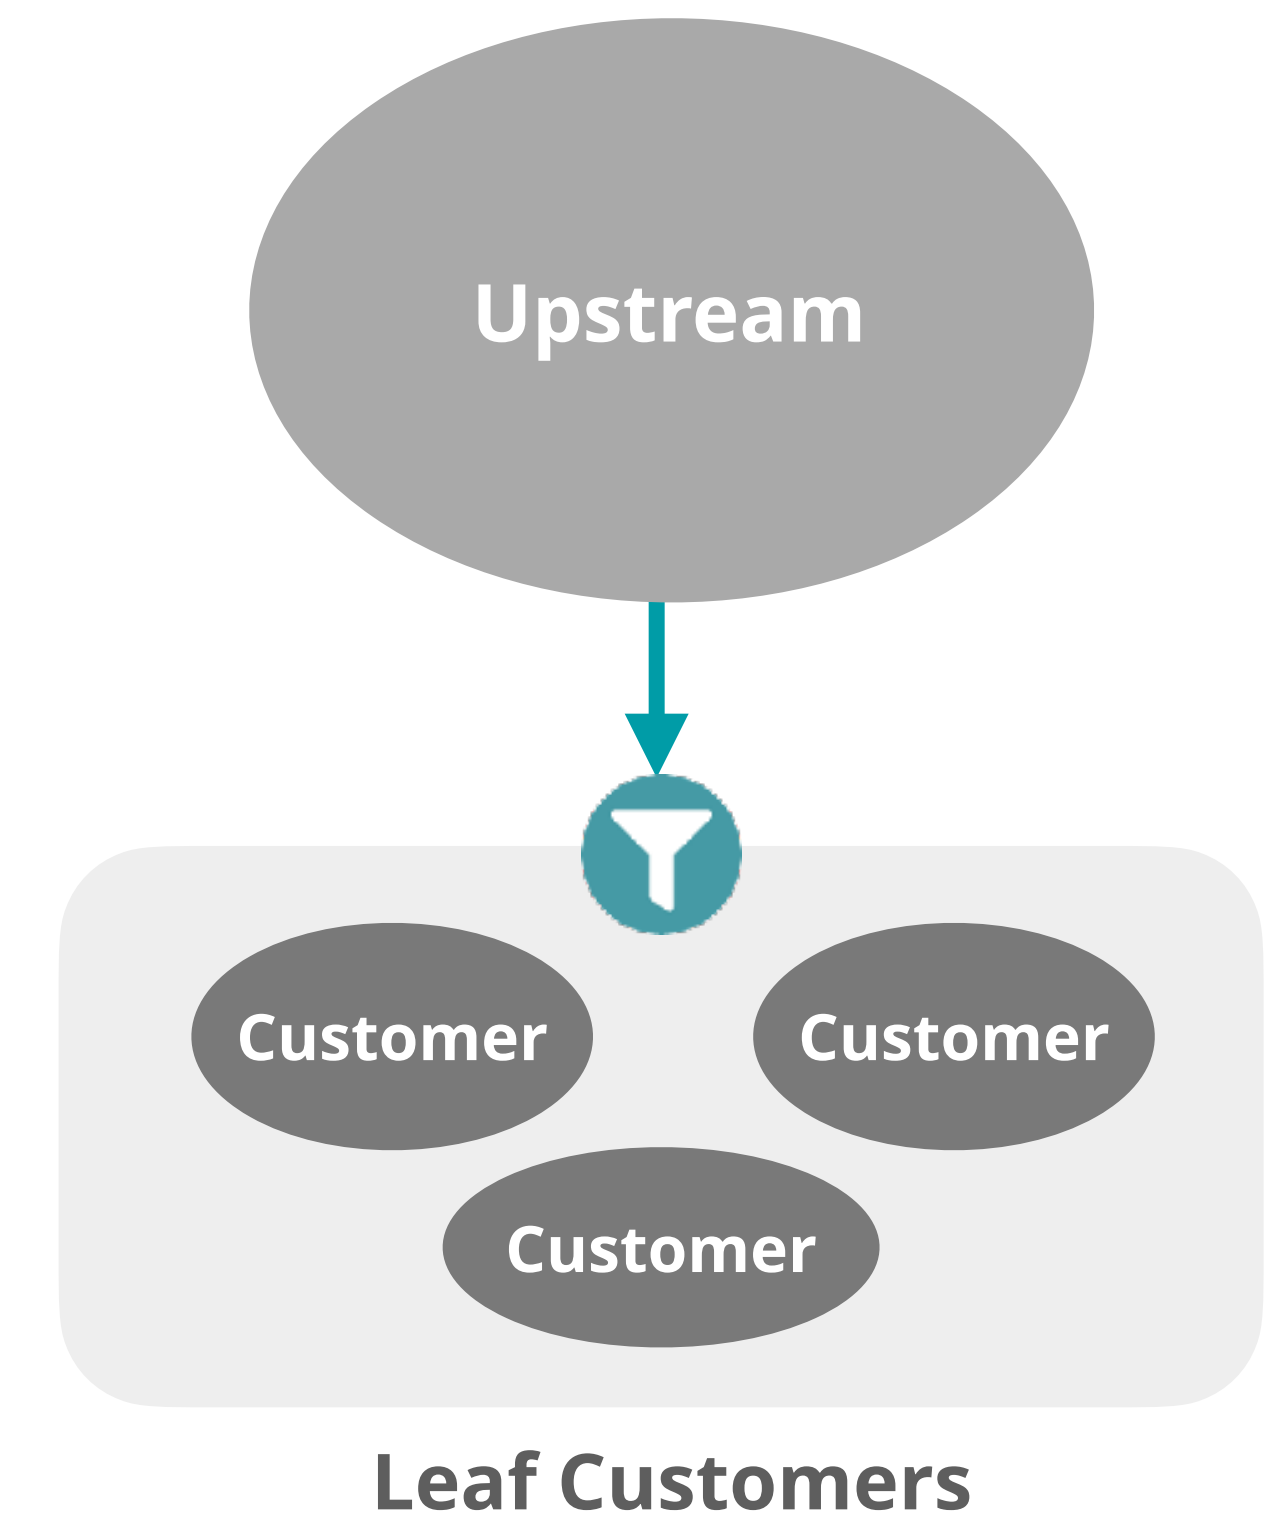
- Filters should be deployed corresponding to the routes requested from upstream
 - If the default route is requested, **accept only the default**
 - If the full route is requested, the followings **should be filtered**:

Prefixes that are not globally routable

Routes that are too specific

Prefixes belonging to the local AS

The default route (depending on whether or not it is requested)





Leaf Customer Network (Outbound)

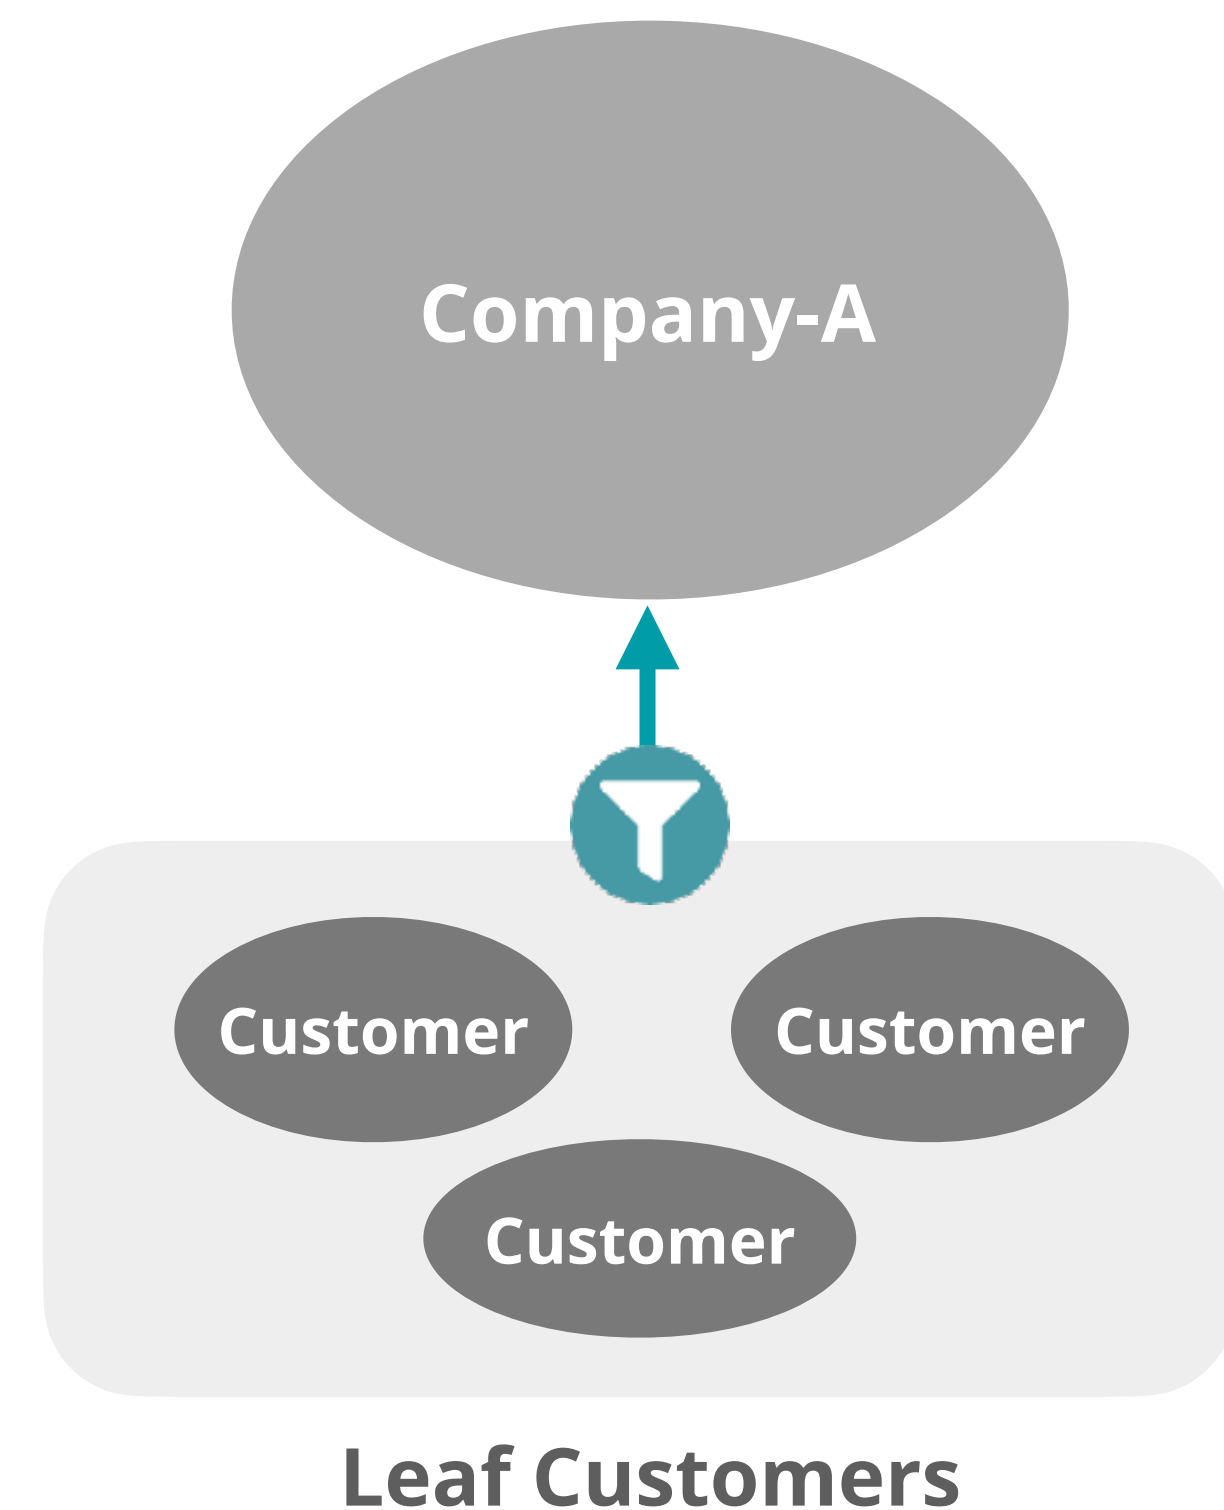
- Outbound policy is very straightforward
 - **Only announce your own prefixes!**
- In order to avoid announcing invalid routes to the upstream, the following **should be filtered**:

Prefixes that are not globally routable

Routes that are too specific

IXP LAN prefixes

The default route





Questions



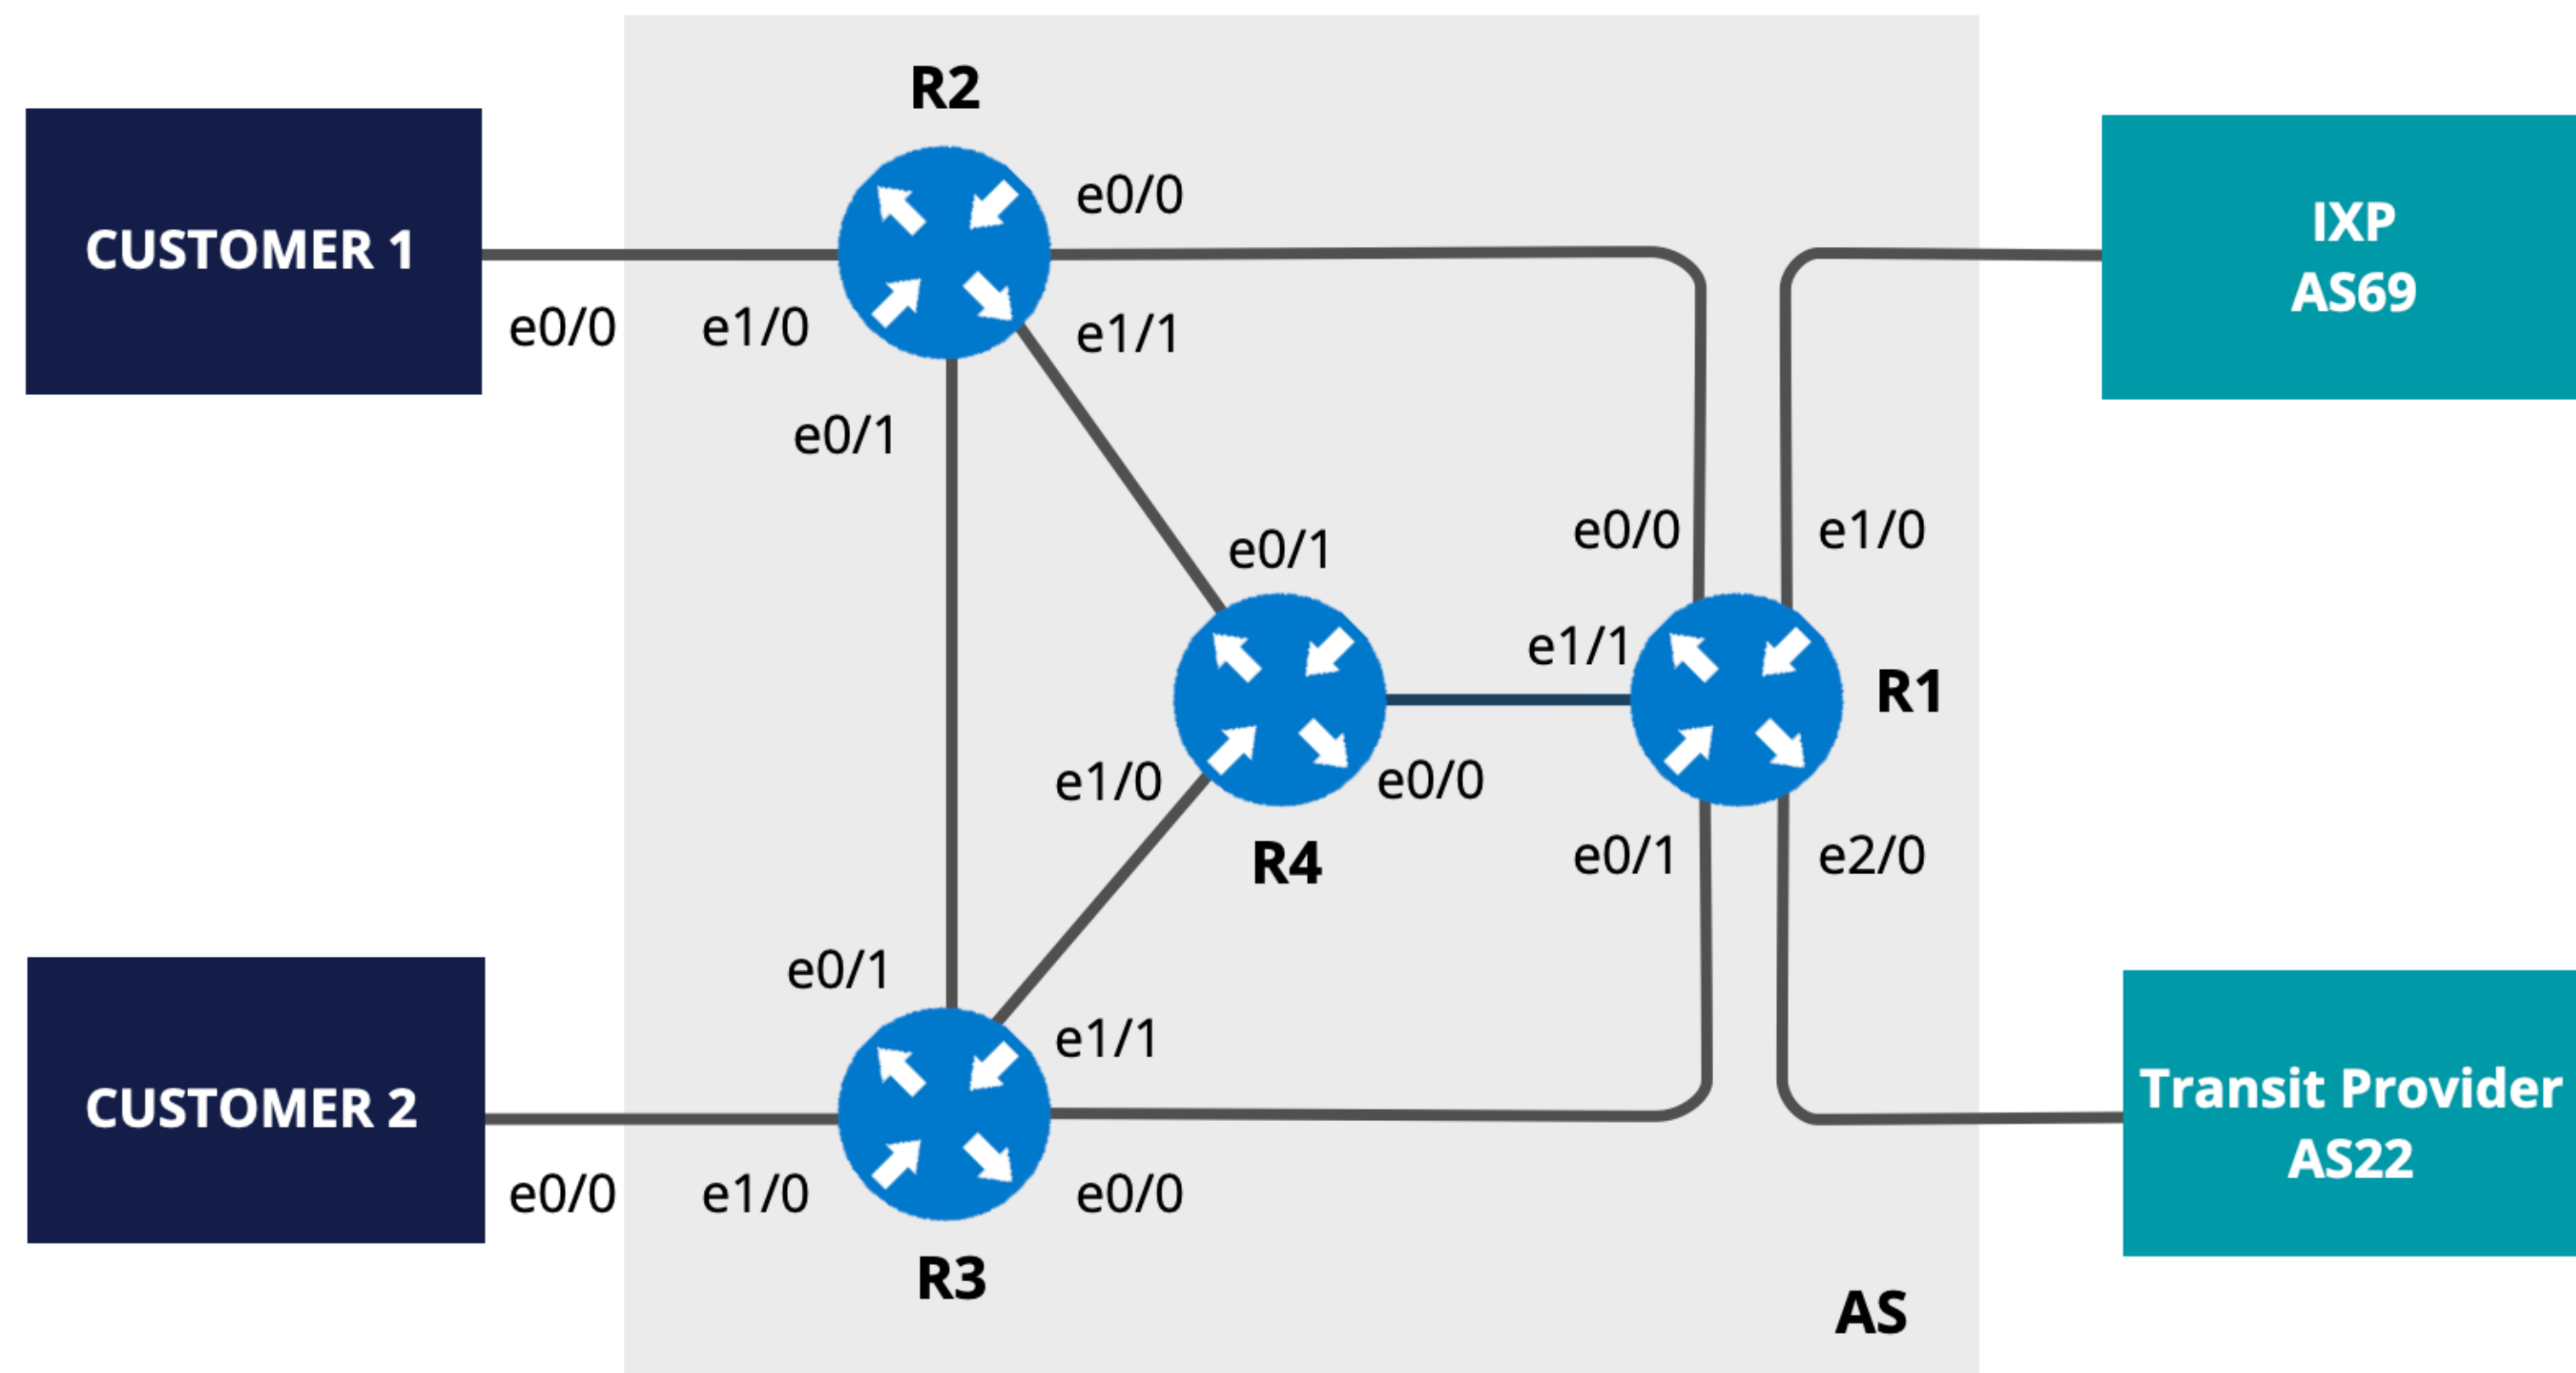


Demo Time!



Let's Filter Too Specific Prefixes

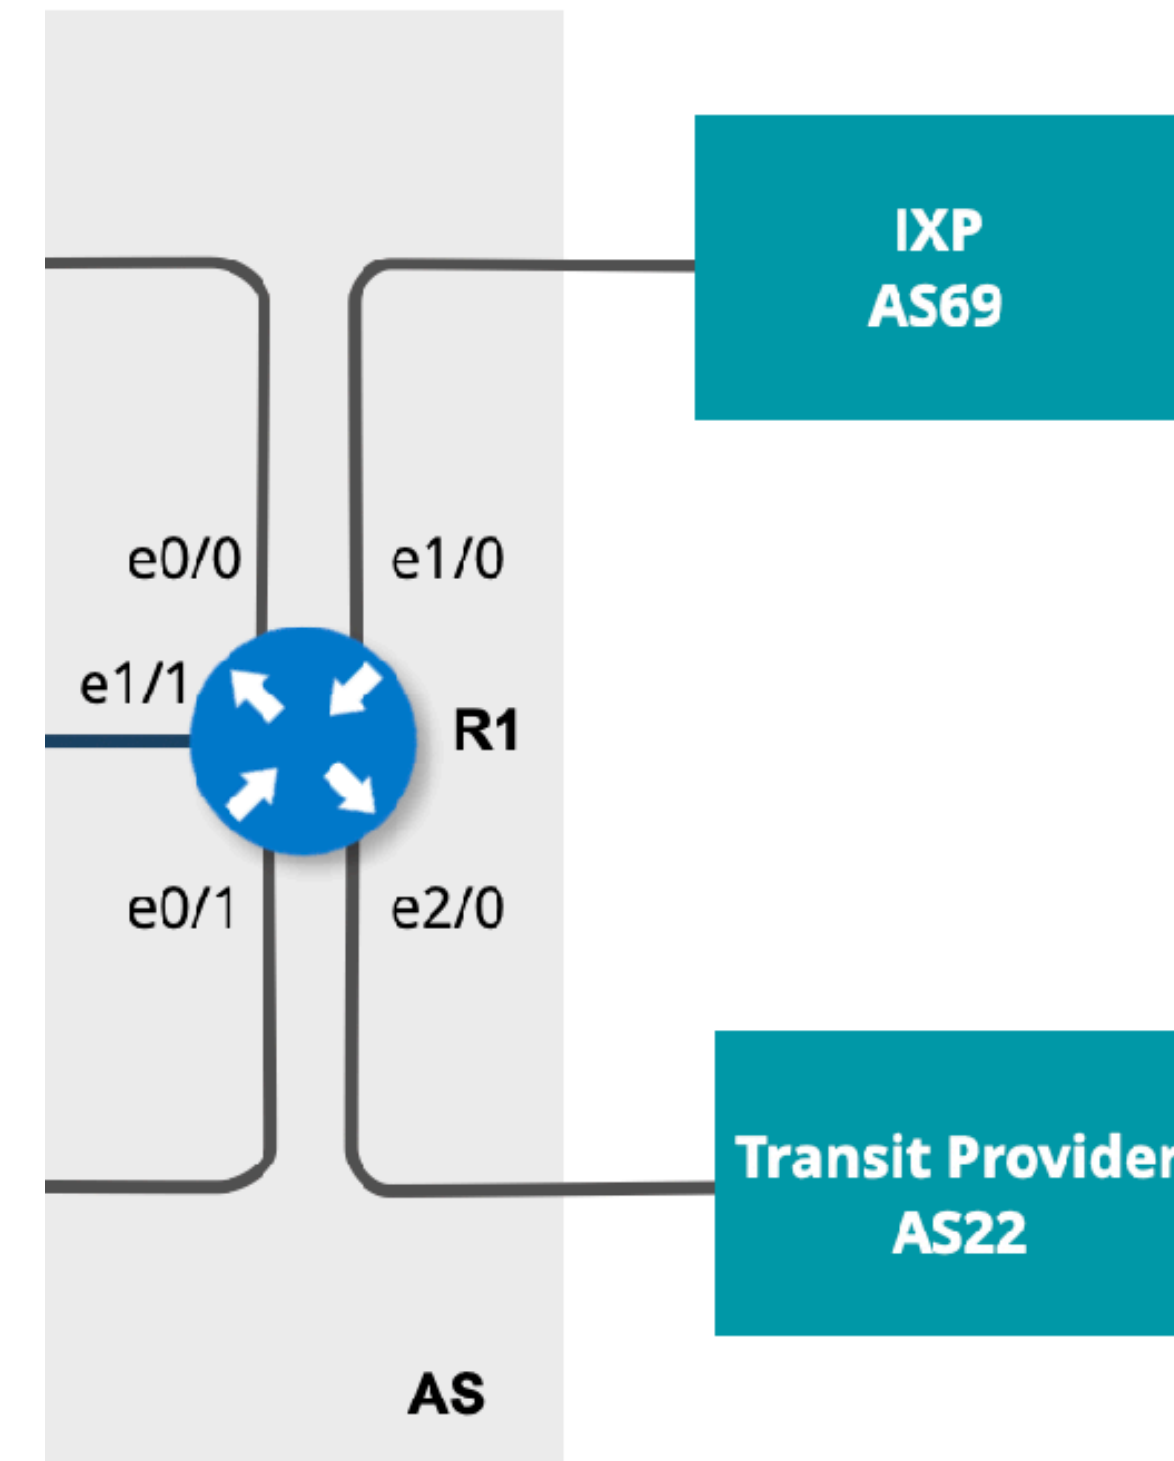
- Smaller prefixes are leaked from a transit provider and an IXP
- And **we need to filter them!**





Preparation (on R1)

- Examine the routing table
- Check for prefixes that are too specific



```
# show ip route bgp
# show ip bgp
# show ipv6 route bgp
# show bgp ipv6
```





```
U1_R1#show ip bgp
BGP table version is 17, local router ID is 172.1.255.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found
```

	Network	Next Hop	Metric	LocPrf	Weight	Path
*	1.1.1.0/24	172.16.0.99		0	69	33 i
*		172.16.0.66		0	69	33 i
*>		10.132.1.1		0	22	33 i
*>		10.132.1.1	0	0	22	i
*	23.226.112.0/20	172.16.0.99		0	69	33 22 i
*		172.16.0.66		0	69	33 22 i
*	65.255.48.0/25	172.16.0.99		0	69	33 22 i
*		172.16.0.66		0	69	33 22 i

--More--



```
U1_R1#sh ip bgp ipv6 unicast
BGP table version is 11, local router ID is 172.1.255.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found
```

	Network	Next Hop	Metric	LocPrf	Weight	Path
*>	2001:65::/64	2001:FF32:0:1::A	0		0 22	i
*>	2001:FF01::/32	::	0		32768	i
*>	2001:FF32::/32	2001:FF32:0:1::A	0		0 22	i
*	2001:FF66::/32	2001:FF69::66			0 69	i
*>		2001:FF69::66	0		0 69	i
*	2001:FF99::/32	2001:FF69::99	0		0 69	i
*>		2001:FF69::99			0 69	i



Filter More Specifics (on R1)

Filtering prefixes that are too specific

```
(config)# ip prefix-list TRANS-IN-V4 seq 10 permit 0.0.0.0/0 le 24
(config)# ip prefix-list IXP-IN-V4 seq 10 permit 0.0.0.0/0 le 24
(config)# ipv6 prefix-list TRANS-IN-V6 seq 10 permit 2000::/3 le 48
(config)# ipv6 prefix-list IXP-IN-V6 seq 10 permit 2000::/3 le 48
```



Filter More Specifics

Apply the inbound policy to the neighbors

```
(config)# router bgp 101
```

```
(config-router)# address-family ipv4
```

```
(config-router-af)# neighbor 10.132.1.1 prefix-list TRANS-IN-V4 in
```

```
(config-router-af)# neighbor 172.16.0.66 prefix-list IXP-IN-V4 in
```

```
(config-router-af)# neighbor 172.16.0.99 prefix-list IXP-IN-V4 in
```

```
(config-router-af)# address-family ipv6
```

```
(config-router-af)# neighbor 2001:ff32:0:01::a prefix-list TRANS-IN-V6 in
```

```
(config-router-af)# neighbor 2001:ff69::66 prefix-list IXP-IN-V6 in
```

```
(config-router-af)# neighbor 2001:ff69::99 prefix-list IXP-IN-V6 in
```



Clear the BGP Sessions (on R1)

```
# clear bgp ipv4 unicast 172.16.0.66 in
# clear bgp ipv4 unicast 172.16.0.99 in
# clear bgp ipv4 unicast 10.132.1.1 in
# clear bgp ipv6 unicast 2001:ff69::66 in
# clear bgp ipv6 unicast 2001:ff69::99 in
# clear bgp ipv6 unicast 2001:ff32:0:01::a in
```


Verify



Check BGP and the routing table

```
# show bgp ipv4 unicast  
# show bgp ipv6 unicast  
# show ip route bgp | include /25  
# show ipv6 route bgp | include /64
```



```
U1_R1#show bgp ipv4 unicast | i /25
U1_R1#
U1_R1#show bgp ipv6 unicast | i /64
U1_R1#
U1_R1#
U1_R1#show ip route bgp | include /25
U1_R1#
U1_R1#
U1_R1#show ipv6 route bgp | include /64
```



Questions



We want your feedback!



What did you think about this session? Take our survey at:

<https://www.ripe.net/feedback/bgp1/>






Learn something new today!
academy.ripe.net





RIPE NCC Certified Professionals



<https://getcertified.ripe.net/>



What's next in BGP?




Webinars

Attend another webinar live wherever you are.

- ❖ BGP Filtering (1 hr)
- ❖ Deploying RPKI (2 hrs)
- ❖ Introduction to RPKI (1 hr)
- ❖ Internet Routing Registry (1 hr)



For more info click the link below



learning.ripe.net



Face-to-face

Meet us at a location near you for a training session delivered in person.

- ❖ BGP Routing Security (8.5 hrs)



E-learning

Learn at your own pace at our online Academy.

- ❖ BGP Security (10 hrs)



For more info click the link below



academy.ripe.net



Examinations

Learnt everything you needed? Get certified!

- ❖ BGP Security Associate



For more info click the link below




getcertified.ripe.net

Have more questions? Ask us!

academy@ripe.net



Ěnn	Соңы	An Críoch	پایان	Ende	Y Diwedd
Vége	Endir	Finvezh	վերջ	Кінець	Koniec
Son	დასასრული	הסוף	Tmíem	Liðugt	Finis
Lõpp	Amaia	Loppu	Slutt	Kraj	
Kraj	Sfârșit	النهاية	Конец	Konec	Fund
Fine	Fin	Einde	Fí	Край	Beigas
					Τέλος
Fim	Slut				Pabaiga



Copyright Statement

[...]

The RIPE NCC Materials may be used for **private purposes, for public non-commercial purpose, for research, for educational or demonstration purposes**, or if the materials in question specifically state that use of the material is permissible, and provided the RIPE NCC Materials are not modified and are properly identified as RIPE NCC documents. Unless authorised by the RIPE NCC in writing, any use of the RIPE NCC Materials for advertising or marketing purposes is strictly forbidden and may be prosecuted. The RIPE NCC should be notified of any such activities or suspicions thereof.

[...]

Find the full copyright statement here:

<https://www.ripe.net/about-us/legal/copyright-statement>

