Large Networks Monitoring

Luca Deri <deri@ntop.org>, @lucaderi

ntop

Goal of This Presentation

- Explain challenges in monitoring large networks.
- Introduce open source AS traffic observability.
- Present the integration work on BGP/BMP.
- Show you that you will now be able to monitor your AS traffic without costly monthly subscriptions to cloud-based products.

Do I "Own" the Monitored Traffic ?

- Yes
  You are monitoring your services (e.g. email. Web etc) so the traffic hitting your servers belongs to you. You can do DPI and store detailed IP information.
  Example: service providers, company, individuals.
- No
  I provide Internet connectivity to my community and my customers. My goal is to keep the network healthy, I can't store/ visualize detailed information.
  Example: IXP Network Operators.

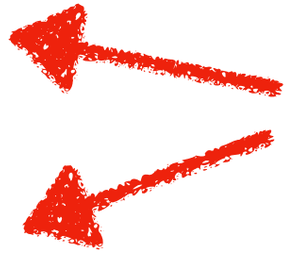

Our use case

- Mixed
  You have a large enterprise where only on selected networks you do DPI and the rest is flow-based.

What's Inside a Flow ?

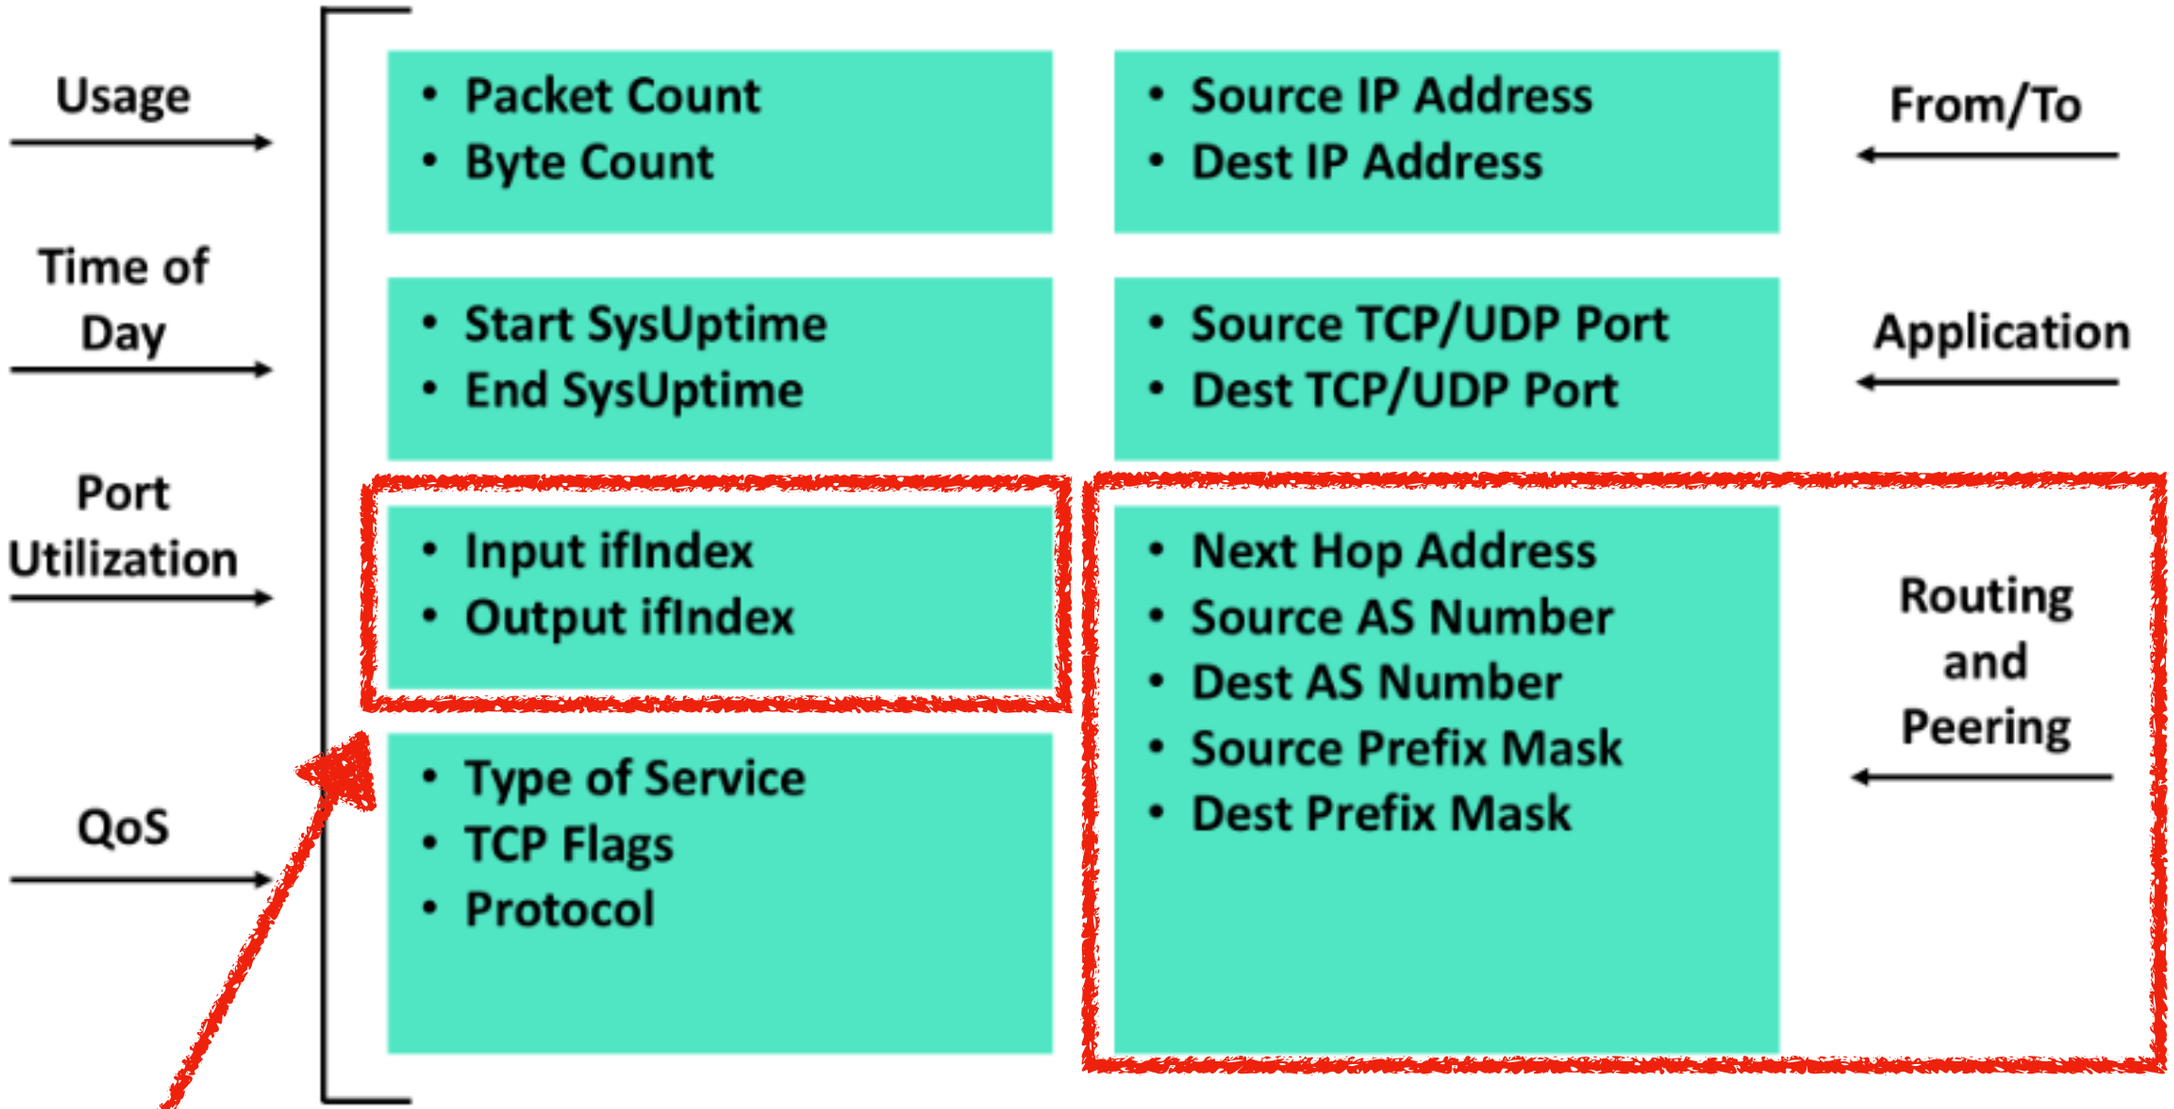

Flow Exporter Information

Flow Collection: nProbe and ntopng

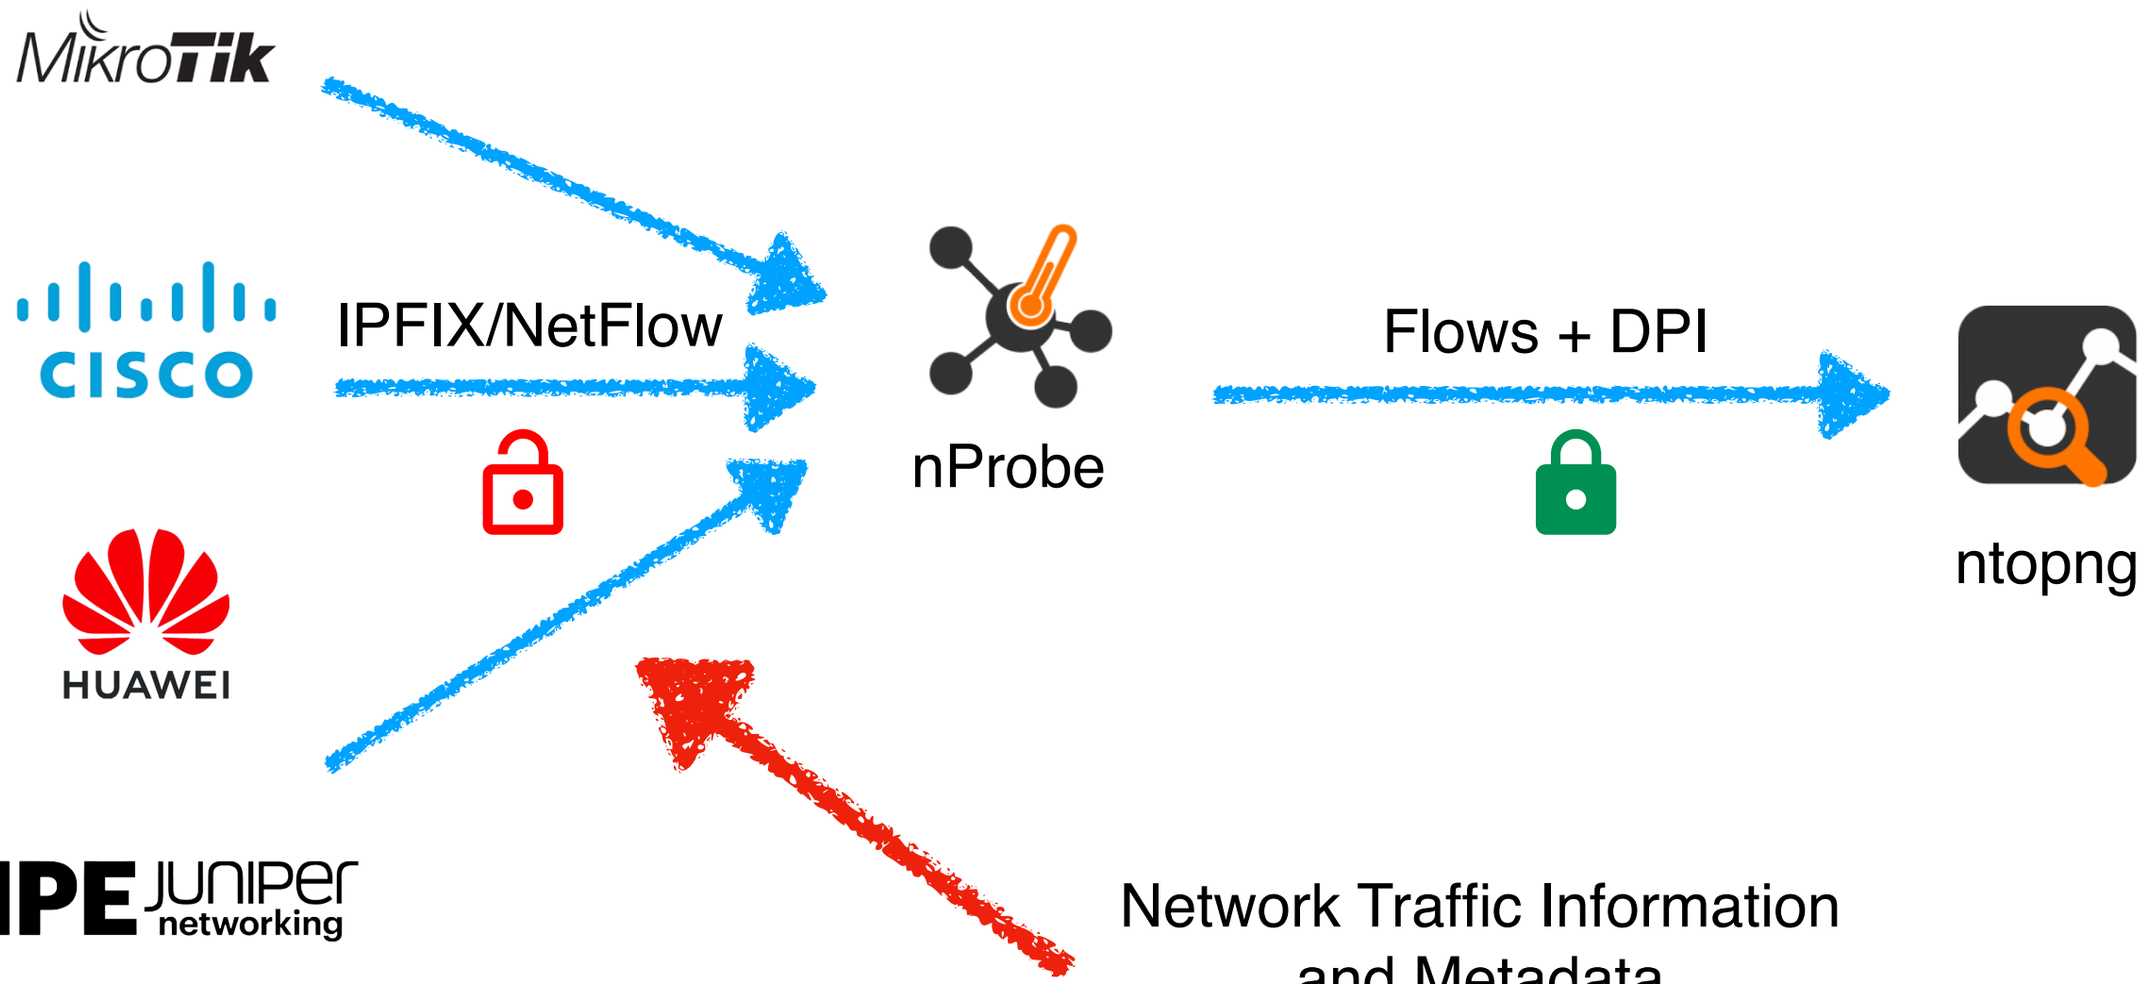

Flow Collection in Large Networks

(Possible) Flow Duplication

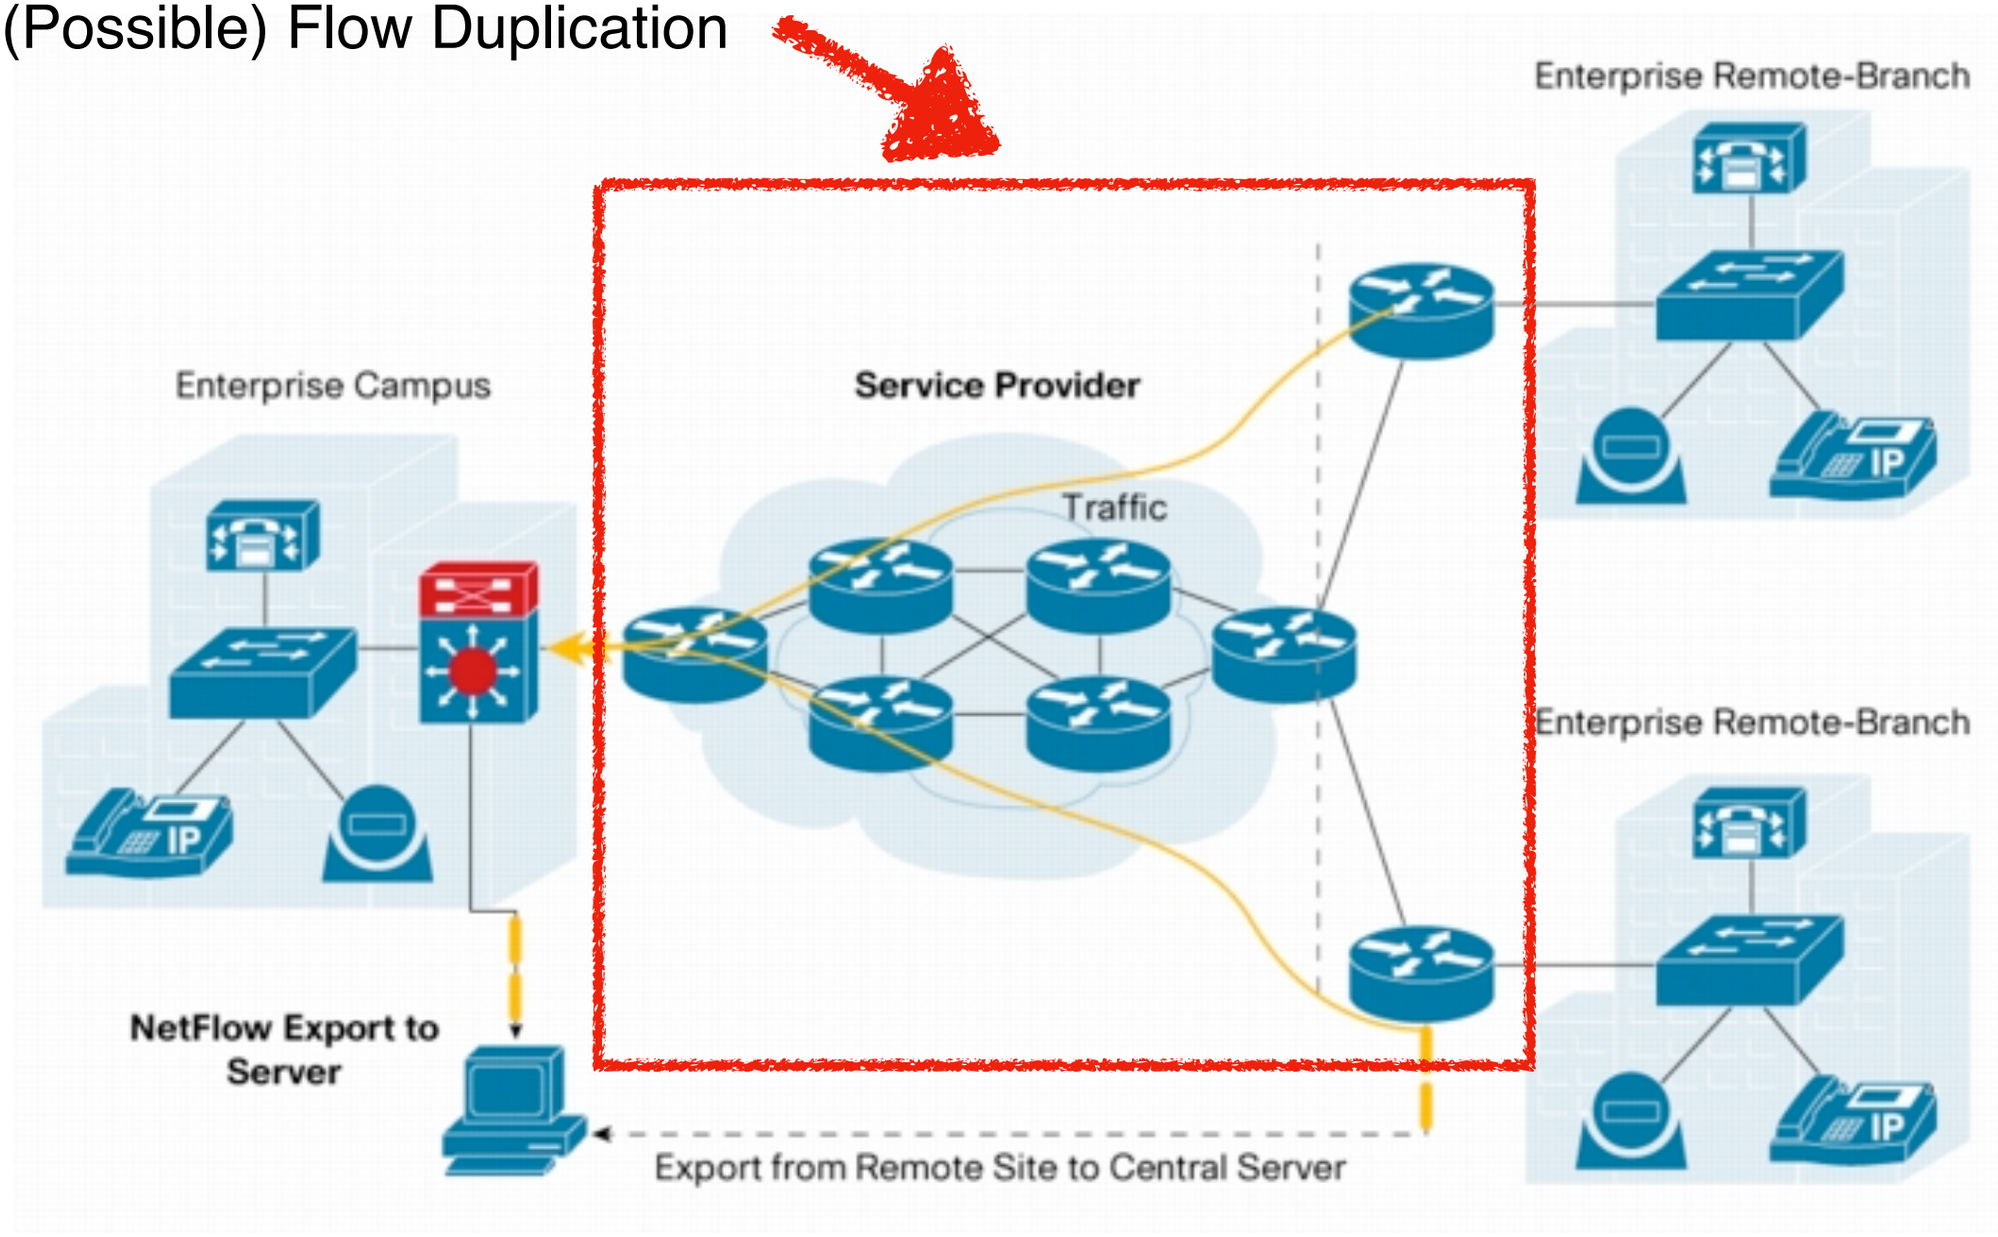

Flow Duplication: Good or Bad? [1/3]

- What is duplicated flow?
  - Same time window
  - Same 5-tuple (proto, IP/Port src/dst)
  - Different flow exporter.
- Flow duplication can create issues when accounting traffic as the same flow can be accounted once, twice... depending on the network.

Flow Duplication: Good or Bad? [2/3]

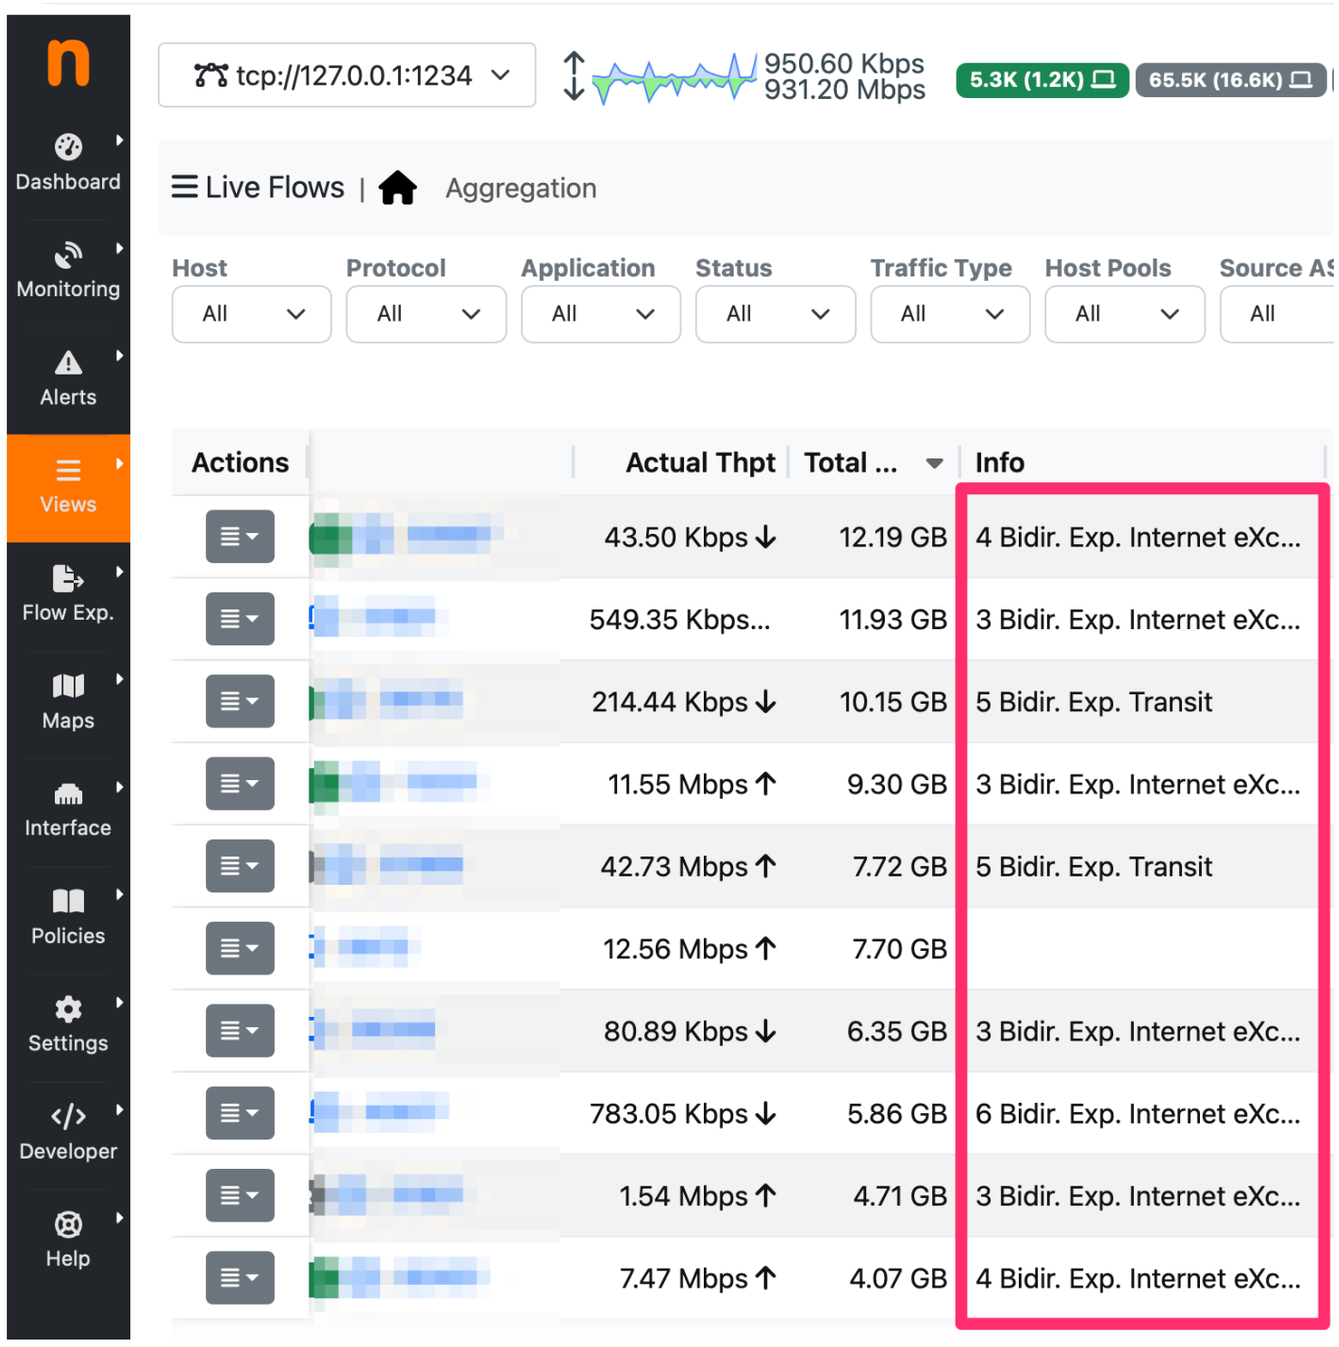

Flow Duplication: Good or Bad? [3/3]

Exporters Information	Flow Exporter / Next Hop	Input Interface / Output Interface
	[redacted] [10.255.255.202] ➜ 10.20.2.10	VLAN1001-edge-01.core-A ➜ VLAN2002-ac101.ac (internal_interface)
	[redacted] [10.255.255.202] ➜ 10.20.2.25 Return Path	VLAN2002-ac101.ac (internal_interface) ➜ VLAN1001-edge-01.core-A
	[redacted] [10.255.255.31] ➜ 10.20.2.9 Return Path	pppoe-B-KINEMAXMON ➜ VLAN2002-edge-02.core (internal_interface)
	[redacted] [10.255.255.31] ➜ 0.0.0.0	VLAN2002-edge-02.core (internal_interface) ➜ pppoe-B-KINEMAXMON
	[redacted] [10.255.255.201] ➜ 10.20.2.26	VLAN100-IPT-FT (transit) ➜ VLAN1001-edge-02.core-A

Flow Deduplication: nProbe vs ntopng

ntopng

ZMQ Statistics				
Collected Flows	817,147,795 [12.7 Kfps] ↑	**Discarded Flows (by ntopng)**	400 [0.0 %] —	
Collected ZMQ Messages	136,254,424 ↑	**Dropped ZMQ Messages**	1,519 [0.0 %] —	**Avg Flows/Msg**: 6
Interface Updates	18,127 ↑	**Counter Updates**	0 —	
Traffic Statistics				
Traffic Anomalies	**Local Hosts Anomalies**	22,241 —	**Remote Hosts Anomalies**	23,117 ↑
Deduplicated Flows	407,756,411 [49.9 %] ↑			

nProbe

```
--flow-collection-deduplication     | Discards duplicated flows (i.e. flows with the same key sent by
                                    | two different flow exporter IP)
```

Best Practices

- Most people can deploy monitoring devices at the edge of the network:
  - No duplication
  - Visibility of north/south traffic.
- You can also monitor internal routers if you want to study the traffic trajectory but this can put significant load on large networks.
- In short: it depends on your monitoring goals and your use case.

Handling Exporters IPs

- Flow exporters often export flows from multiple interfaces hence with multiple IP addresses.
- In order to bind all IPs to the same exporter, we have enhanced SNMP to poll IPv4/v6 addresses.

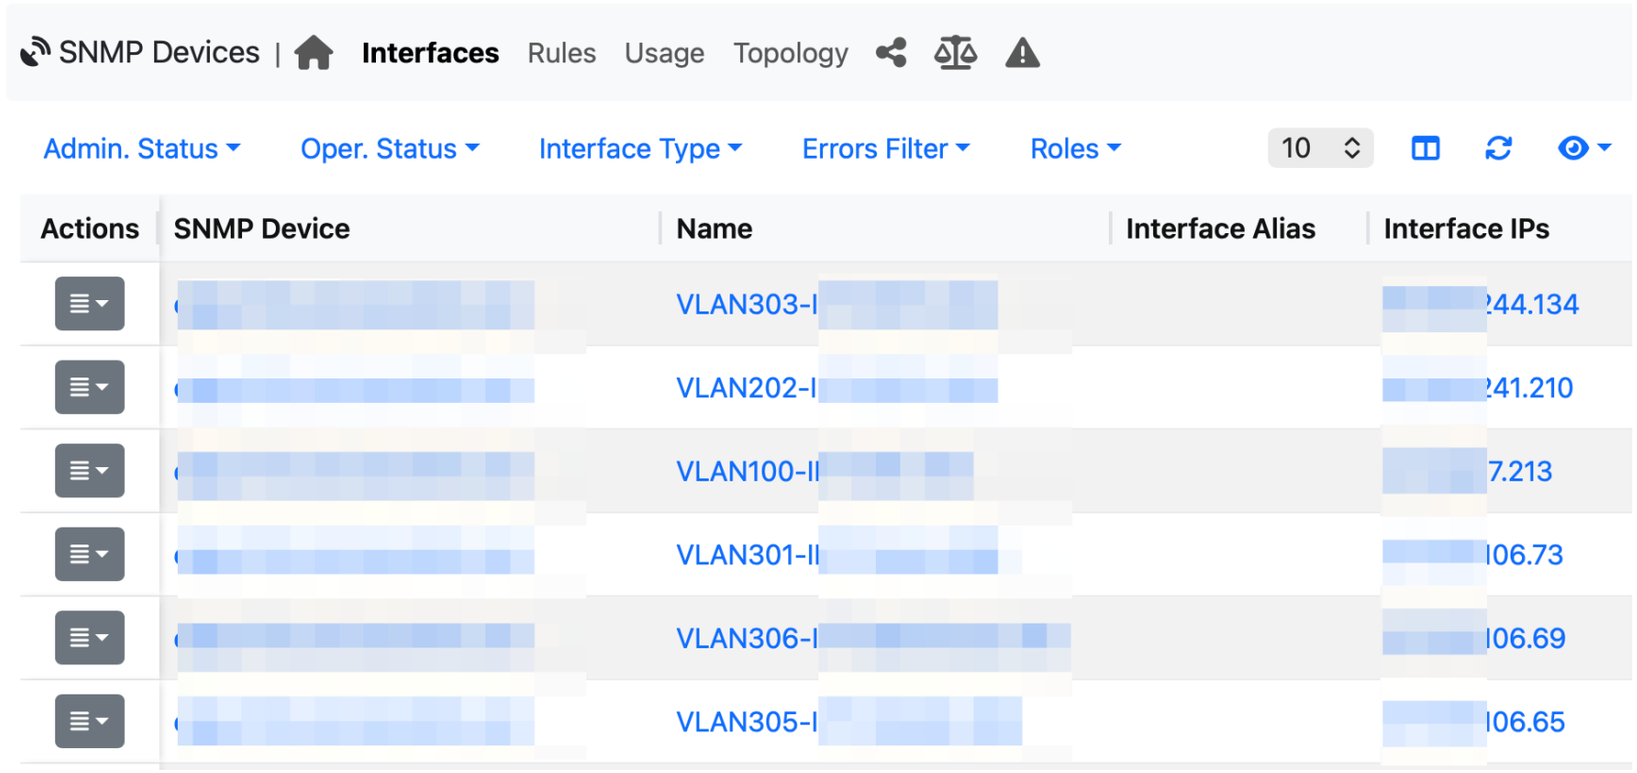

nProbe vs Exporters

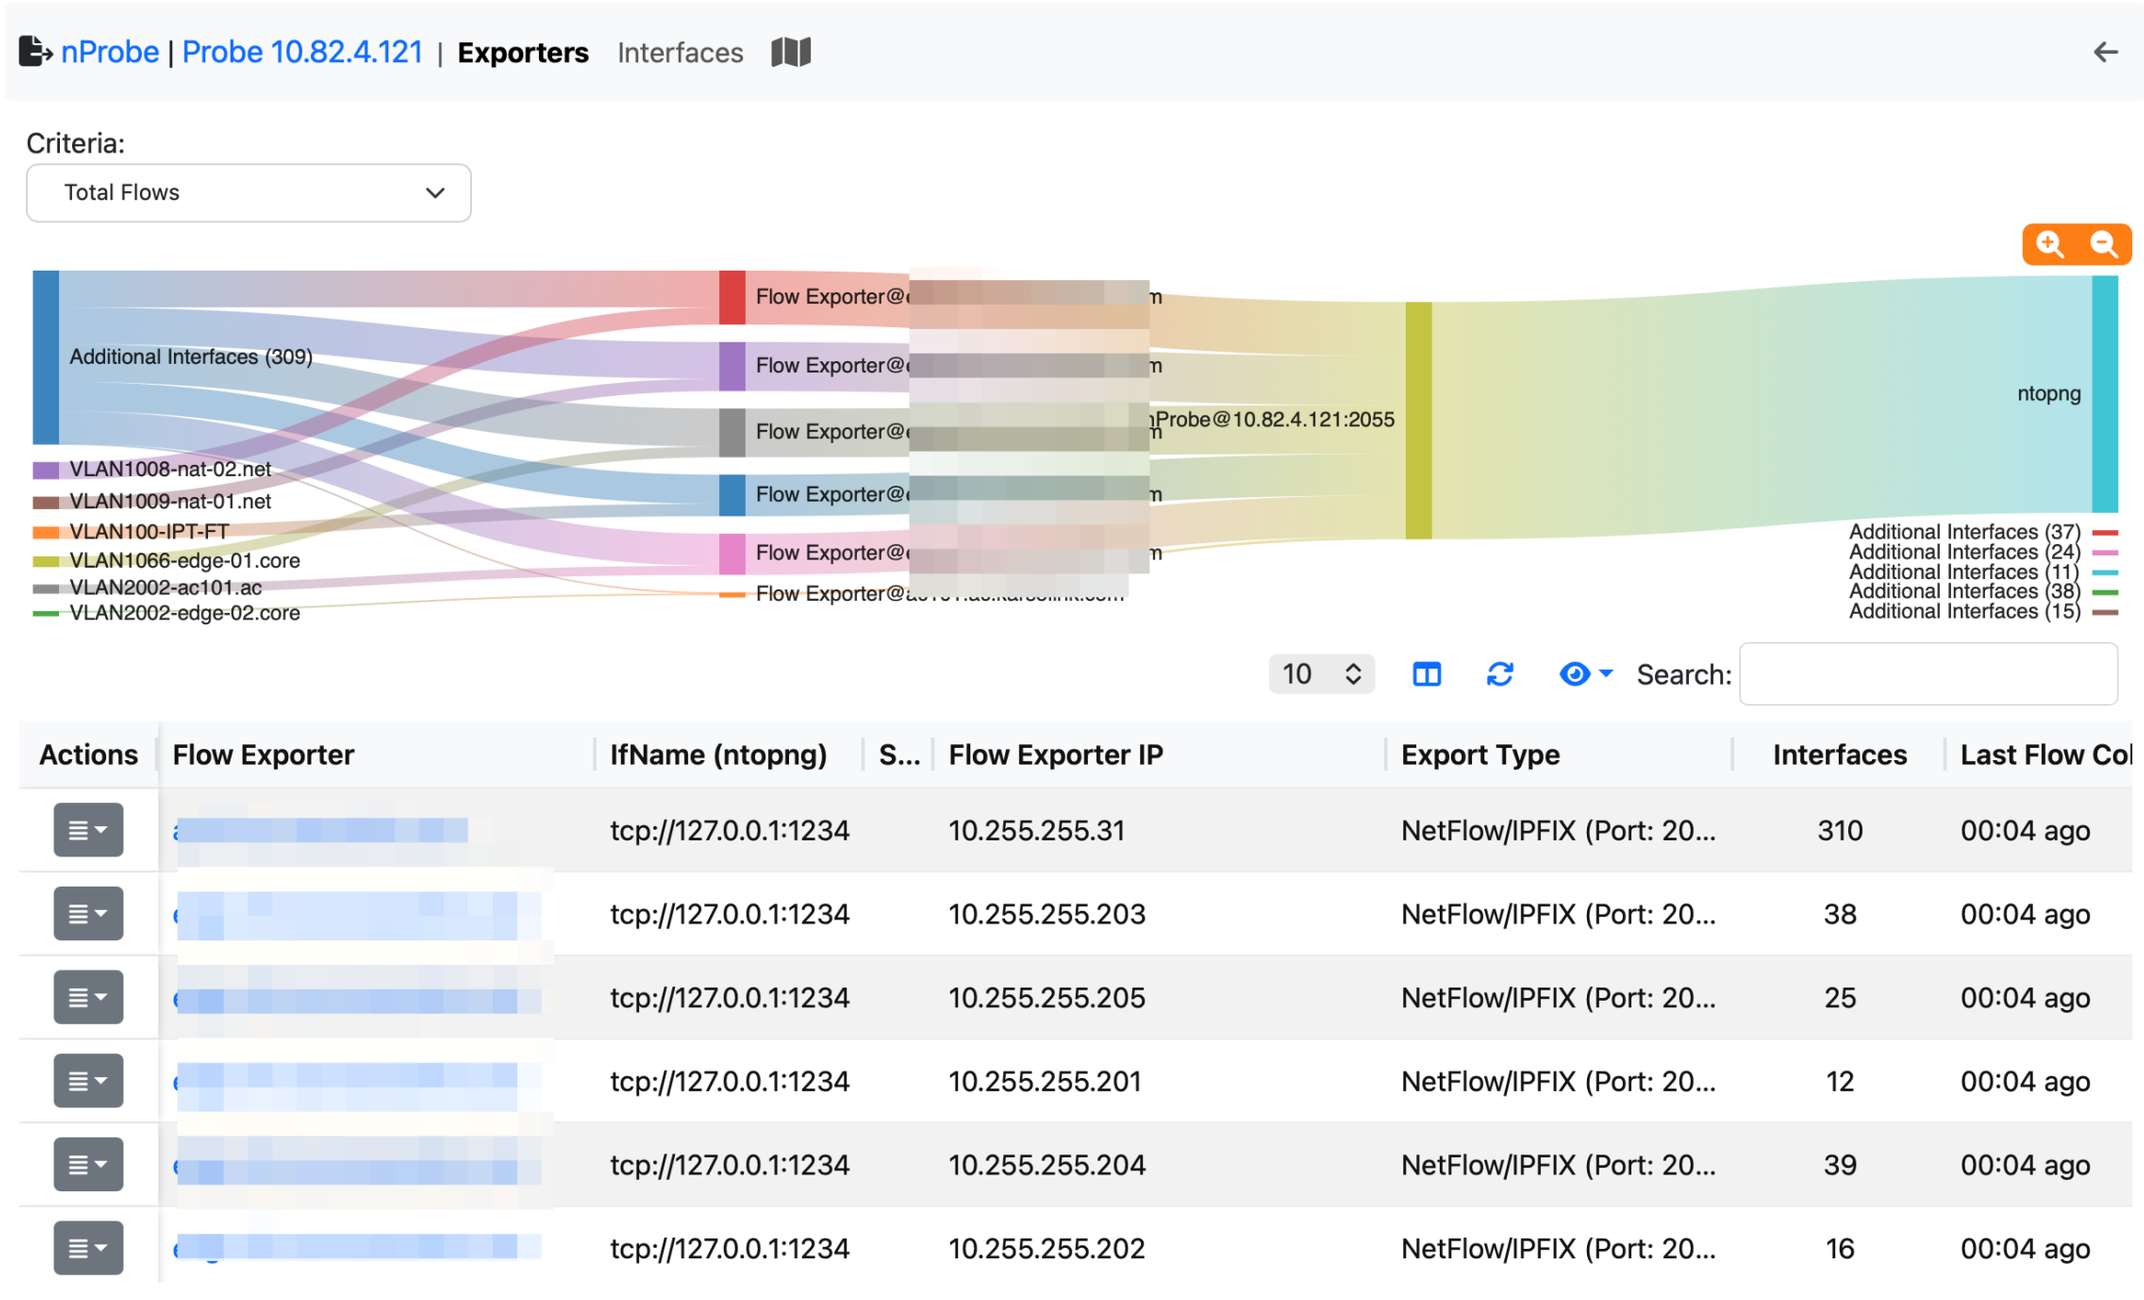

Actions	Flow Exporter	IfName (ntopng)	S...	Flow Exporter IP	Export Type	Interfaces	Last Flow Col
		tcp://127.0.0.1:1234		10.255.255.31	NetFlow/IPFIX (Port: 20...	310	00:04 ago
		tcp://127.0.0.1:1234		10.255.255.203	NetFlow/IPFIX (Port: 20...	38	00:04 ago
		tcp://127.0.0.1:1234		10.255.255.205	NetFlow/IPFIX (Port: 20...	25	00:04 ago
		tcp://127.0.0.1:1234		10.255.255.201	NetFlow/IPFIX (Port: 20...	12	00:04 ago
		tcp://127.0.0.1:1234		10.255.255.204	NetFlow/IPFIX (Port: 20...	39	00:04 ago
		tcp://127.0.0.1:1234		10.255.255.202	NetFlow/IPFIX (Port: 20...	16	00:04 ago

Combining Exporters with Flow NextHop

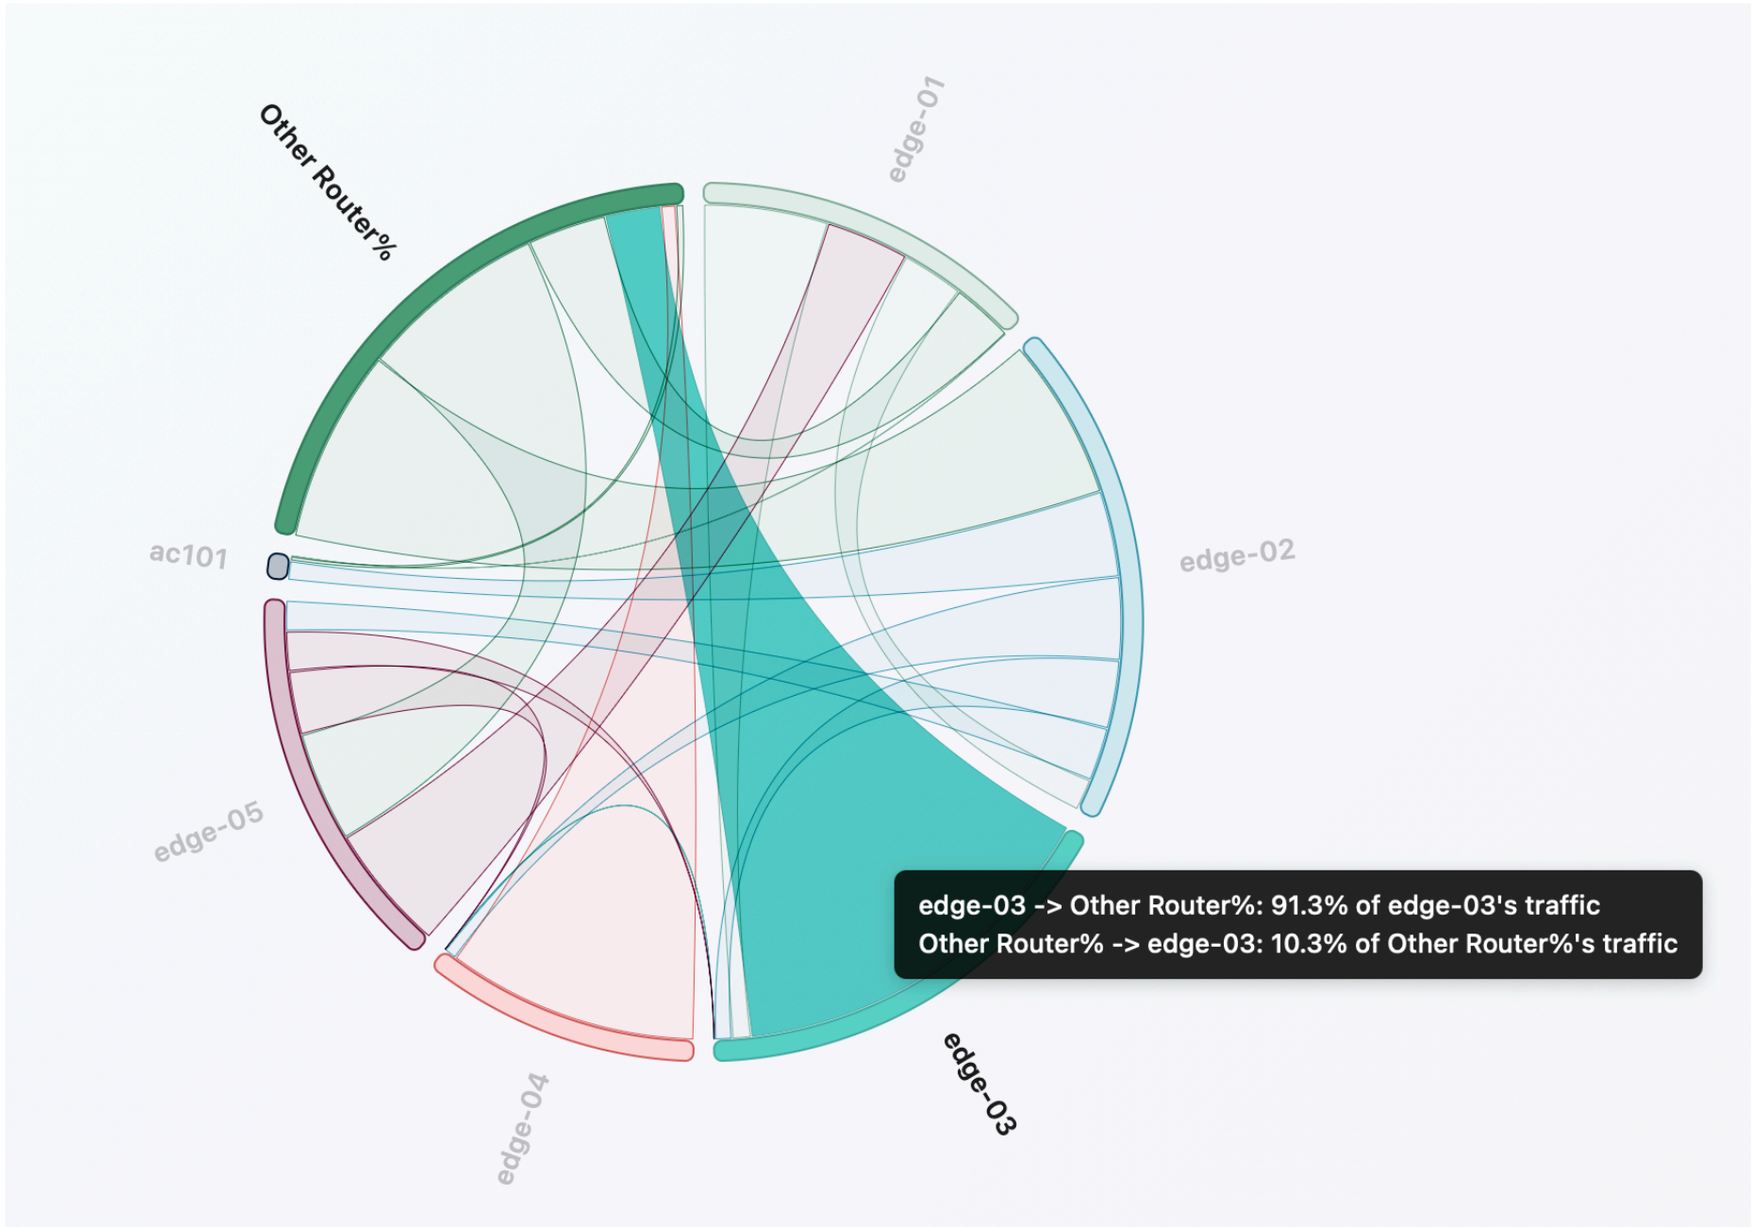

Networks vs Sites [1/2]

- On large networks, IPs might be physically deployed across multiple locations.
- Sites allow you to cluster IPs across locations.

Networks | Sites

10 + Search:

Actions	Site	Description	Location
	Default		
	Milano	Milan site	45.138464° N, 10.260198° E
	Padova	Università Padova	45.739169° N, 11.380144° E
	Trieste	Main site	45.919325° N, 13.487195° E

Networks vs Sites [2/2]

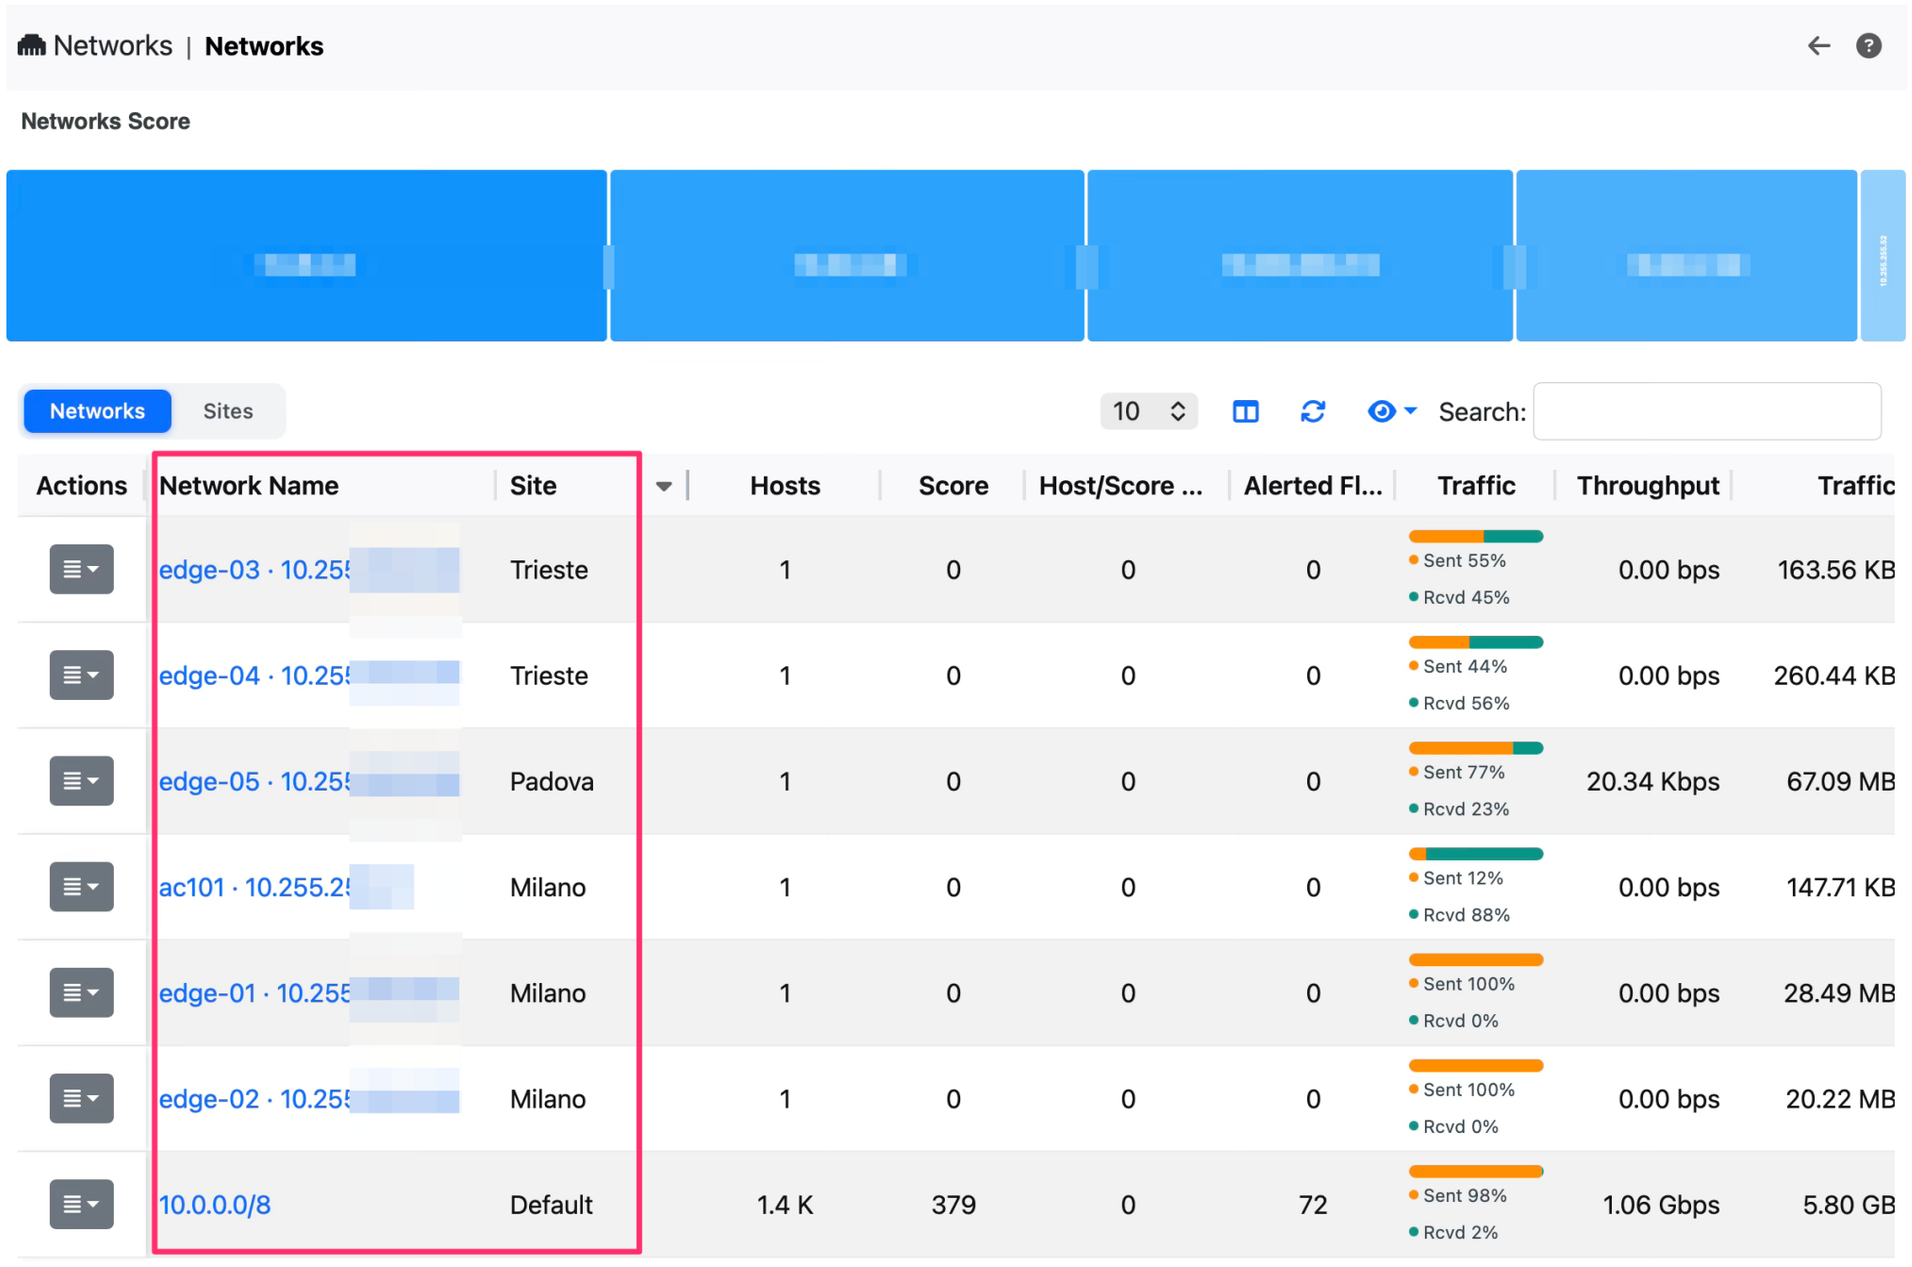

Sites Traffic Distribution

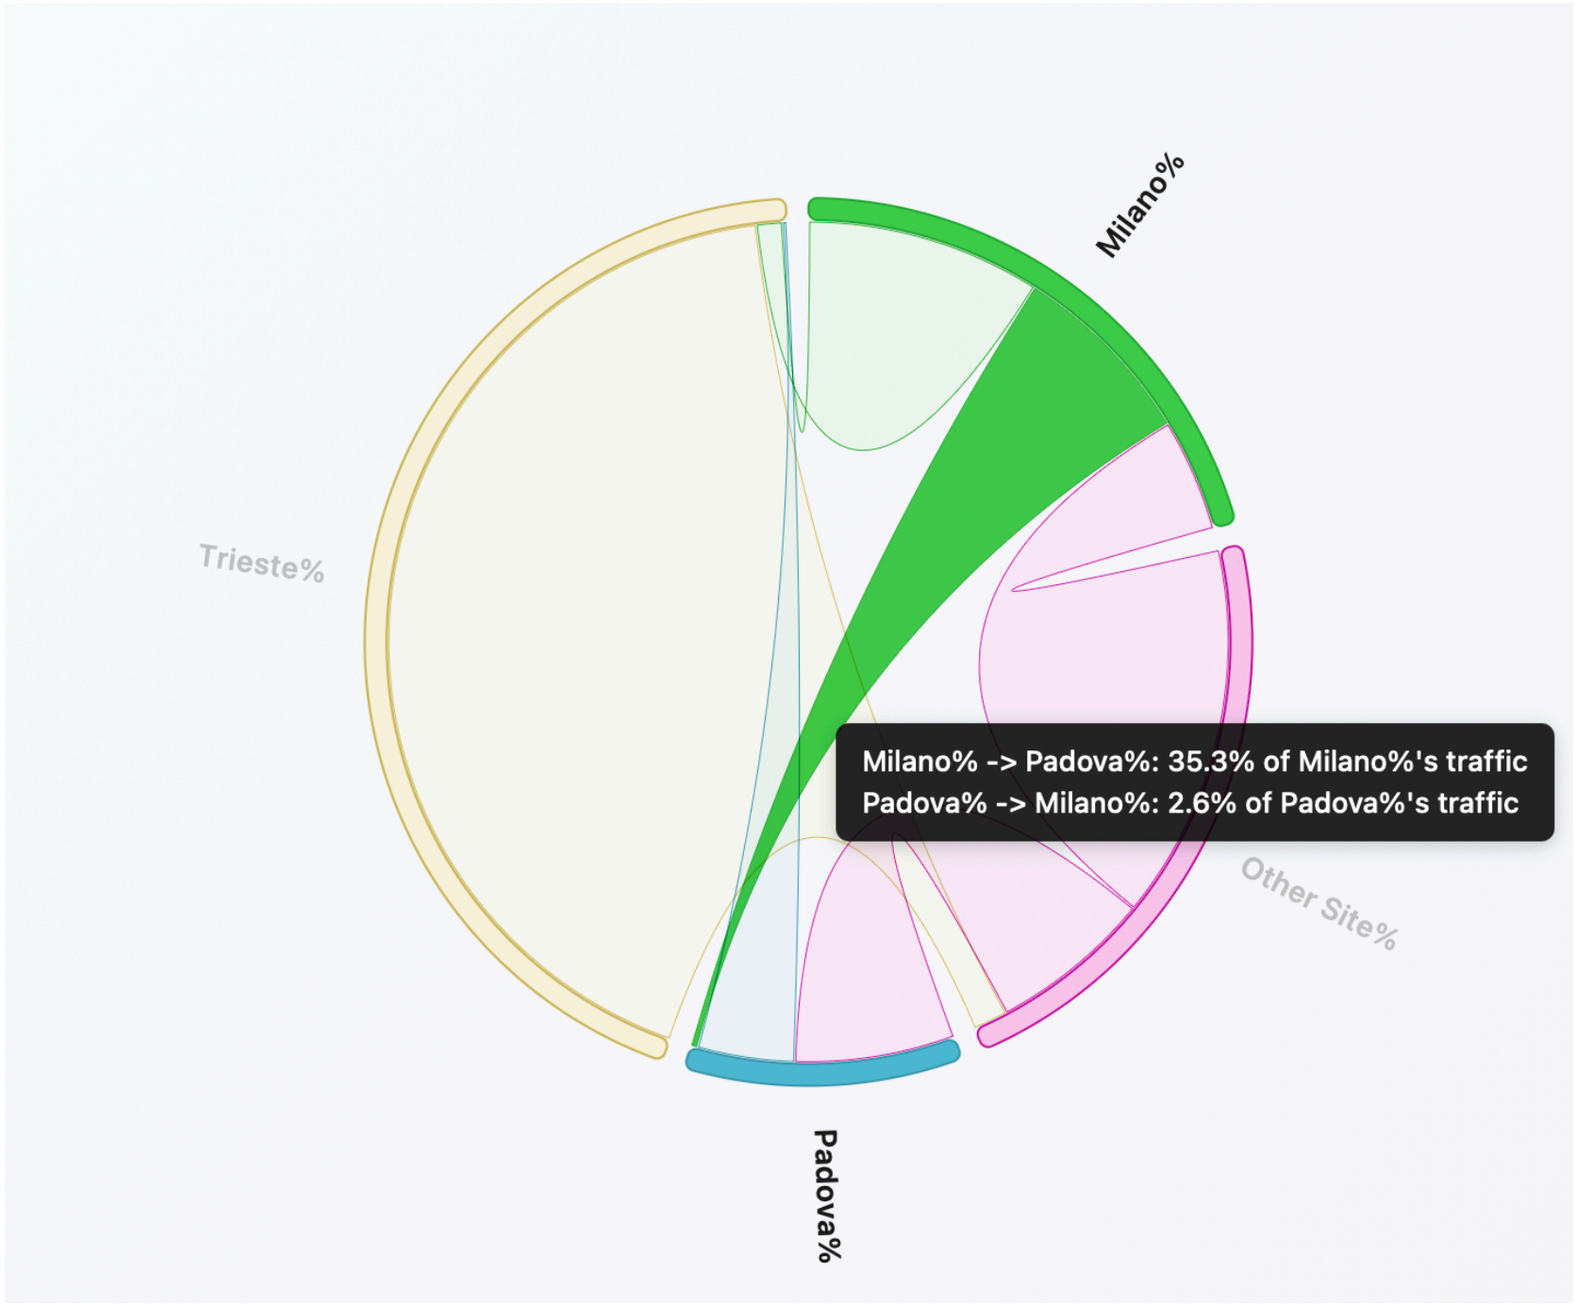

Binding Interfaces to Roles

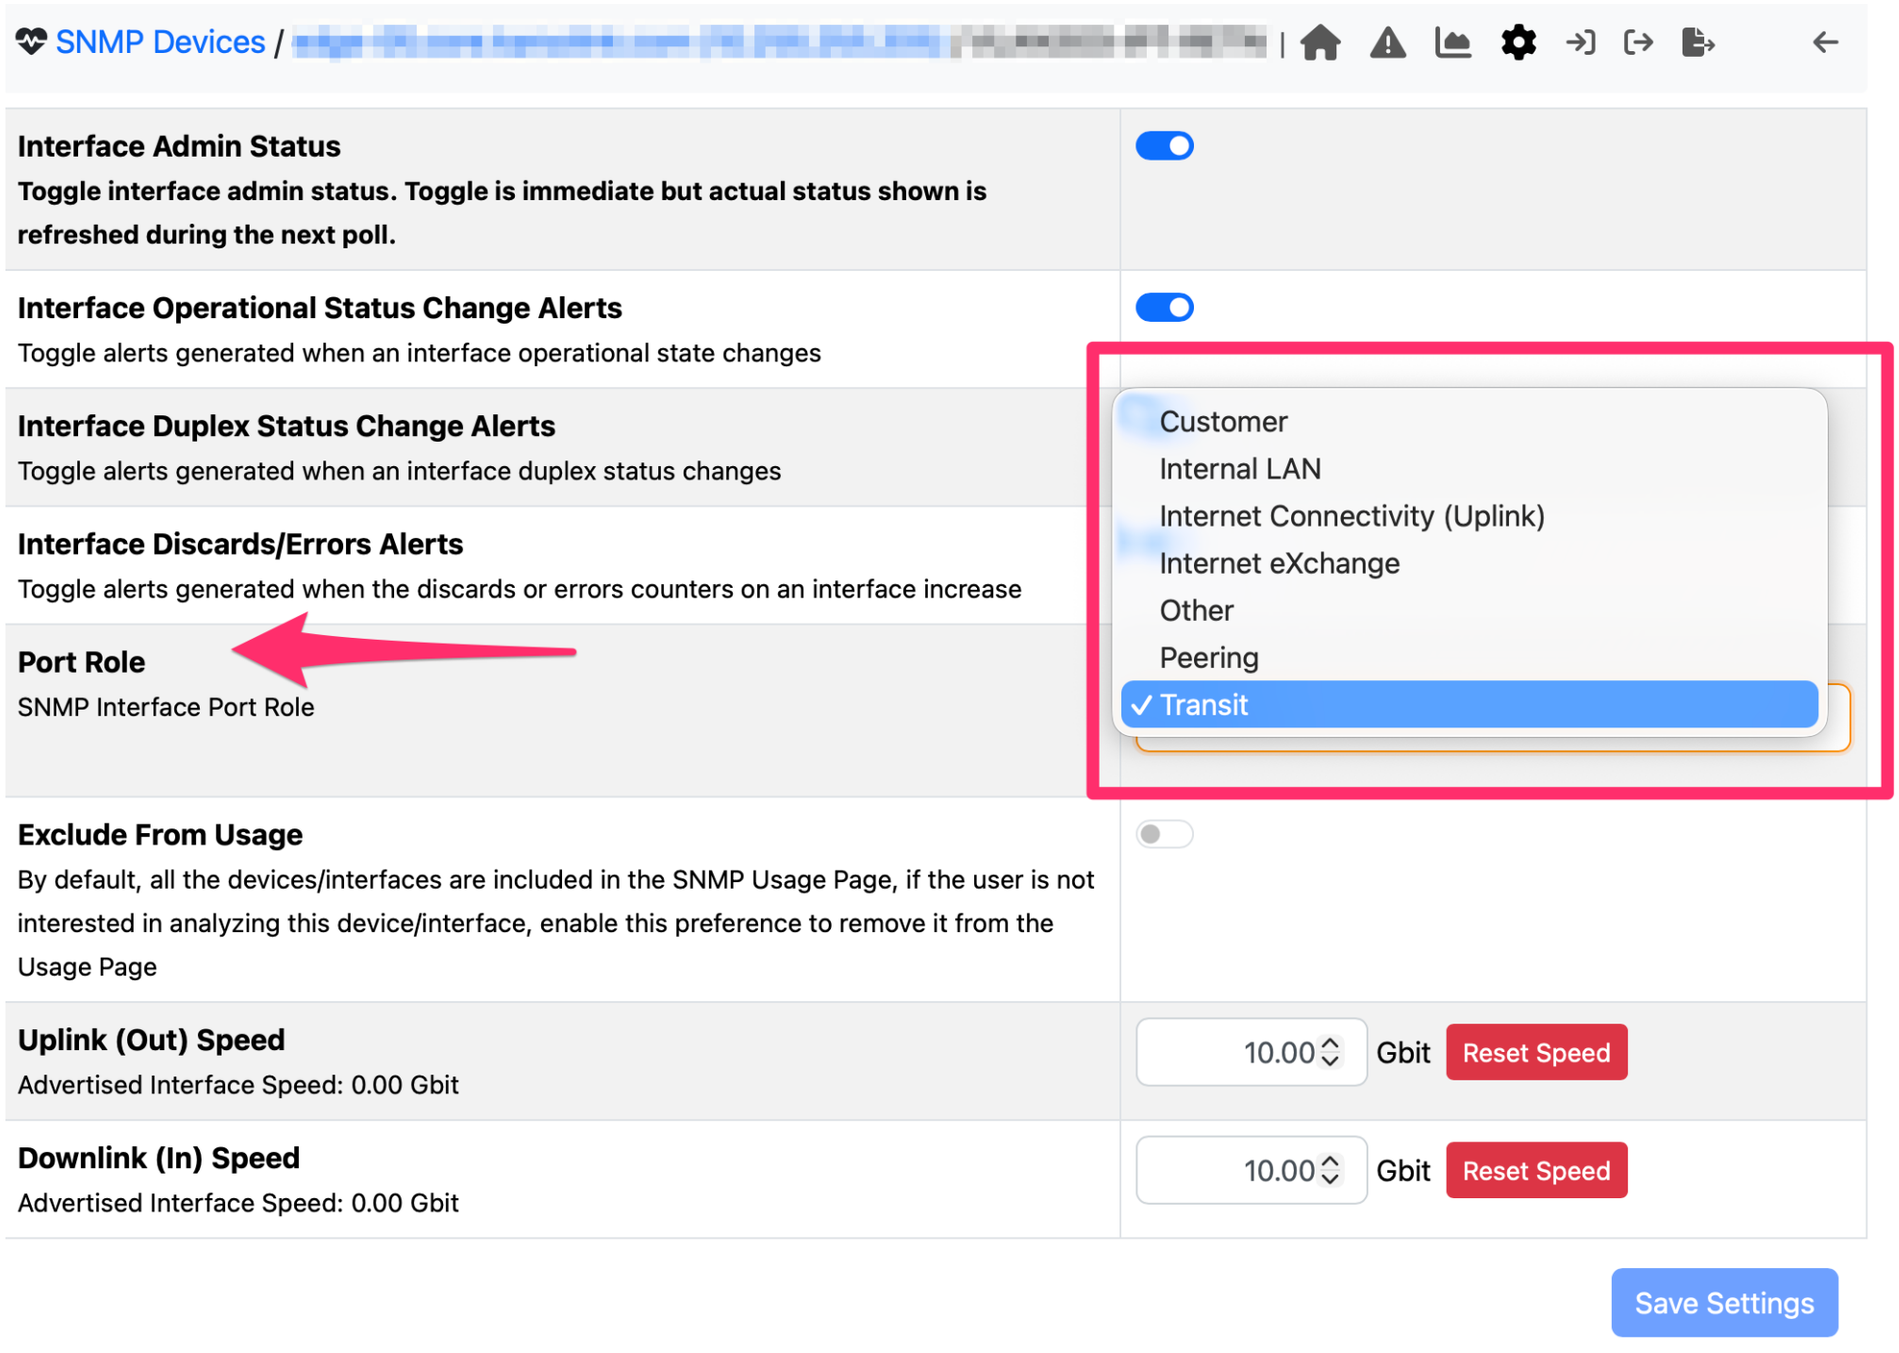

Binding Exporters to Sites

- Step 1: define a site name
- Step 2: bind the exporter network to the site

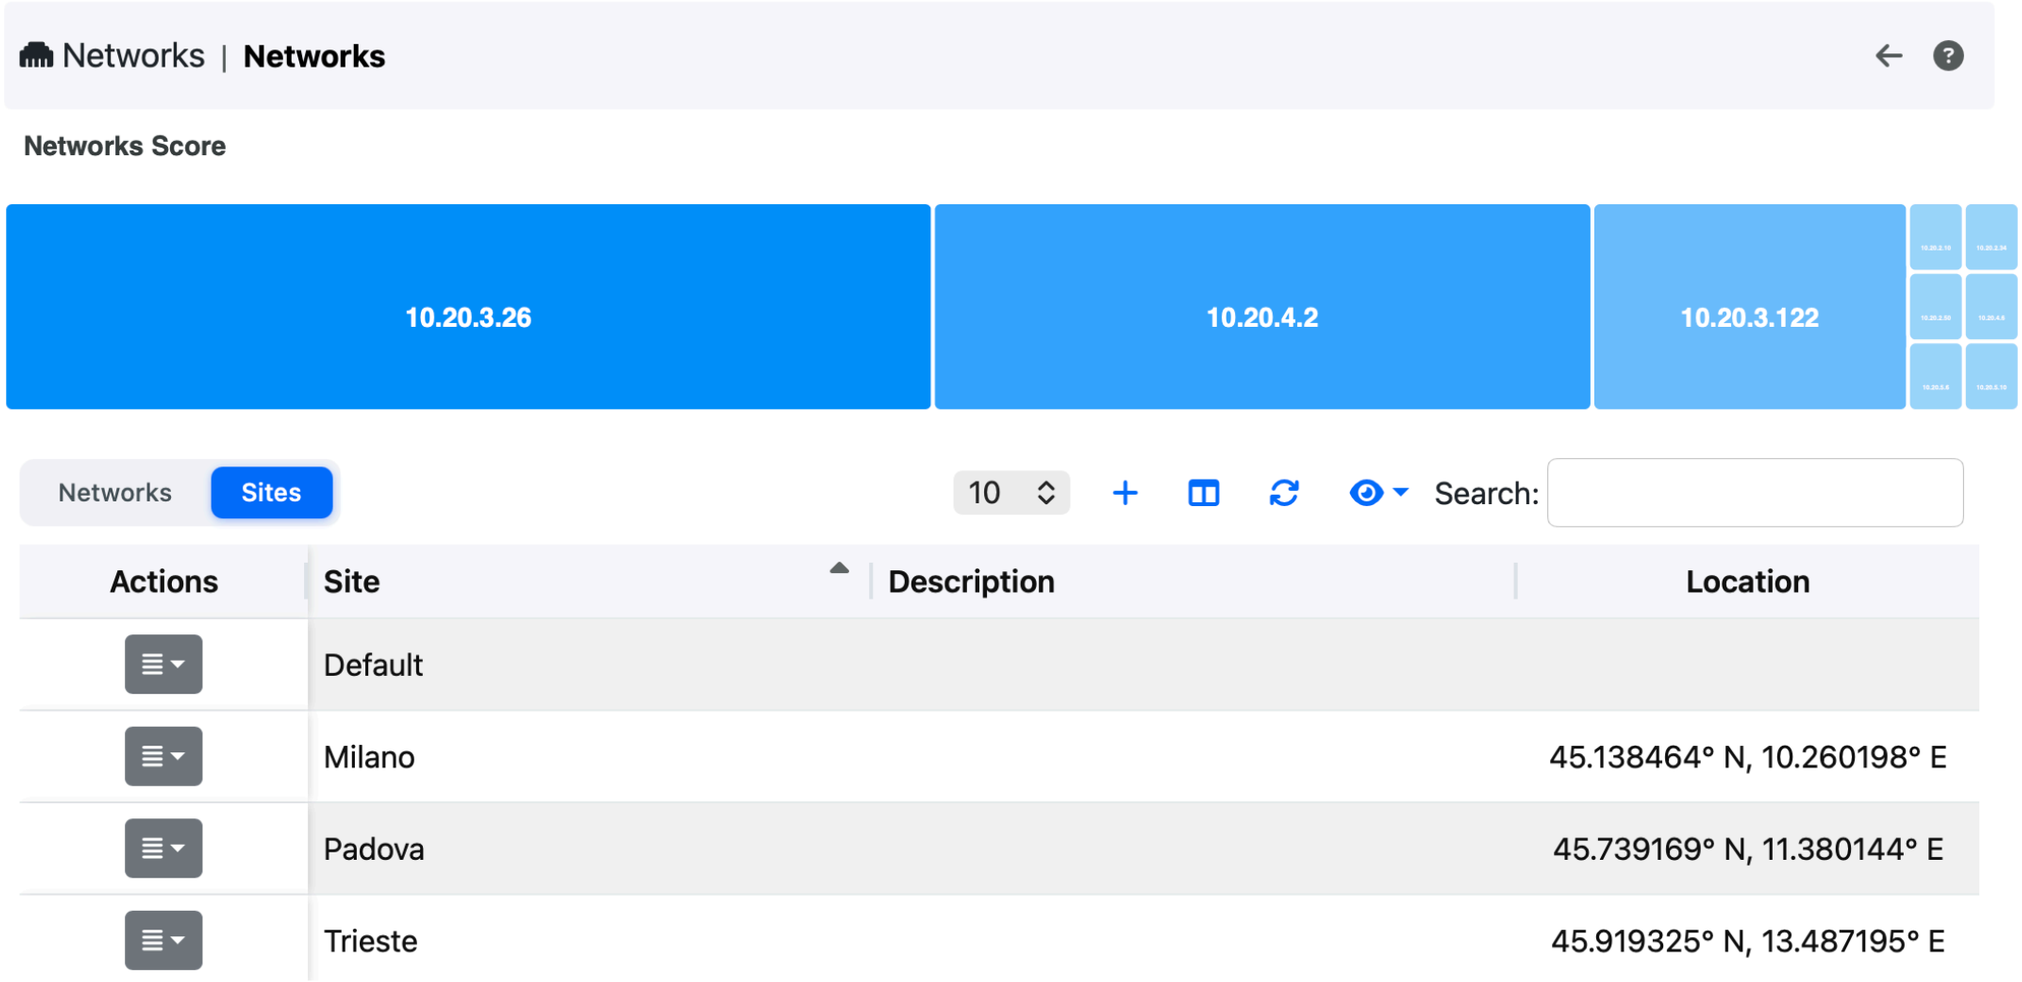

Merging All Together

Role	Internet eXchange	
CommunityId	1:A92Wj45BN9euVc6bRY+bJ1ciby8=	
Actual / Peak / Average Throughput	402.18 kbps / 9.97 Mbps / 6.64 Mbps	
ASN [Client / Server]	[redacted]	[redacted]
Exporters Information	**Flow Exporter / Next Hop**	**Input Interface / Output Interface**
	[redacted] (Default) [10.255.255.203] ➜ 10.20.4.21 Return Path	VLAN1062-edge-02.core (internal_interface) ➜ VLAN1007-nat-02.net
	[redacted] (Default) [10.255.255.202] ➜ 217.29.67.15	VLAN1063-edge-04.core (internal_interface) ➜ VLAN100-IPT-MIX (ix)

Exploring Traffic Flow Paths

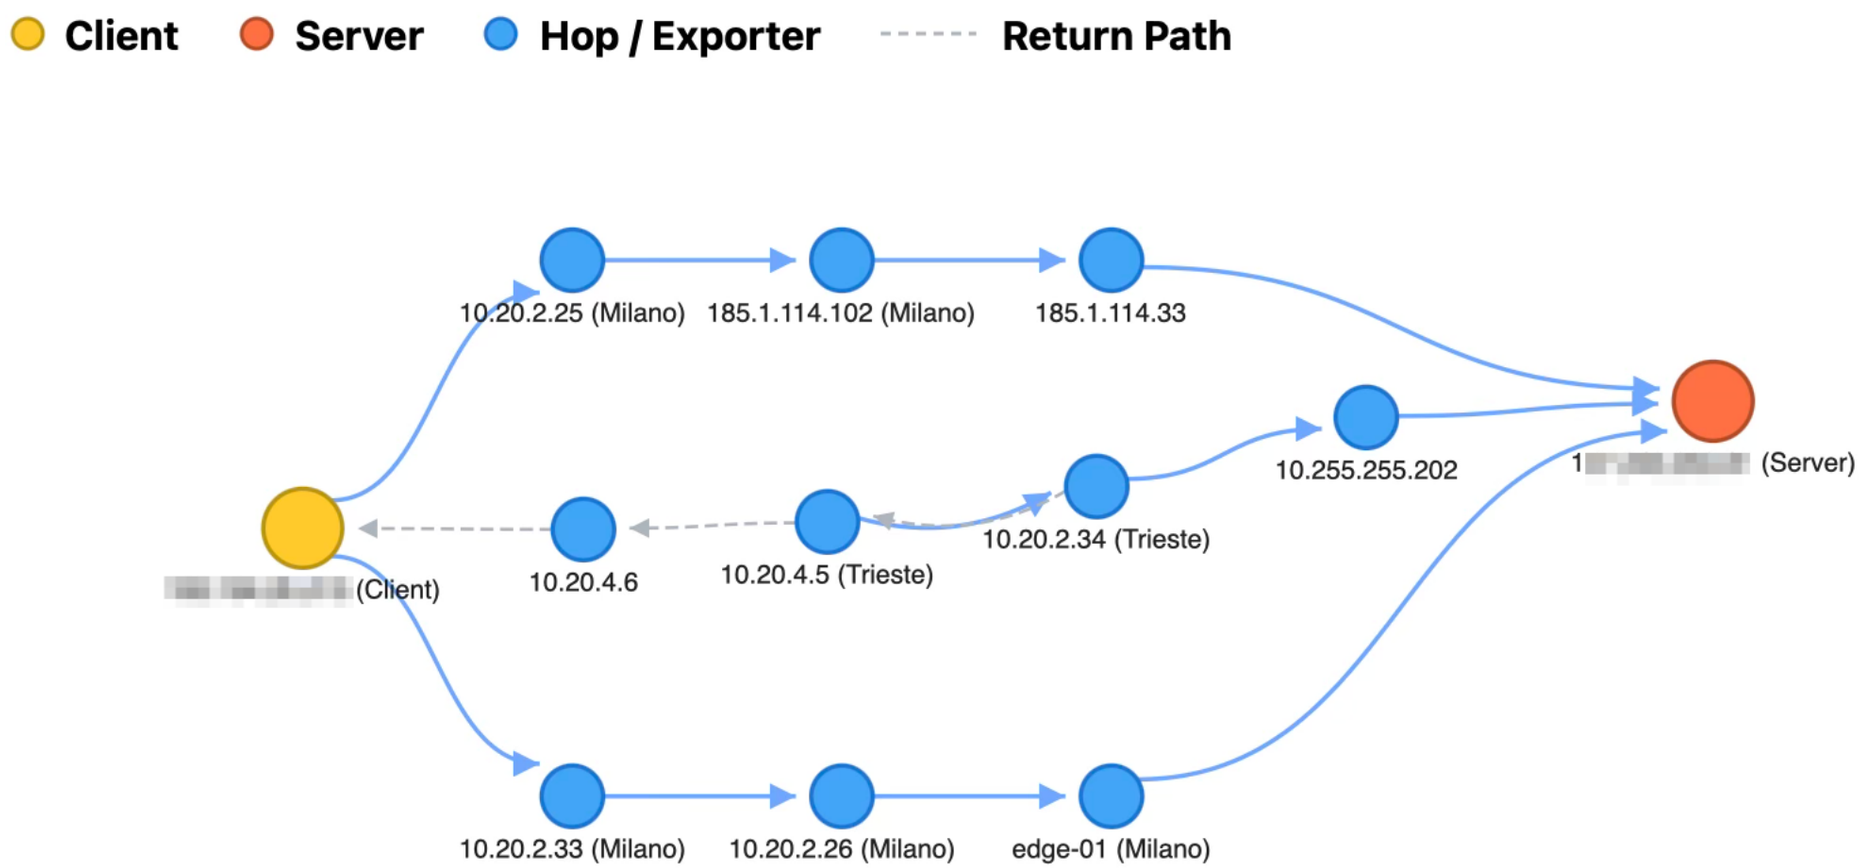

Note
Paths are depicted based on (duplicated) collected flows.

Role Traffic Breakdown

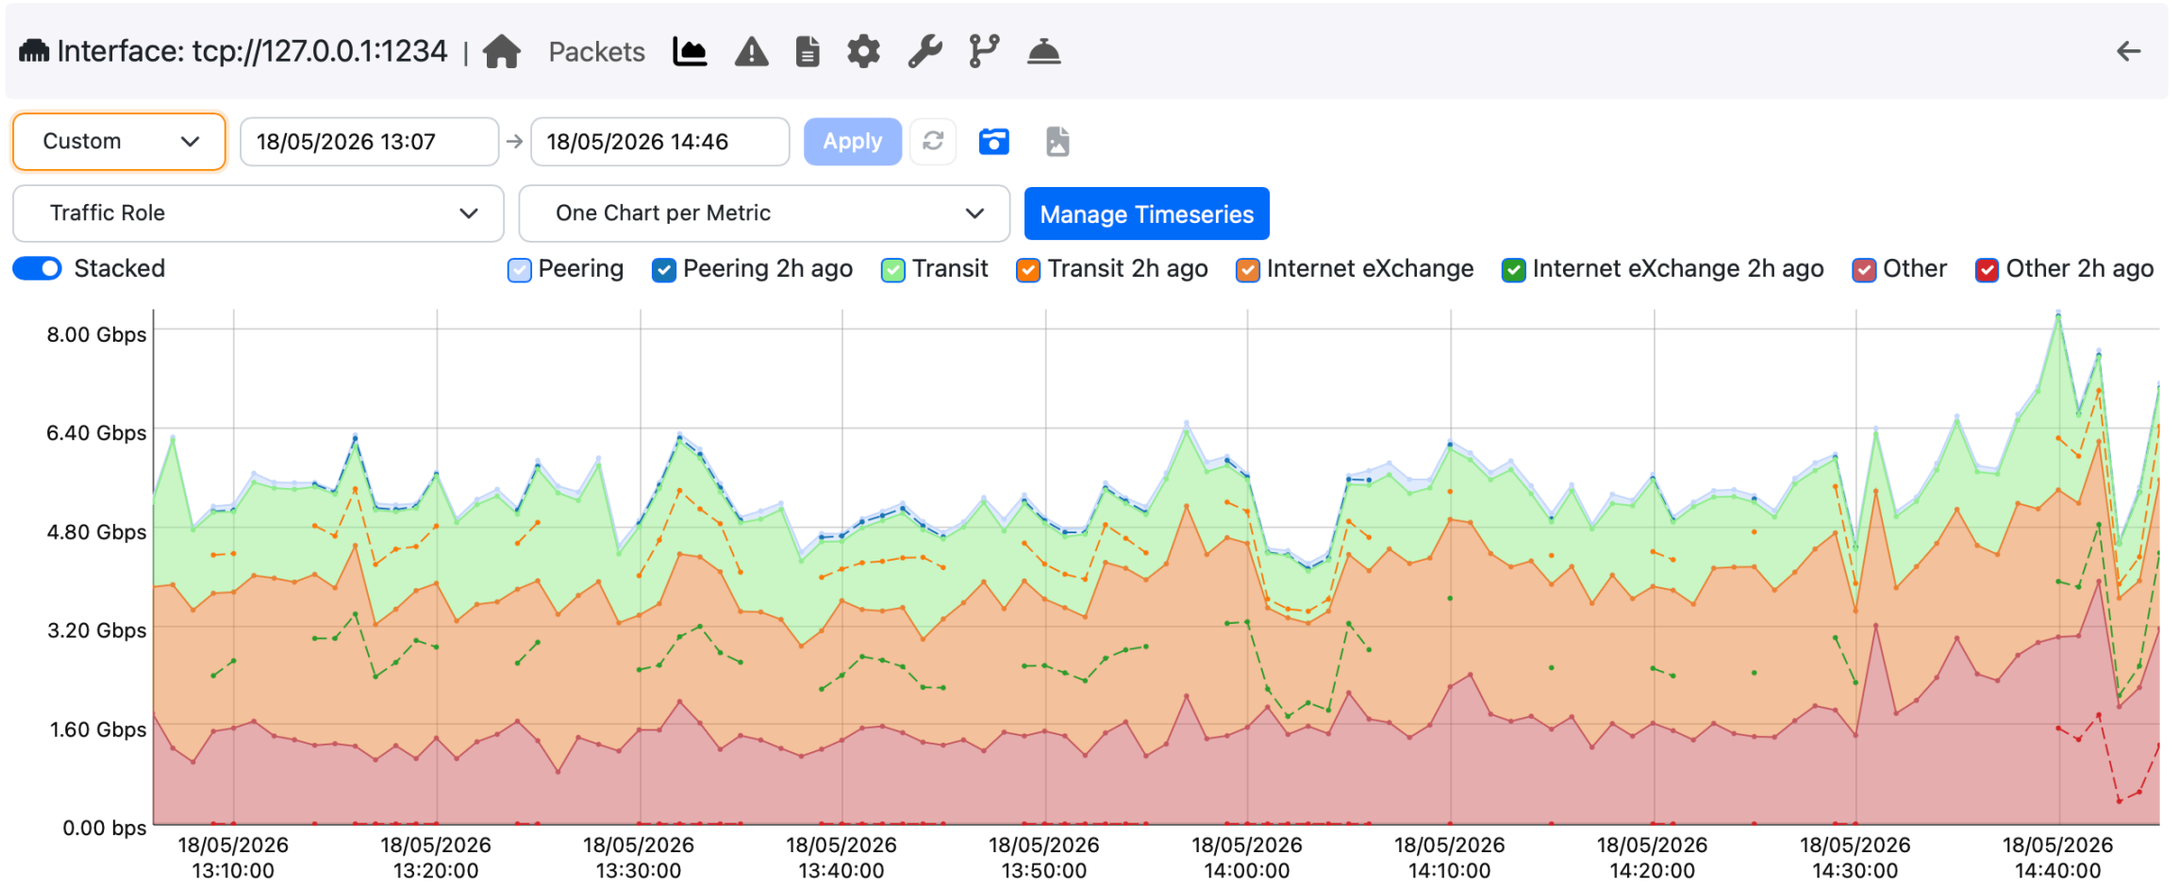

Metric	Average	95th Percentile	Max	Min	Total
Peering	68.37 Mbps	99.56 Mbps	151.46 Mbps	27.19 Mbps	47.76 GB
Transit	714.27 Mbps	1.08 Gbps	1.95 Gbps	417.46 Mbps	498.91 GB
Internet eXchange	1.23 Gbps	1.61 Gbps	1.75 Gbps	758.23 Mbps	855.90 GB
Other	1.52 Gbps	2.07 Gbps	3.06 Gbps	843.53 Mbps	1.04 TB

Interface Traffic Breakdown

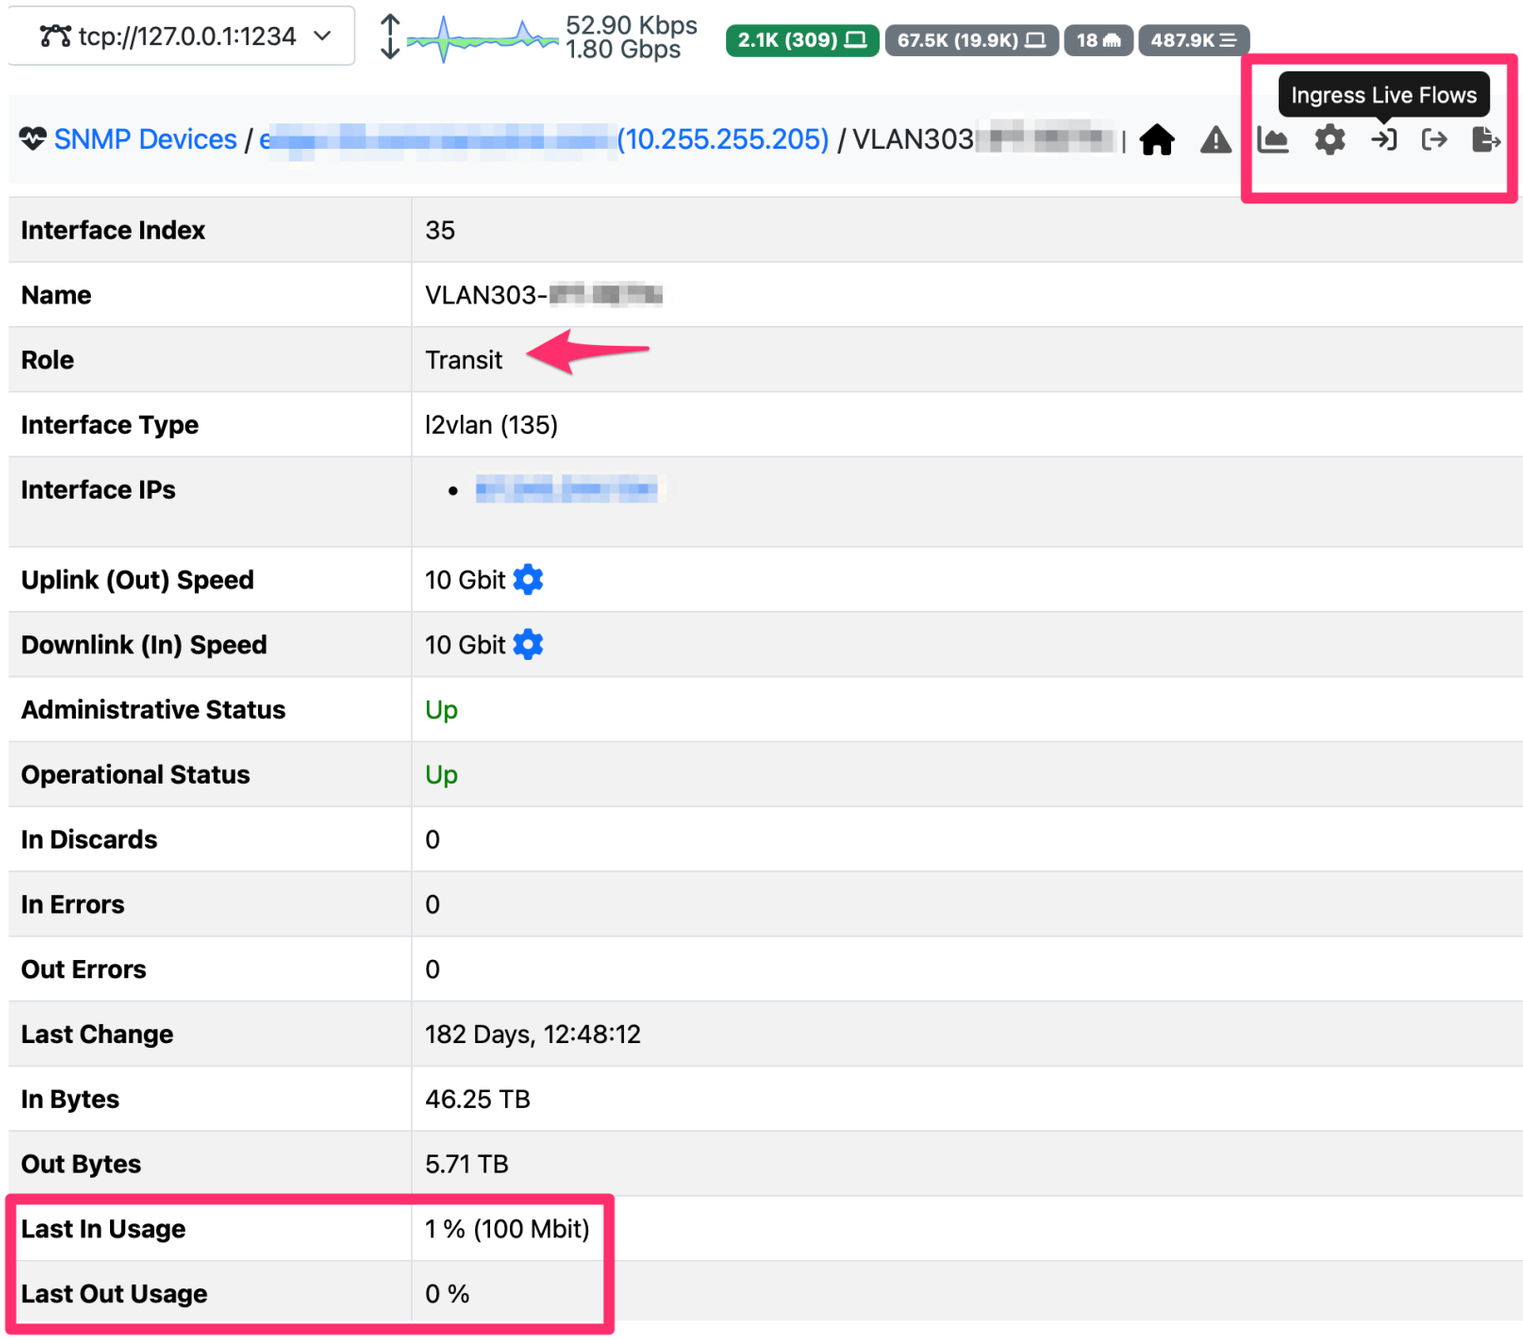

Enhanced Flow Filtering

- In addition to the "legacy" host/protocol/etc. filters it is now possible to filter flows according to:
  - ASN
  - Exporter
  - Exporter Interfaces (id and roles).

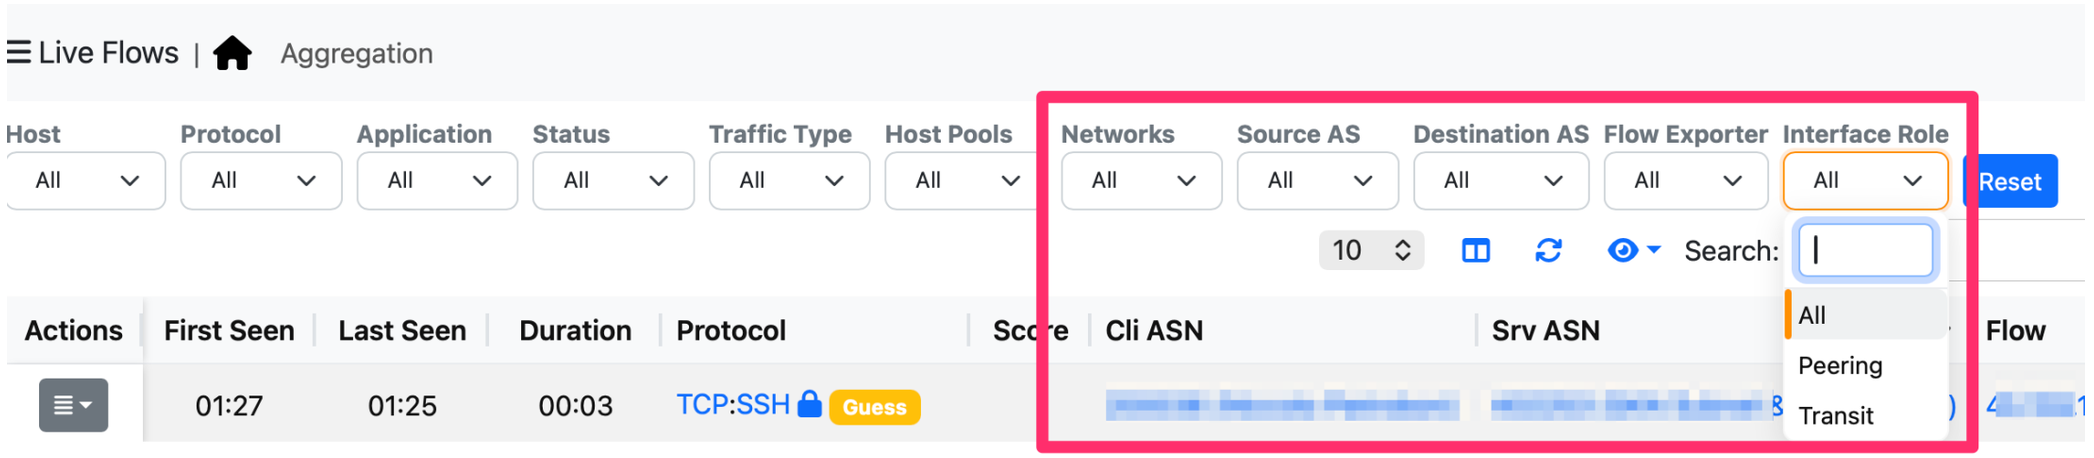

ASN Classification

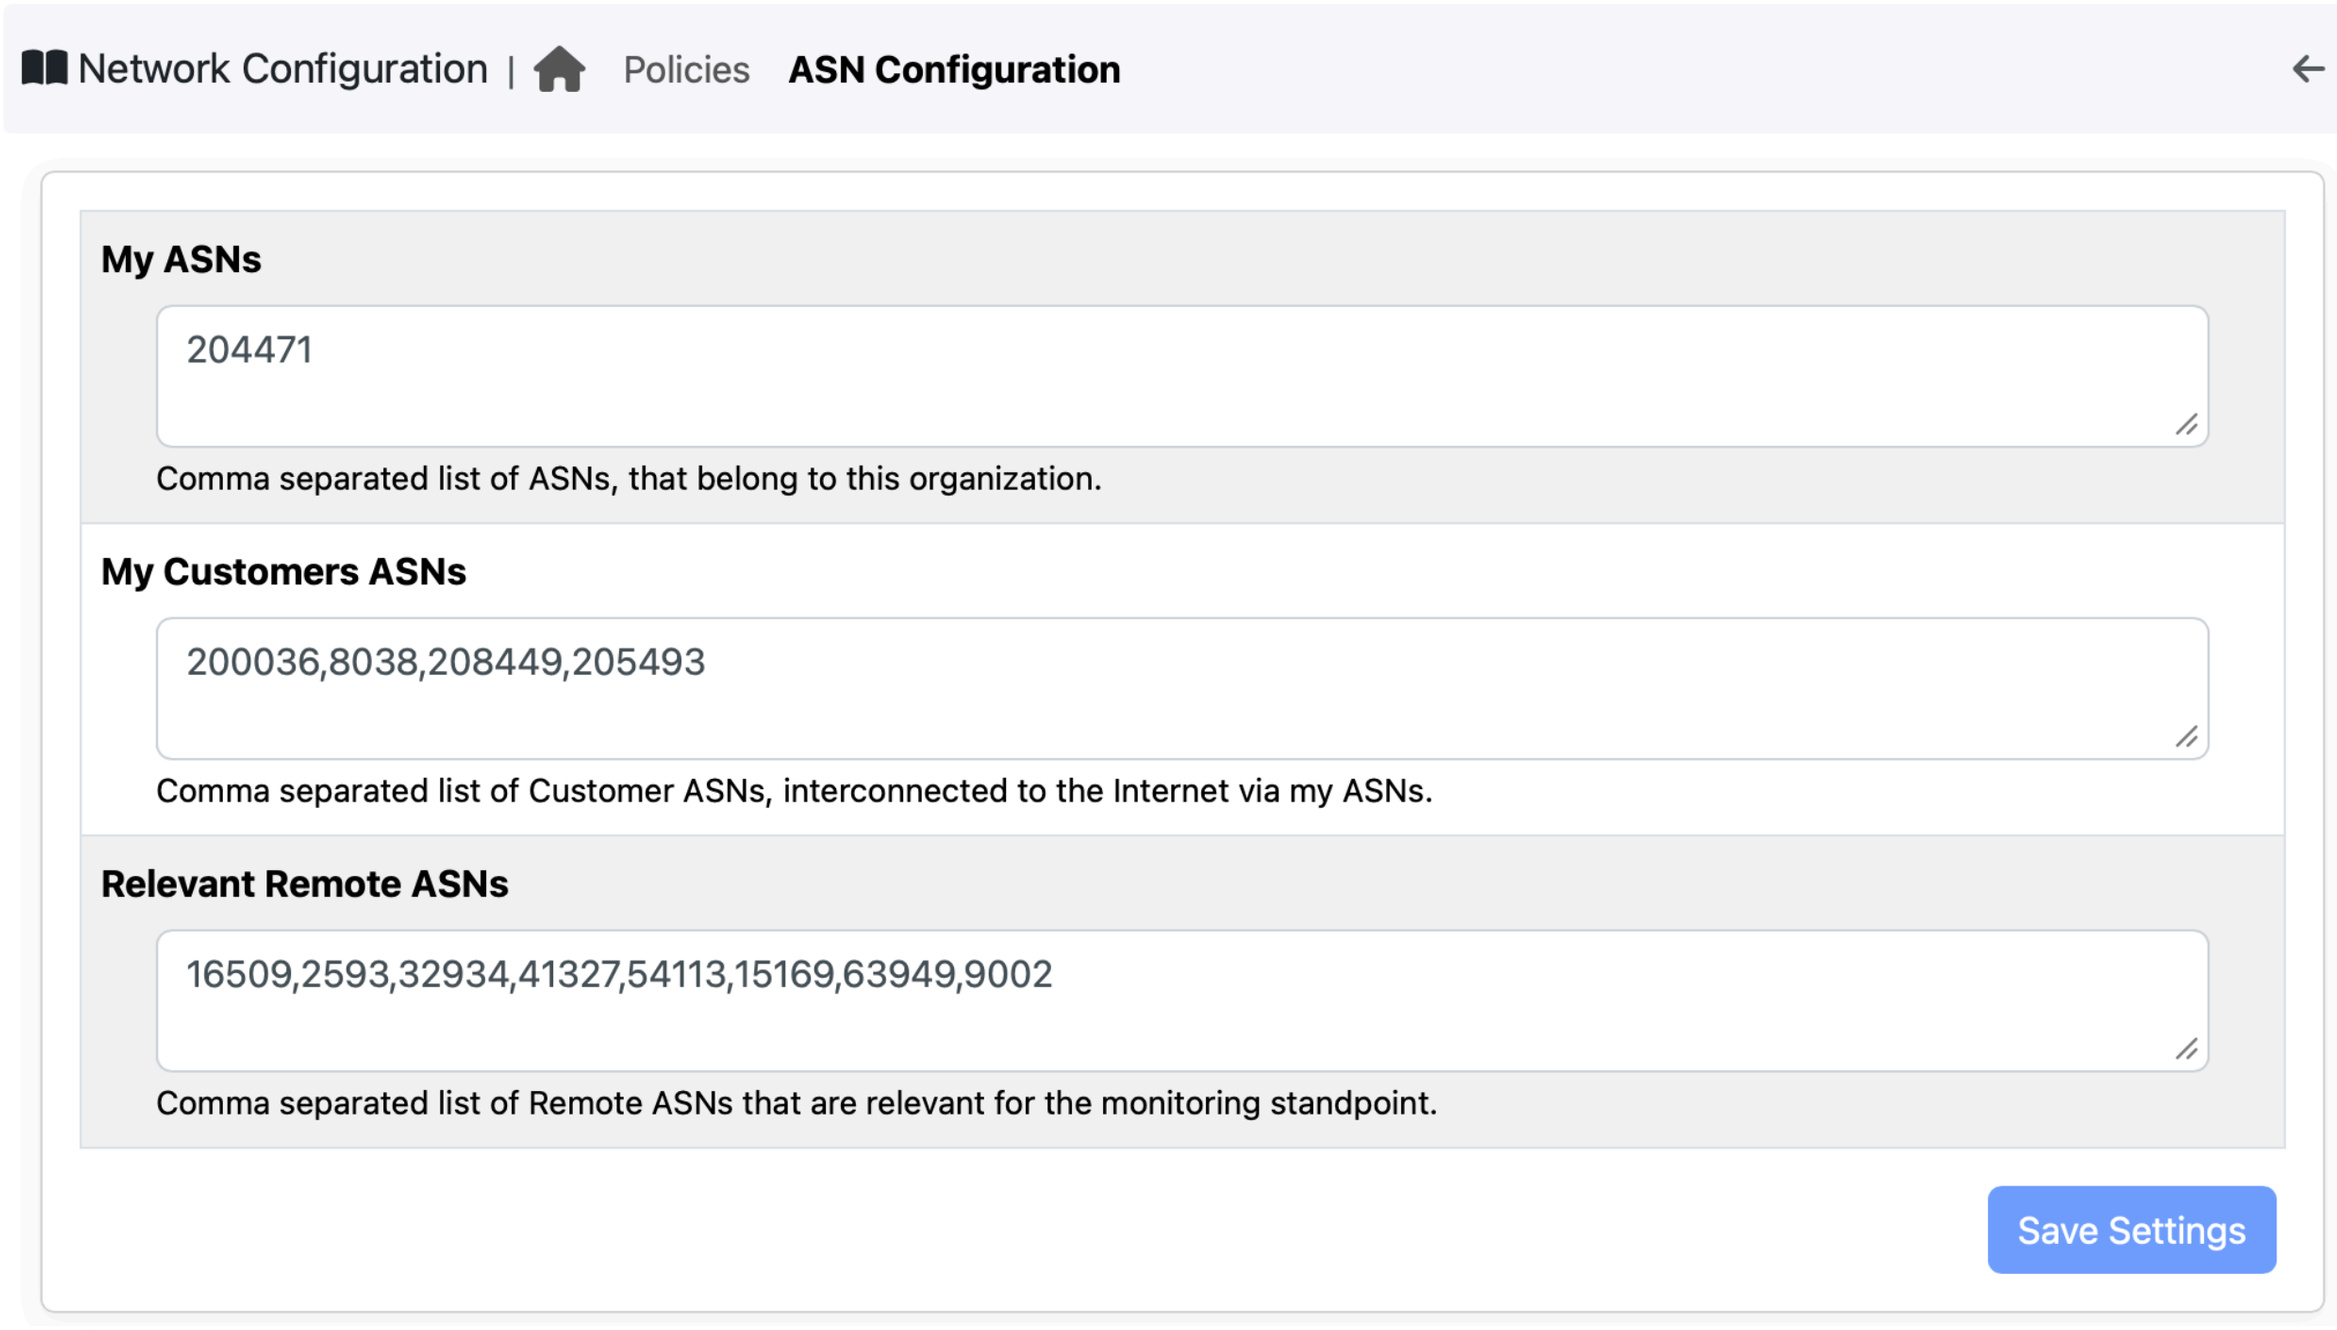

ASN Page Revamped

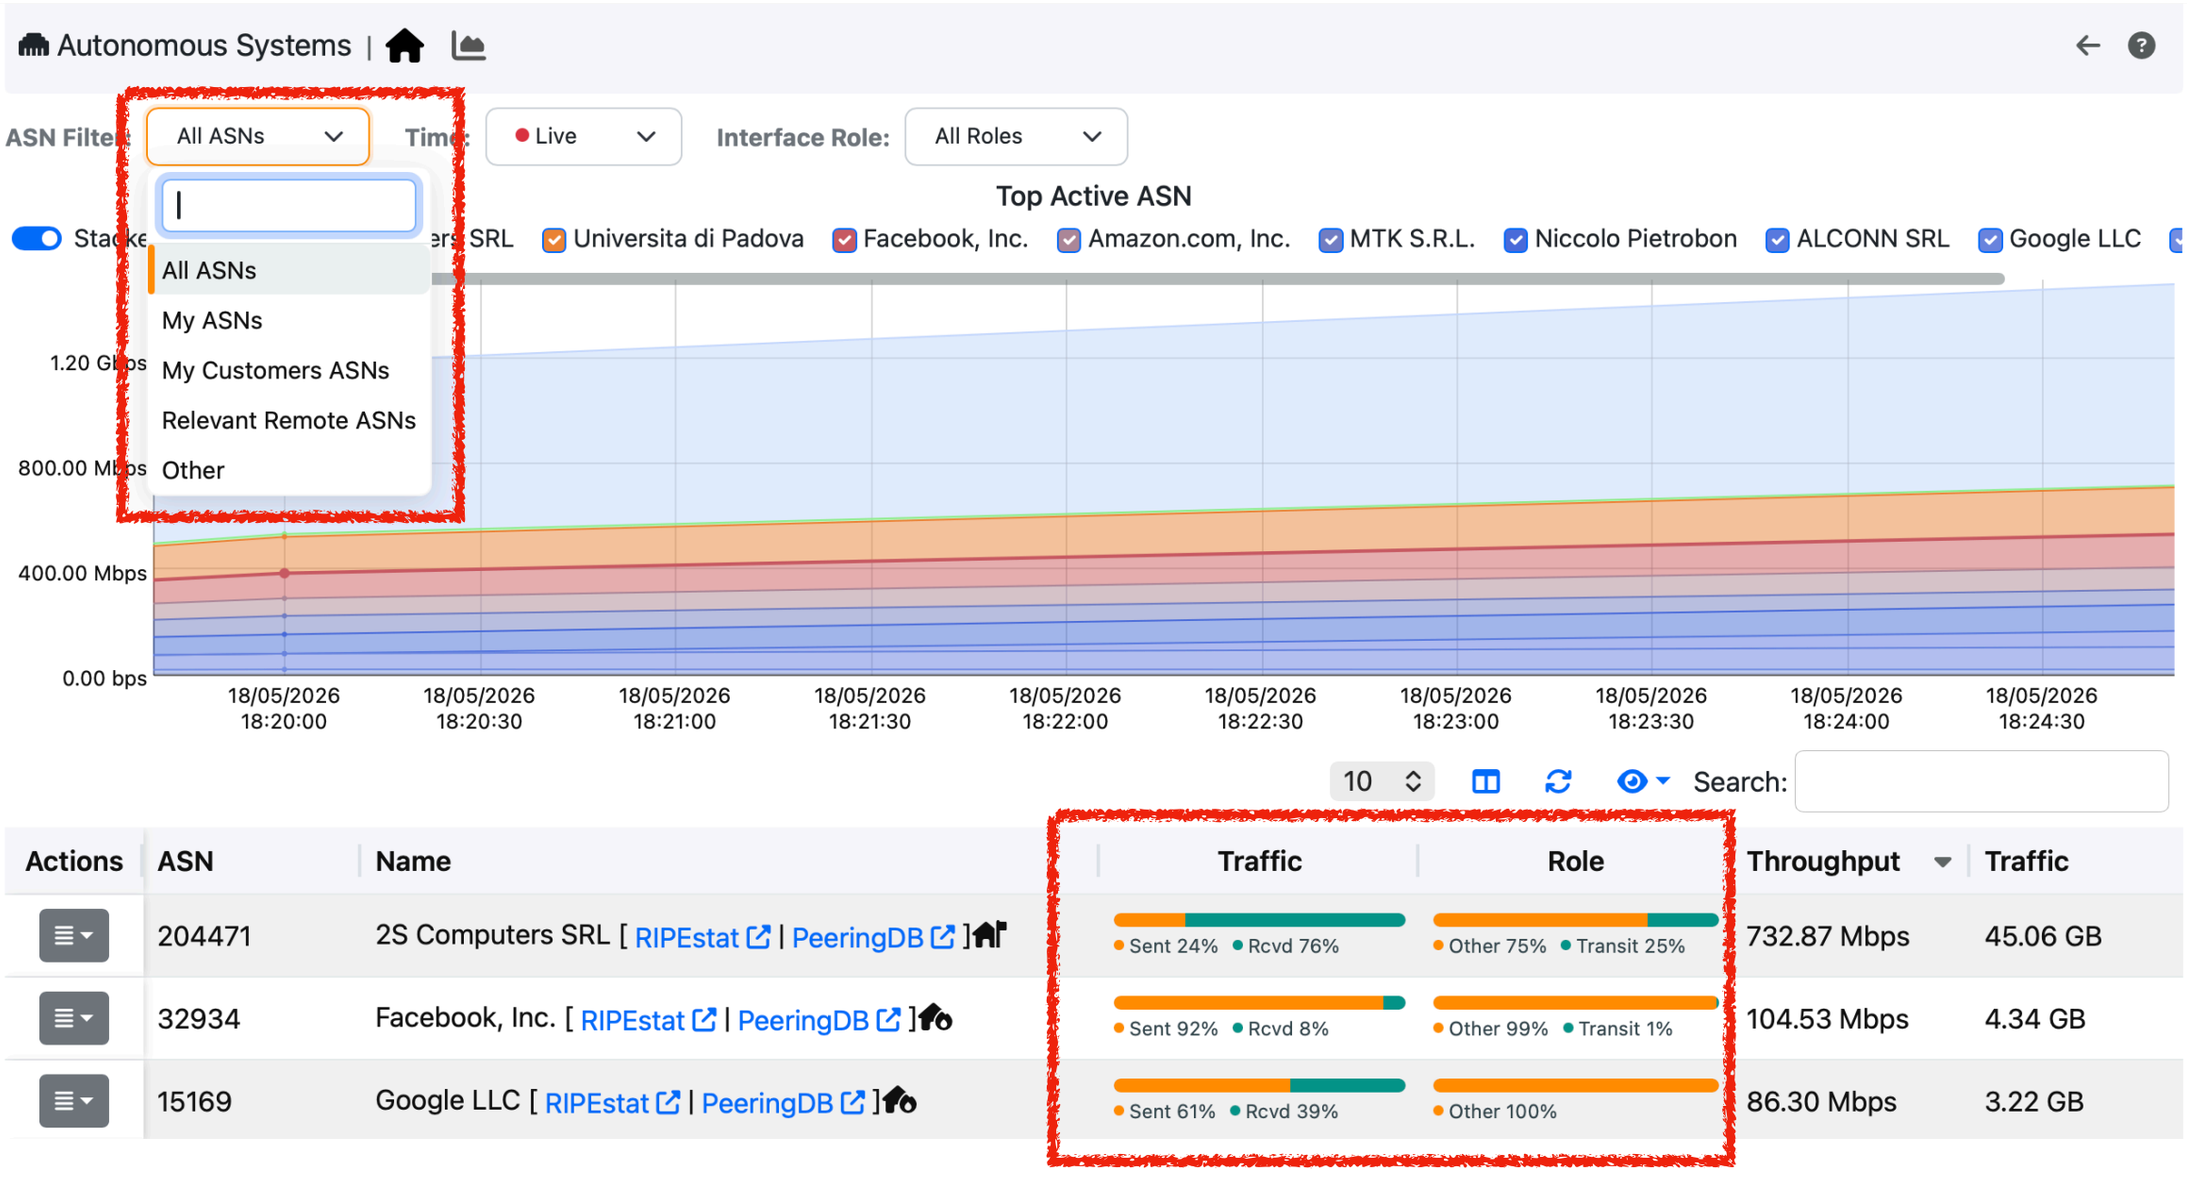

ASN: Ingress/Egress Interfaces

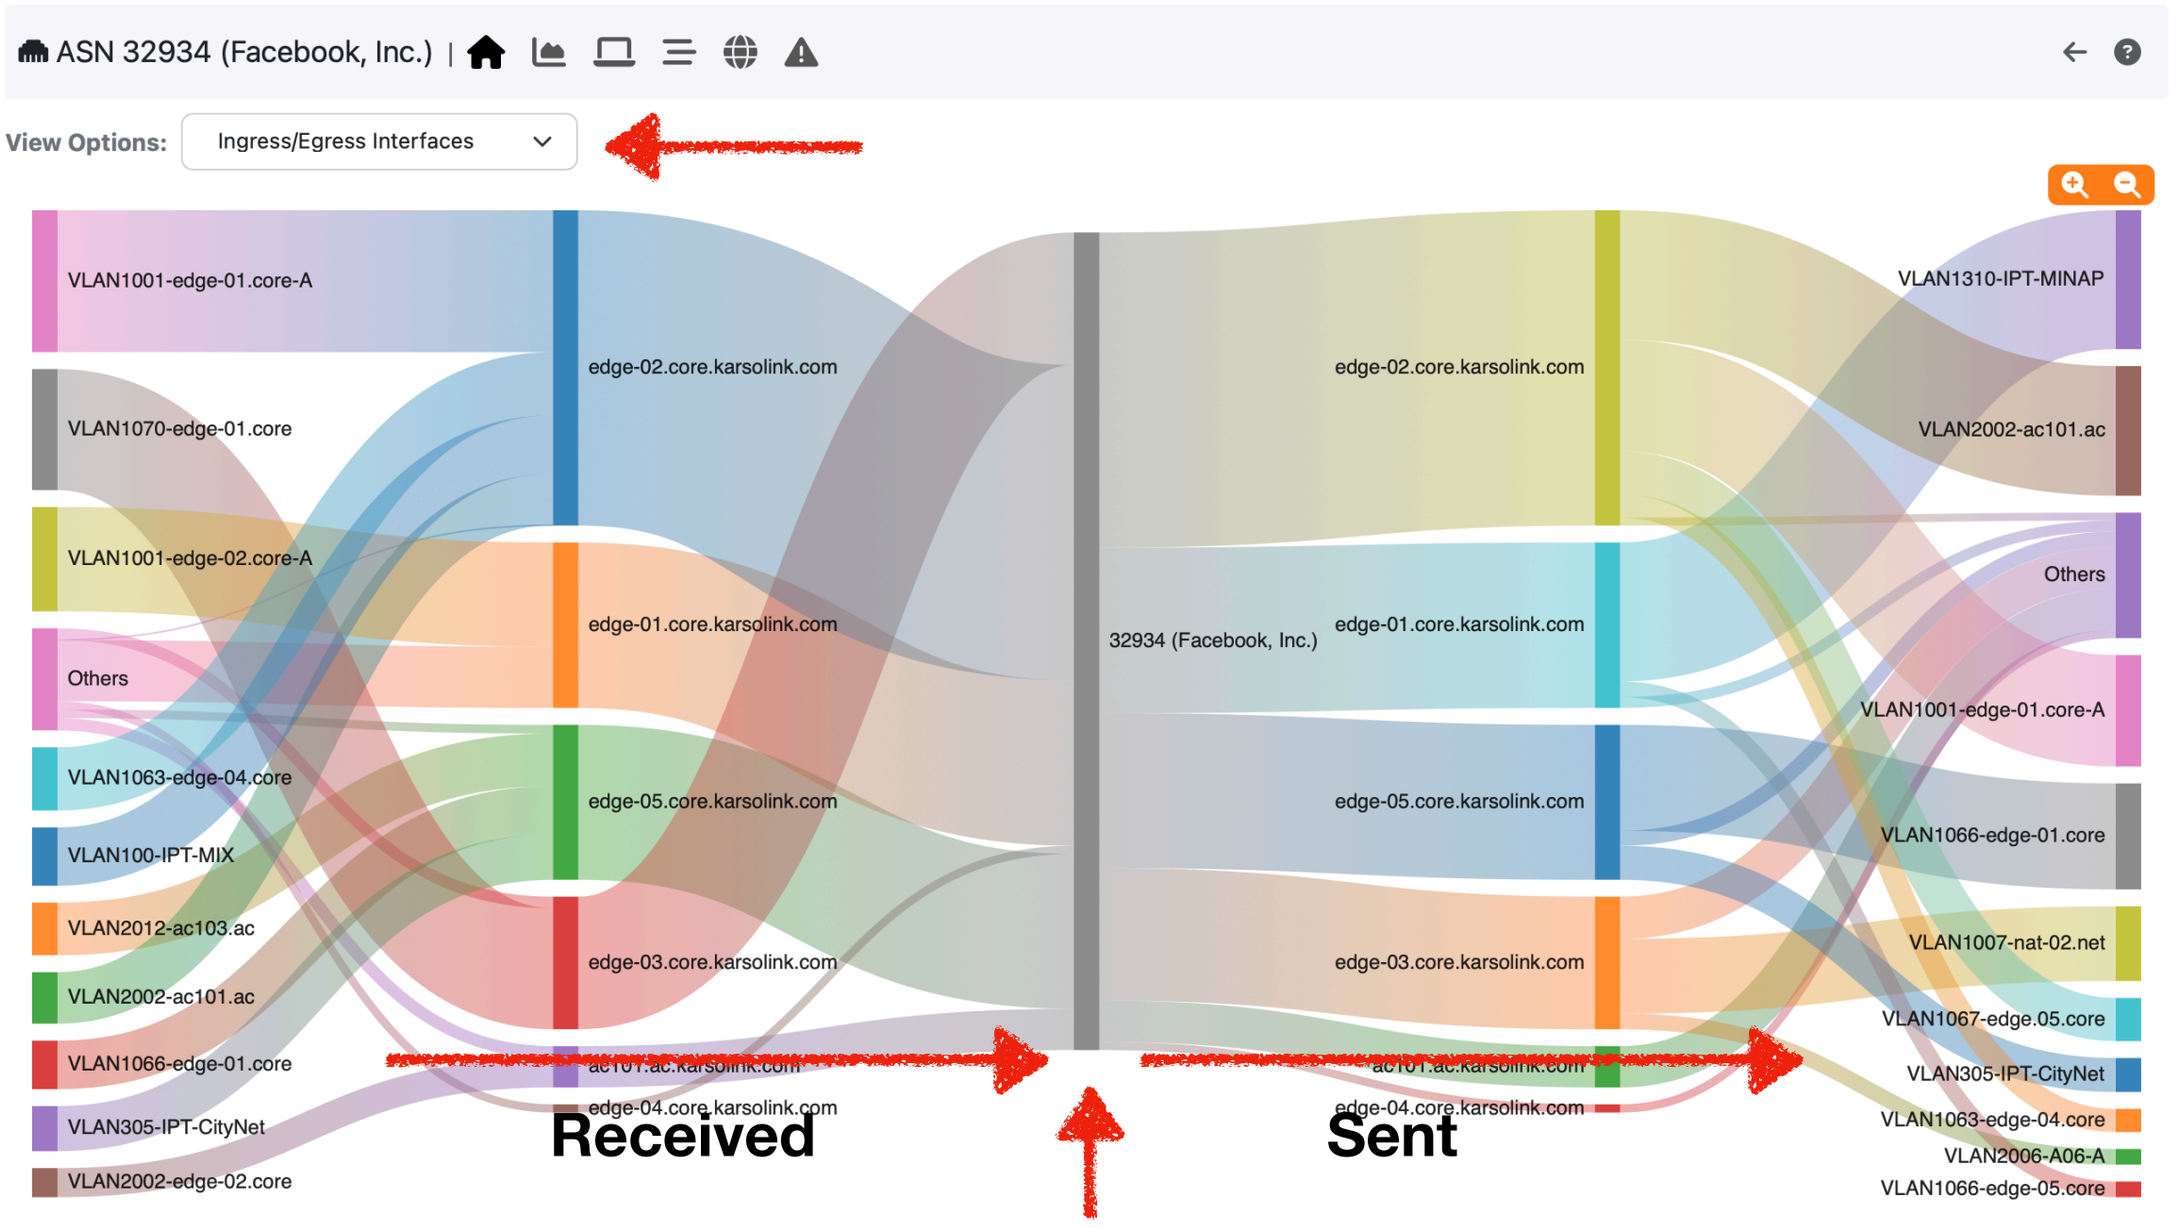

ASN Traffic View

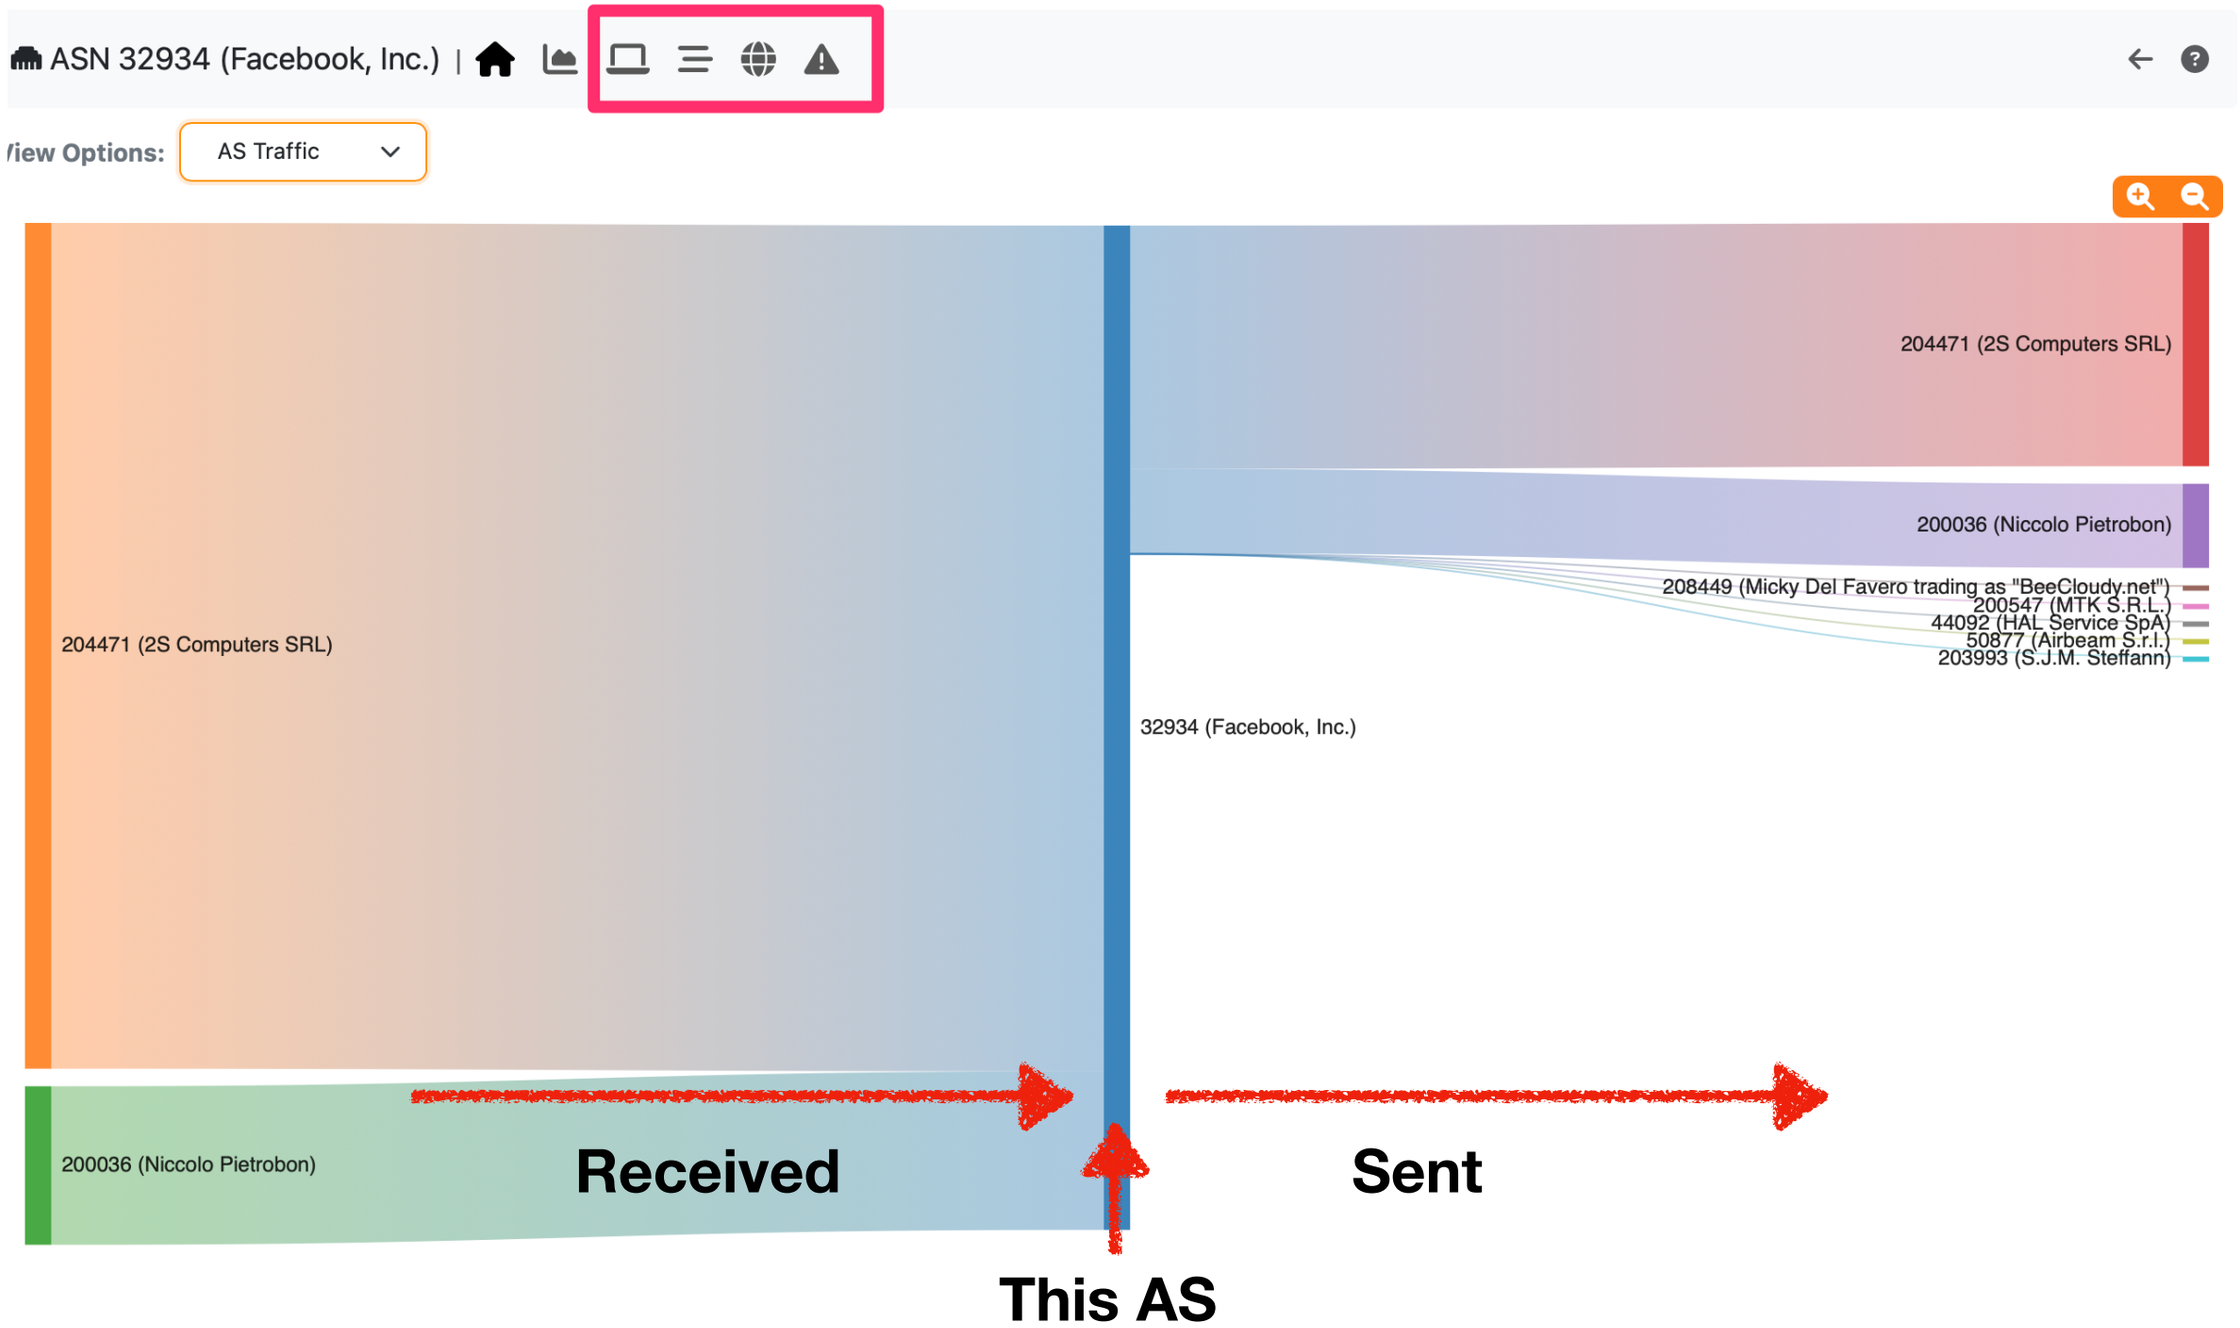

AS View: My Customers Breakdown

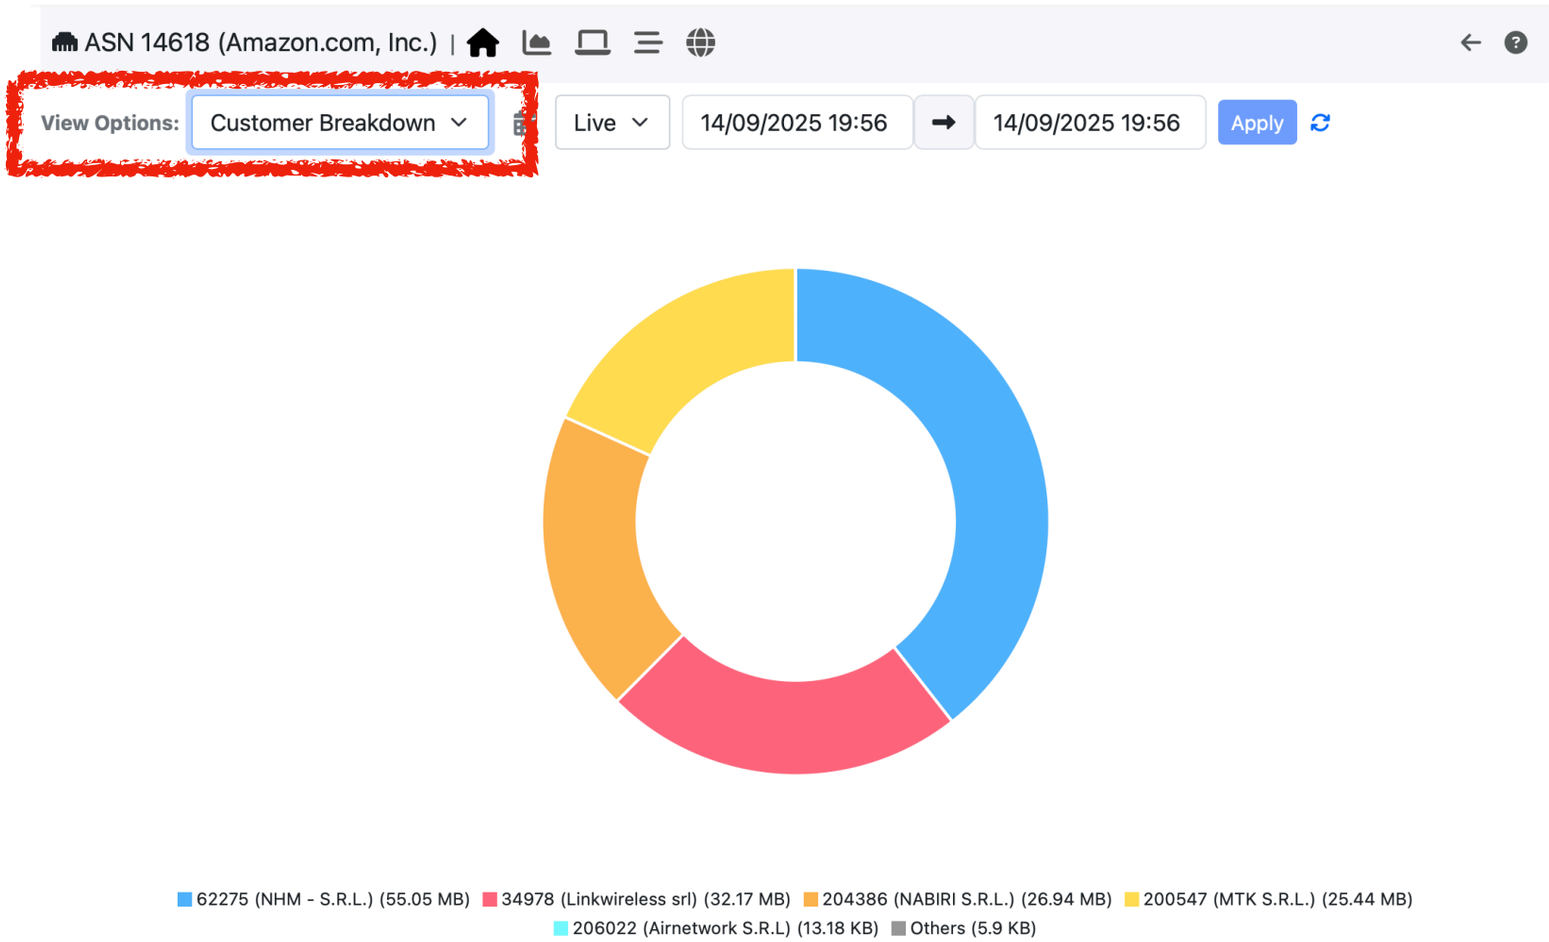

10 Search:

Customer ASNs	Bytes Sent	Bytes Rcvd	Tota
62275 (NHM - S.R.L.)	31.18 MB	23.52 MB	
34978 (Linkwireless srl)	23.38 MB	7.81 MB	

Enabling ASN Mode: ntopng

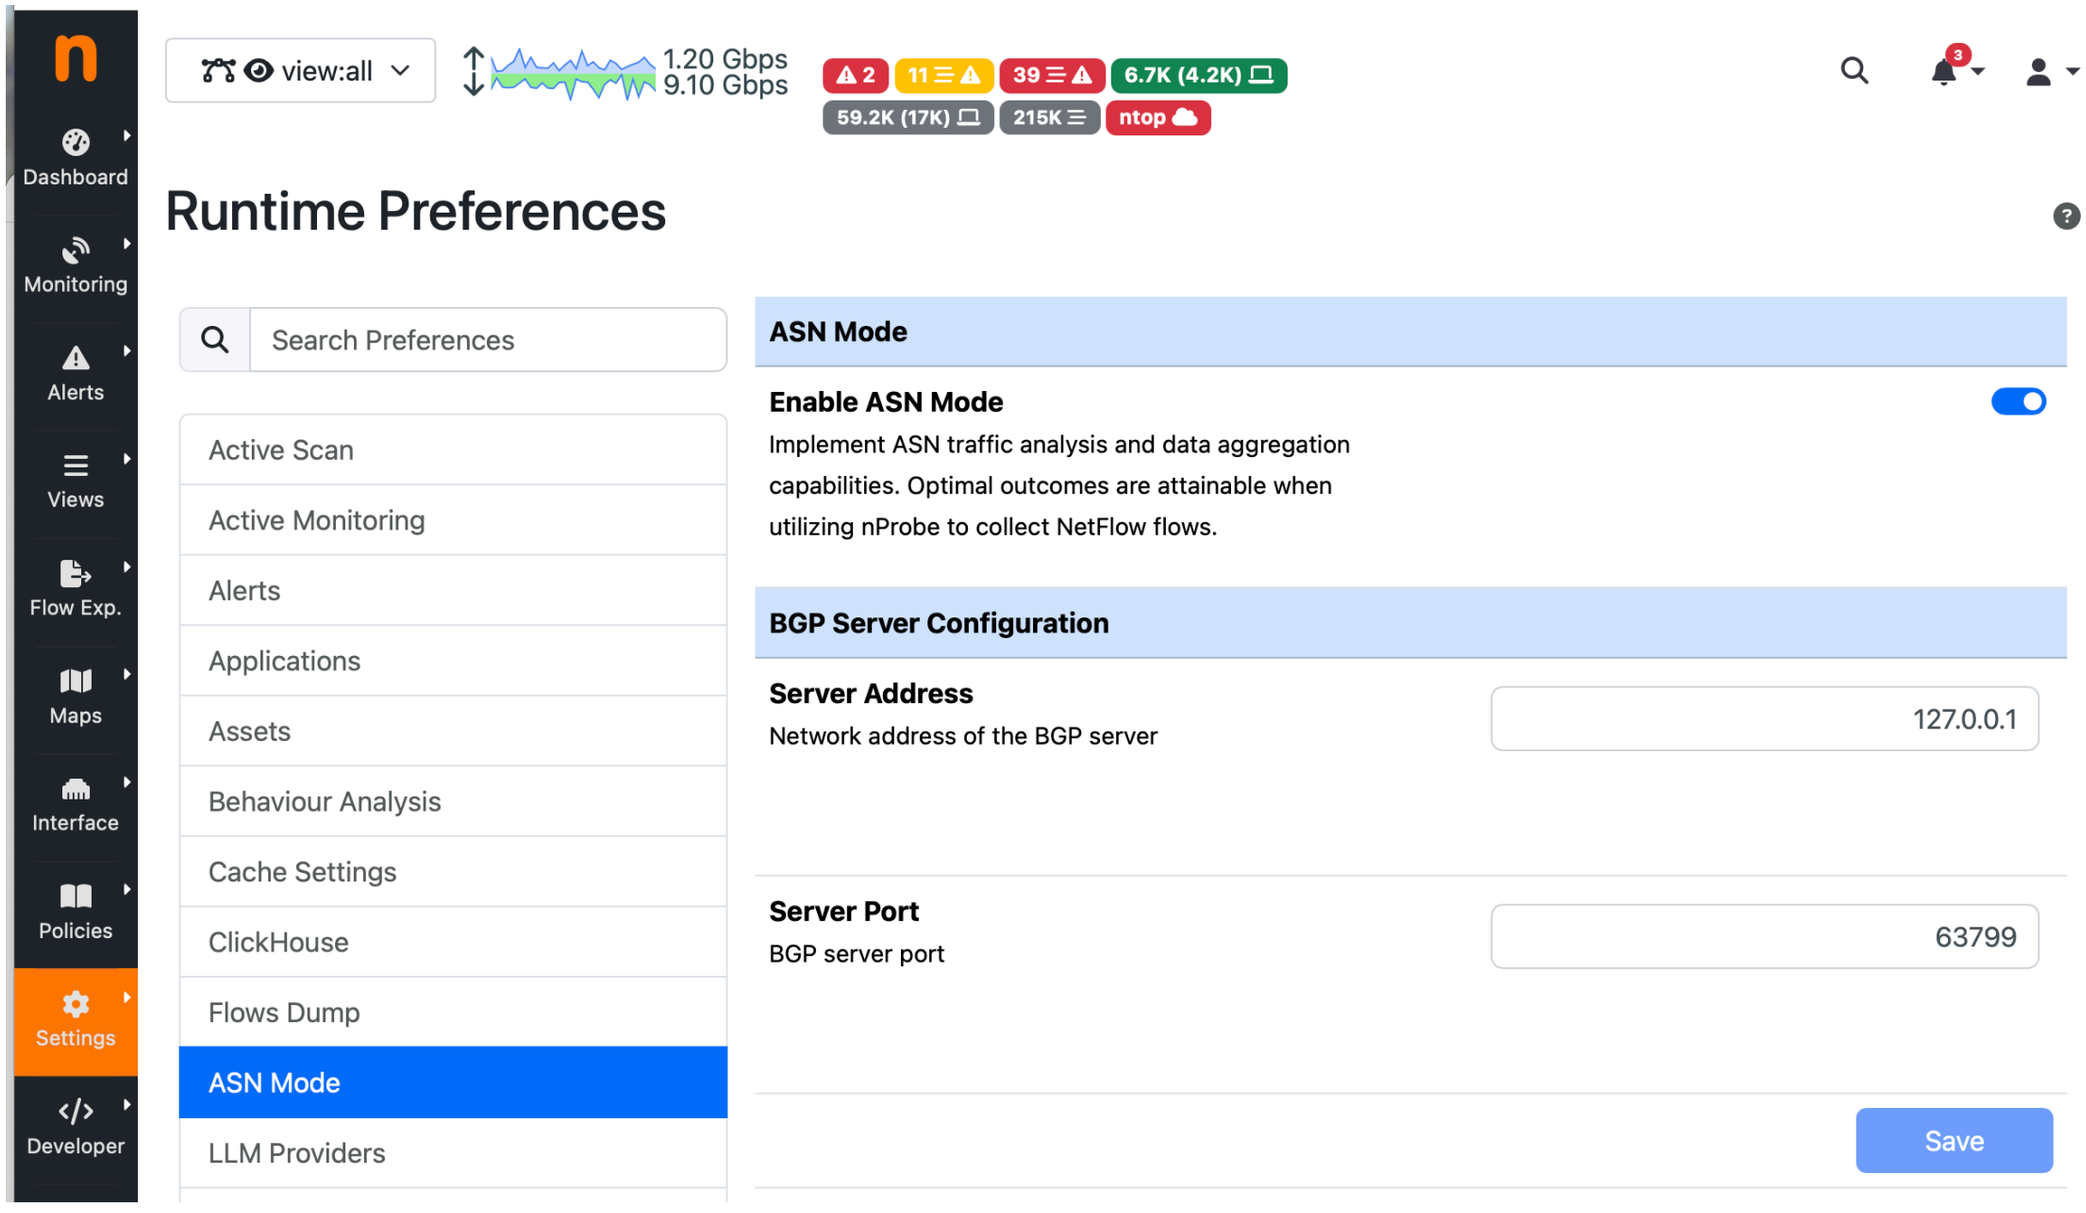

Enabling ASN Mode: nProbe

- You have the option to:
  - Collect flows as they are received (i.e. with full IP information).
  - [Optional] Mask IP addresses (according to the flow netmask) in order to hide the exact IP address.

```
--asn-mode                    | Collect flows and optimize export for AS traffic analysis.
                              | This CLI option has no effect in packet mode
```

- Note: DPI in flow collection operates partially (no packets) using IP addresses (e.g. the Office365 IP range) and protocol+ports.

BGP/BMP Collection [1/2]

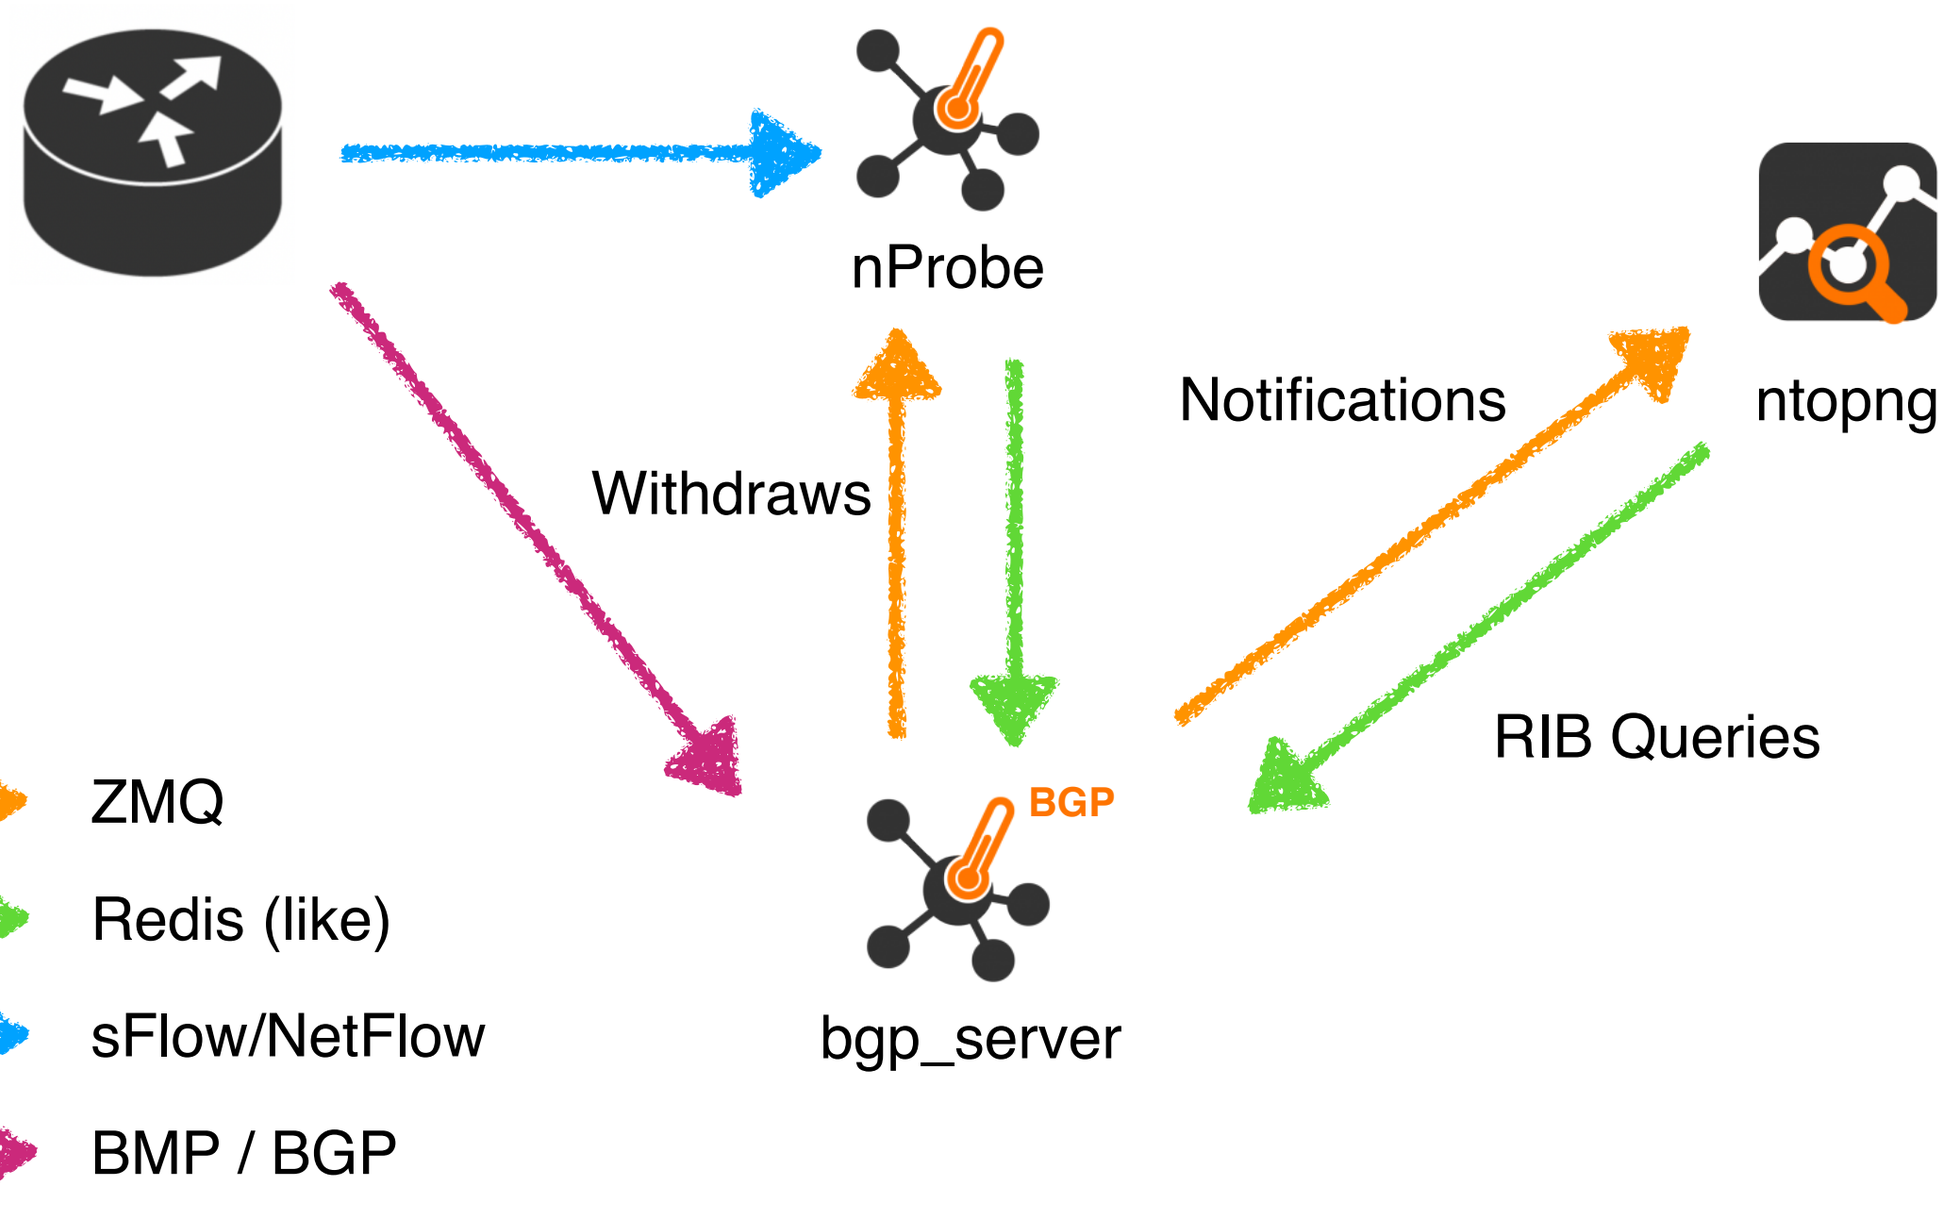

BGP/BMP Collection [2/2]

```
Welcome to bgp_server v.11.1.260406 for macOS 26.4
Copyright 2026 - ntop.org

Usage: bgp_server -z <URL> [-p <BMP port>] [-r <redis port>]
                  [-b <BGP port>] [-a <local ASN>] [-i <BGP router-ID>]
                  [-n <path>] [-v]

  -z <URL>    ZMQ publish URL
  -p <port>   BMP listen port           (default 11019)
  -r <port>   Redis query port          (default 63799)
  -b <port>   BGP passive listen port   (0 = disabled, default BGP port is 179)
  -a <ASN>    Local AS number for BGP OPEN  (default 65000)
  -i <IP>     Local BGP router-ID (dotted decimal, default 1.1.1.1)
  -n <path>   List of IPv4/v6 prefixes to be notified (add/withdraw)
  -v          Verbose: print every parsed message

Example:
[BMP] bgp_server -z tcp://127.0.0.1:11059 -p 11019 -n prefixes.txt
[BGP] bgp_server -z tcp://127.0.0.1:11059 -i 10.82.4.121 -p 179 -a 65000 -n prefixes.txt
```

Merging Routing with Flows

Flow: [redacted] → [redacted] | BGP

BGP Client Info

Prefix	45.159.195.0/24
BGP Peer Id	217.197.106.201
BGP Peer ASN	204471 (KARSOLINK - 2S Computers SRL)
Origin	IGP
Next Hop	10.255.255.201
AS Path	204471 (KARSOLINK - 2S Computers SRL) 200036 (CITYNET - Niccolo Pietrobon)
MED (Multi-Exit Discri...	
Local Pref.	
Communities	

BGP Server Info

Prefix	209.99.184.0/21
BGP Peer Id	217.197.106.201
BGP Peer ASN	204471 (KARSOLINK - 2S Computers SRL)
Origin	IGP
Next Hop	10.255.255.201
AS Path	204471 (KARSOLINK - 2S Computers SRL) 9002 (RETN-AS - RETN Limited) 3257 (GTT-BACKBONE GTT Communications Inc.) 30823 (AUROLOGIC - aurologic GmbH) 402253 (SKN-NETWORK-1 - SKN Subnet & Teleco...
MED (Multi-Exit Discri...	
Local Pref.	
Communities	9002:9002 9002:64641

BGP Looking Glass

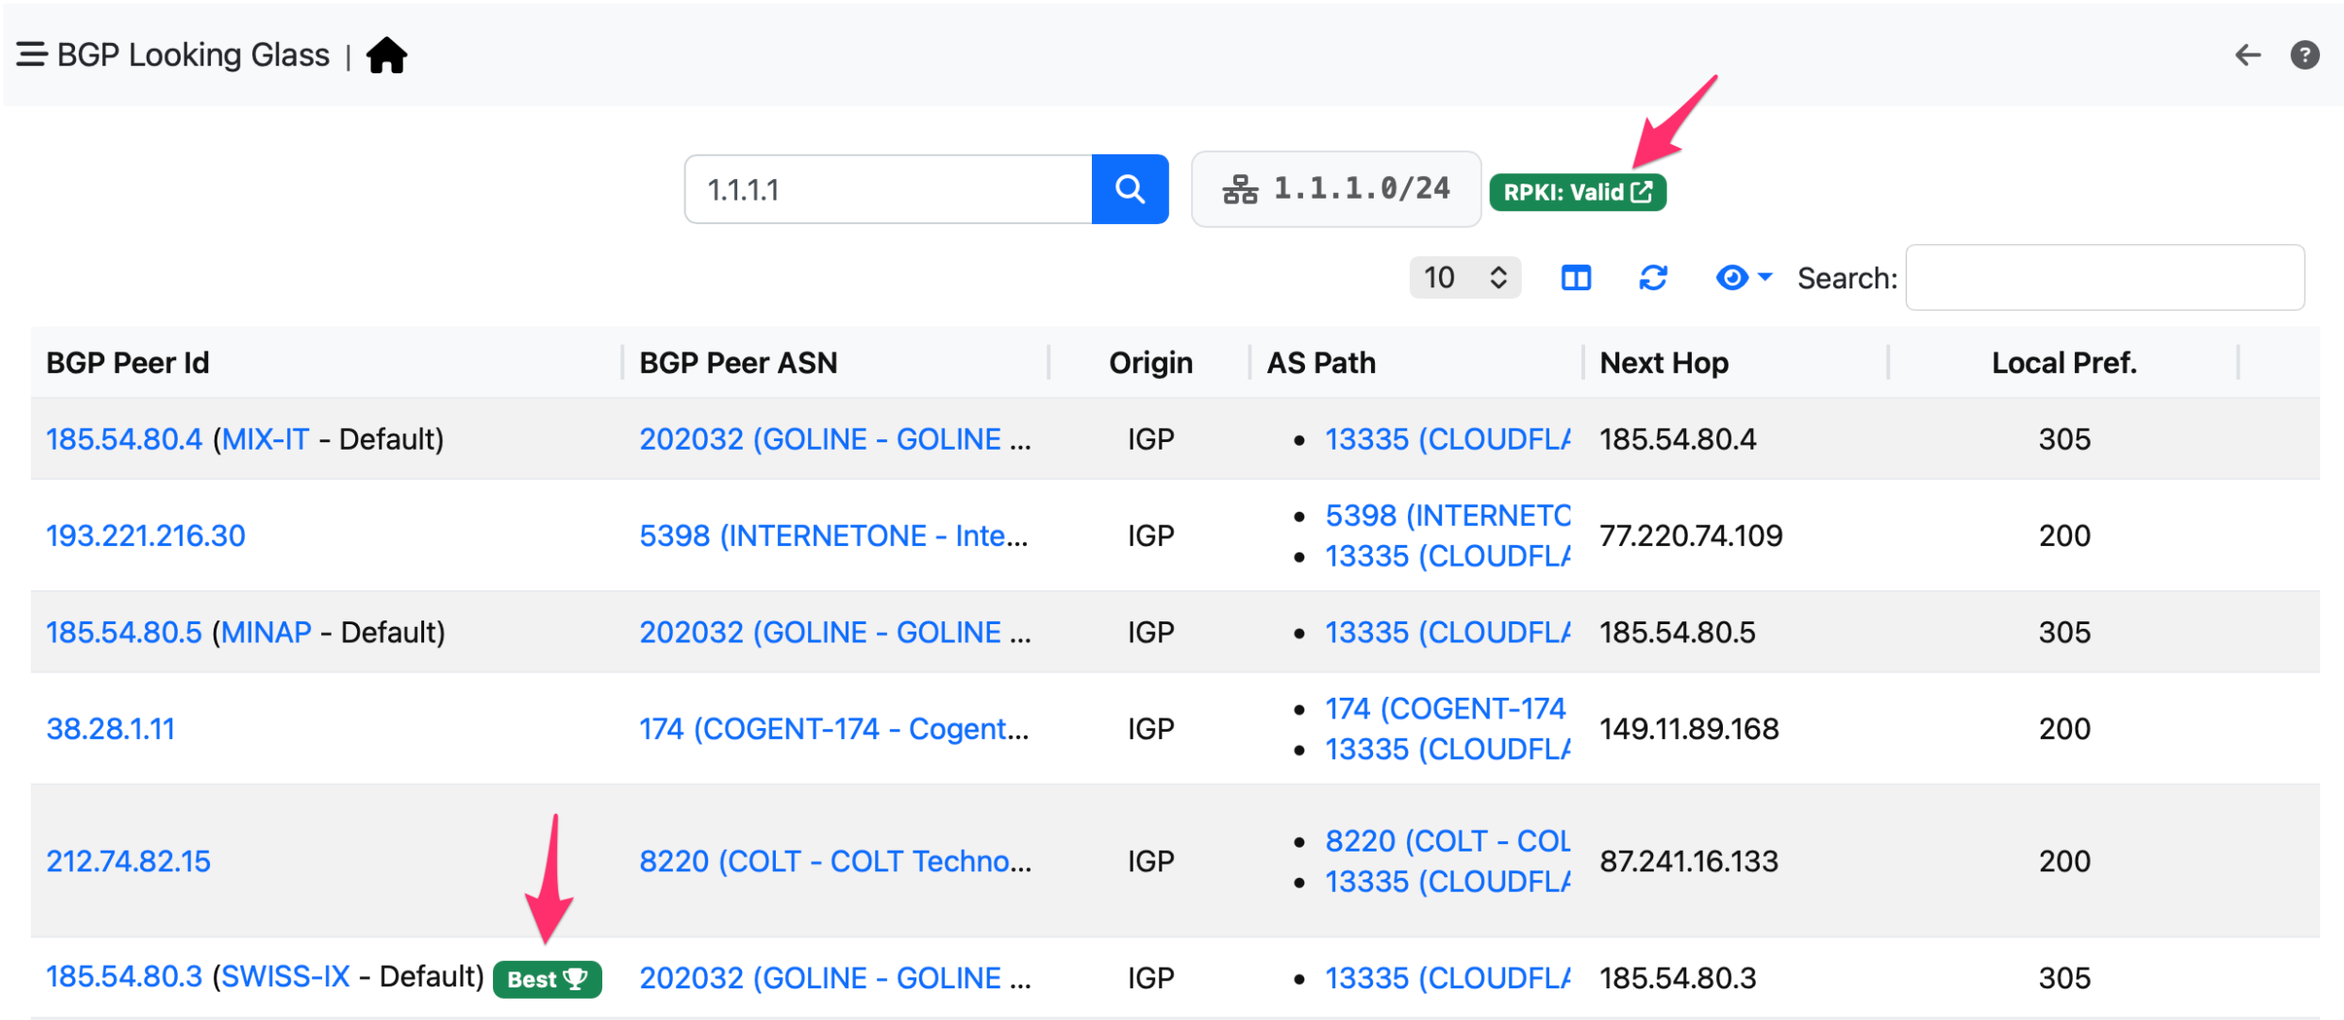

Alerts [1/2]

Alert |

AS	15169 (Google LLC)
Date / Time	09:06:19
Alert	Threshold Crossed
Description	[Metric: Traffic (RX + TX)] [Condition: 409.37 Mbps > 400 Mbps] [Check Frequency: 5 Minutes]

Alerts [2/2]

Alert | 

AS	23344 (Disney Worldwide Services, Inc.)
Date / Time	20:00:35
Alert	AS Exporter Ranking Changed
Description	Ingress ranking changed to [rank 1] NE8K_F1A_1_NAMEX:NAMEX_PEERING_LAN (17.72 GB) [rank 2] NE8K_M14_MIX:MINAP_via_SEEWEB-TGE1/0/39 (944.45 MB) [rank 3] NE8K_M14_MIX:MIX_PEERING (125.9 MB) [rank 4] NE8K_M14_MIX:BSO_IX_REACH_DECIX_FRANKFURT - ID_2288 (1.5 MB) from [rank 1] NE8K_F1A_1_NAMEX:NAMEX_PEERING_LAN (18.32 GB) [rank 2] NE8K_M14_MIX:BSO_IX_REACH_DECIX_FRANKFURT - ID_2288 (515.33 KB)

Availability

- What presented today is available in the dev branch. They will be included in the next stable.
- The bgp_server does not require an additional license other than nProbe (M and above), and it is based on open source software we developed (URL at the end of this presentation).
- AS/Routing information is available in all ntopng versions if available.

Future Work Items

- AI will be integrated into ntopng to ease data analysis.
- Additional alerts (e.g. DDoS, BGP peers state...).
- Detection of traffic spikes not due to a DDoS (e.g. soccer match).
- Provide more insight about billing costs per customer (peering exposed), in order to better tune the monthly fees based on the current usage.
- What else ?

What's Next ?

ntop AI Webinar June 9th [announce preview]

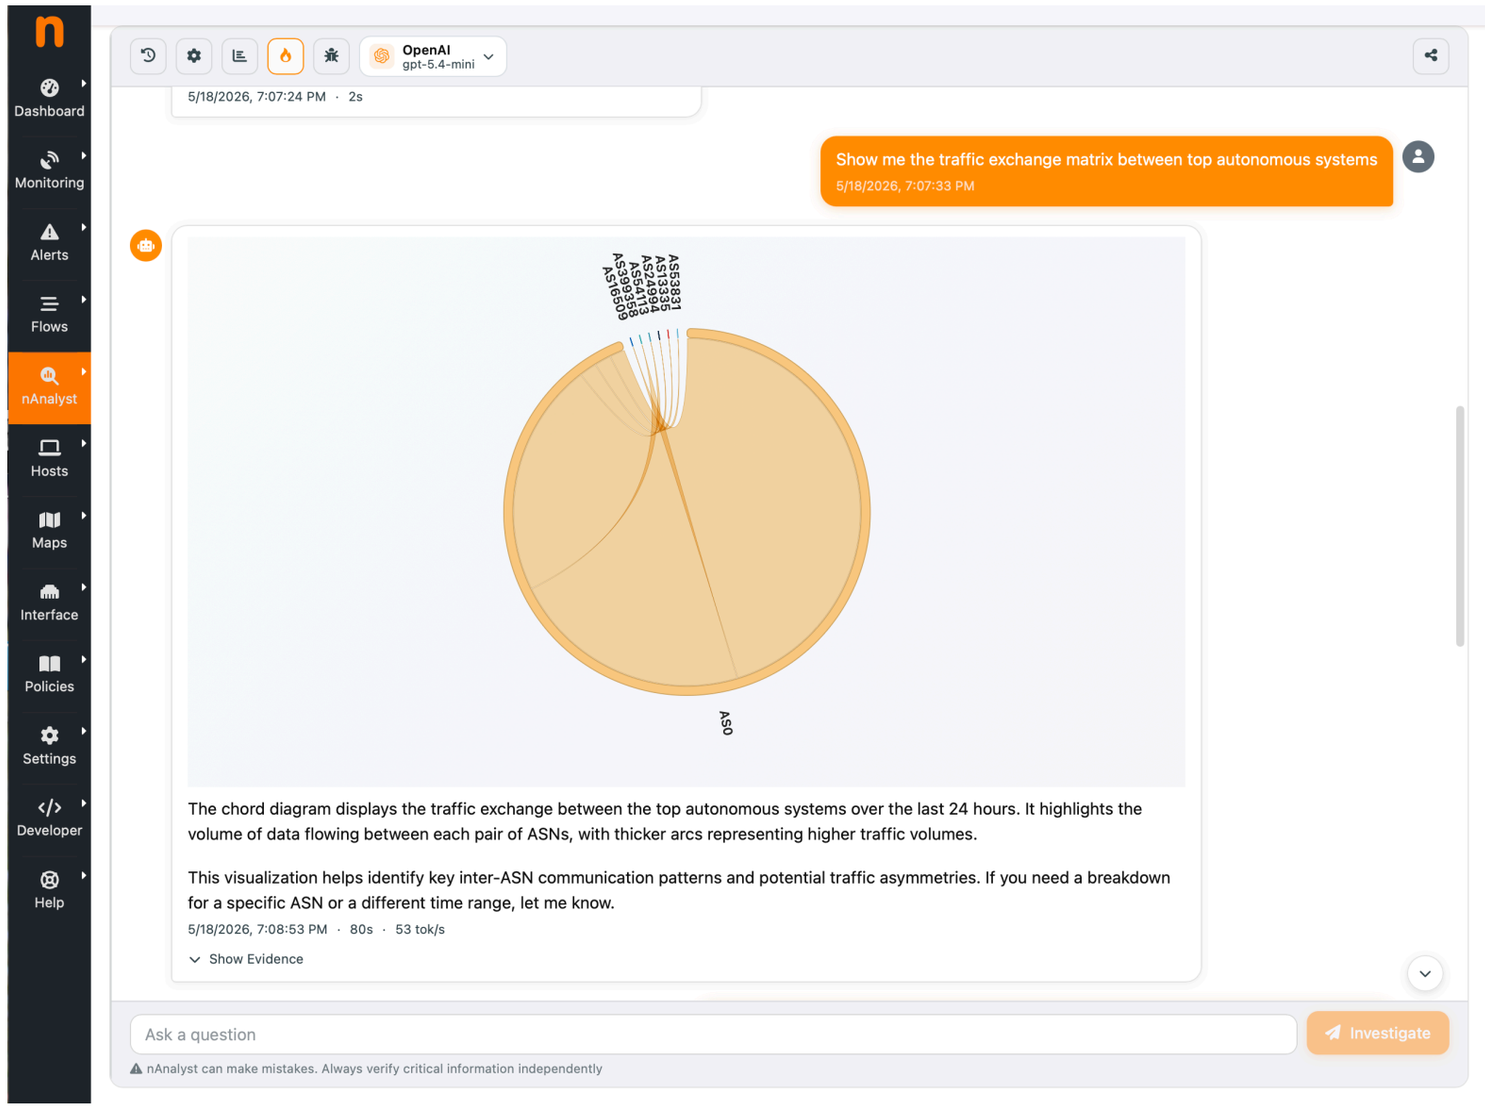

https://github.com/ntop/ntopng

https://github.com/ntop/nProbe/tree/master/bgp

Credits

- Federico Santulli - NHM (AS 62275)
- Vasja Krizmancic - Karsolink (AS 204471)
- Paolo Caparrelli - Goline (AS 202032)