Infusing Heterogeneous Data to Troubleshoot & Improve Peering Performance and Security

Some Use Cases

Julie Liu

March 2024

Some “BGP Routing Tasks”

BGP Coordinators, Network Engineers for BGP Operation……

Peering Coordination: peering evaluation suiting the peering policy

Traffic Route Management: traffic monitoring for traffic route optimization, anomaly identification and troubleshooting

Route Health Monitoring: BGP route message and RPKI status analysis for routing health monitoring

Route Anomaly Detection: it’d be even better if someone could notify me whenever abnormal route behaviors happened

➡ “*Anything you need to quantify can be measured in some way that is superior to not measuring it at all.*” —Gilb's Law:
Collect and fuse network data from multiple sources

Peering Coordination

To Peer or Not to Peer, That's the Question

Your Peering Policy: No, Open, Selective, Restricted?

- **No:** To buy transit cost efficiently. So need to identify the best candidate to buy transit from
- **Open:** To peer with as many networks as possible should it can save transit costs. So need to identify the right candidate to peer with
- **Selective:** To peer only with those who have significant values to us. So need to verify whether a peering request makes sense
- **Restrictive:** To better run our transit business operations. So need to understand 'transit prospective customers' traffic behavior for compelling business case building

➡ **Multiple network datasets:**
Flow and BGP data

Peering Coordination

Use Case

Open Peering Policy

Know who you wanna negotiate with: Create a list of networks with whom you exchange traffic but aren't peering with yet, and rank them by traffic volume. Then learn other traffic behavior patterns of these networks

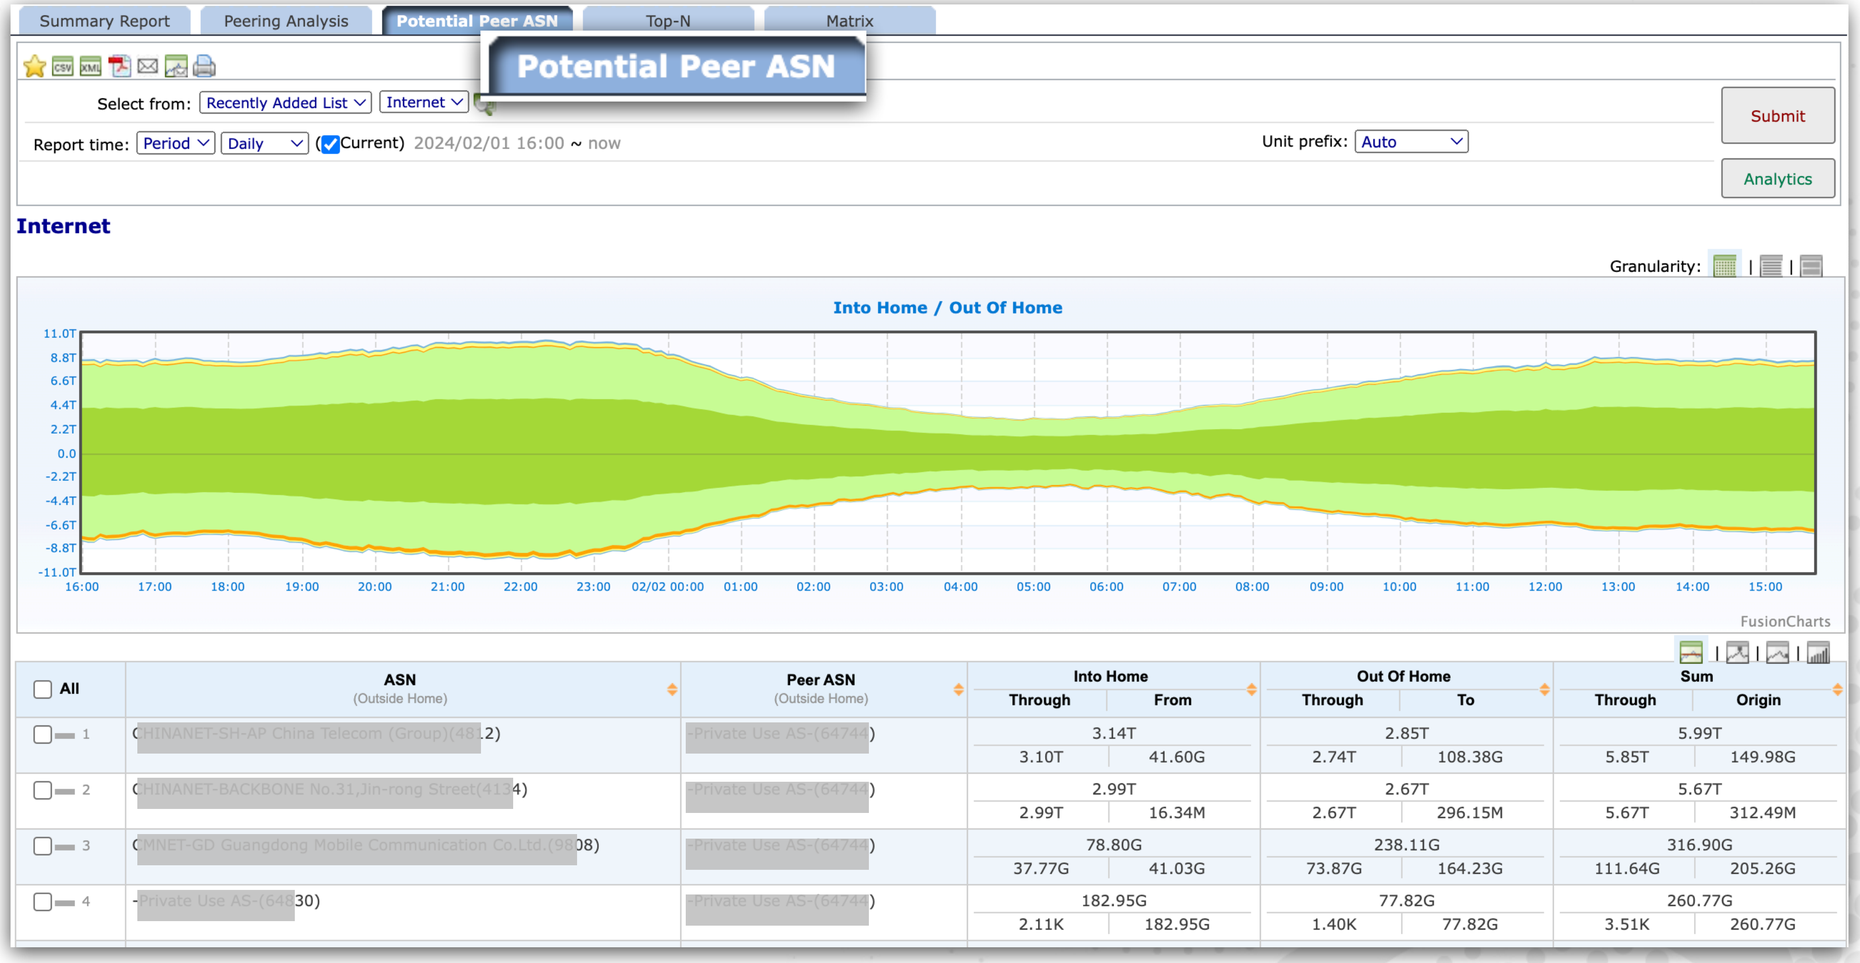

All	ASN (Outside Home)	Peer ASN (Outside Home)	Into Home		Out Of Home		Sum	
			Through	From	Through	To	Through	Origin
1	CHINANET-SH-AP China Telecom (Group)(4812)	-Private Use AS-(64744)	3.14T		2.85T		5.99T	
			3.10T	41.60G	2.74T	108.38G	5.85T	149.98G
2	CHINANET-BACKBONE No.31,Jin-rong Street(4134)	-Private Use AS-(64744)	2.99T		2.67T		5.67T	
			2.99T	16.34M	2.67T	296.15M	5.67T	312.49M
3	CMNET-GD Guangdong Mobile Communication Co.Ltd.(9808)	-Private Use AS-(64744)	78.80G		238.11G		316.90G	
			37.77G	41.03G	73.87G	164.23G	111.64G	205.26G
4	-Private Use AS-(64830)	-Private Use AS-(64744)	182.95G		77.82G		260.77G	
			2.11K	182.95G	1.40K	77.82G	3.51K	260.77G

Peering Coordination

Use Case

Open Peering Policy

Know who you wanna negotiate with: Create a list of networks with whom you exchange traffic but aren't peering with yet, and rank them by traffic volume. Then learn other traffic behavior patterns of these networks

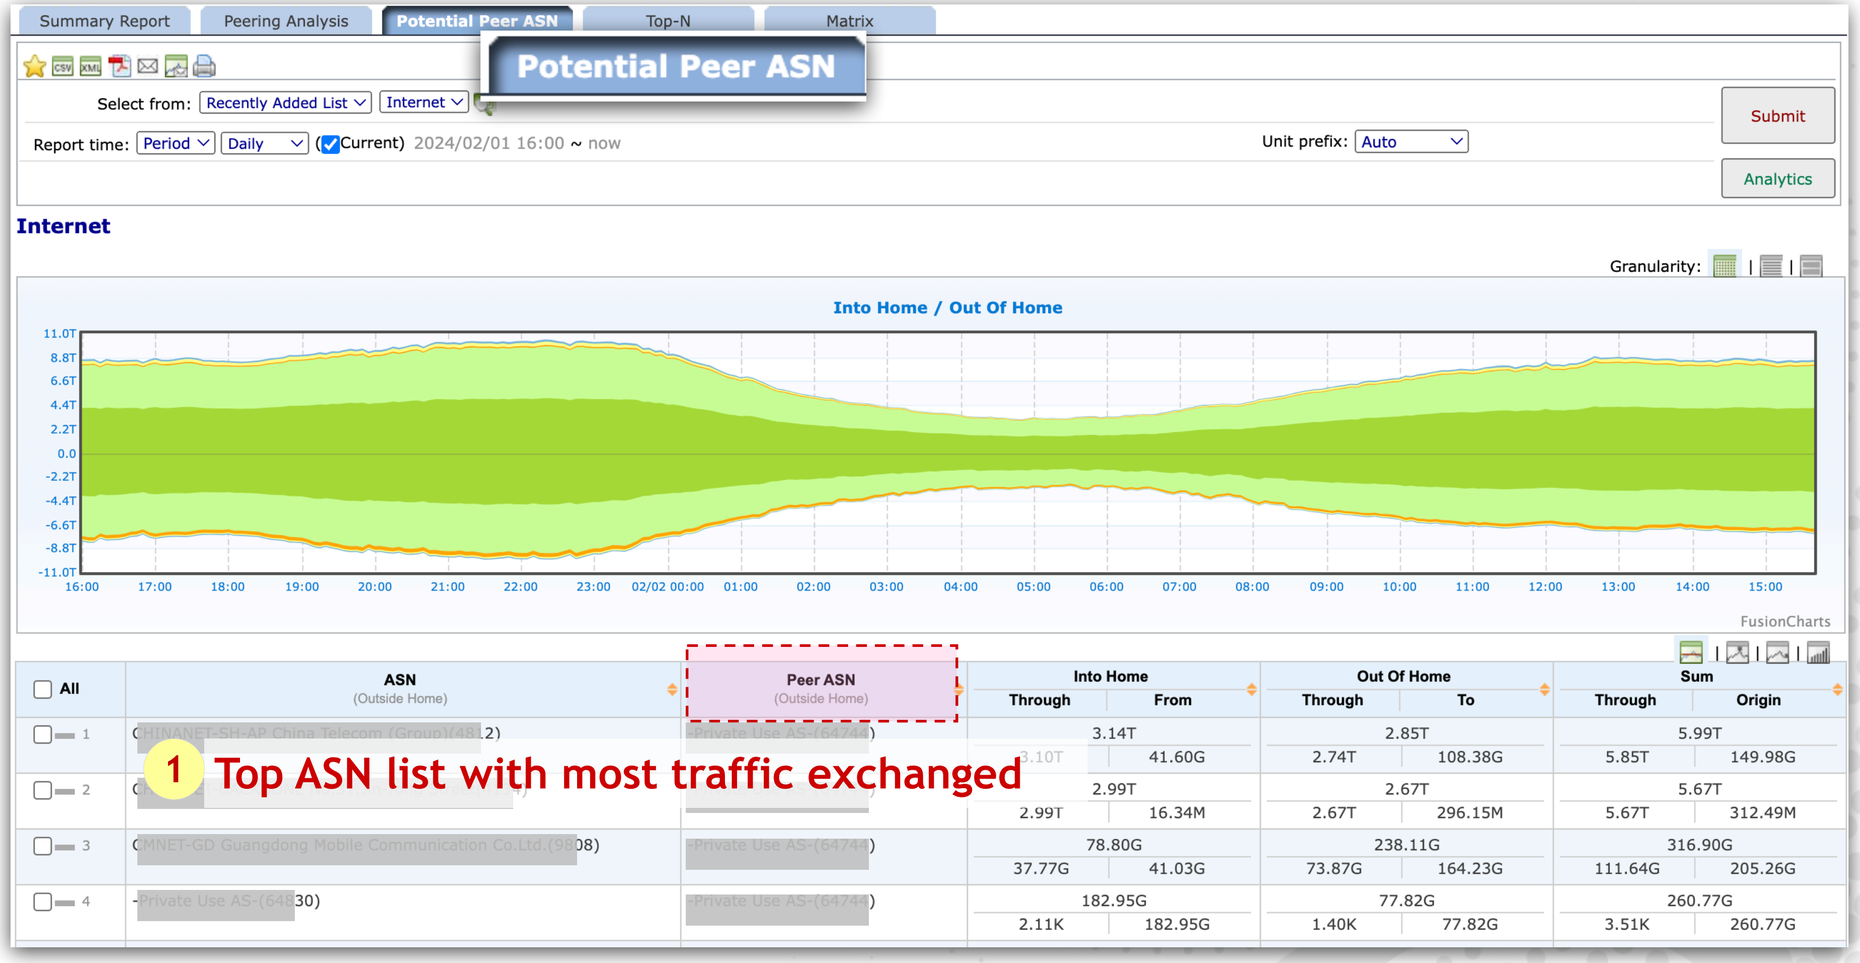

Peering Coordination

Use Case

Open Peering Policy

Know who you wanna negotiate with: Create a list of networks with whom you exchange traffic but aren't peering with yet, and rank them by traffic volume. Then learn other traffic behavior patterns of these networks

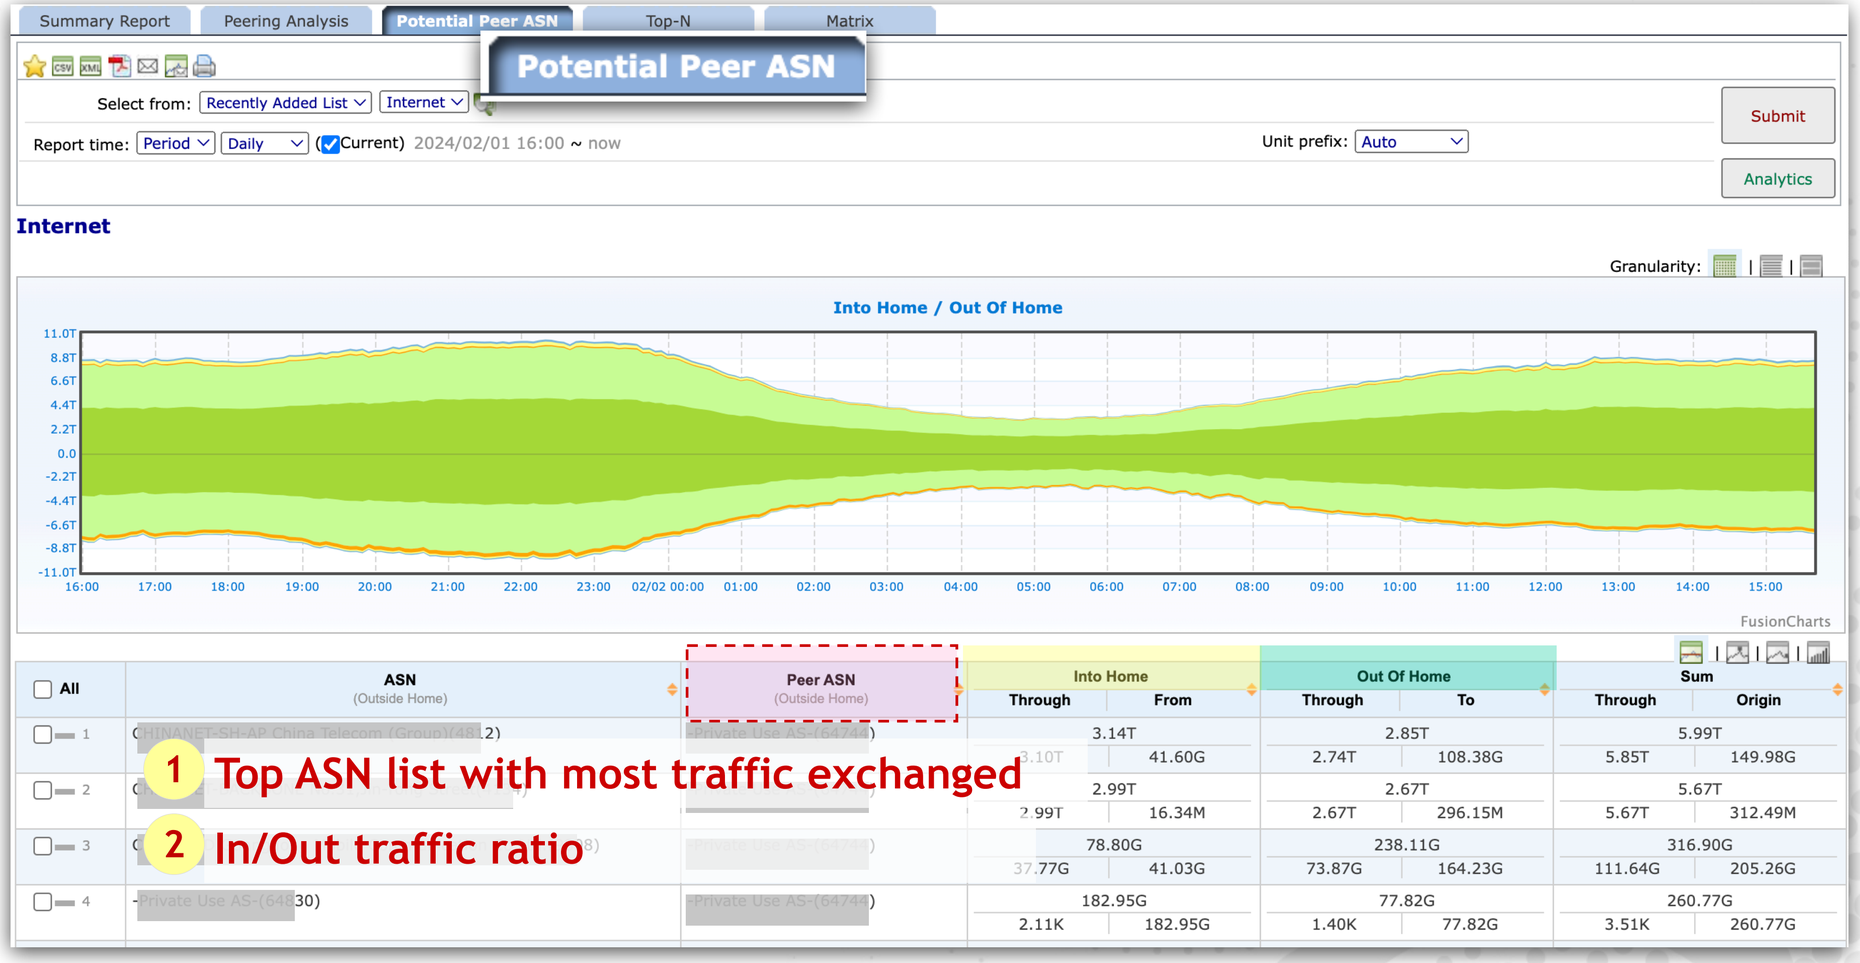

Peering Coordination

Use Case

Open Peering Policy

Know who you wanna negotiate with: Create a list of networks with whom you exchange traffic but aren't peering with yet, and rank them by traffic volume. Then learn other traffic behavior patterns of these networks

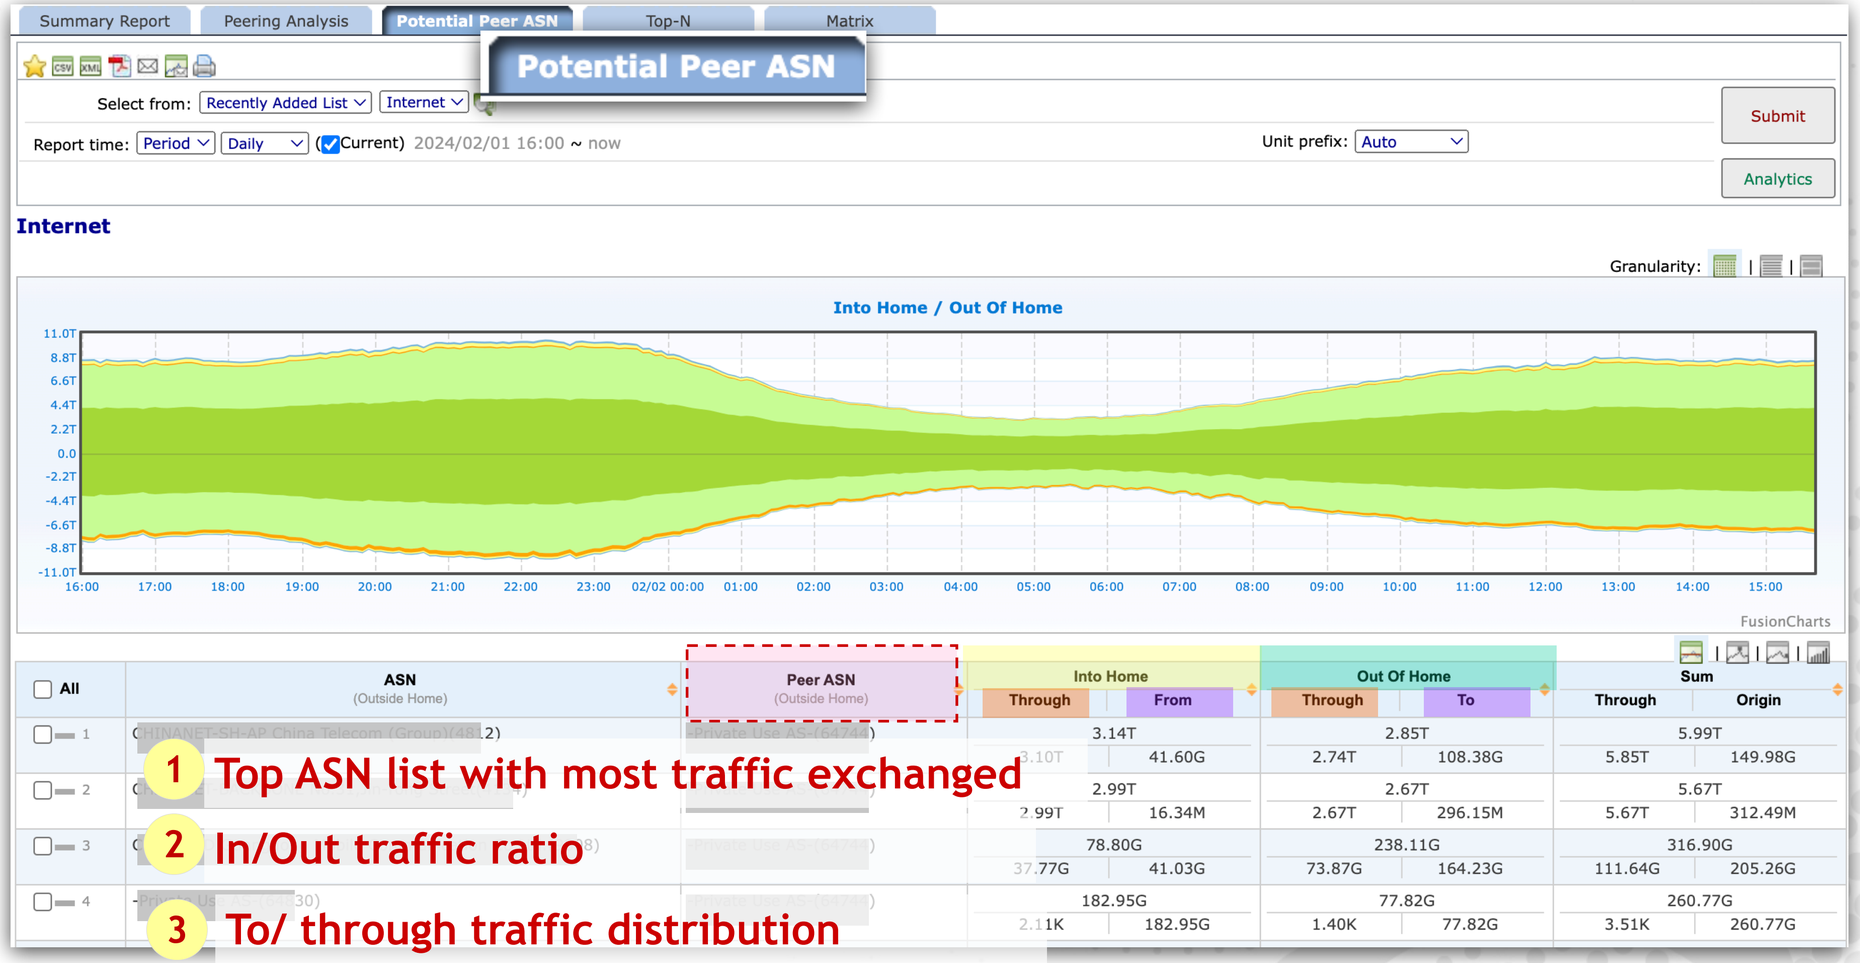

Peering Coordination

Use Case

Cost Analysis

Does it make sense to expand your network in order to peer.

Get quotes for port access, transport and transit from the providers.

Estimate how much traffic you will be able to peer. Compute how much it's going to cost to build a new peer.

Is cost lower than how much you pay for transiting the traffic?

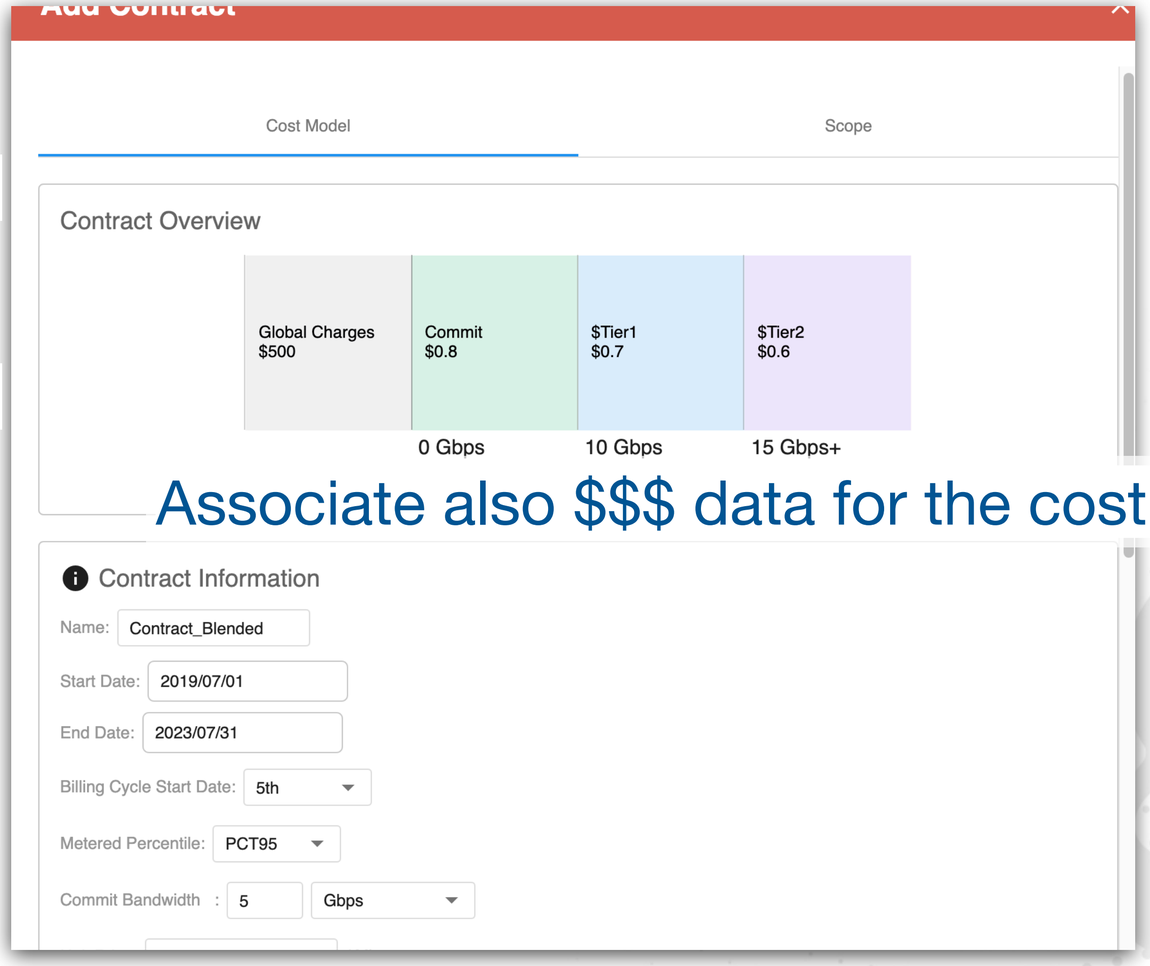

Associate also $$$ data for the cost analysis

Peering Coordination

Use Case

Cost Analysis

Does it make sense to expand your network in order to peer.

Get quotes for port access, transport and transit from the providers.

Estimate how much traffic you will be able to peer. Compute how much it's going to cost to build a new peer.

Is cost lower than how much you pay for transiting the traffic?

Associate also $$$ data for the cost analysis

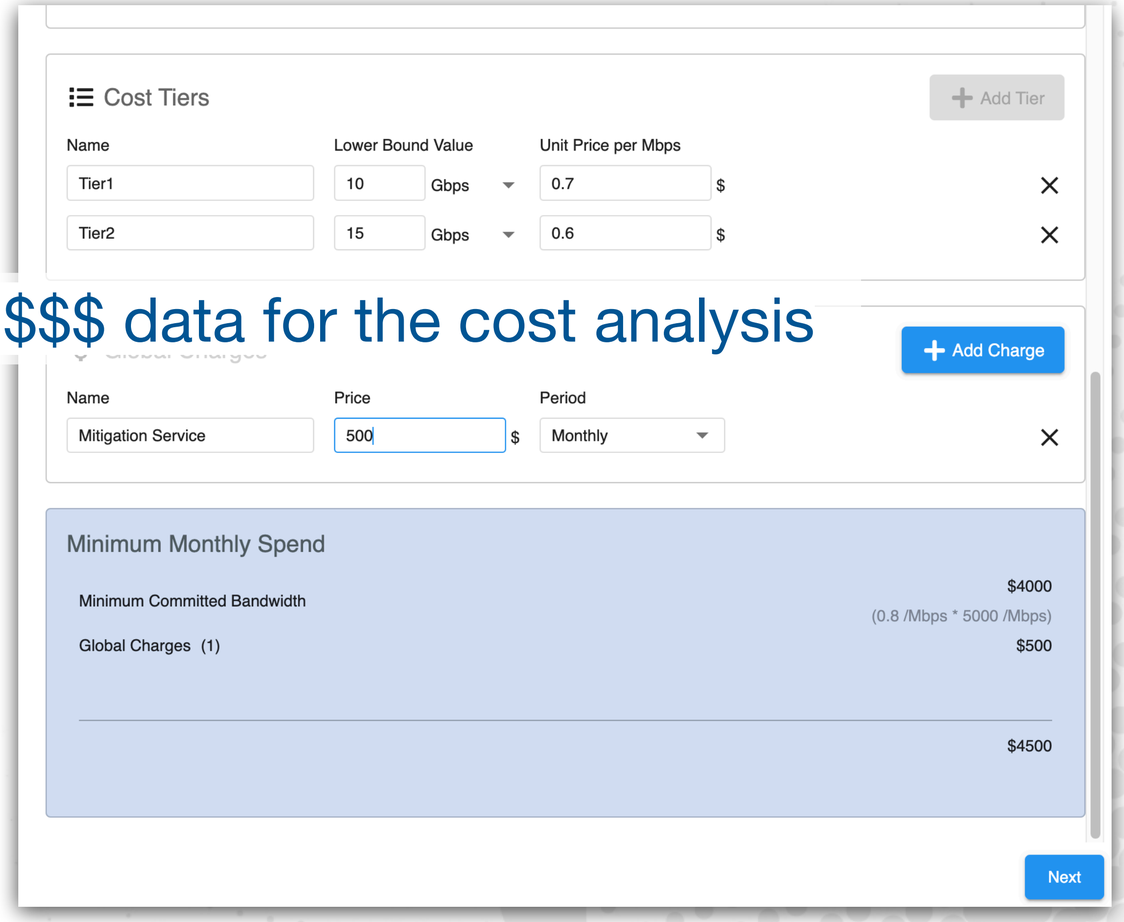

Peering Coordination

Use Case

Cost Analysis

Does it make sense to expand your network in order to peer.

Get quotes for port access, transport and transit from the providers.

Estimate how much traffic you will be able to peer. Compute how much it's going to cost to build a new peer.

Is cost lower than how much you pay for transiting the traffic?

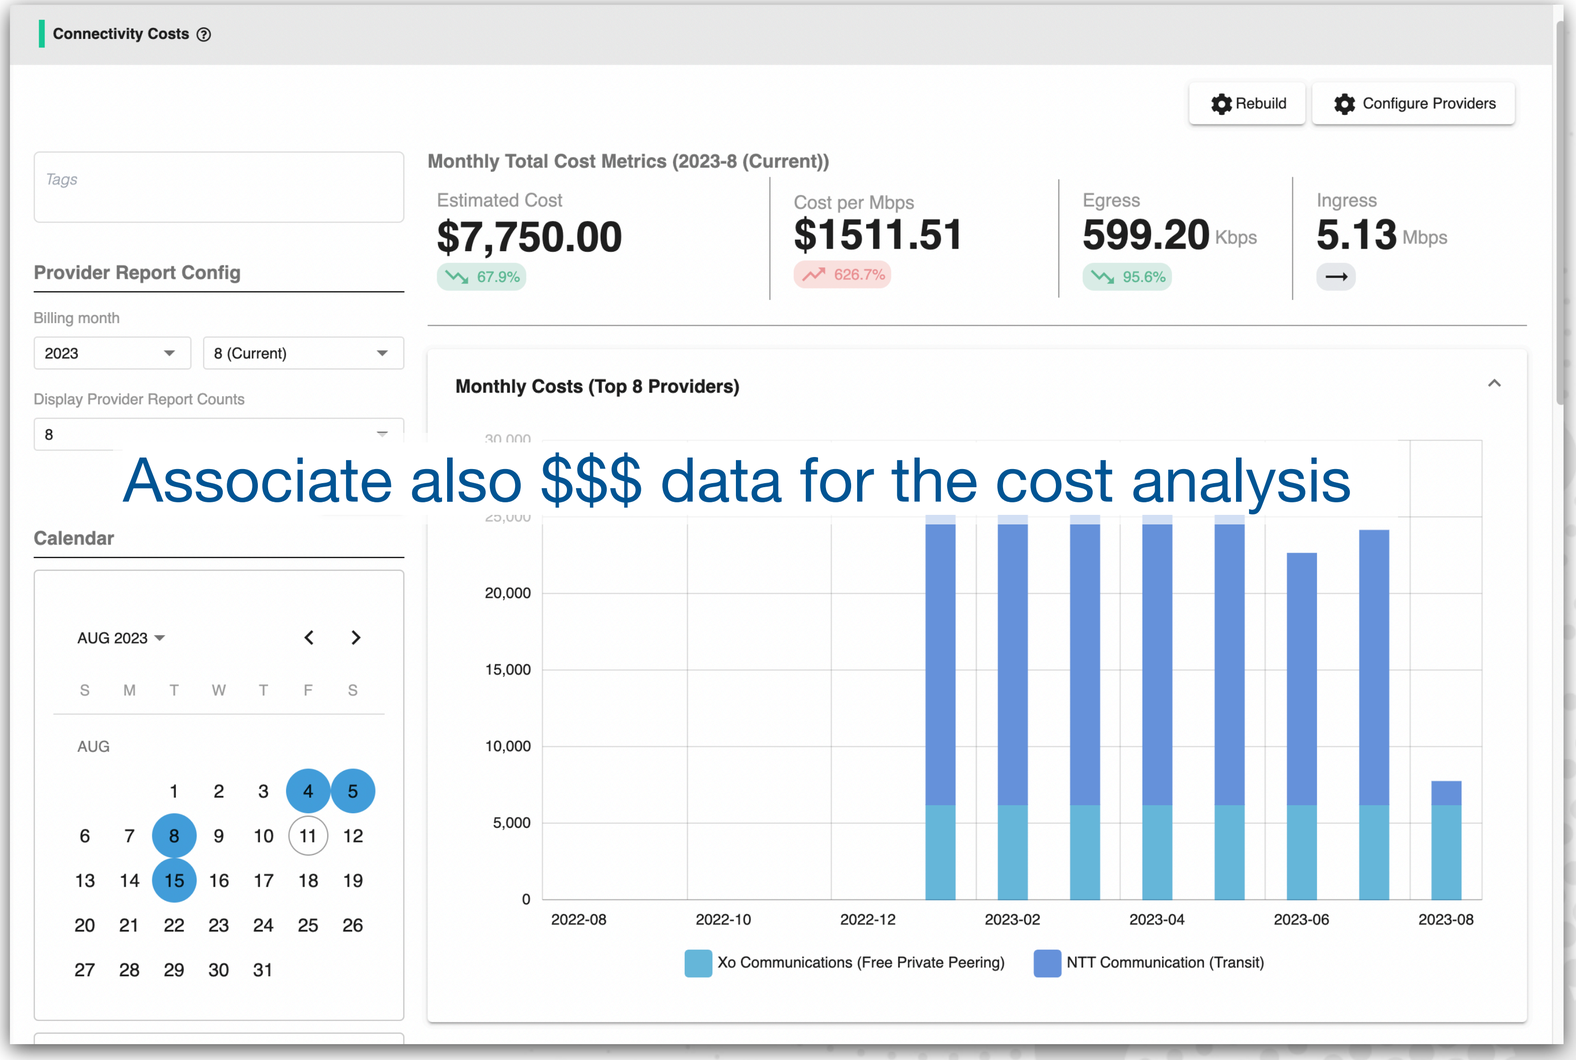

Associate also $$$ data for the cost analysis

Traffic Route Monitoring

Monitoring for Optimization

Traffic engineering: for better performance, resource utilization or congestion avoidance

- Knowing **What traffic** is **leaving/entering your network @where** is helpful for adjusting how the traffic going across the network
- E.g., How much traffic is going to an AS, an AS Path, a BGP Community, or a Next-hop?
- E.g., How much traffic is going through an interface/link or a set of interfaces/links?
- E.g., What's or who's traffic it is? Top talkers/listeners, services, etc.

  Have facts and figures then may help act to deterministically move the traffic around

➡ **Multiple network datasets:**

Flow, BGP and network device/interface (SNMP) data

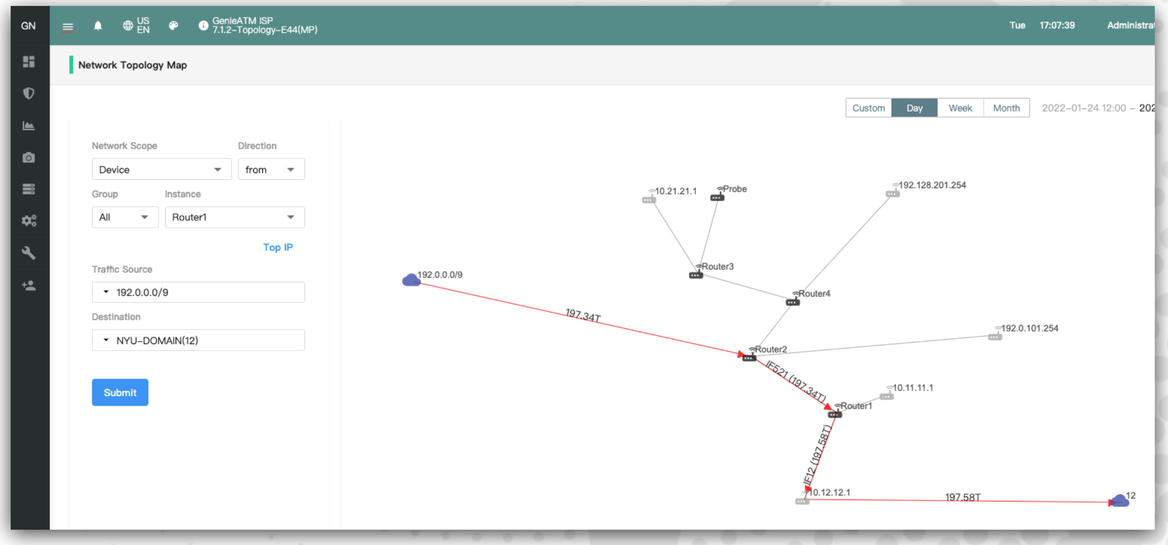

Traffic Route Monitoring

Use Case

Apply route policy with actual measurements

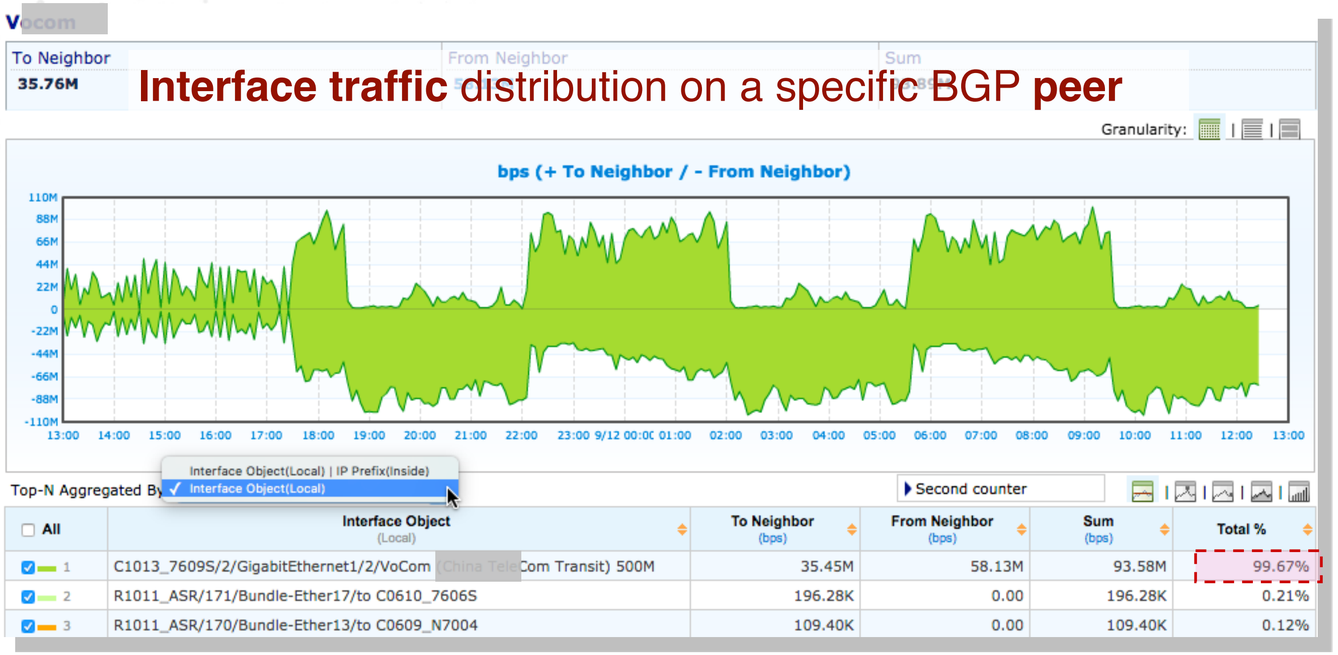

- 3 potential exit points to reach ISP-V
- Currently most thru router C1's link
- Want to shift some traffic to other links
- What traffic shall be shifted?

Traffic Route Monitoring

Use Case

Apply route policy with actual measurements

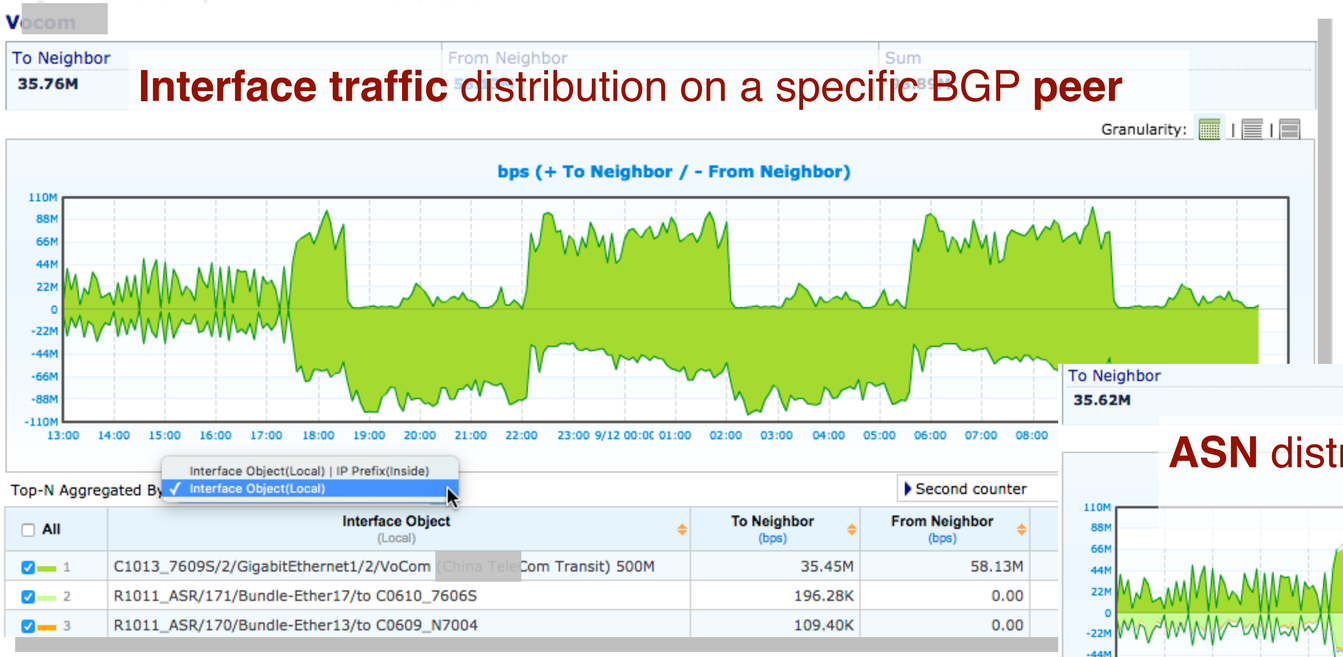

- 3 potential exit points to reach ISP-V
- Currently most thru router C1's link
- Want to shift some traffic to other links
- What traffic shall be shifted?

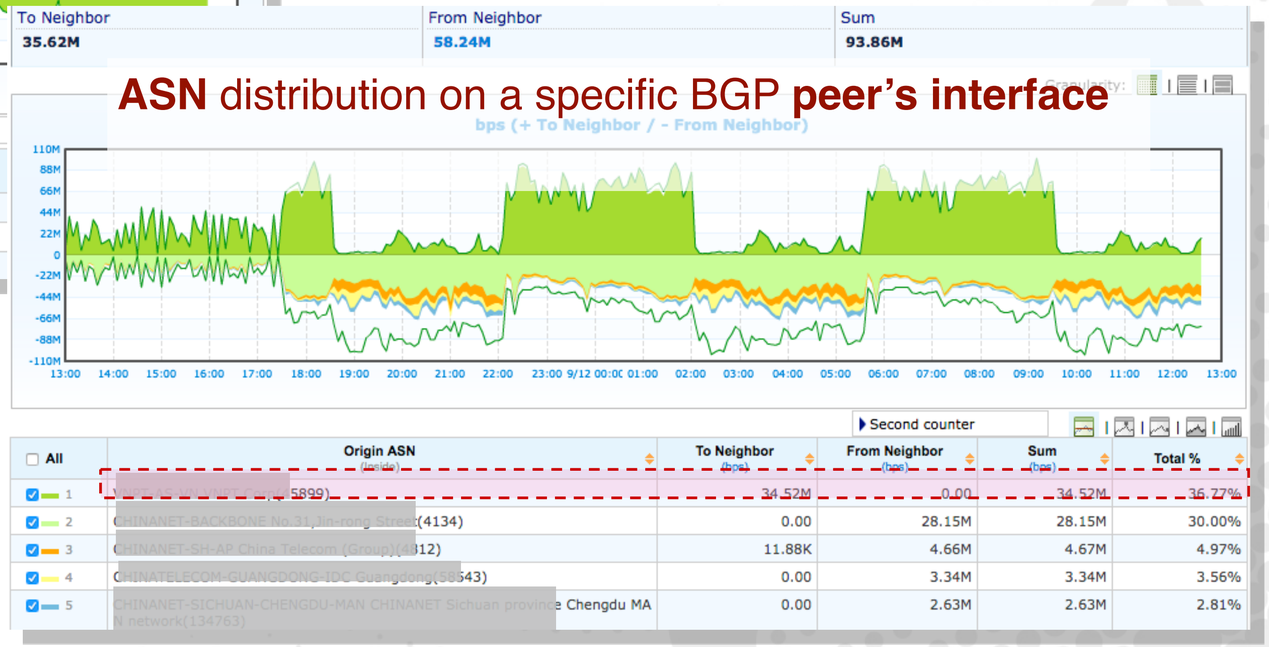

- We may like to move certain share of traffic (e.g. 36%, traffic going to an ASN) from one interface (C's1) to another by BGP methods (e.g., tuning LOCAL_PRF)

Traffic Route Troubleshooting

Use Case

Identify Unreasonable Routes

BGP setting verification: verify whether we have made all configuration changes as we expected to facilitate the expected traffic route changes

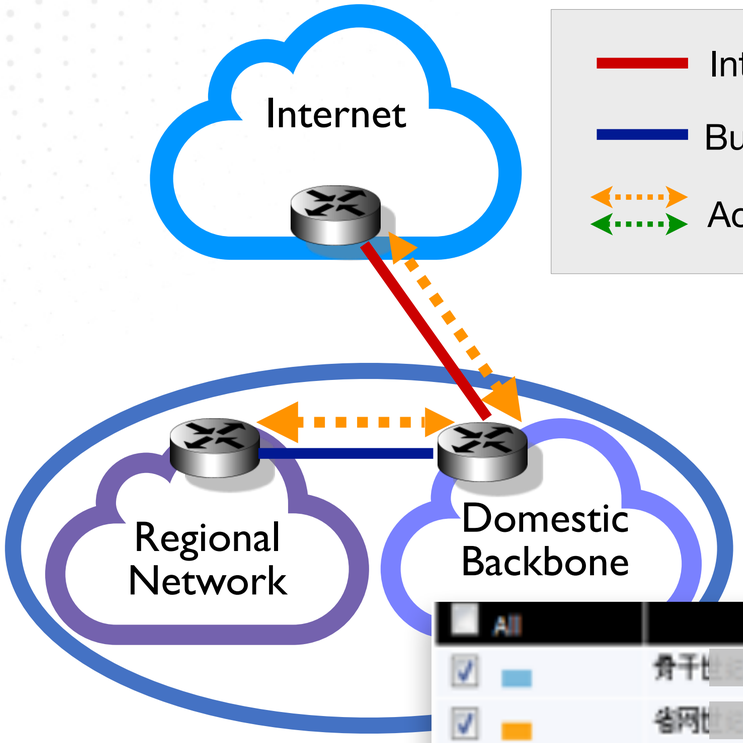

Internet Exit

Busy Link to offload

Actual Traffic routes

All	Filter Name	Filter Direction	Opposite Direction	Sum	Total %
☑	骨干[redacted]	7,855.30K	95.28M	102.95M	78.44%
☑	省网[redacted]	20.82M	7,659.53K	28.30M	21.56%
☑	Total	28.49M	102.76M	131.25M	100.00%

Traffic Route Troubleshooting

Identify Unreasonable Routes

BGP setting verification: verify whether we have made all configuration changes as we expected to facilitate the expected traffic route changes

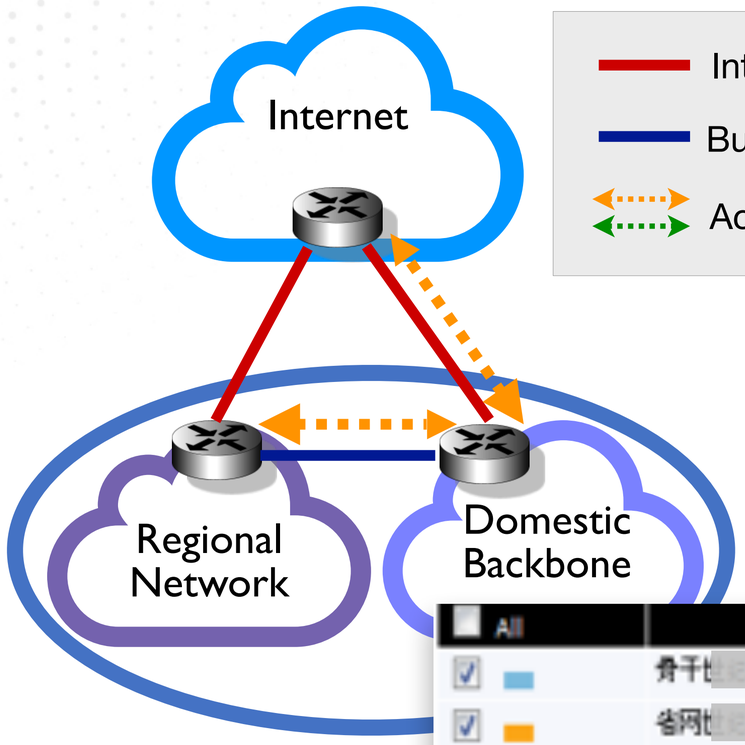

Internet Exit

Busy Link to offload

Actual Traffic routes

In order to offload the regional network's Internet traffic from the busy link connecting to its domestic backbone, the SP has added a direct exit link from the regional network to the Internet.

All	Filter Name	Filter Direction	Opposite Direction	Sum	Total %
☑	骨干[redacted]	7,855.30K	95.28M	102.95M	78.44%
☑	省网[redacted]	20.82M	7,659.53K	28.30M	21.56%
☑	Total	28.49M	102.76M	131.25M	100.00%

Traffic Route Troubleshooting

Identify Unreasonable Routes

BGP setting verification: verify whether we have made all configuration changes as we expected to facilitate the expected traffic route changes

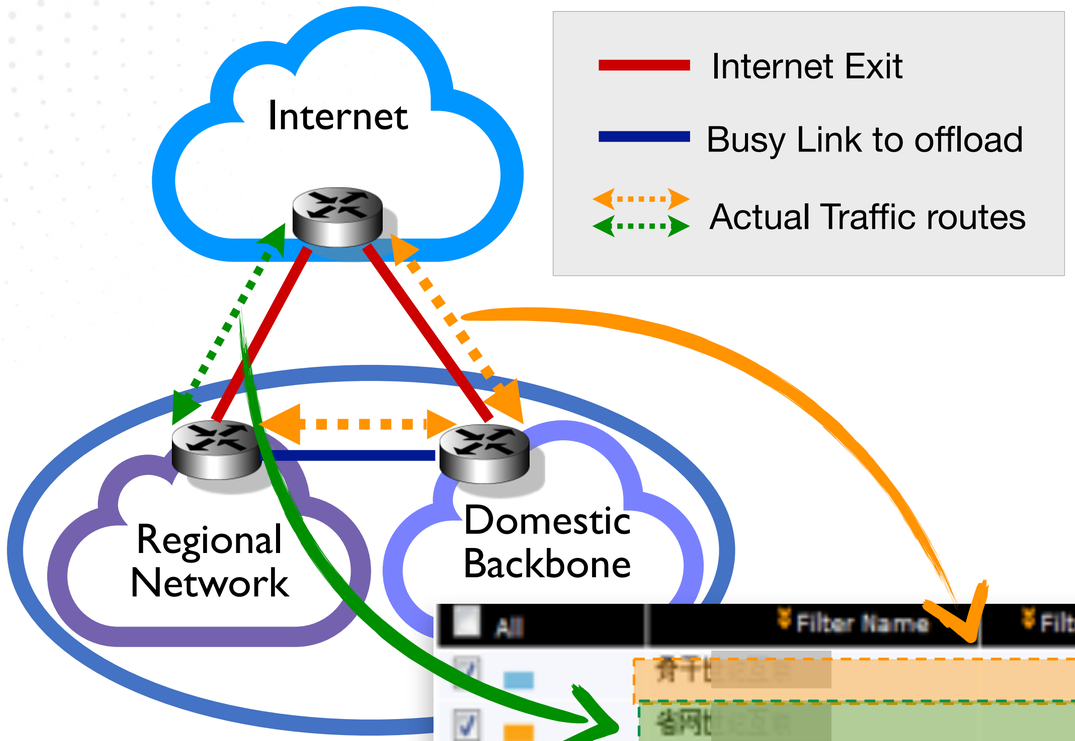

In order to offload the regional network's Internet traffic from the busy link connecting to its domestic backbone, the SP has added a direct exit link from the regional network to the Internet.

Despite completing the corresponding configuration changes, the regional network is found continues routing most of its Internet traffic through the link to the domestic backbone, wasting the network resources.
It turned out that some BGP policies were not changed along correctly…

All	Filter Name	Filter Direction	Opposite Direction	Sum	Total %
	青干[redacted]	7.855.30K	95.28M	102.95M	78.44%
	省网[redacted]	20.82M	7,659.53K	28.30M	21.56%
	Total	28.49M	102.76M	131.25M	100.00%

Traffic Route Troubleshooting

Identify Unreasonable Routes

Example: peers dumping traffic at you for routes you didn't send them

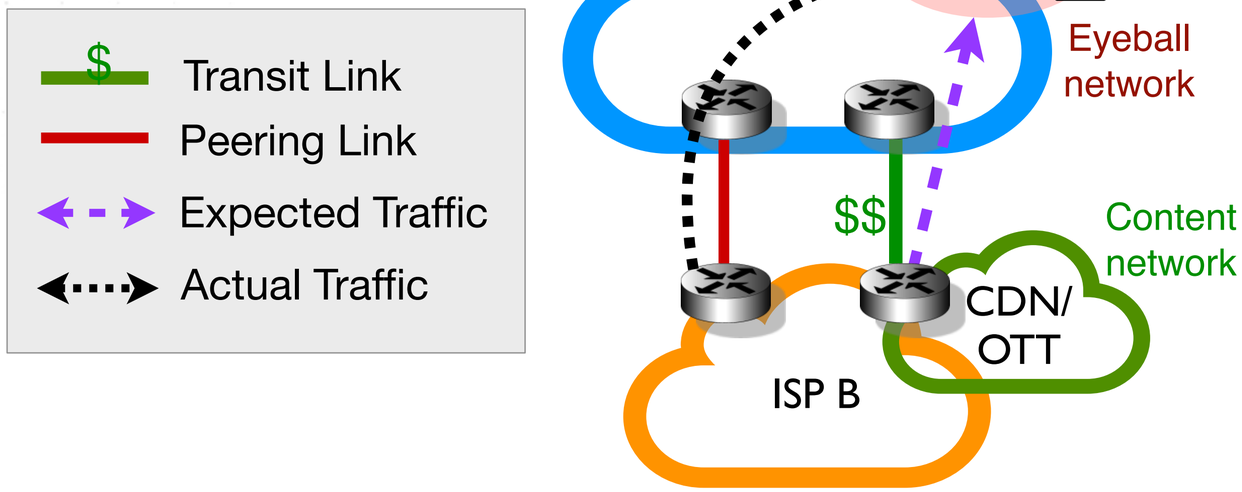

Instead of diverting the traffic through paid transit links, the content traffic is dumped at ISP A through the free-peering arrangement without knowing agreement

ISP A considers this undermined the terms of peering arrangement and unfairly exploits ISP A's resources

➡ **Multiple network datasets:**
Flow, BGP and DNS (for identifying CDN or OTT service providers)

Route Health Monitoring

Route Diagnostics

BGP route instability: know the BGP instability and its source (peer, prefix)

- Analyze the BGP routes received
- Route Events/Selection Events: AAdup, AAdiffBetter, AAdiffWorse, Wdown, Wnull, Ainit, WAdup, WAdiff, Tup, Tbetter, Tworse, Tdup, TW, etc.

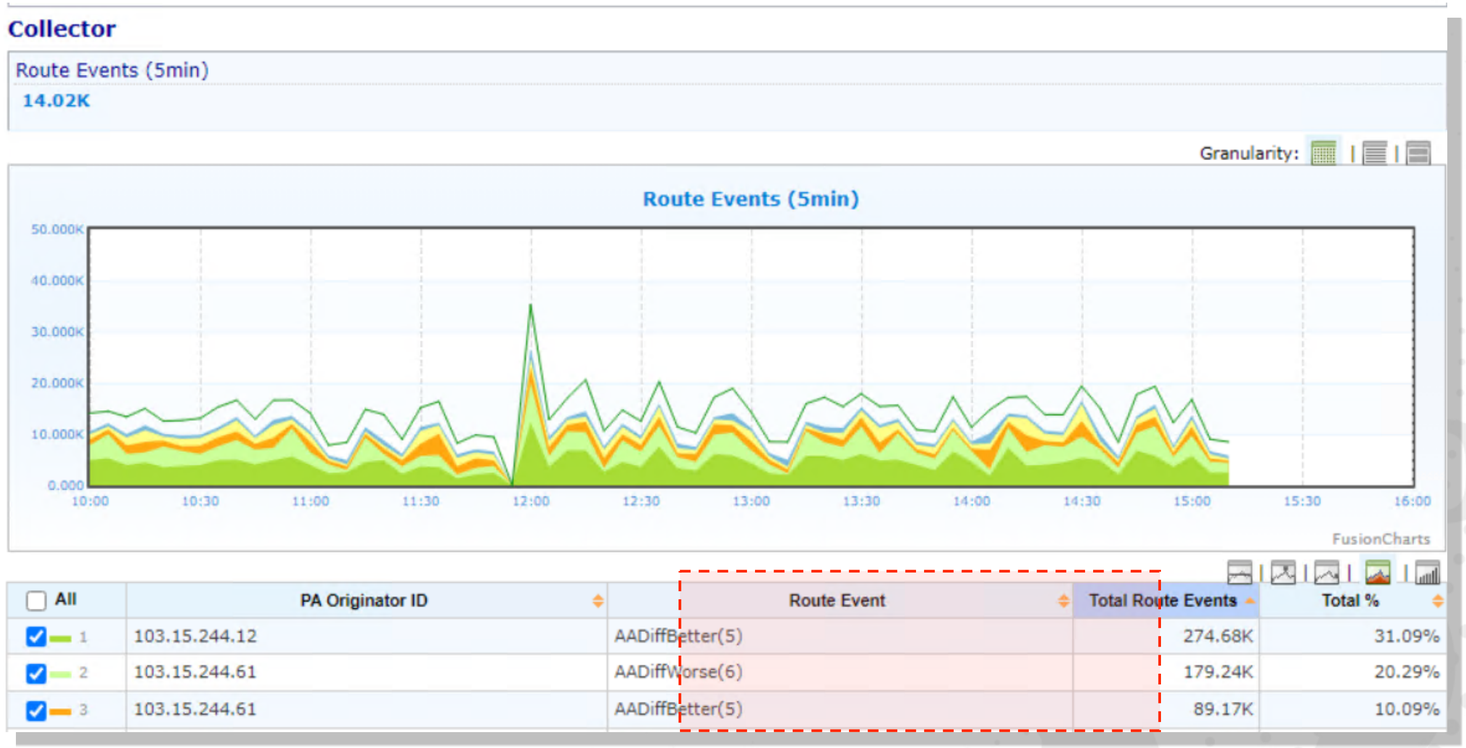

➡ **Datasets:**
BGP messages

Route Health Monitoring

Route Diagnostics

BGP route instability: know the BGP instability and its source (peer, prefix)

- Analyze the BGP routes received
- Route Events/Selection Events: AAdup, AAdiffBetter, AAdiffWorse, Wdown, Wnull, Ainit, WAdup, WAdiff, Tup, Tbetter, Tworse, Tdup, TW, etc.

➡ **Datasets:**
BGP messages

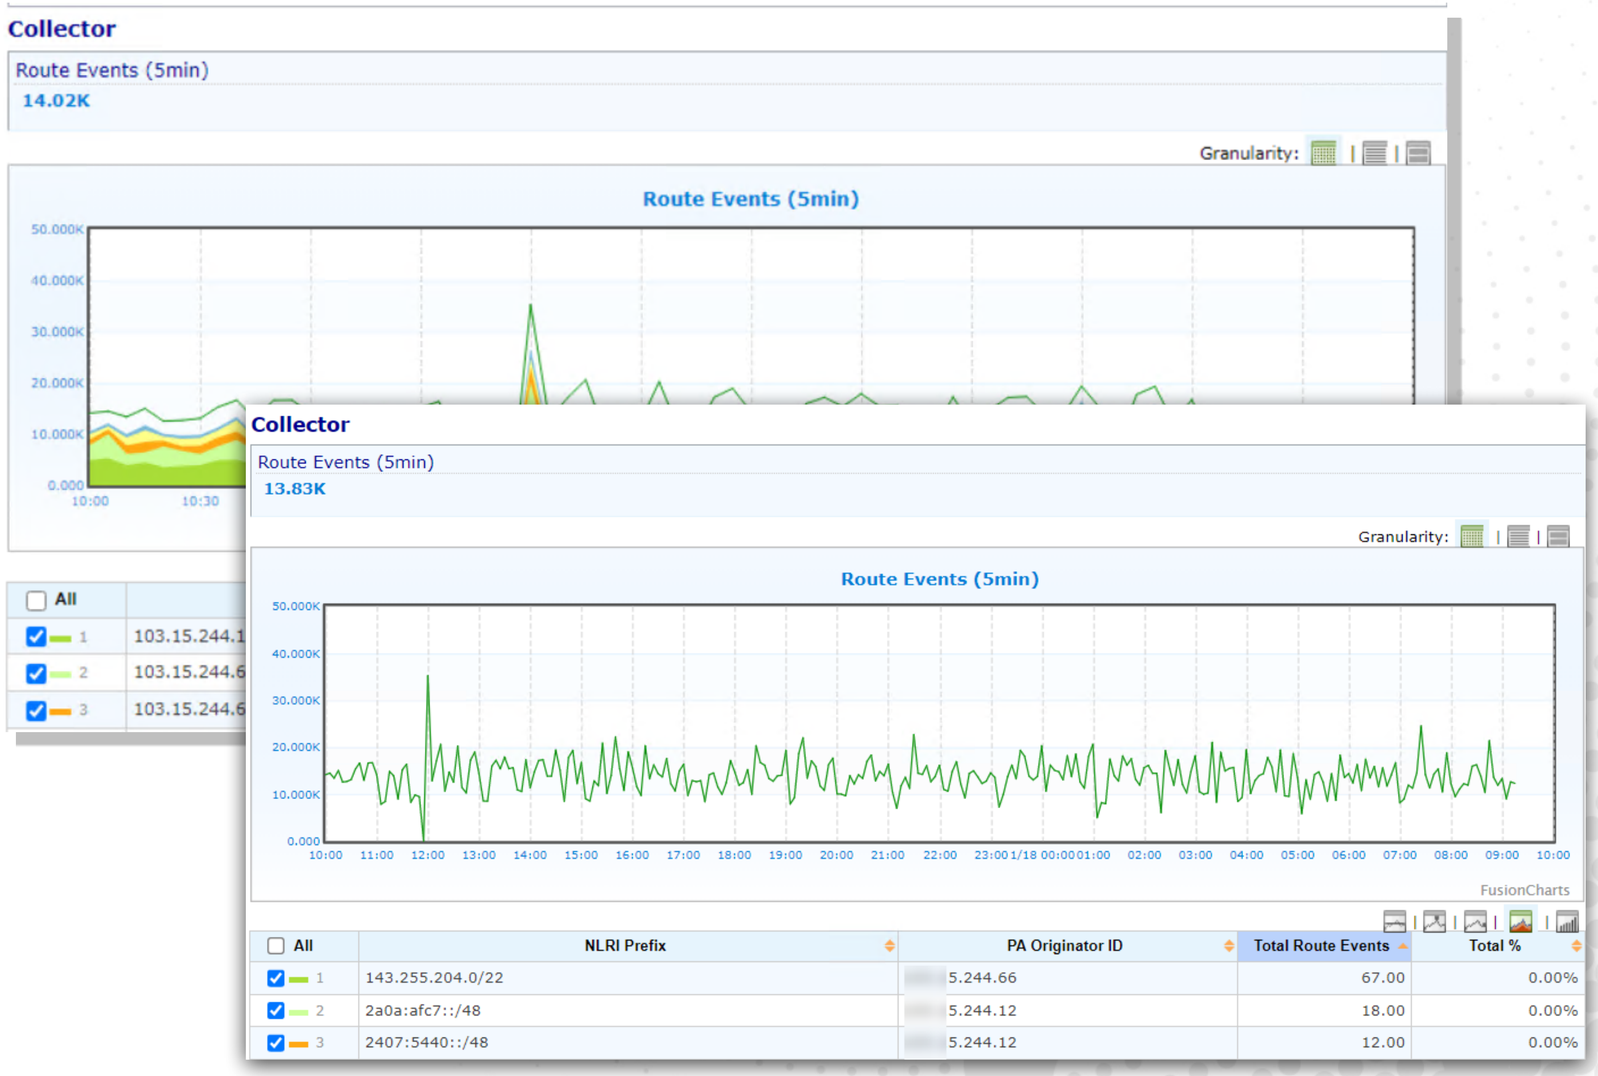

Route Health Monitoring

Route Diagnostics

BGP RPKI status: know the RPKI validation status of the BGP routes received

- RPKI status: Valid, Invalid ASN, Invalid length, Unknown
- Can also analyze the RPKI status of traffic flows

➡ **Multiple datasets:** BGP, RPKI and traffic Flow data

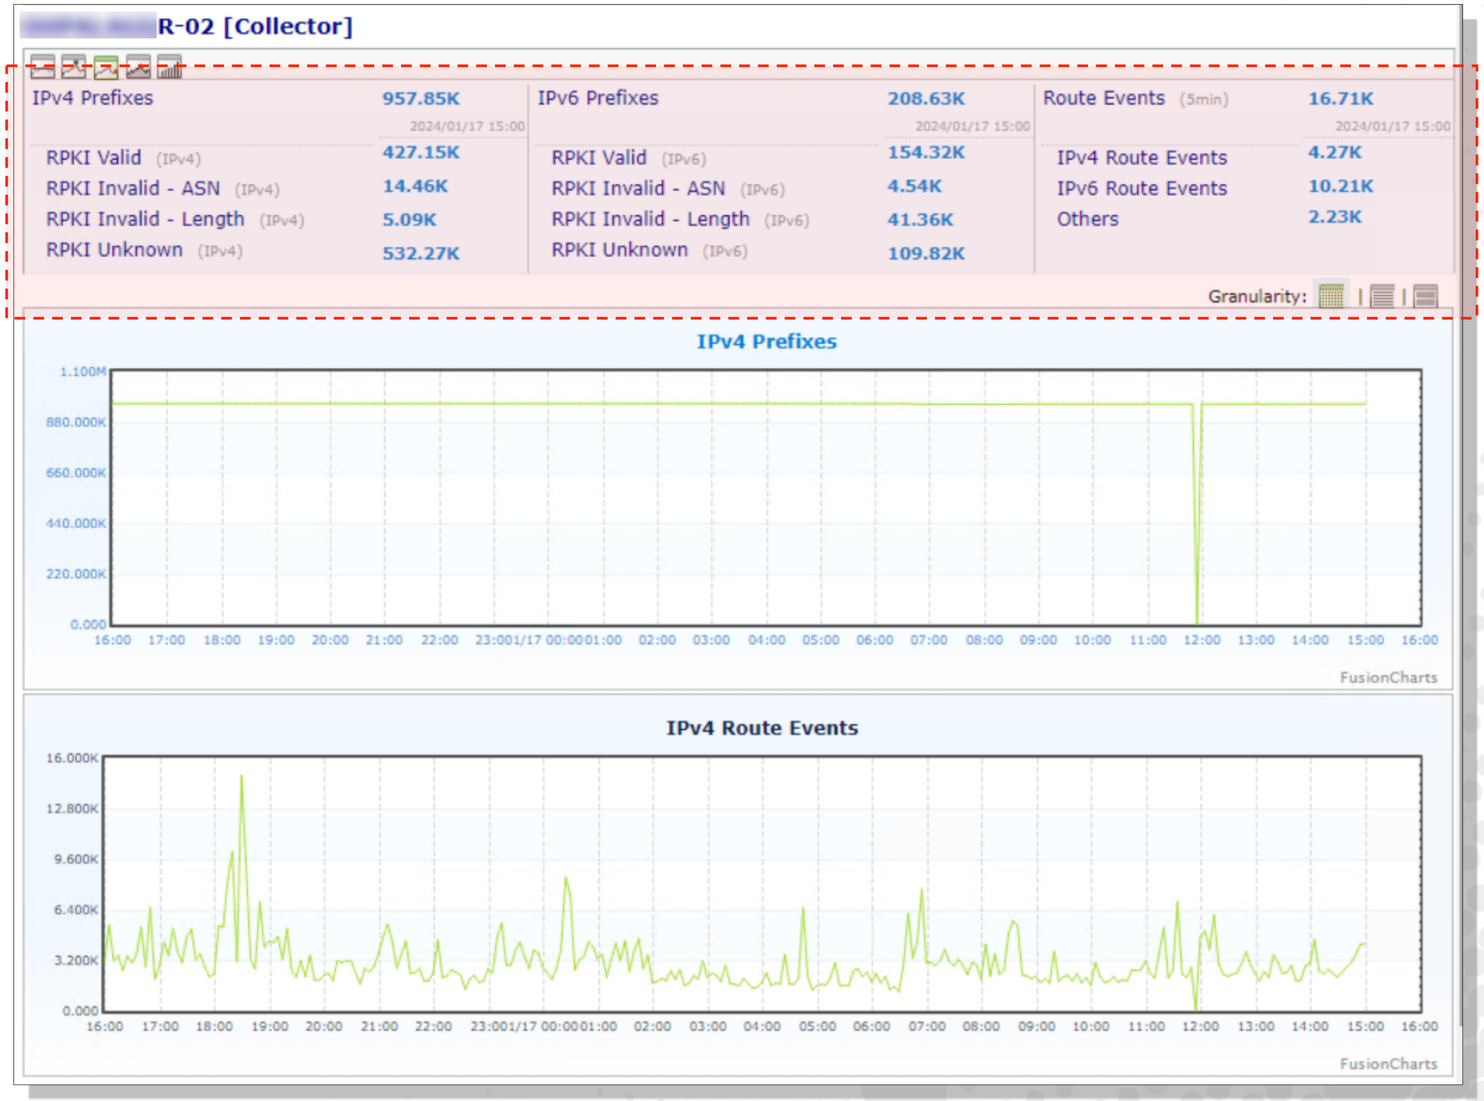

Route Health Monitoring

Route Diagnostics

BGP RPKI status: know the RPKI validation status of the BGP routes received

- RPKI status: Valid, Invalid ASN, Invalid length, Unknown
- Can also analyze the RPKI status of traffic flows

➡ **Multiple datasets:** BGP, RPKI and traffic Flow data

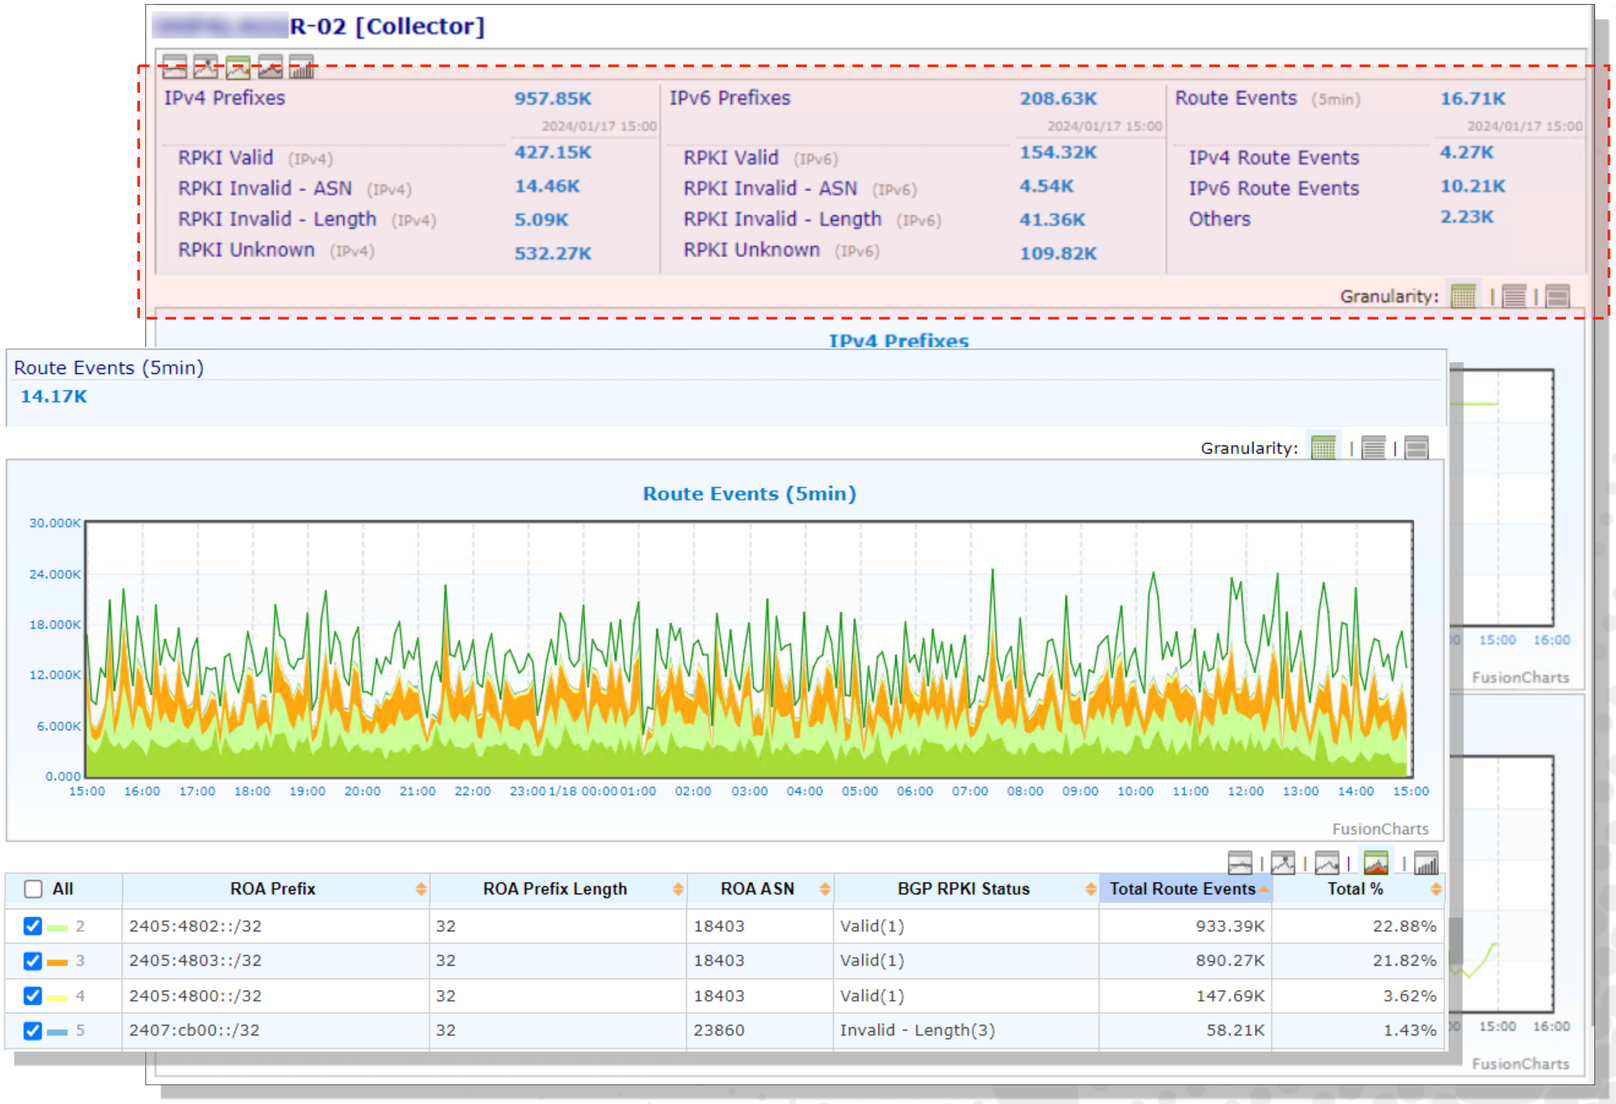

Route Anomaly Detection

BGP Route Anomalies

Alert when abnormal route behaviors detected:

- A BGP peer monitored by BMP goes up and down
- Route changes while RPKI invalid
- Unstable routes due to too many route change events
- Too frequent route announcements from a BGP router

➡ **Multiple datasets:**
BGP, RPKI and BMP data

Edit Baseline

Baseline

Detection Base: per-router

***Name:** BGP Route Anomaly Baseline

No.	ID	Name	No. of BGP update message		Remarks
			Status	Threshold	
1	1	BGP peer up/down	Enabled		
2	2	Route change (instable routes)	Enabled	1 events/5min	(Replace / Fail-over events)
3	3	Route change while RPKI invalid	Enabled	1 events/5min	(Add / Delete / Replace Route or Fail-over routes, while RPKI is invalid)
4	4	Unusual route announcements	Enabled	50 events/5min	(All route events)

Route Anomaly Detection

BGP Route Anomalies

Alert when abnormal route behaviors detected:

- A BGP peer monitored by BMP goes up and down
- Route changes while RPKI invalid
- Unstable routes due to too many route change events
- Too frequent route announcements from a BGP router

➡ **Multiple datasets:**
BGP, RPKI and BMP data

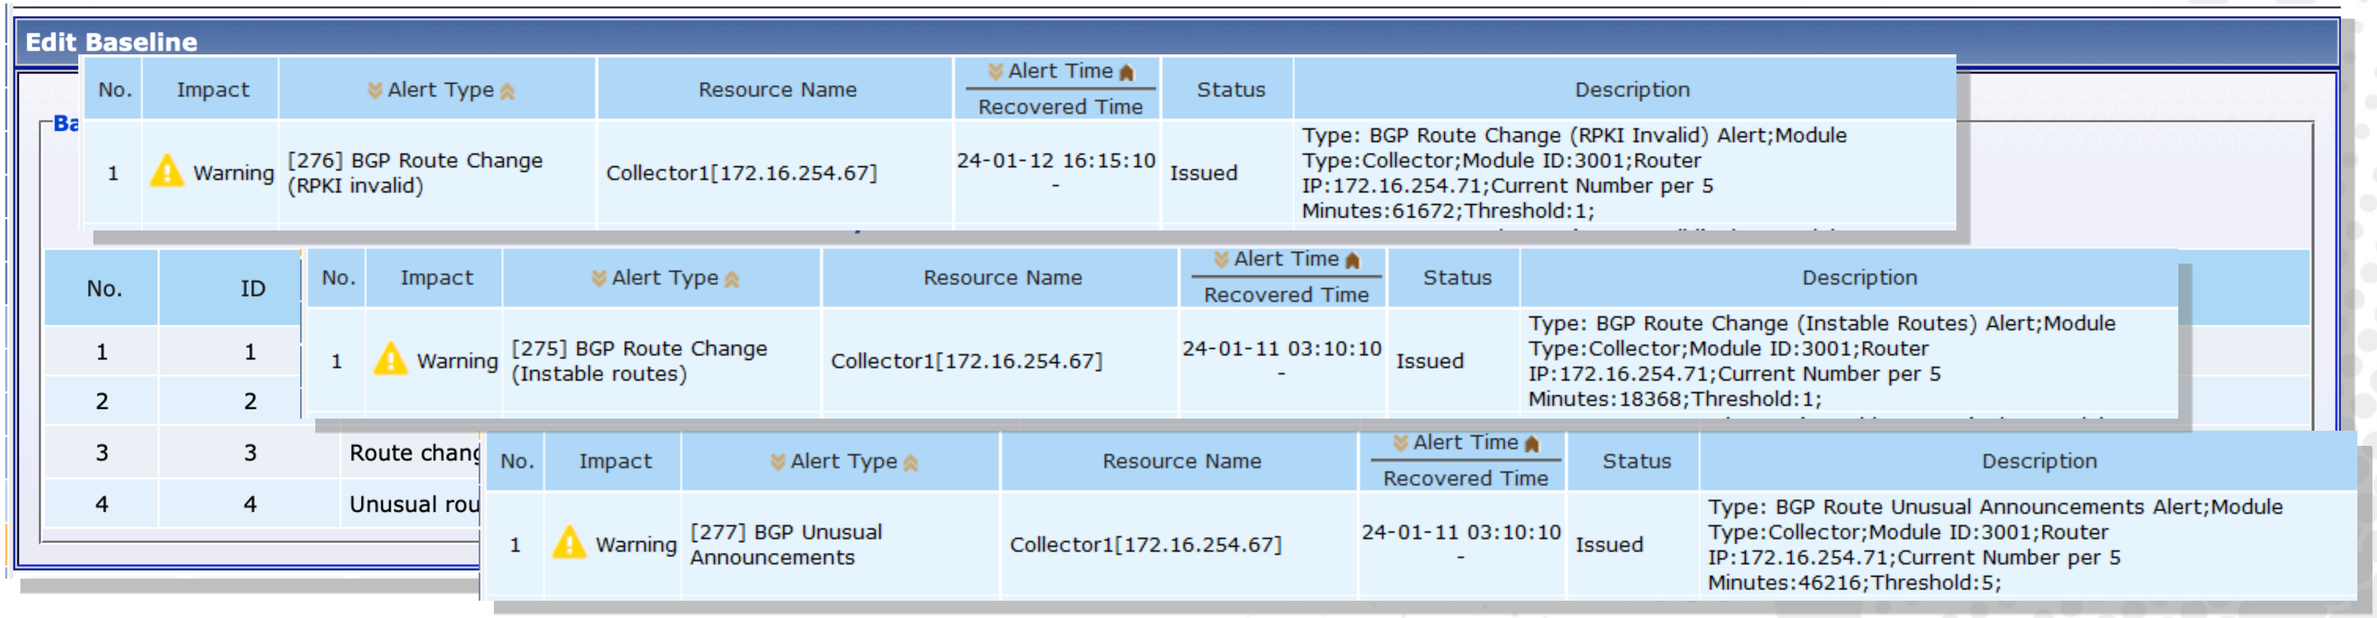

Key Takeaways

Some BGP Routing Tasks

Peering Coordination: candidate identification, peering evaluation with cost analysis, etc.

Traffic Route Monitoring: traffic route optimization, unreasonable route identification, traffic route troubleshooting, etc.

Route Health Monitoring: route event monitoring, RPKI status monitoring, etc.

Route Anomaly Identification: abnormal route behavior alerting

➡ **Correlate traffic Flow, BGP route, device SNMP, Service (DNS), RPKI and BMP data will help the tasks done more effectively and efficiently!**

Genie Networks at a Glance

Founded in **2000**
Headquartered in **Taipei Taiwan**

Solutions:
Carrier-grade Traffic Analytics & DDoS Protection

Over 450 carrier customers worldwide

NTT
KDDI
SoftBank
YAHOO! JAPAN
Singtel
OPTUS yes
Jio
Telin by Telkom Indonesia
Garena
NBIP nationale beheersorganisatie internet providers
POST LUXEMBOURG
中国电信 CHINA TELECOM
中国移动 China Mobile
China unicom中国联通

THANK YOU!

Julie Liu

www.genie-networks.com

julie@genie-networks.com