Infusing Heterogeneous Data to Troubleshoot & Improve Peering Performance and Security

Some Use Cases

Siarhei Matashuk
CCIE #27340
May 2025

“BGP Routing Tasks”

Peering Coordinators, Network Engineers for BGP Operation……

Peering Evaluation: Align peering decisions with your policy
Traffic Route Management: Monitor traffic paths to optimize routing, detect anomalies, and troubleshoot issues.
Route Health Monitoring: Use BGP updates and RPKI validation to assess route stability and integrity.
Route Anomaly Detection: Enable proactive alerts for abnormal BGP behavior.

“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

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

Align with your peering policy:

- No Peering: Focus on choosing the best transit providers for cost-efficiency.
  Pick providers who are best for your specific traffic
- Open: Peer with as many networks as possible to reduce transit costs.
  To decide if you should peer with a new network
  To convince others to peer with you
- Selective: Only peer with networks that offer significant mutual value.
  To decide if you should peer with a new network
- Restrictive: Assess potential customer traffic for transit business revenue opportunities.
  To understand 'transit prospective customers’ traffic behavior for compelling business case building

Open Peering Policy Use Case

Goal: **Find and evaluate new peering candidates**

- **Identify the Candidates:** Settlement-free peering reduces transit costs by enabling direct exchange of traffic, bypassing third-party transit providers.
- **Cost Analysis:** While the peering arrangement itself is settlement-free, it still incurs infrastructure costs such as router port utilization and colocation expenses at the peering facility.

Steps:

1)Identify ASNs with significant traffic volume not yet peered with: Rank the ASNs based on **the volume** of exchanged traffic

2)Identify ASNs with a **balanced traffic ratio**: where the inbound and outbound traffic volumes are approximately equal. A roughly equal exchange makes the peering relationship more sustainable and likely to be accepted

3)Distinguish between **direct traffic** (**sent to/from**) and **through traffic** (**through,** transit paths) (**sent to/from or through**): If the ASN is just a transit provider or intermediary (i.e., you’re sending traffic *through* it to reach others),

Open Peering Policy Use Case (contd)

Estimate cost savings vs. infrastructure cost:

Cost Avoidance from Transit: how much transit cost would be saved if the traffic is offloaded

- Example: If you're paying $X and $Y per Mbps for transit A and B, and this peering would offload m and n Mbps, you can estimate monthly savings [$(m·X + n·X)]

Peering Intra Cost: Calculate the port & colocation cost for establishing the new peering

- Do you already have infrastructure (routers, ports, power, rack space) at the same IX where the candidate ASN peers?

  If yes, the incremental cost is lower

  If not, factor in additional port costs (e.g., 10G, 100G) and rack/power/cross-connect fees.

A simplified cost-benefit formula like:.

Net Benefit = (Transit Cost Saved) – (Peering Infra Cost + Operational

Peering Evaluation

Not just for Open Peering. Evaluation helps across all policy types:

- No Peering: Cost analysis can also help Choose the most efficient transit providers for your specific traffic
- Selective: Assess whether peering requests make sense based on traffic volume, traffic ratio and traffic patterns
- Restrictive: Understand potential customer traffic to build a business case

Use data to support decisions:

Flow data: Understand traffic volume and directionality

BGP data: Analyze prefix visibility, AS paths, and routing dynamics (build your own virtual RIB)

Traffic Route Management

Traffic Engineering Use Cases

Knowing What traffic is leaving/entering your network @where is helpful for adjusting how the traffic going across the network.

- **Congestion mitigation**: Identify which links or peers are overloaded
- **Exit point balancing**: Shift traffic using BGP policies (e.g., LOCAL_PREF)
- **Policy verification**: Ensure that configuration changes are effective
- **Route integrity**: Detect route leaks or peers violating traffic agreements

Key Data Sources:

- Flow records: traffic volume, interface mapping, etc.
- BGP updates: prefix paths, next-hop, etc.
- SNMP: interface utilization, etc.

Congestion Mitigation by Exit Point Balancing

Use Case

A Use Case to mitigate congestion and improve resource utilization

Example: A peer (ISP V) is complained “being saturated.” So let’s see what is going on there:

- Knowing only how much total traffic there is through ISP-V is not enough

Need to know the traffic through ISP-V

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

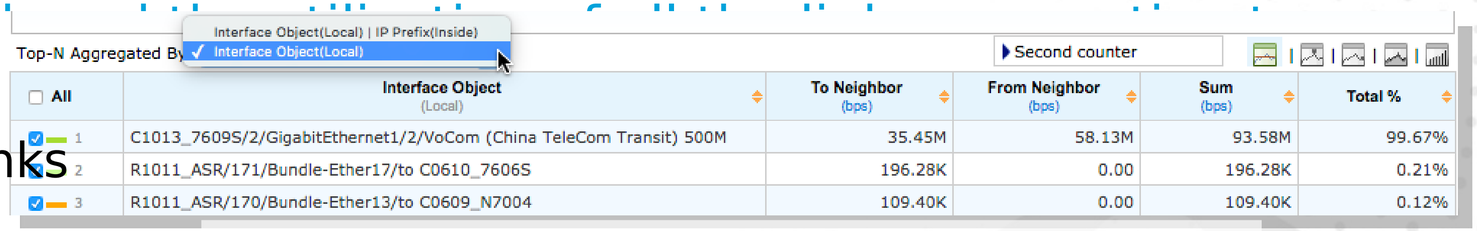

All	Interface Object (Local)	To Neighbor (bps)	From Neighbor (bps)	Sum (bps)	Total %
1	C1013_7609S/2/GigabitEthernet1/2/VoCom (China TeleCom Transit) 500M	35.45M	58.13M	93.58M	99.67%
2	R1011_ASR/171/Bundle-Ether17/to C0610_7606S	196.28K	0.00	196.28K	0.21%
3	R1011_ASR/170/Bundle-Ether13/to C0609_N7004	109.40K	0.00	109.40K	0.12%

Interface traffic distribution on a specific BGP **peer**

Apply route policy with actual measurements

- We may move certain share of traf (e.g. 30%, which is pretty much a specific ASN’s traffic) from one interface (C’s1) to another (ASR’s) BGP methods (e.g., LOCAL_PRF)

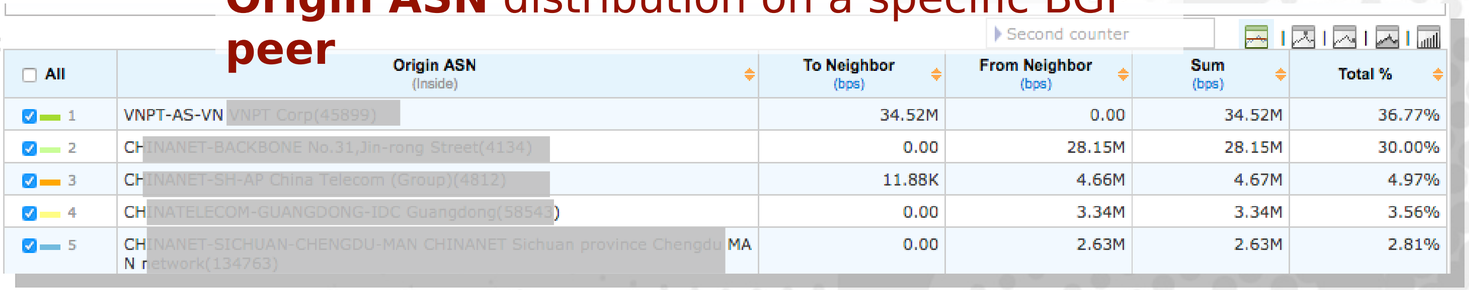

All	Origin ASN (Inside)	To Neighbor (bps)	From Neighbor (bps)	Sum (bps)	Total %
1	VNPT-AS-VN VNPT Corp(45899)	34.52M	0.00	34.52M	36.77%
2	CHINANET-BACKBONE No.31,Jin-rong Street(4134)	0.00	28.15M	28.15M	30.00%
3	CHINANET-SH-AP China Telecom (Group)(4812)	11.88K	4.66M	4.67M	4.97%
4	CHINATELECOM-GUANGDONG-IDC Guangdong(58543)	0.00	3.34M	3.34M	3.56%
5	CHINANET-SICHUAN-CHENGDU-MAN CHINANET Sichuan province Chengdu MAN network(134763)	0.00	2.63M	2.63M	2.81%

Origin ASN distribution on a specific BGP **peer**

BGP Policy Validation

Use Case

Troubleshooting — Expected route shifts don’t happen

Suboptimal BGP policy setting: for verifying whether we have made all configuration changes as we expected to facilitate the expected traffic route changes

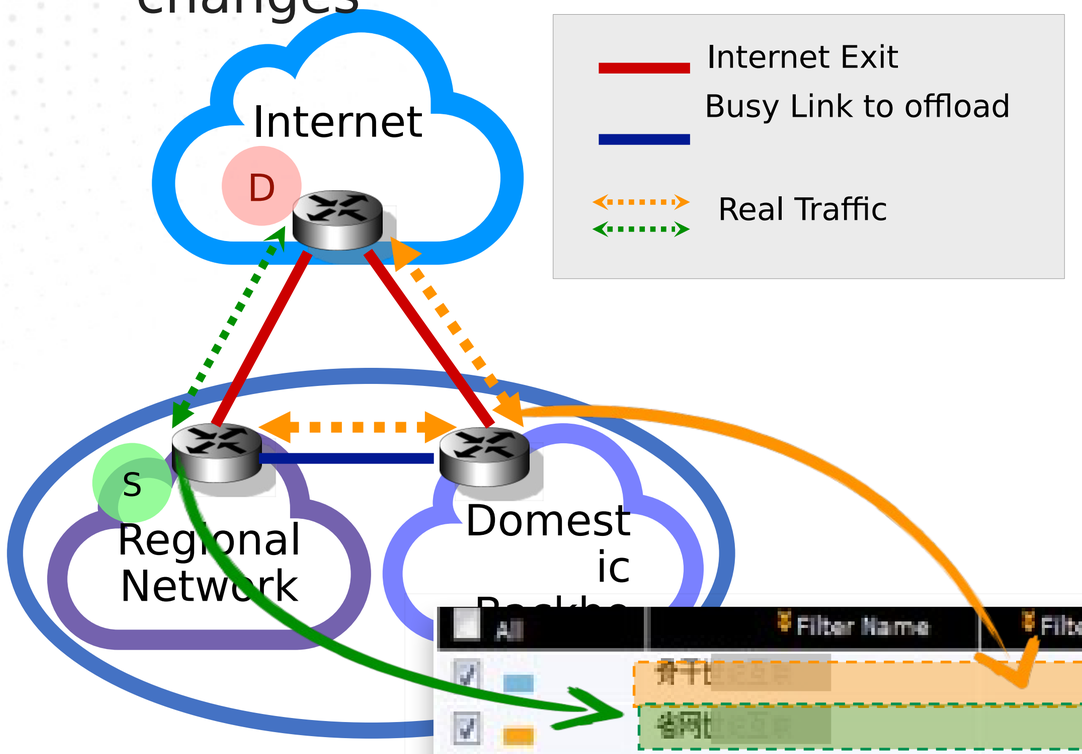

In order to offload the regional network’s Internet traffic from the busy link connecting the regional network with its domestic backbone, the SP has added an exit link from the regional network to the Internet directly.

Even though the corresponding configuration changes were considered done, the regional network is found still routing most of its Internet traffic through the link to the domestic backbone, wasting the domestic backbone resources.

It turned out that some BGP policies were not changed along correctly…

All	Filter Name	Filter Direction	Opposite Direction	Sum	Total %
	骨干t	7.855.30K	95.28M	102.95M	78.44%
	省网t	20.82M	7.659.53K	28.30M	21.56%
	Total	28.49M	102.76M	131.25M	100.00%

Route Integrity Verification

Use Case

Troubleshooting — Unreasonable peer behavior

Example: Peers dumping traffic at you for routes they didn’t receive from you (CDN/service providers abusing peering terms)

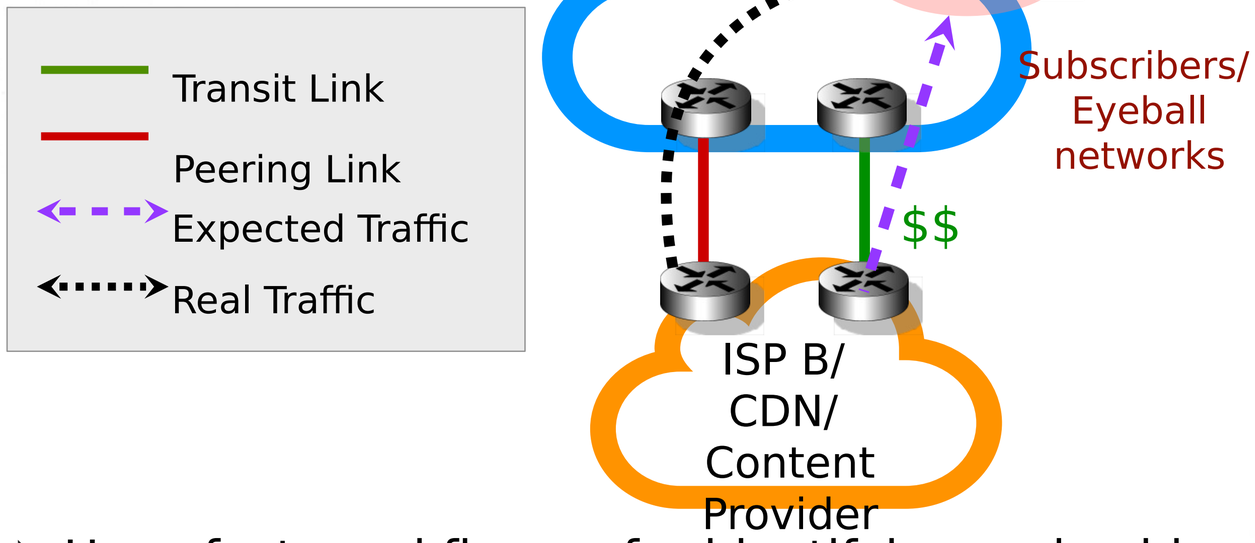

Instead of diverting the traffic through paid transit links, ISP B dumps traffic at ISP A for routes ISP A didn’t send it trough the free-peering agreement.

ISP B is stealing resources from ISP A, by violating the peering agreements with ISP A.

➡ Have facts and figures for identifying and evidencing these unreasonable routes:
Correlate service (CDN, OTT, etc.) identity information can be also helpful, in the cases that ISP B is a CDN network or service provider who’s traffic cannot be identified via origin ASN…

Route Health

Analyze BGP Route Instability

By examining BGP route status changes received from peers, BMP data, and correlation with RPKI information, we can analyze BGP route instability. The route status events reflect the internal BGP decision-making process. Analyzing them may help:

- Detect flapping routes or unstable prefixes,
- Identify problematic peers
- Trace the root cause behind frequent route changes (e.g., upstream instability, policy change, implementation bugs)

Possible Data Sources

- **BGP routing messages** (updates from peers): Such as `UPDATE` messages with `NLRI`, Withdrawn Routes, associated BGP attributes (AS_PATH, NEXT_HOP, Communities, etc.)
- **BMP messages** (from the routers' perspective): per-peer per-prefix events like AAdiff, Wdown, Status transitions (`Tbetter`, `Wdown`, etc.)

Route Health

Monitoring Examples

BGP: Monitor the Top-N prefixes with the highest event count (especially Wdown/TW), etc.

By Peer: Track which peers cause the most TW/Wdown events, etc.

Time-based: Are the events clustered in short bursts (bursty behavior)? Do they happen at regular times (e.g., scheduled policy

All	NLRI Prefix	PA Originator ID	Total Route Events
1	143.255.204.0/22	5.244.66	67.00
2	2a0a:afc7::/48	5.244.12	18.00
3	2407:5440::/48	5.244.12	12.00

All	PA Originator ID	Route Event	Total Route Events
1	103.15.244.12	AADiffBetter(5)	274.68K
2	103.15.244.61	AADiffWorse(6)	179.24K
3	103.15.244.61	AADiffBetter(5)	89.17K

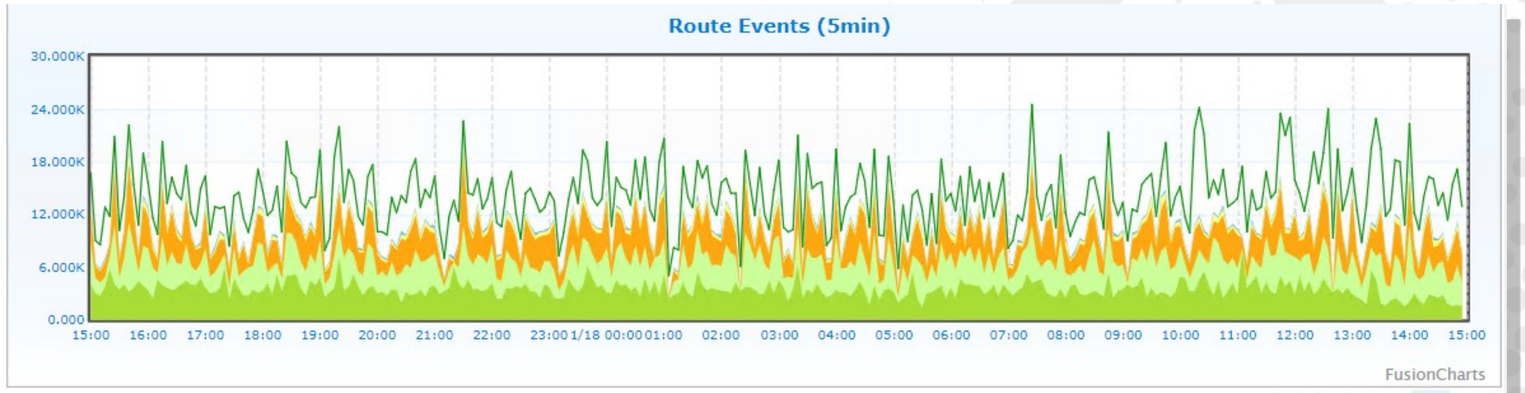

All	ROA Prefix	ROA Prefix Length	ROA ASN	BGP RPKI Status	Total Route Events	Total %

Route Anomaly Detection

Some BGP Route Anomaly Alerting Ideas

- **A BGP peer monitored by BMP goes up and down**: Indicate BGP peer flapping, which triggers frequent route-convergence and generates a massive churn of BGP messages.
  Example: detect > N times Peer flaps (Up/Down) in M minutes
- **BGP route state changes while RPKI invalid:** Invalid routes should be rejected (depending on policy.) If invalid routes cause replacement of valid paths, hijack or misconfig is likely.
  Example: detect > N route change events (Add, Withdraw, Replace, Fail Over) while RPKI invalid (INVALID_ASN or INVALID_LENGTH) in M minutes
- **Unstable routes:** For a prefix, if too many route change events in short time window, it may suggest route flapping (instability) due to upstream routing churn, BGP convergence churn due to remote outages, Intermittent RPKI status flips (valid <-> invalid), etc.
  Example: detect > N times route flaps (Replace, Fail Over) in M minutes
- **Too frequent route announcements from a BGP router:** It could be a route leak or policy Misconfiguration, a BGP Speaker Misbehaving, a Prefix Hijack or

Key Takeaways

Key Tasks Empowered by BGP Telemetry Analysis

Peering Evaluation: Identify valuable candidates; assess cost-benefit

Traffic Route Management: Balance, optimize, and validate routing paths

Route Health Monitoring: Detect flapping, problematic peers, misconfigurations

Route Anomaly Detection: Alert on abnormal route behaviors

Use heterogeneous data: Flow records, BGP routes, SNMP, service layer info (DNS info), RPKI validation and BMP telemetry to help the tasks done more effectively and efficiently!

DDoS attacks up 358% year-

KrebsOnSecurity last week was hit by a near record distributed denial-of-service (DDoS) attack that clocked in at **more than 6.3 Tbps**.

Google Security Engineer **Damian Menscher** told KrebsOnSecurity the May 12 attack was the largest Google has ever handled.

After comparing notes with Cloudflare, Menscher said the botnet that launched both attacks bears the fingerprints of **Aisuru**, a digital siege machine that first surfaced less than a year ago.

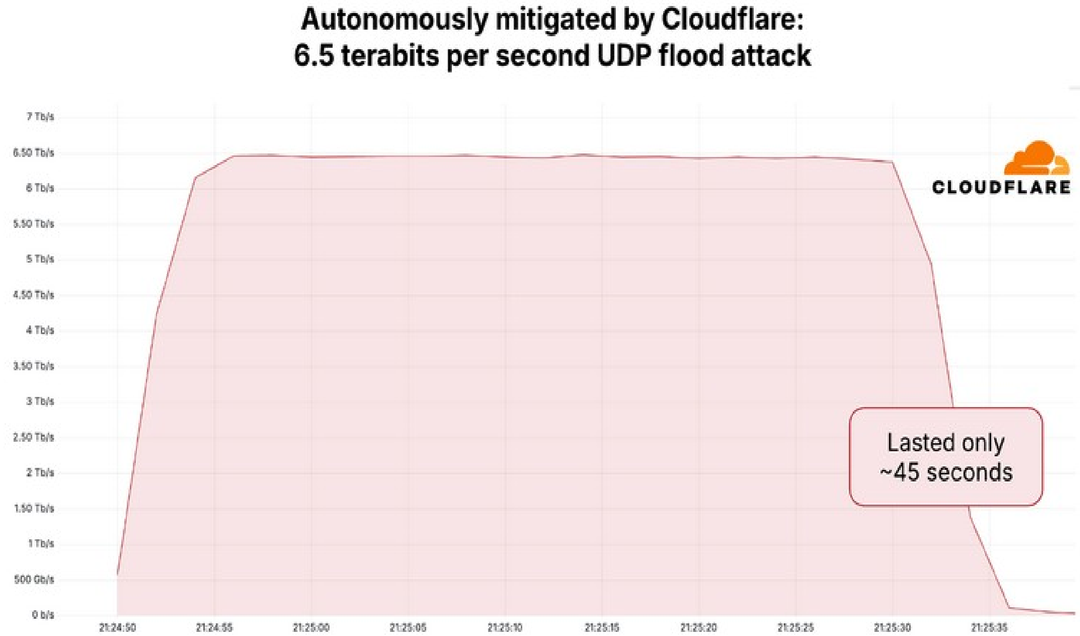

Sub-minute 6.5Tbps attacks using UDP originated from 147 countries and targeted multiple IP addresses and ports of a hosting provider.

What can you

- Use flow tools to ensure traffic visibility
  you can't afford to be blind.
- Know your upstream/IX AntiDDoS capabilities:
  - RTBH
  - Advanced RTBH
  - FlowSpec
  - Upstream scrubbing
  - Clean Pipes
- Maintain a list of emergency contacts for quick response.
- Communicate, share, and collaborate with your peers.

THANK YOU!

Siarhei Matashuk

www.genie-networks.com

s.matashuk
@genie-networks.com