Egress Monitoring at Scale

Alexander Azimov, Yango

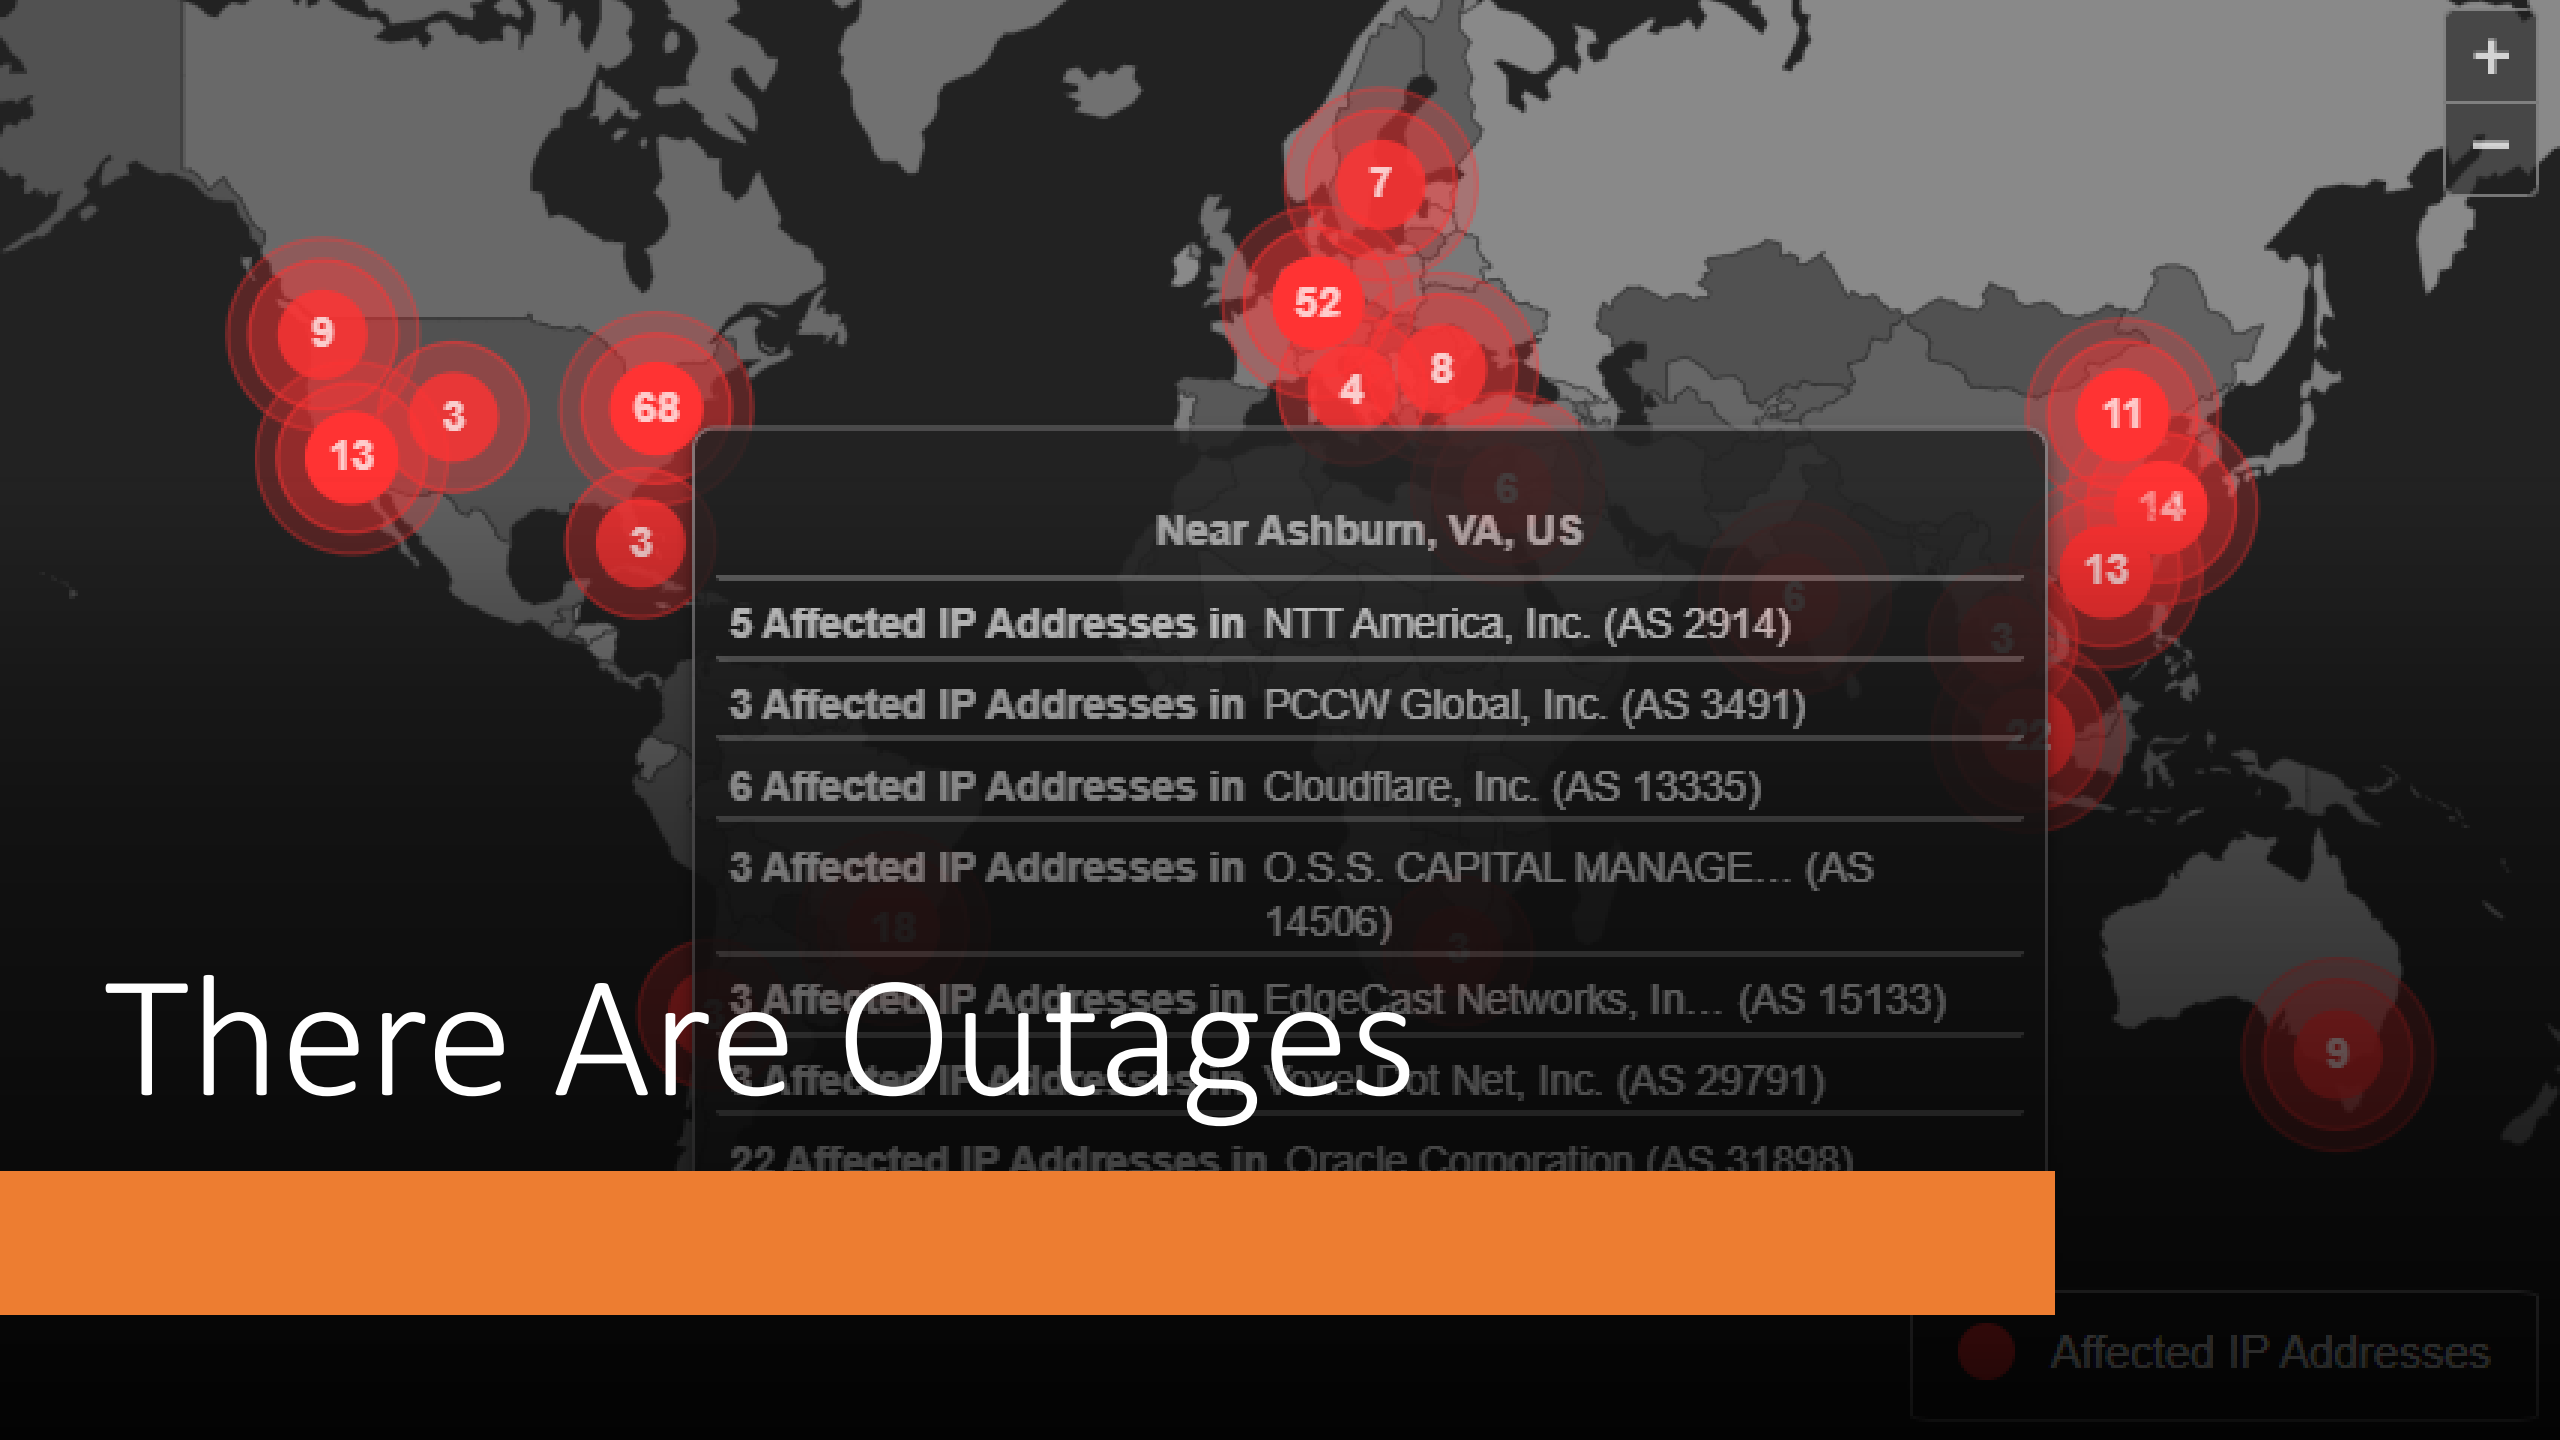

There Are Outages

The MTTI is how long it takes for the networking organization to prove it is not the network causing the degradation. Once that task is accomplished, it is common to assume some other component of IT such as the servers must be at fault.

Network Error Logging

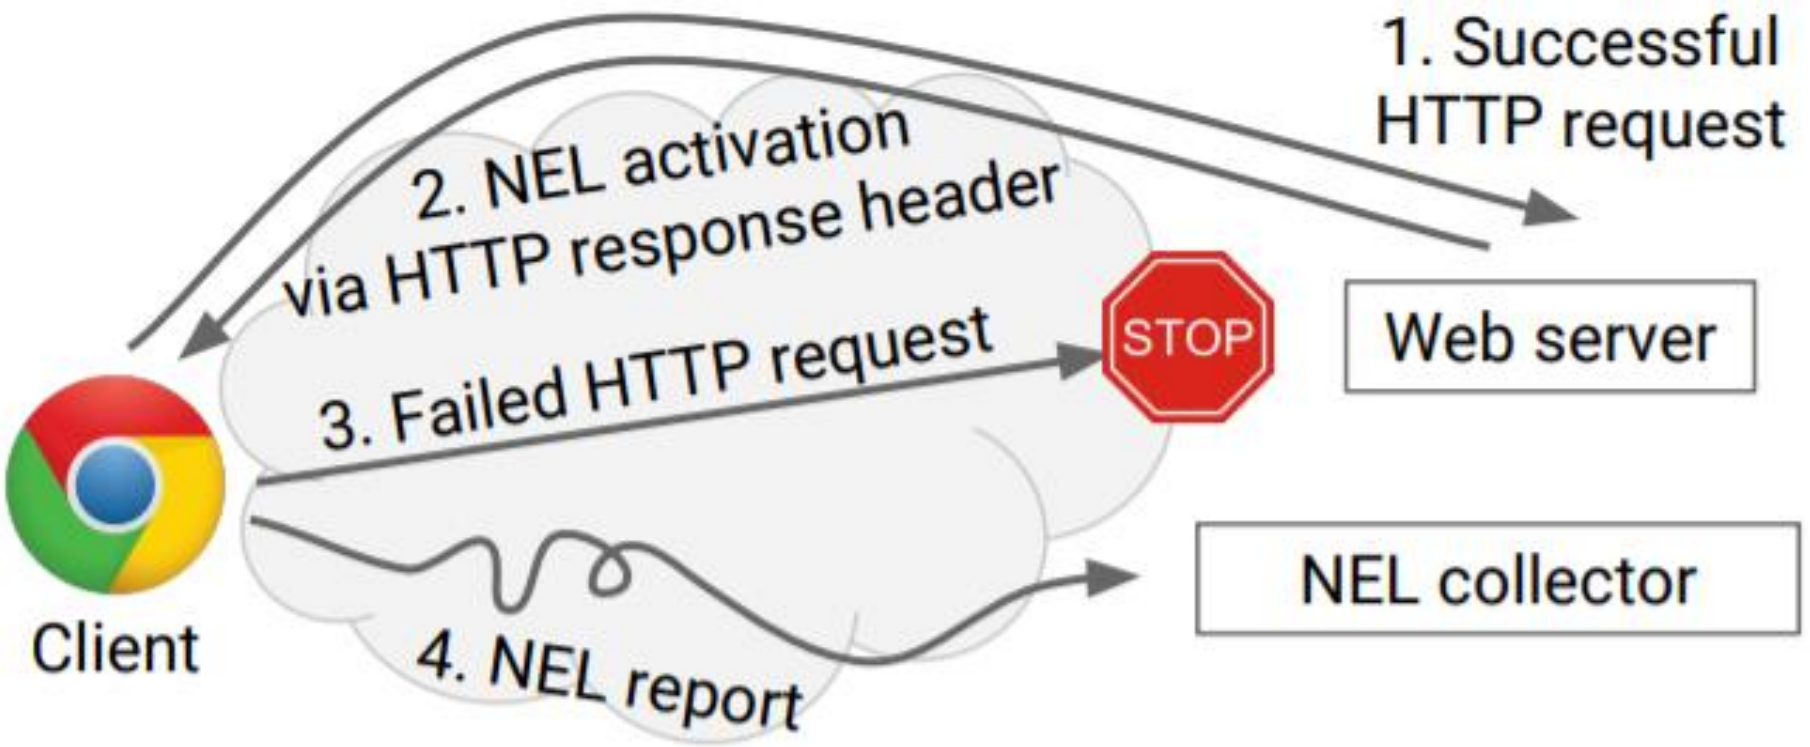

Source: https://www.usenix.org/system/files/nsdi20-paper-burnett.pdf

Network Error Logging: Report


```
[
  {
    "age": 666,
    "body": {
      "elapsed_time": 37,
      "method": "GET",
      "phase": "connection",
      "protocol": "http/1.1",
      "referrer": "https://www.example.com/",
      "sampling_fraction": 1,
      "server_ip": "1.2.3.4",
      "status_code": 0,
      "type": "tcp.reset"
    },
    "type": "network-error",
    "url": "https://www.example.com/image.png",
    "user_agent": "Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like
  }
]
```

Network Error Logging: Support

	Chrome	Edge	Firefox	Opera	Safari	Chrome Android	Firefox for Android	Opera Android	Safari on iOS	Samsung Internet	WebView Android
NEL	✓ 71	✓ 79	✗ No	✓ 58	✗ No	✓ 71	✗ No	✓ 50	✗ No	✓ 10.2	✓ 71

Source: https://developer.mozilla.org/en-US/docs/Web/HTTP/Network_Error_Logging#browser_compatibility

NEL: Good Example

Outage of peering patner

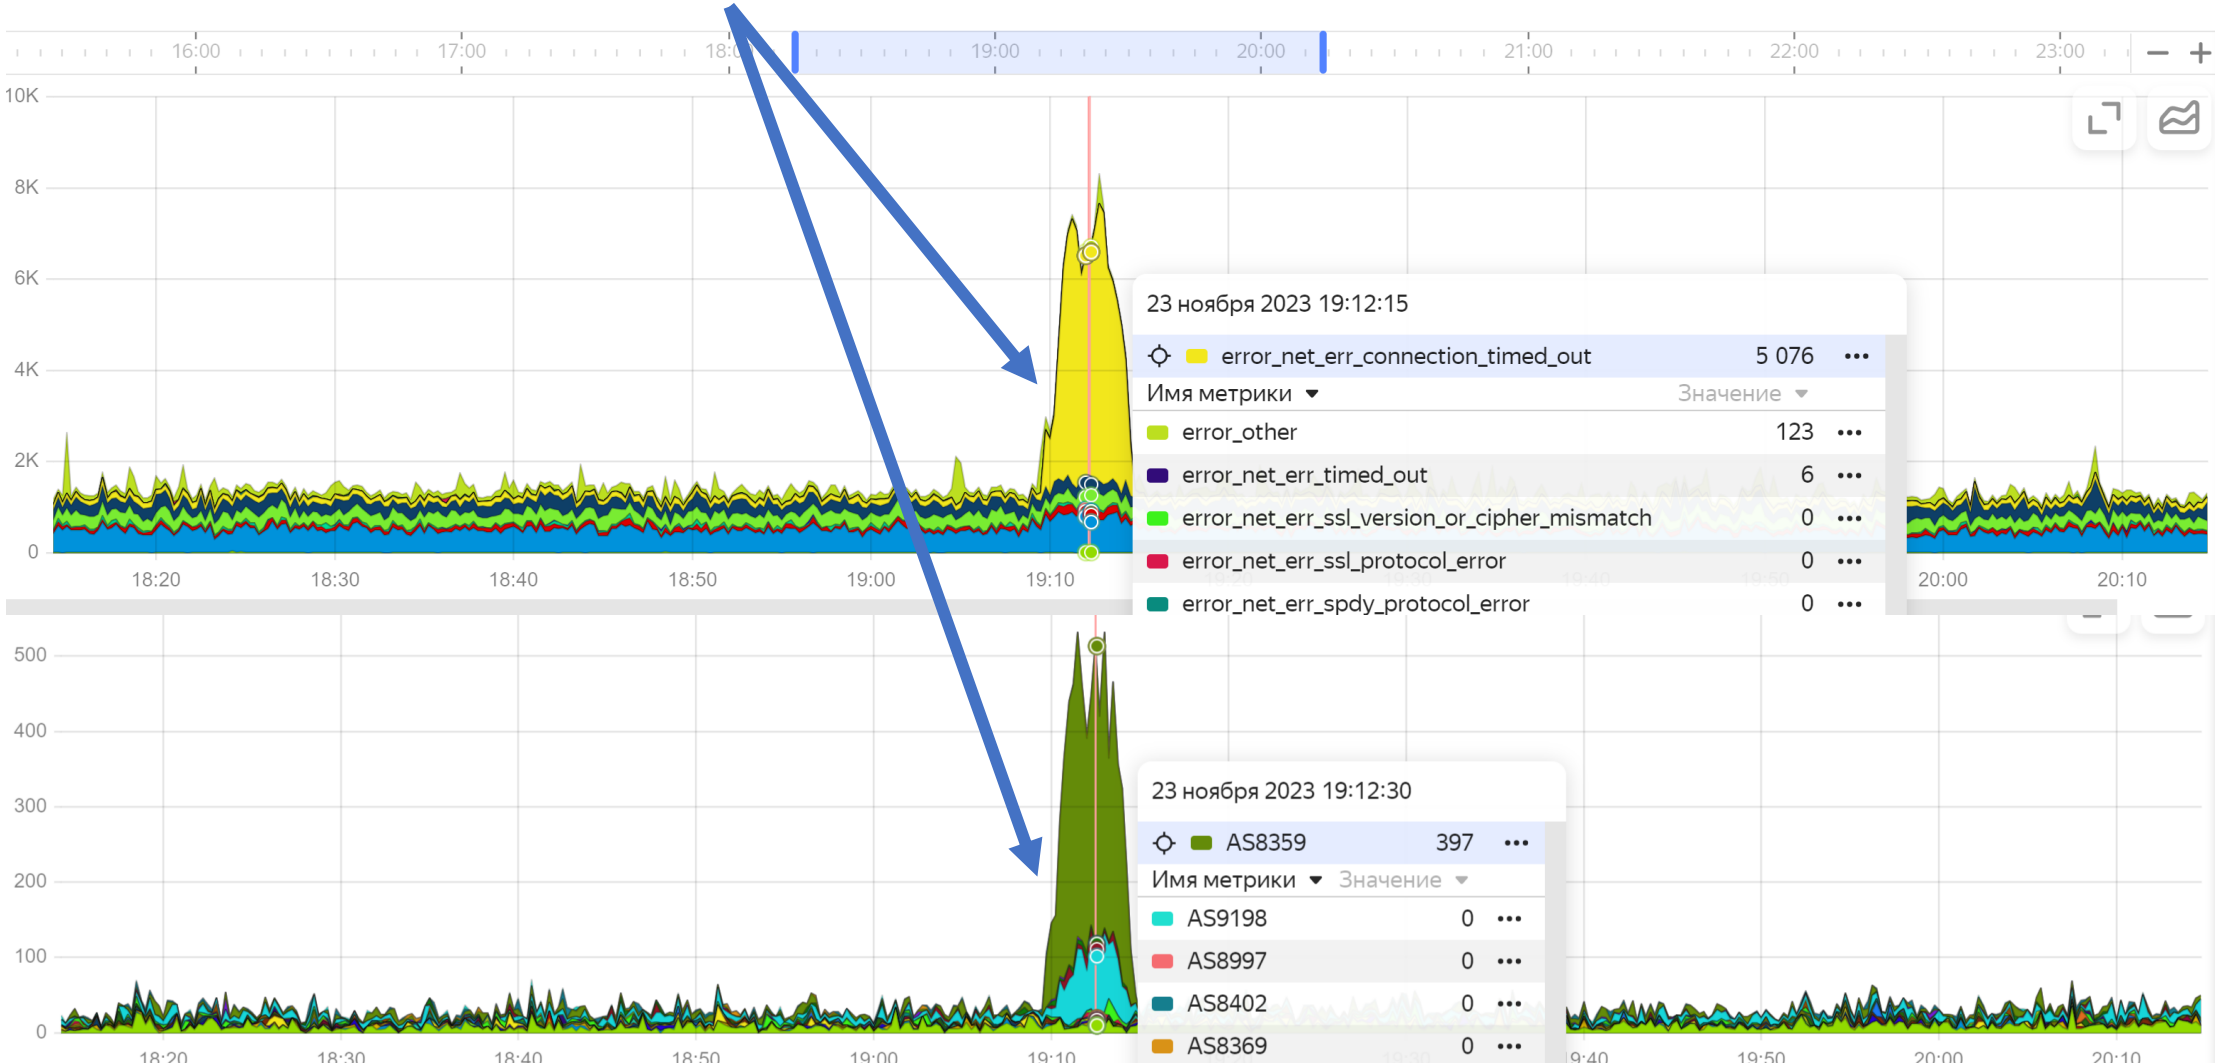

NEL: Not So Good Example

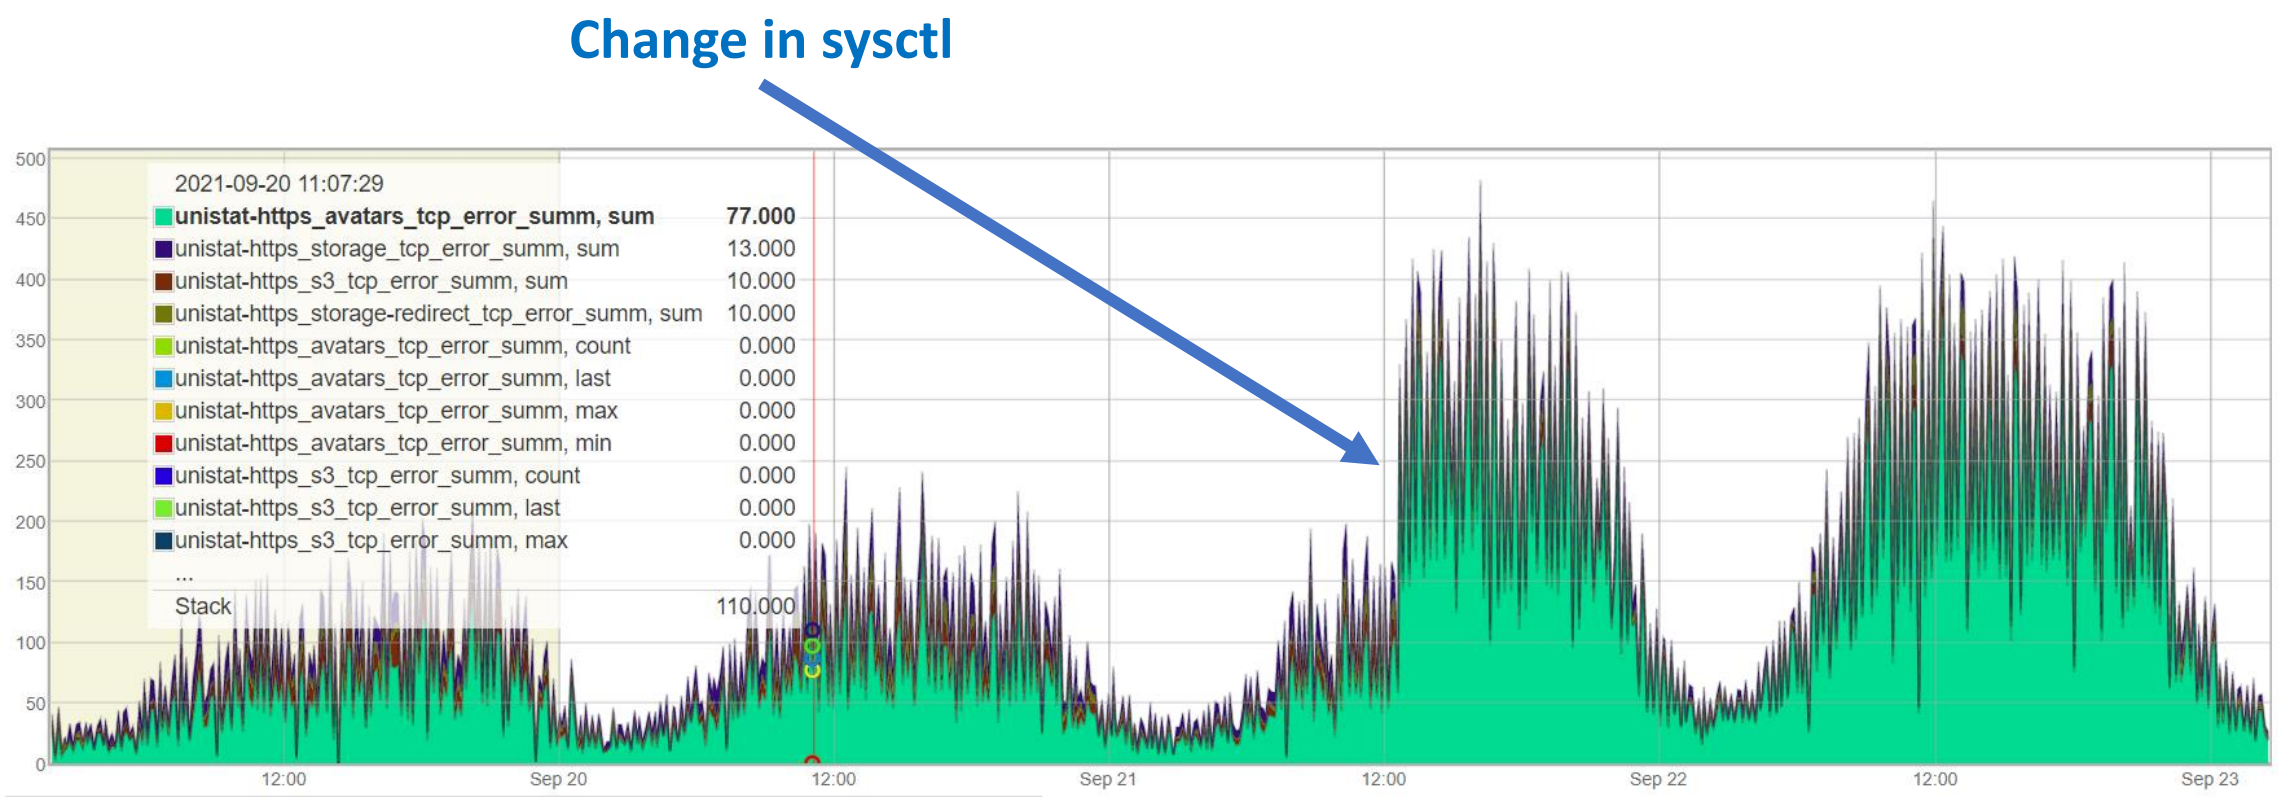

Network Error Logging: Summary

What NEL can do:

- Distinguish between different anomaly types;

What NEL can’t do:

- Distinguish between TCP and network anomalies;
- Distinguish between egress and ingress anomalies;
- Detect failing link;

Remote Probes

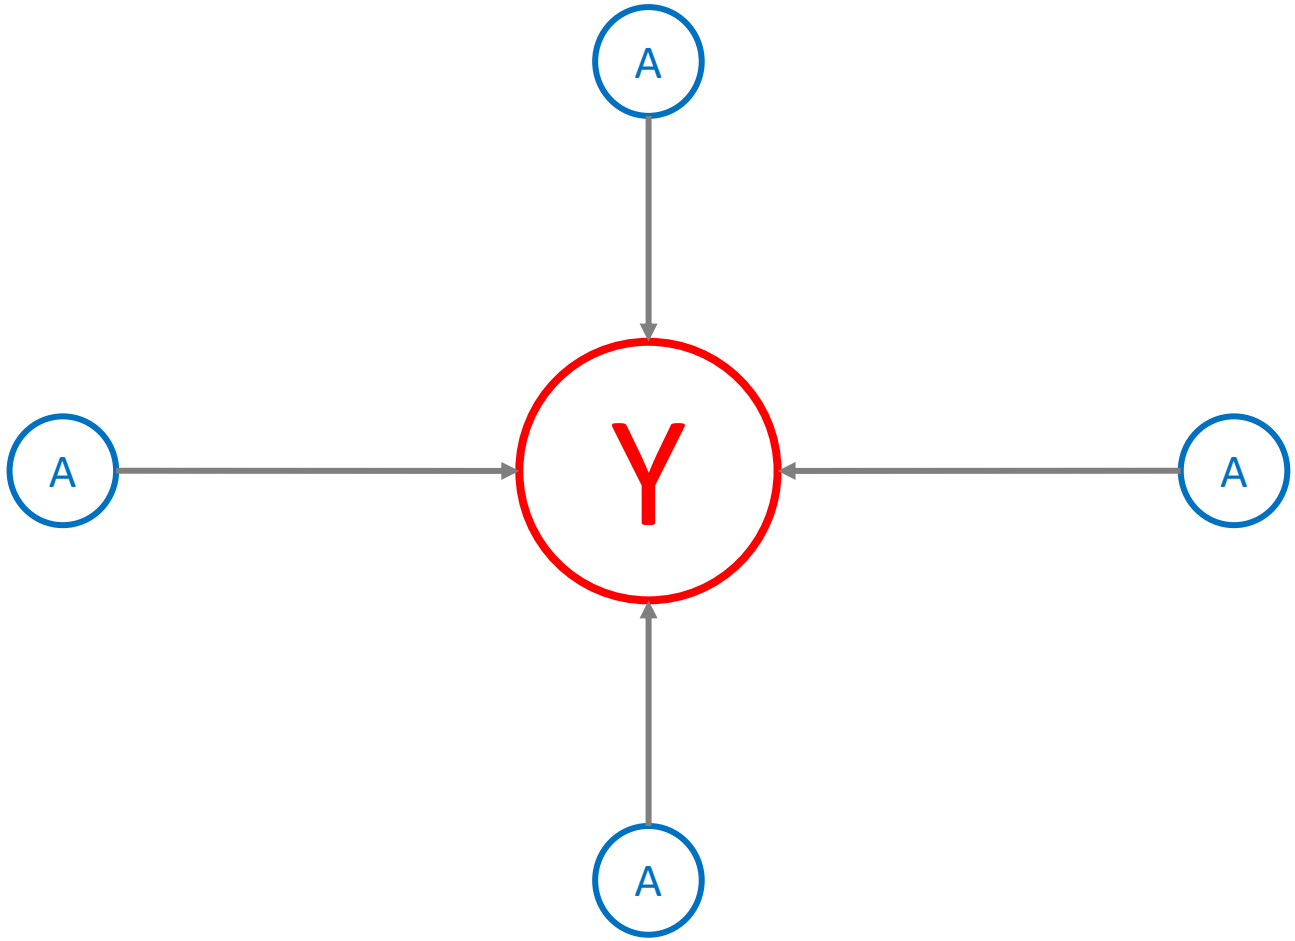

Need More Probes!

Remote Probes: Summary

What remote probes can do:

- Detect network anomalies according to limited coverage;

What remote probes can’t do:

- Detect network anomalies outside coverage;
- Distinguish between egress and ingress anomalies;
- Detect failing link;


Temet Nosce

Upside-down Probing

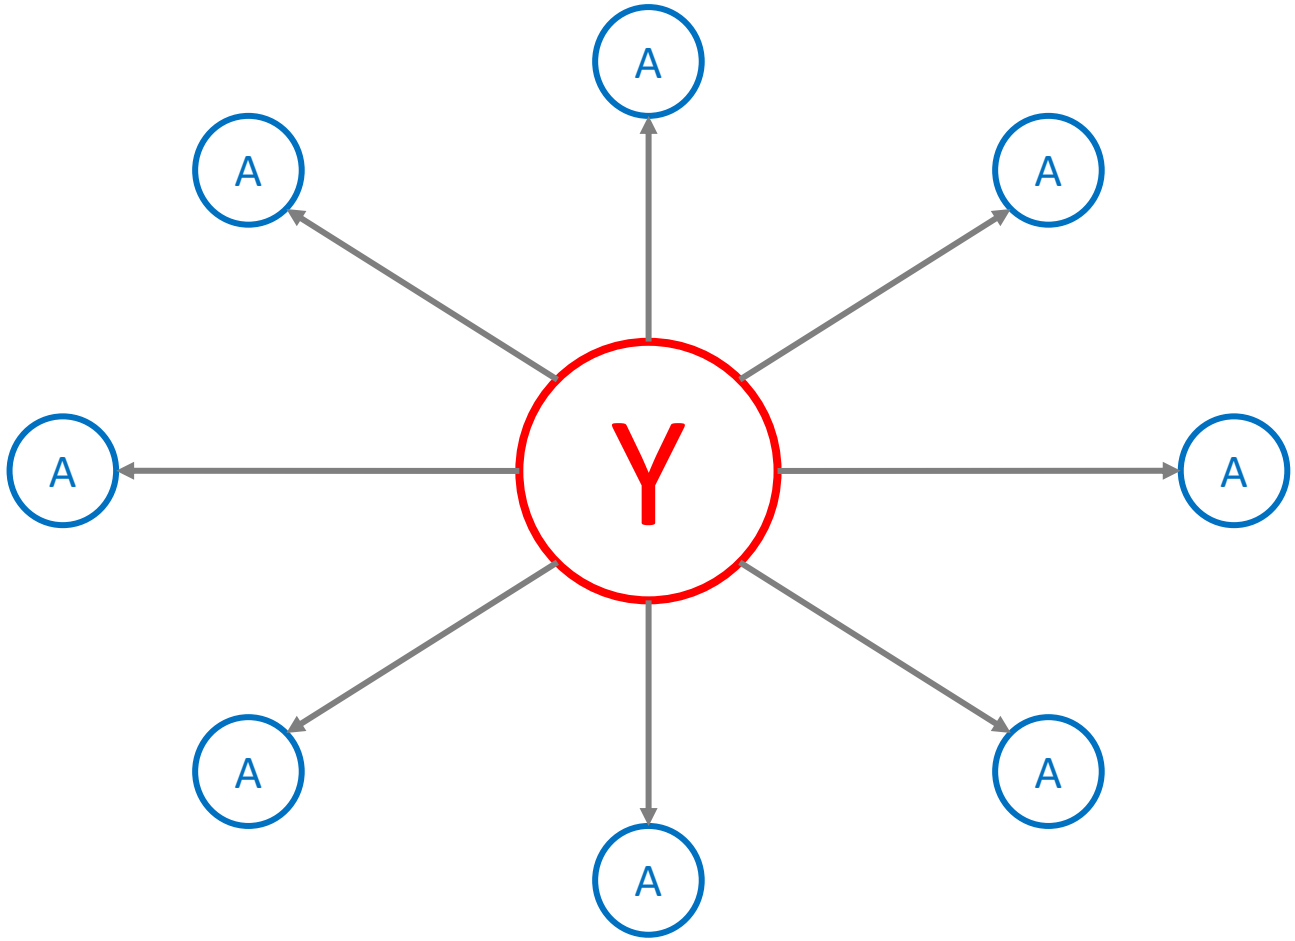

How It Works

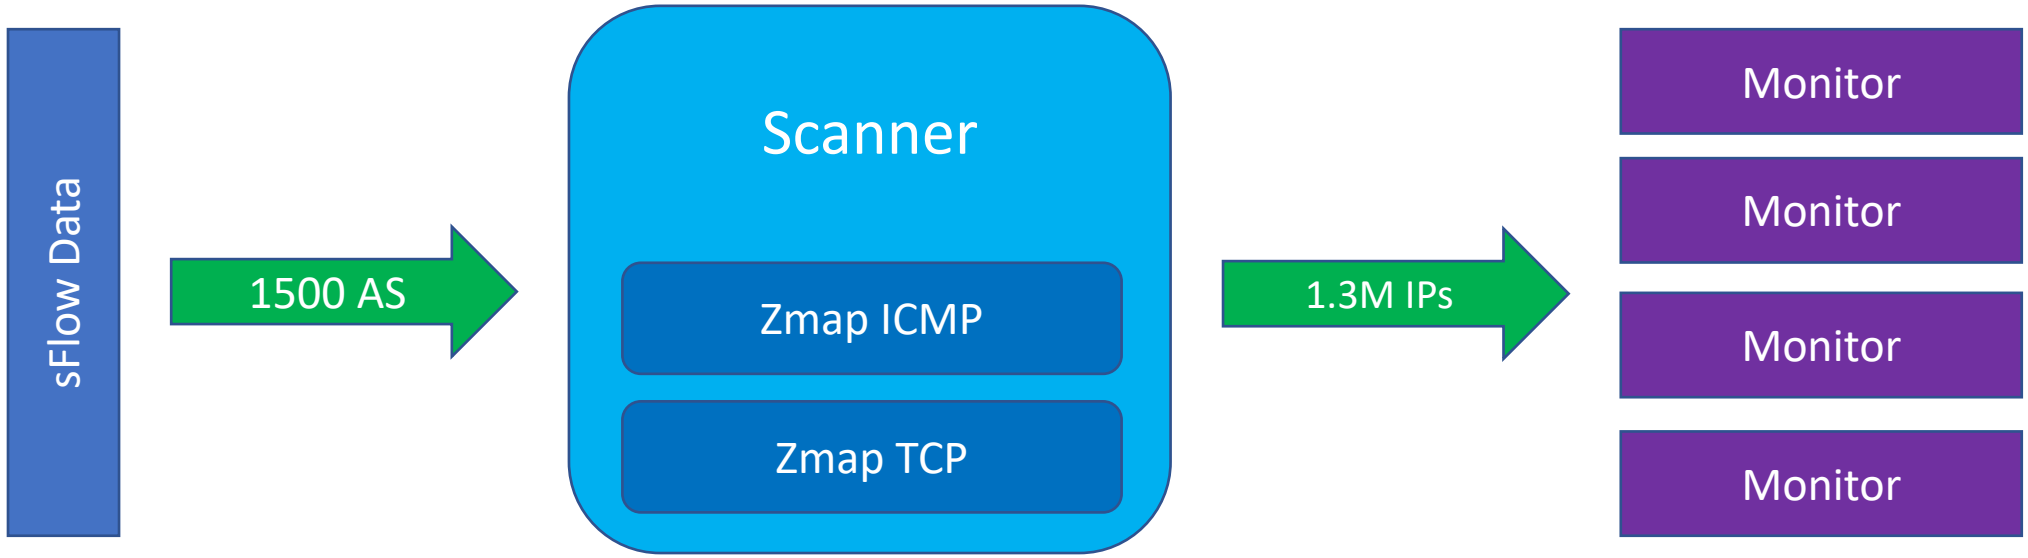
sFlow Data
1500 AS
Scanner
Zmap ICMP
Zmap TCP
1.3M IPs
Monitor
Monitor
Monitor
Monitor

How It Works

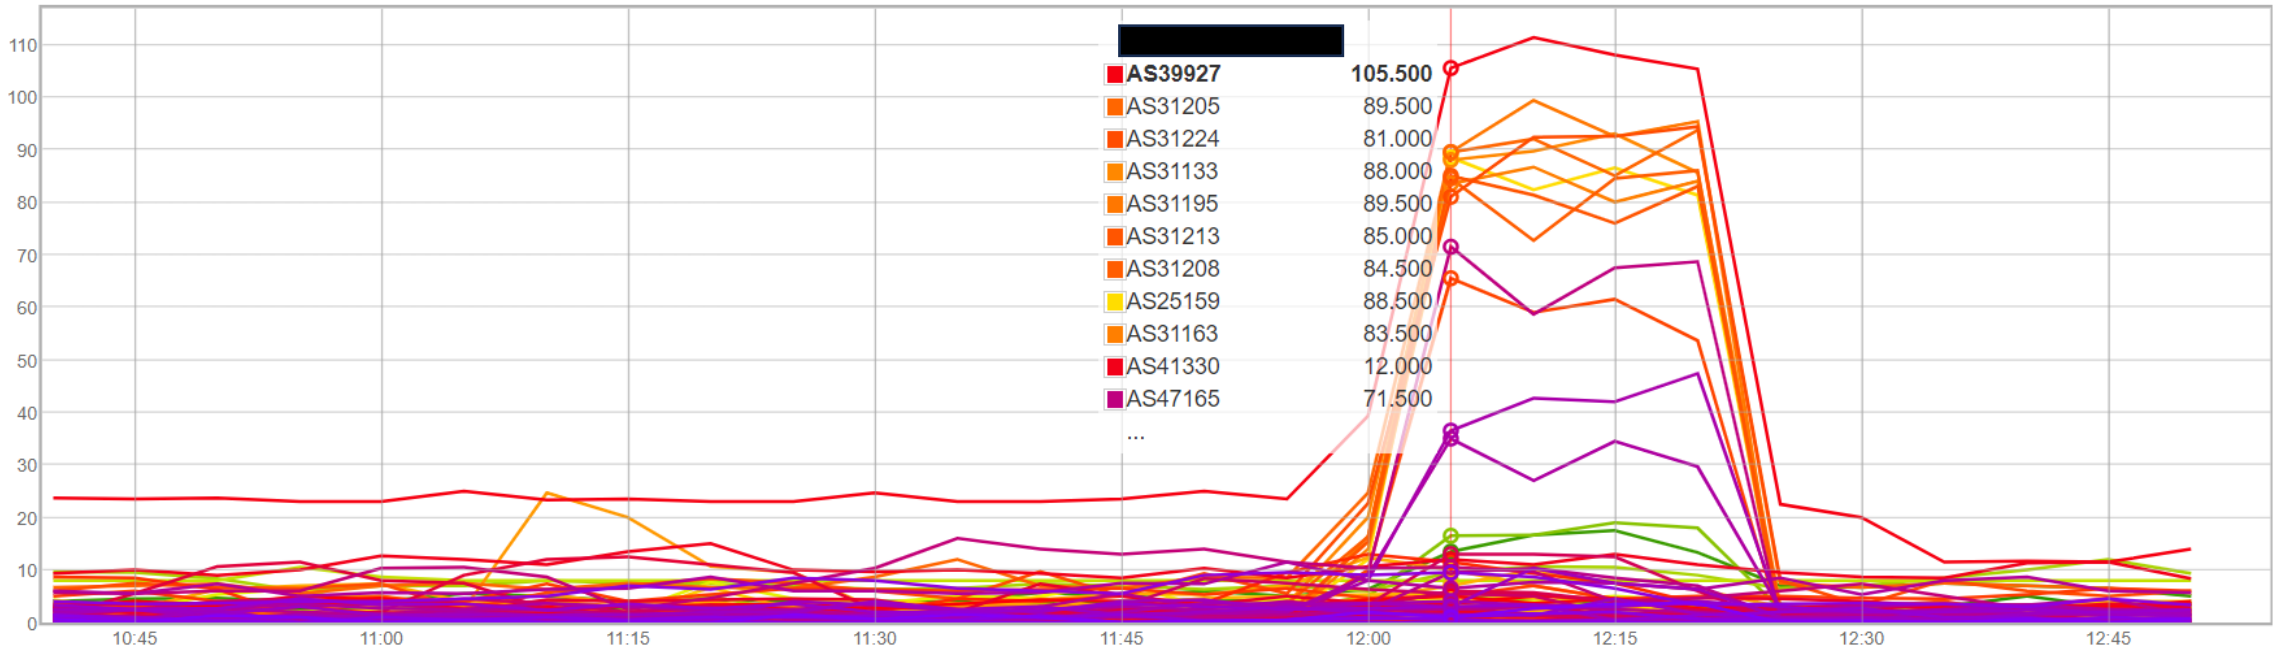
AS39927 105.500
AS31205 89.500
AS31224 81.000
AS31133 88.000
AS31195 89.500
AS31213 85.000
AS31208 84.500
AS25159 88.500
AS31163 83.500
AS41330 12.000
AS47165 71.500
...
110
100
90
80
70
60
50
40
30
20
10
0
10:45
11:00
11:15
11:30
11:45
12:00
12:15
12:30
12:45

In the philosophy of language, a **proper name** – examples include a name of a specific person or place – is a name which ordinarily is taken to uniquely identify its referent in the world.


Megaping

Megaping: Summary

What Megaping can do:

- Detect network anomalies, easily scales;

What Megaping can’t do:

- Detect failing link;
- Hardly applicable to IPv6;

Egress Traffic Pipeline

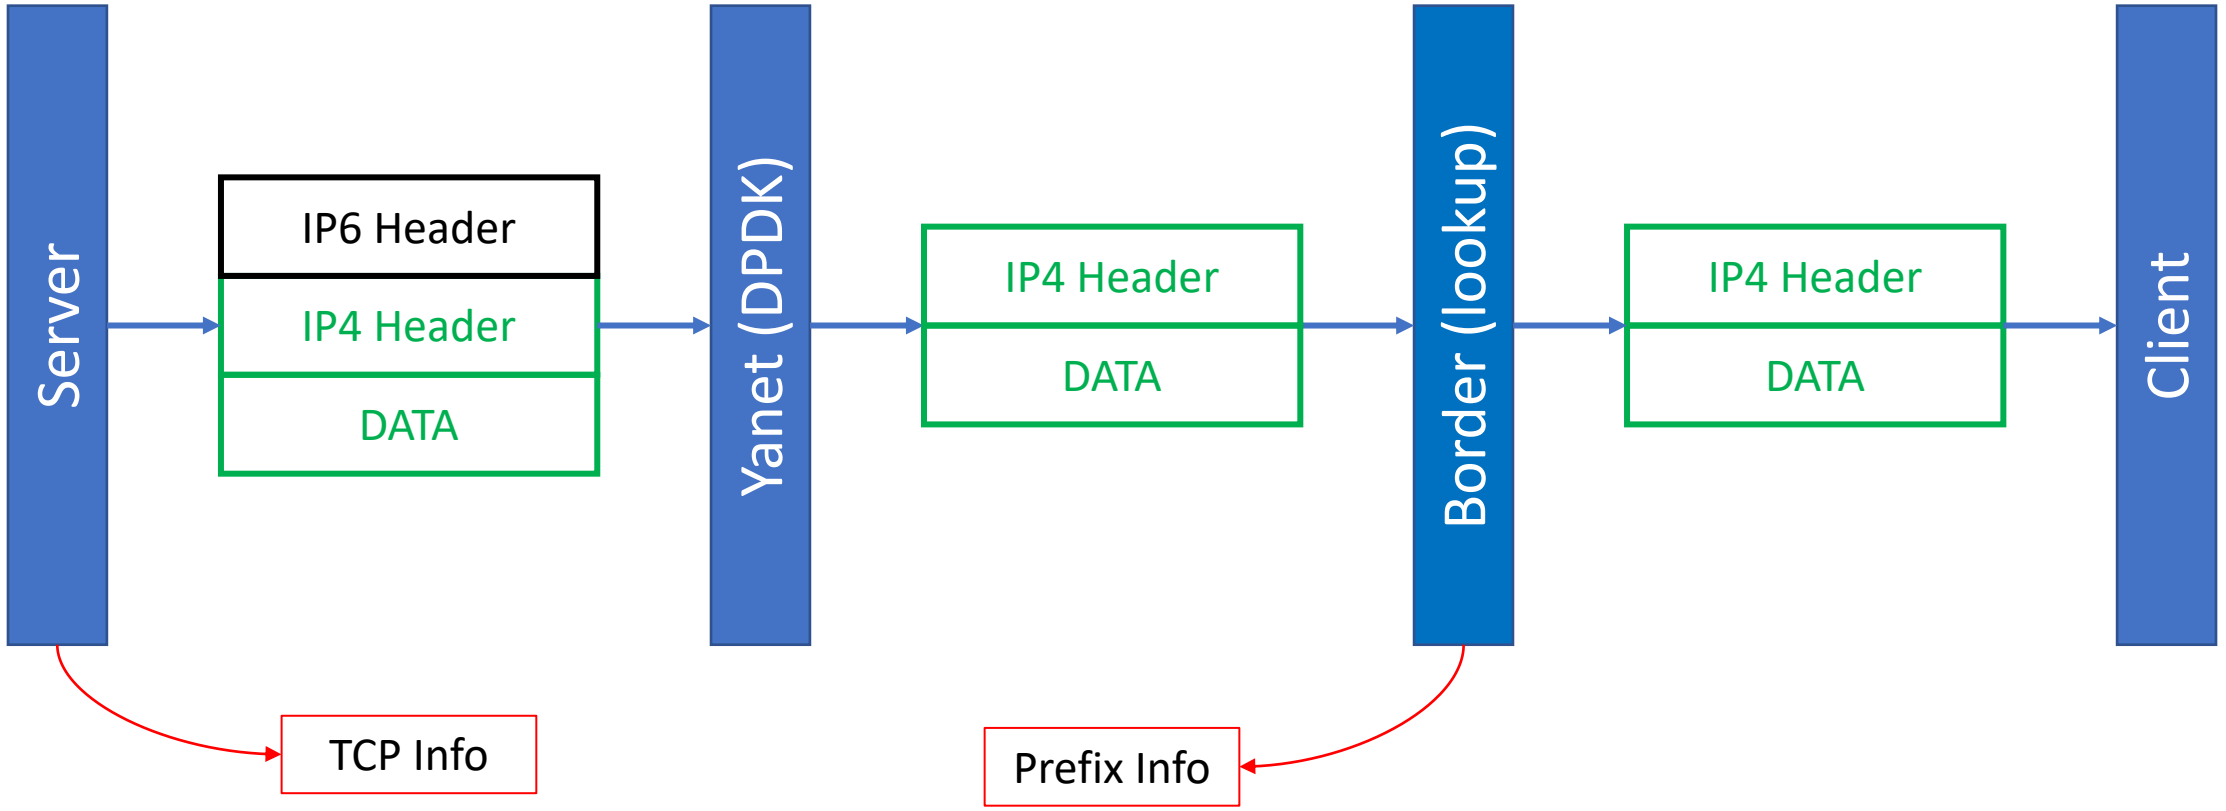

Shared Use of Experimental TCP Options (RFC6994)

Abstract

This document describes how the experimental TCP option codepoints can concurrently support multiple TCP extensions, even within the same connection, using a new IANA TCP experiment identifier. This approach is robust to experiments that are not registered and to those that do not use this sharing mechanism. It is recommended for all new TCP options that use these codepoints.

Kernel Hackers! ☺

```
[TCP Segment Len: 271]
Sequence number: 1    (relative sequence number)
Sequence number (raw): 3331852610
[Next sequence number: 272    (relative sequence number)]
Acknowledgment number: 1    (relative ack number)
Acknowledgment number (raw): 3427112418
1100 .... = Header Length: 48 bytes (12)
▸ Flags: 0x018 (PSH, ACK)
Window size value: 9
[Calculated window size: 9]
[Window size scaling factor: -1 (unknown)]
Checksum: 0xa2d0 [unverified]
[Checksum Status: Unverified]
Urgent pointer: 0
▾ Options: (28 bytes), No-Operation (NOP), No-Operation (NOP), Timestamps, No-Operation (NOP), No-Operation (NOP), Experimental
  ▸ TCP Option - No-Operation (NOP)
  ▸ TCP Option - No-Operation (NOP)
  ▸ TCP Option - Timestamps: TSval 3301259616, TSecr 35314146
  ▸ TCP Option - No-Operation (NOP)
  ▸ TCP Option - No-Operation (NOP)
  ▾ TCP Option - Experimental
      Kind: RFC3692-style Experiment 1 (253)
      Length: 14
      Magic Number: 0x7961
```

Encodes number of acked and retransmitted packets plus RTT

Egress Traffic Pipeline

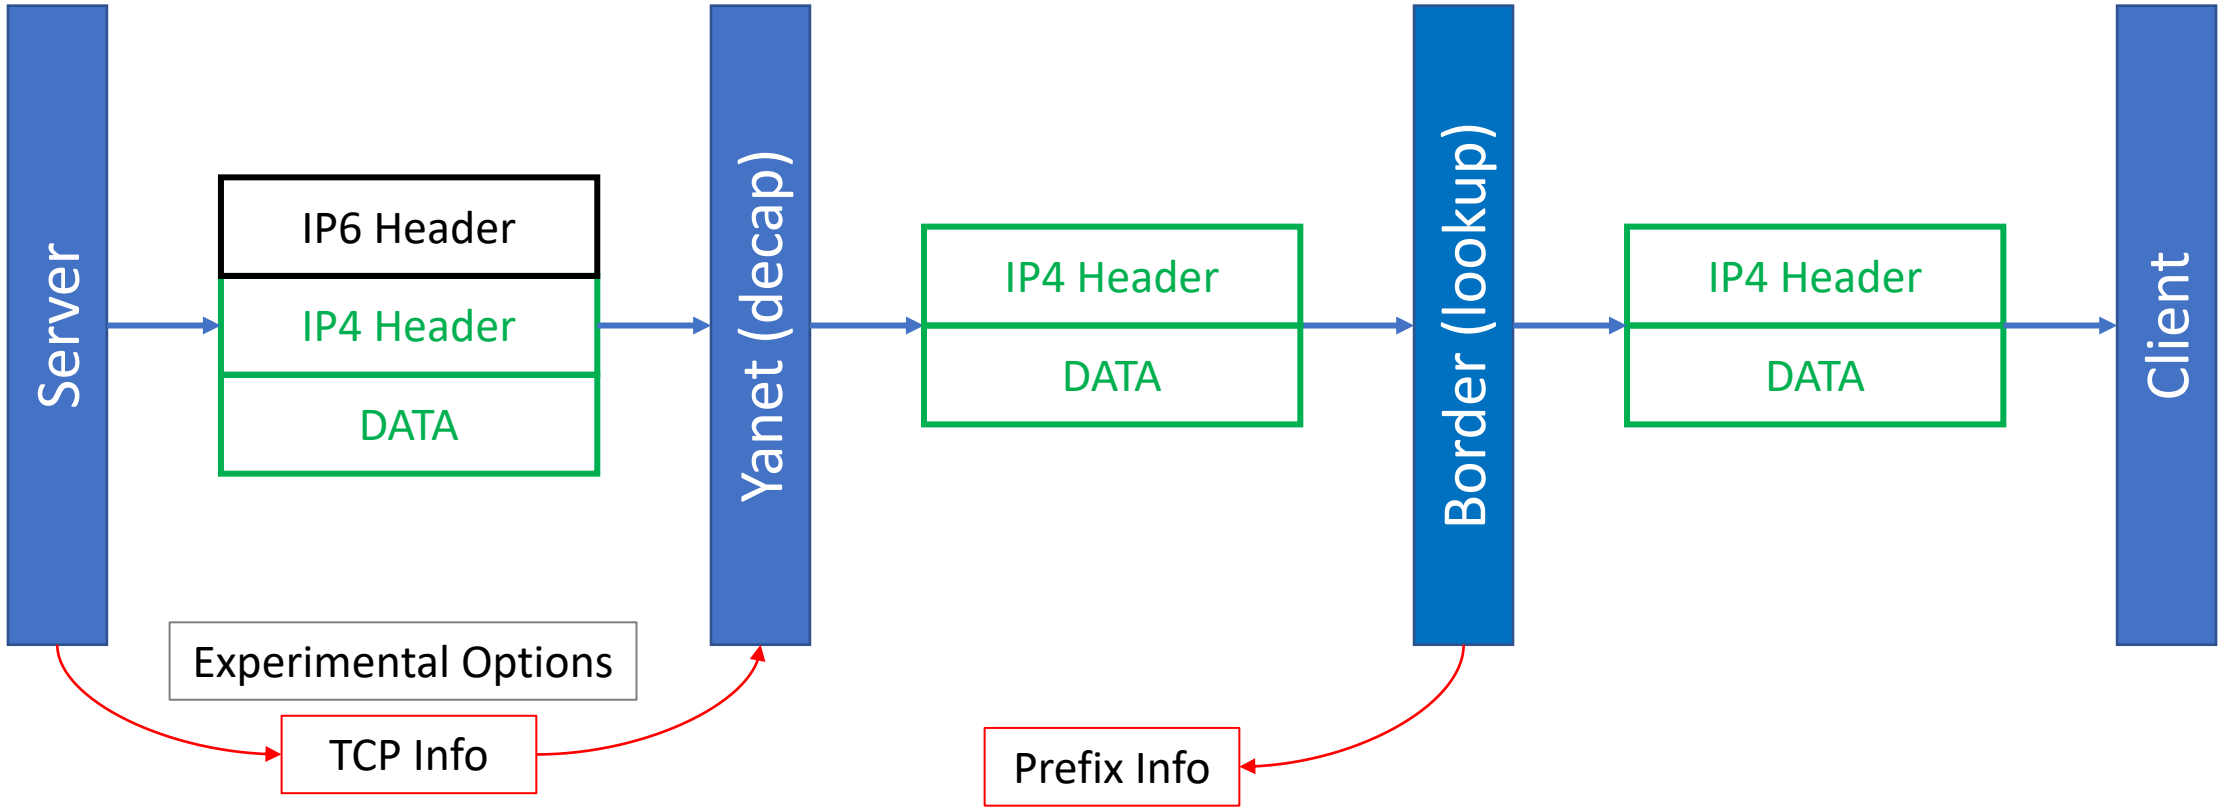

Each Peering Becomes a VRF

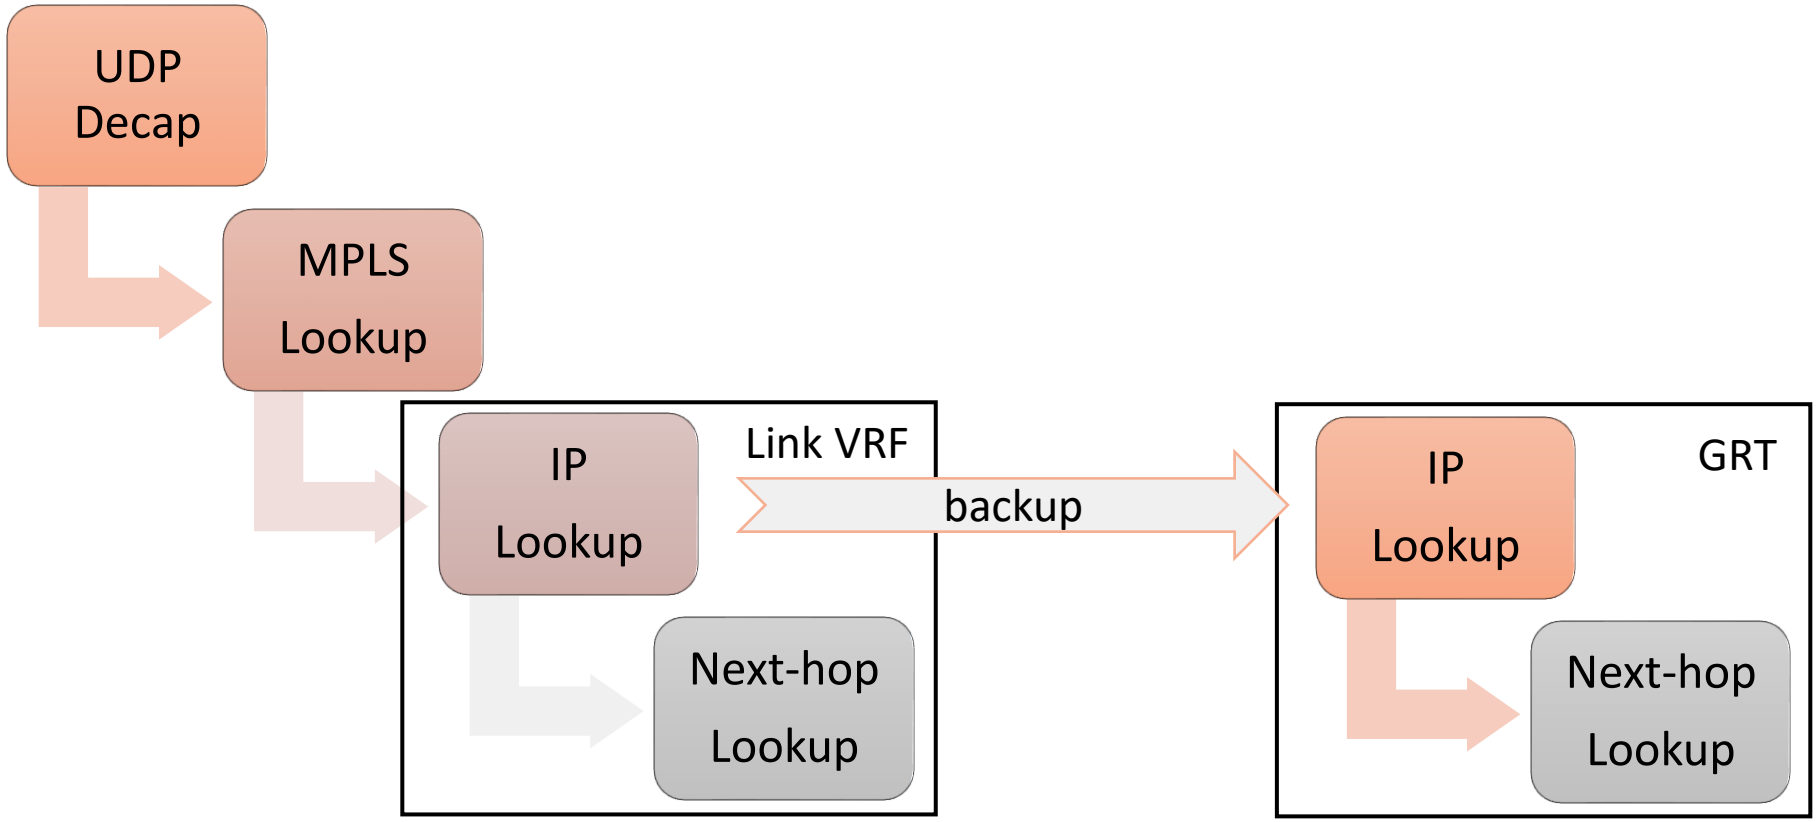

Monitoring Pipeline: IPv4

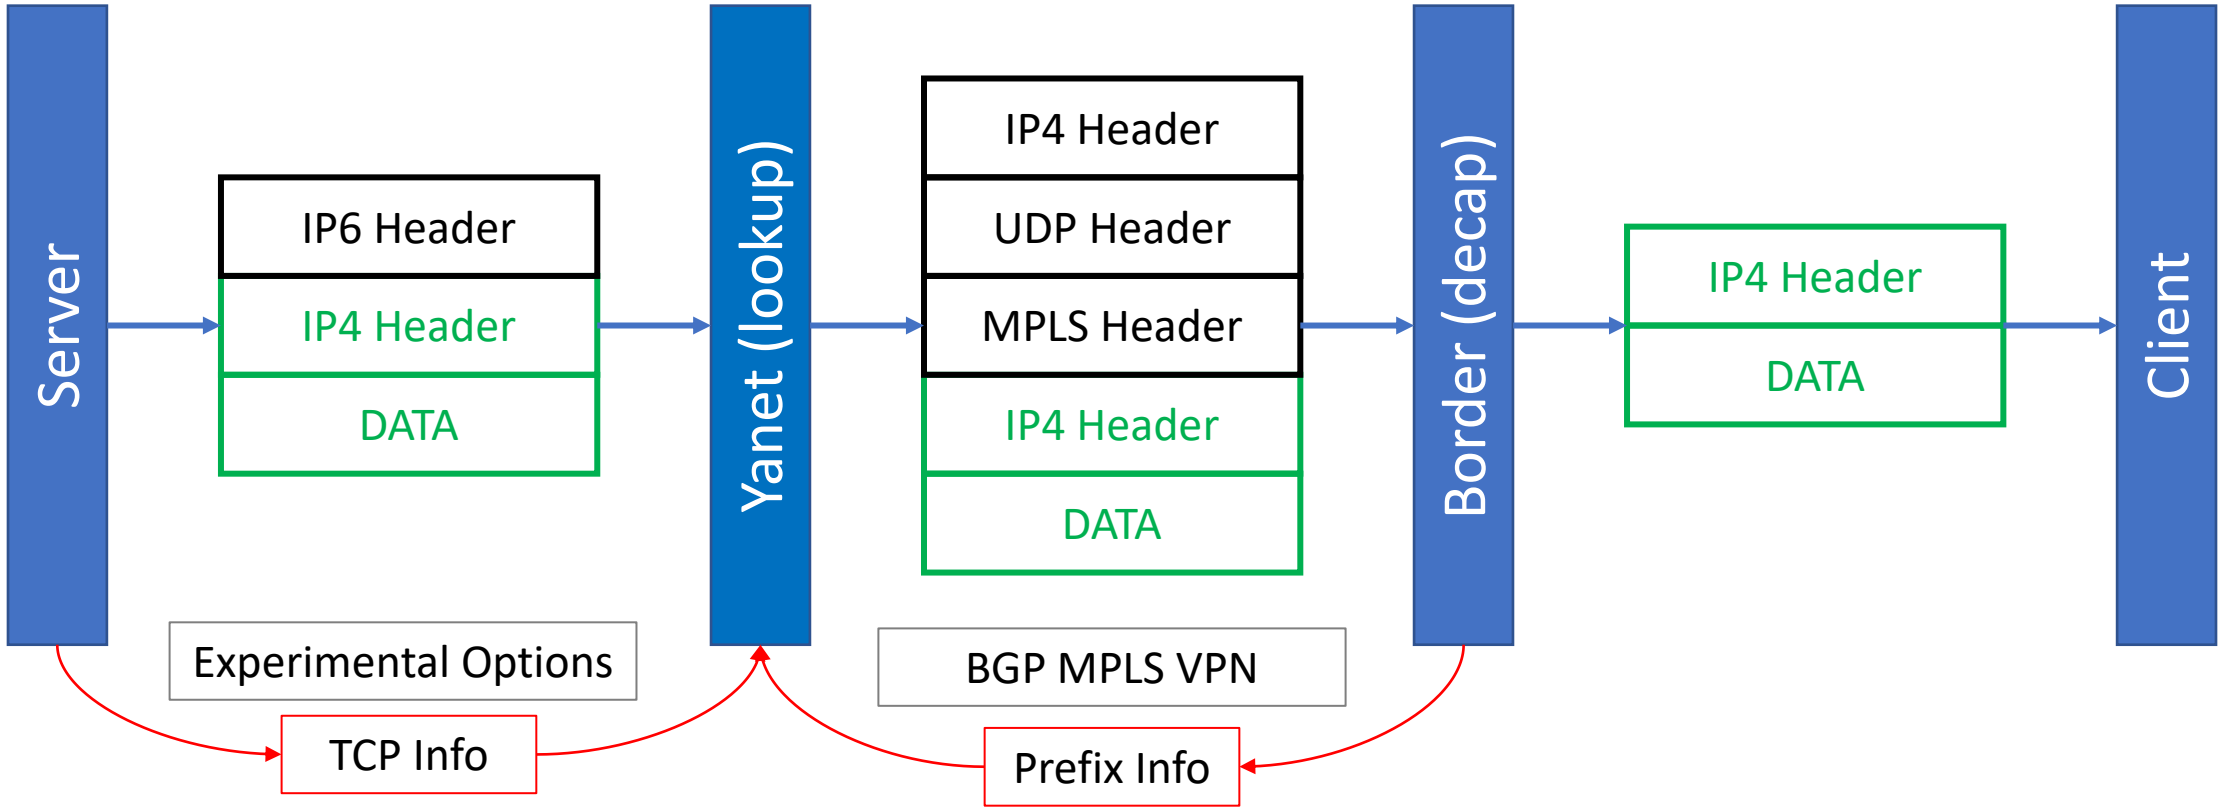

Monitoring Pipeline: IPv6

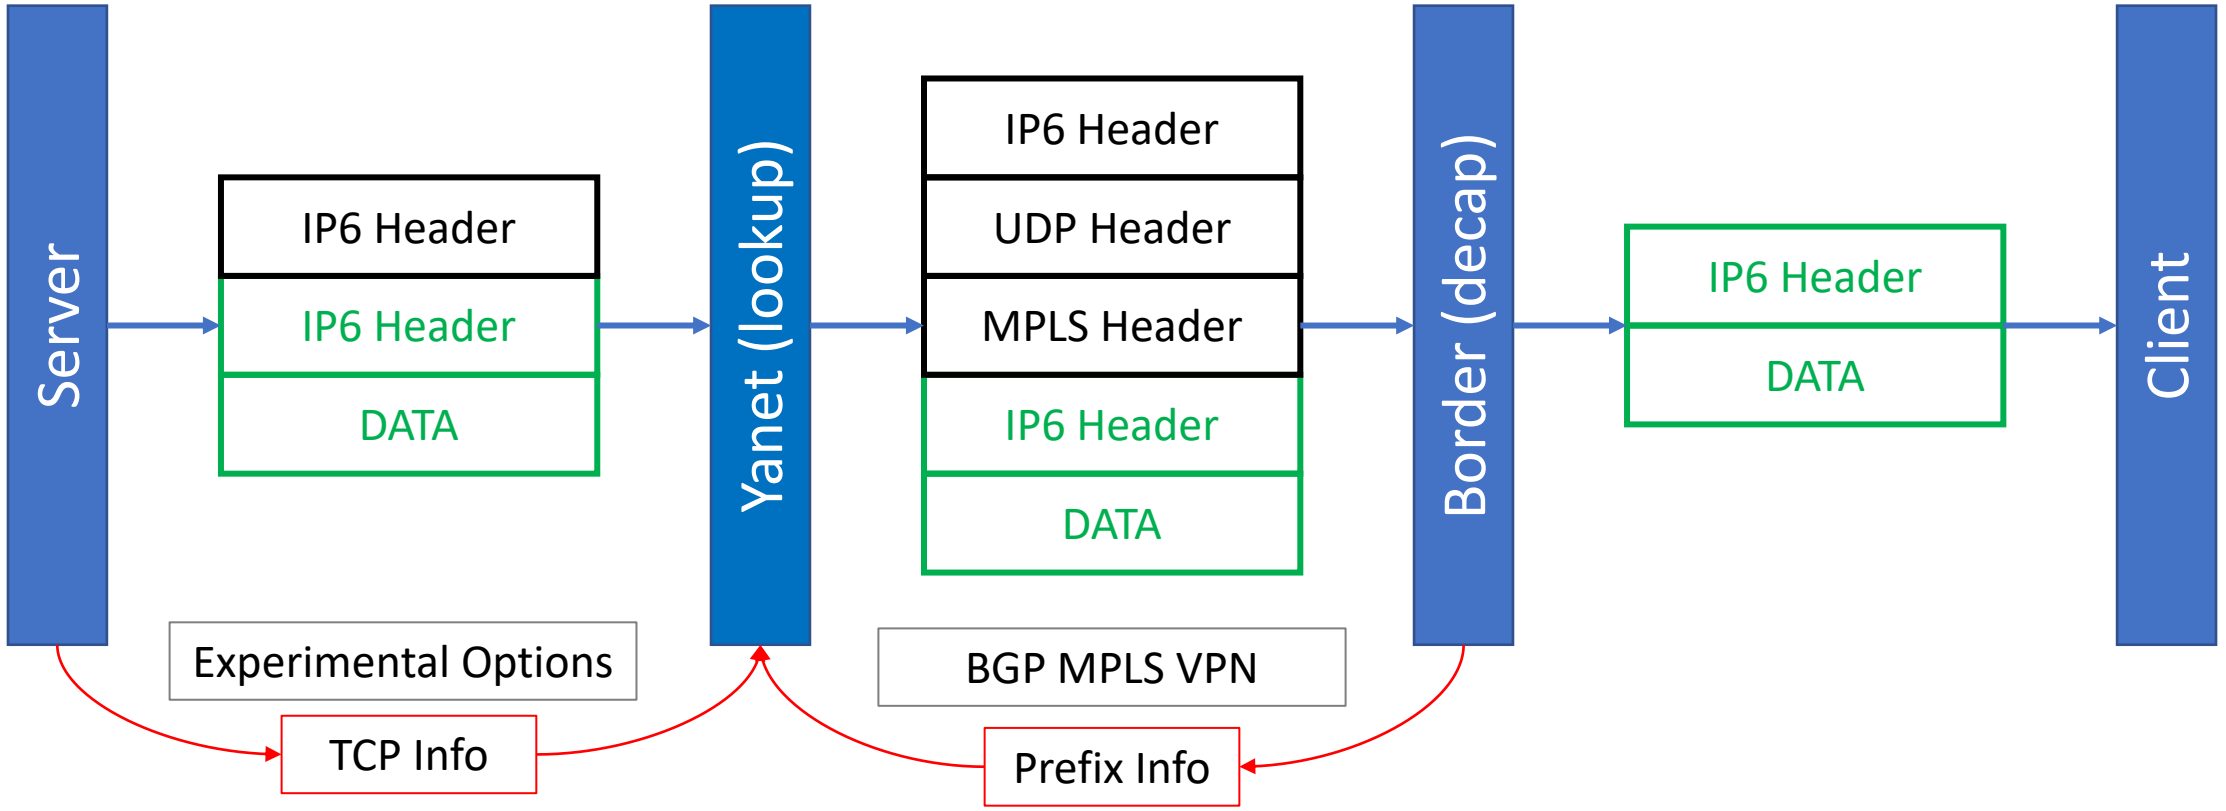

Example: Different Paths to China Telecom

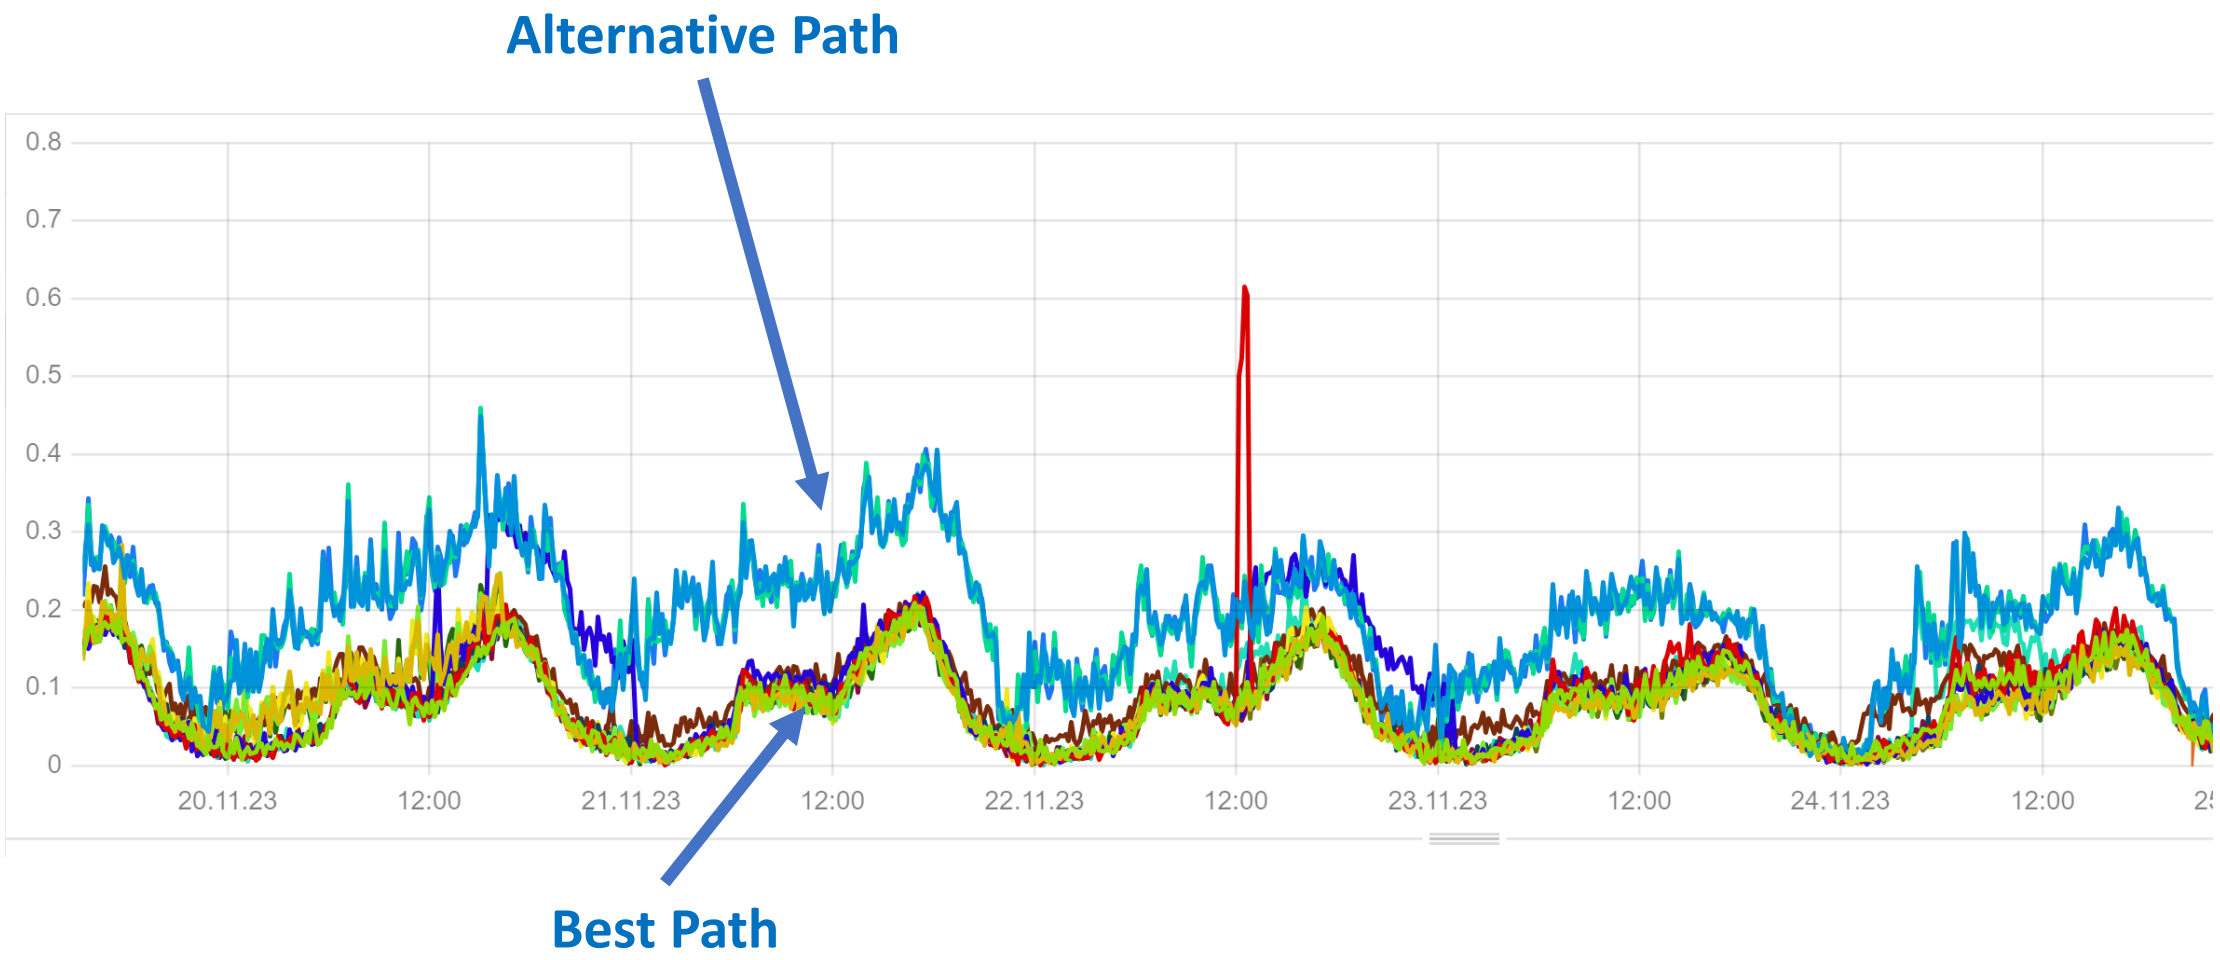

Example: One Peer, Two Paths

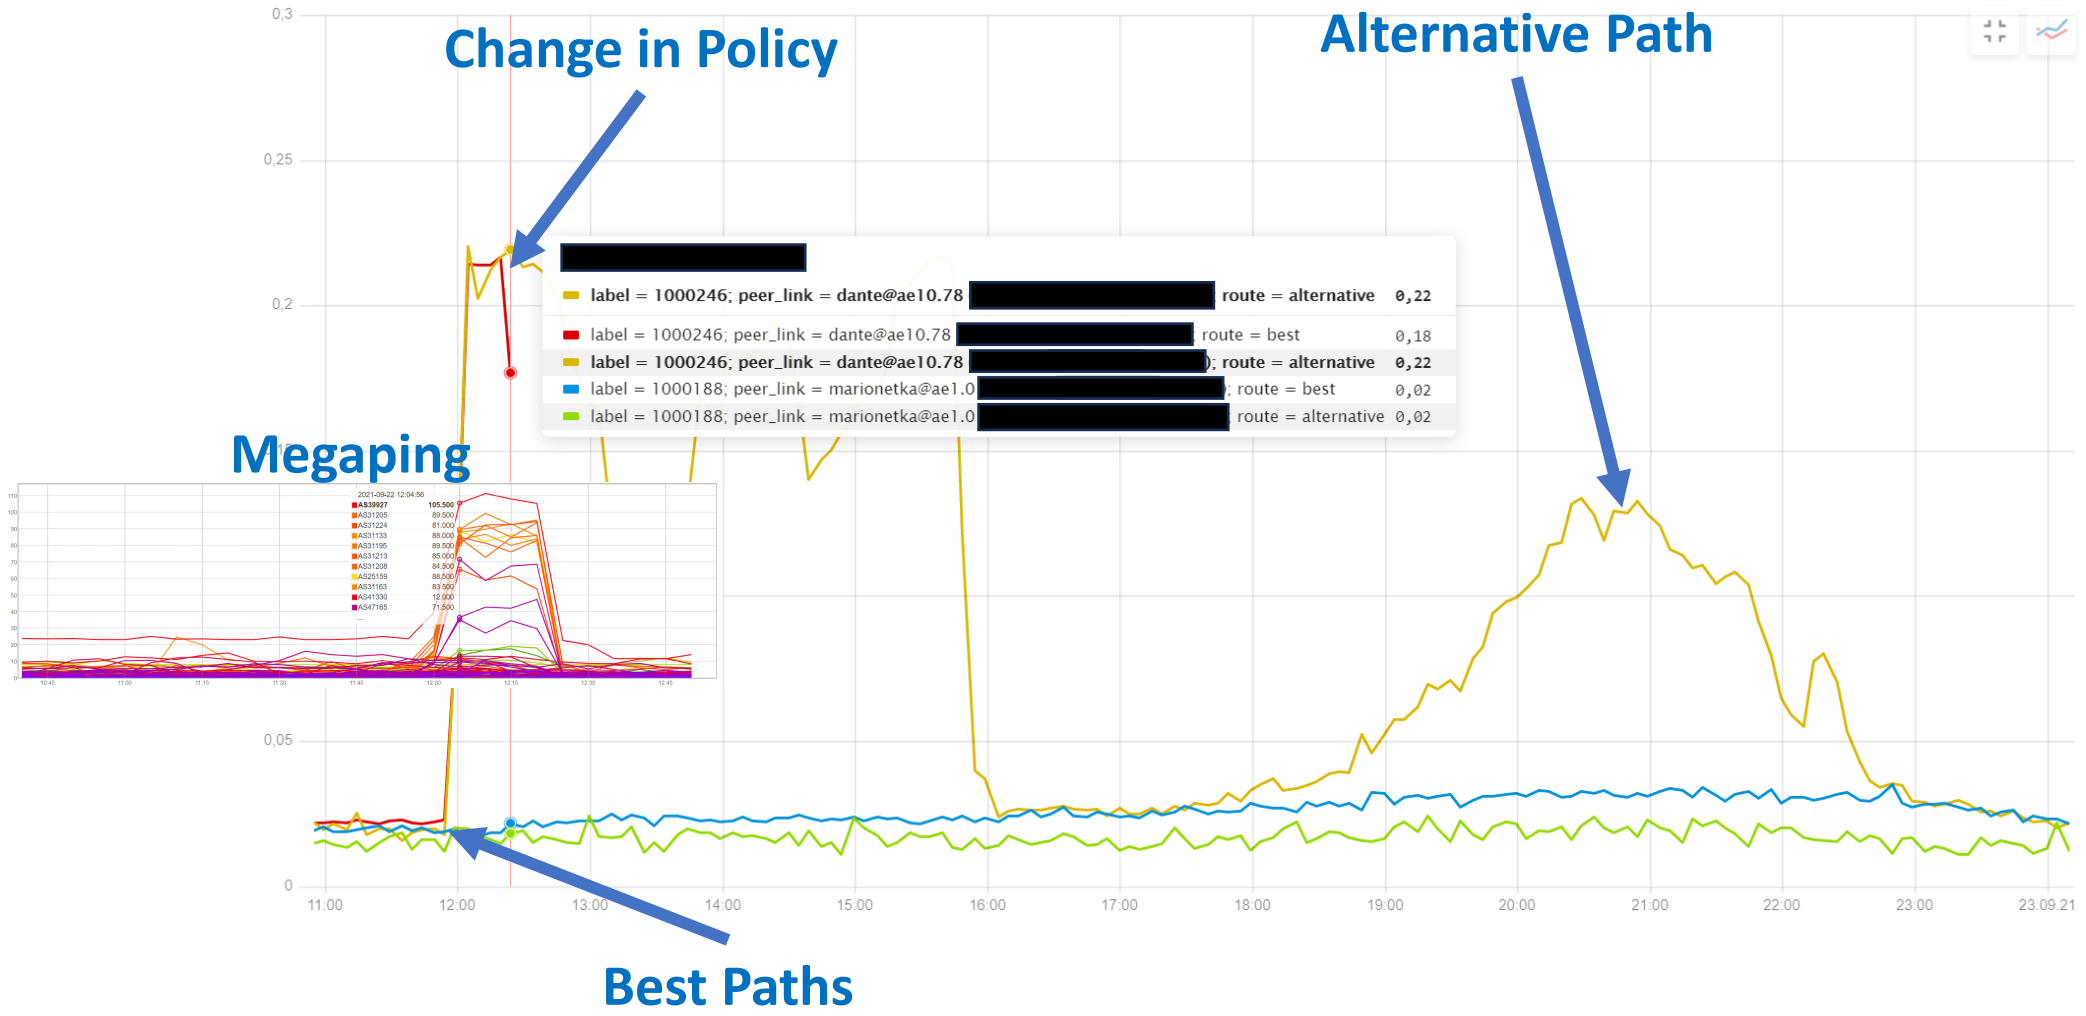


Dr. Egress

Dr. Egress: Summary

What Dr. Egress can do:

- Detect failing links/routes;
- Works both for IPv4 and IPv6;
- Make your RIB great again…

What Dr. Egress can’t do:

- Send you email… Wait for it. ☺
- Ingress monitoring is tricky;

Multilayer Egress Monitoring

NEL

Browser as a monitoring agent

Megaping

Ingress and egress network monitor

Dr. Egress

Per-link egress network monitor