Find The needle in a Haystack

Tenant-level network impact analysis in a Global-scale Infrastructure

Giscard Fernandes Faria
SRE Engineer
Cloud Reliability Lab, Huawei Ireland

Before: NEC, Senior Sistemas, AWS

SRECON EUROPE MIDDLE EAST AFRICA

HUAWEI

How a Packet Flows in a Cloud Network?

HTTP
IP

HTTP
EIP

CDN

Imagine troubleshooting it in a billion request and millions of device scale.

IP

L4

HTTP
LB-VPC
VXLAN
LB-VIP

L7

HTTP
LB-VPC
VXLAN
LBVIP
LB7-IP

HTTP
VPC-IP
VXLAN
BM-IP

HTTP
VPC-IP

VM

HTTP

OBS

How Virtual Private Cloud Works?

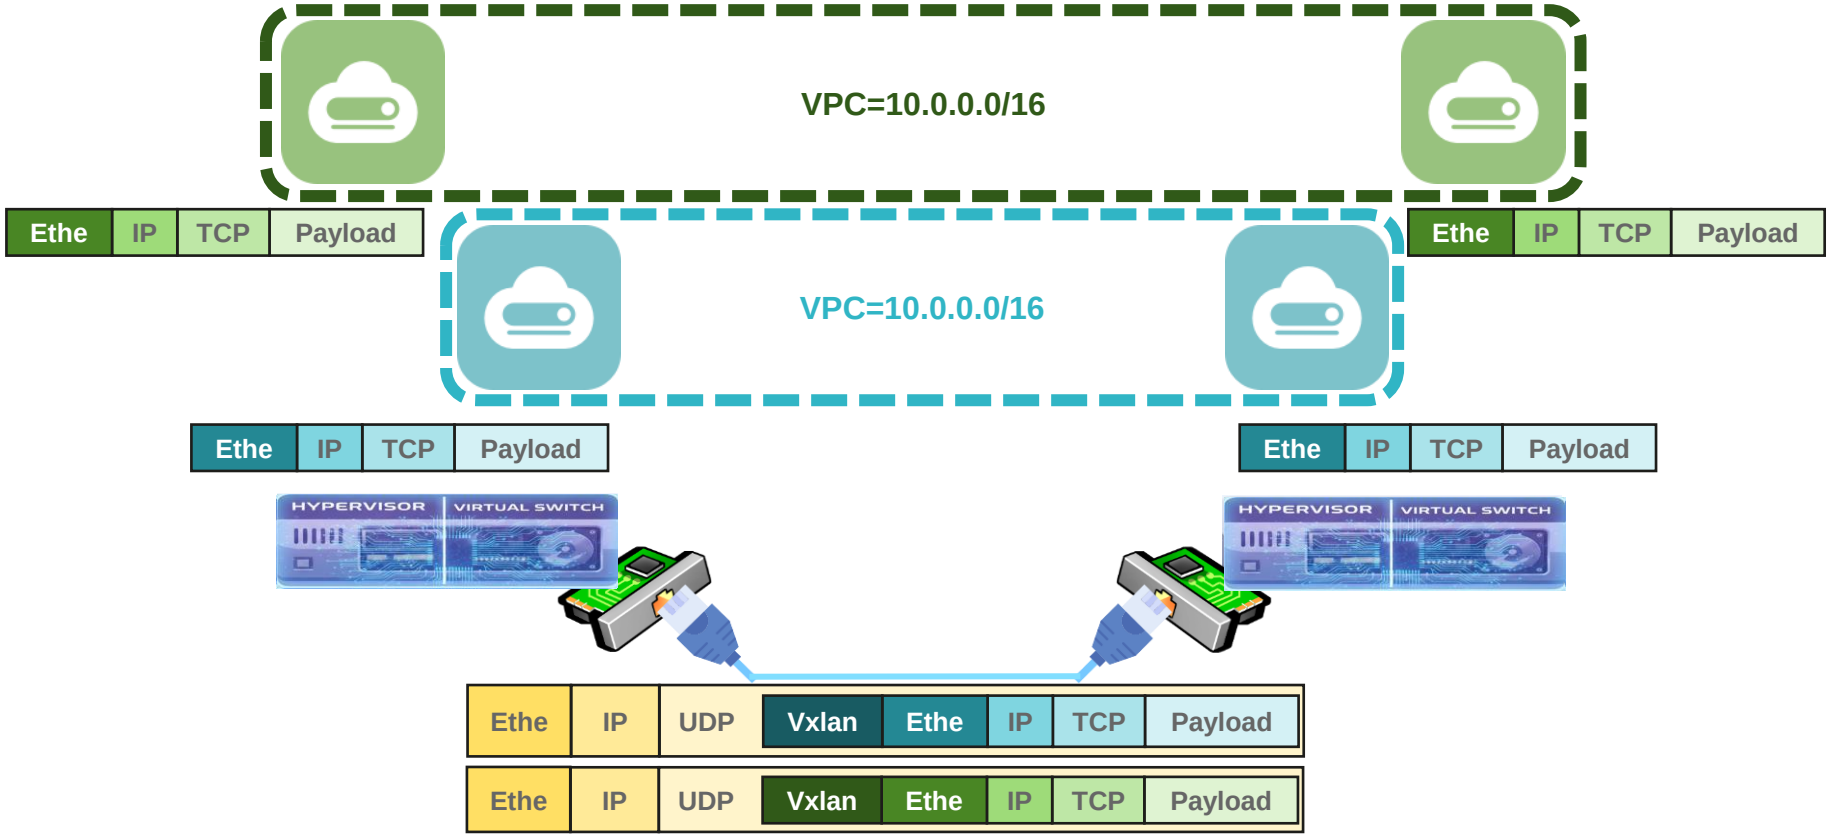

Outer vs Inner
Physical vs Virtual
Underlay vs Overlay

Building Blocks of a Cloud Infrastructure

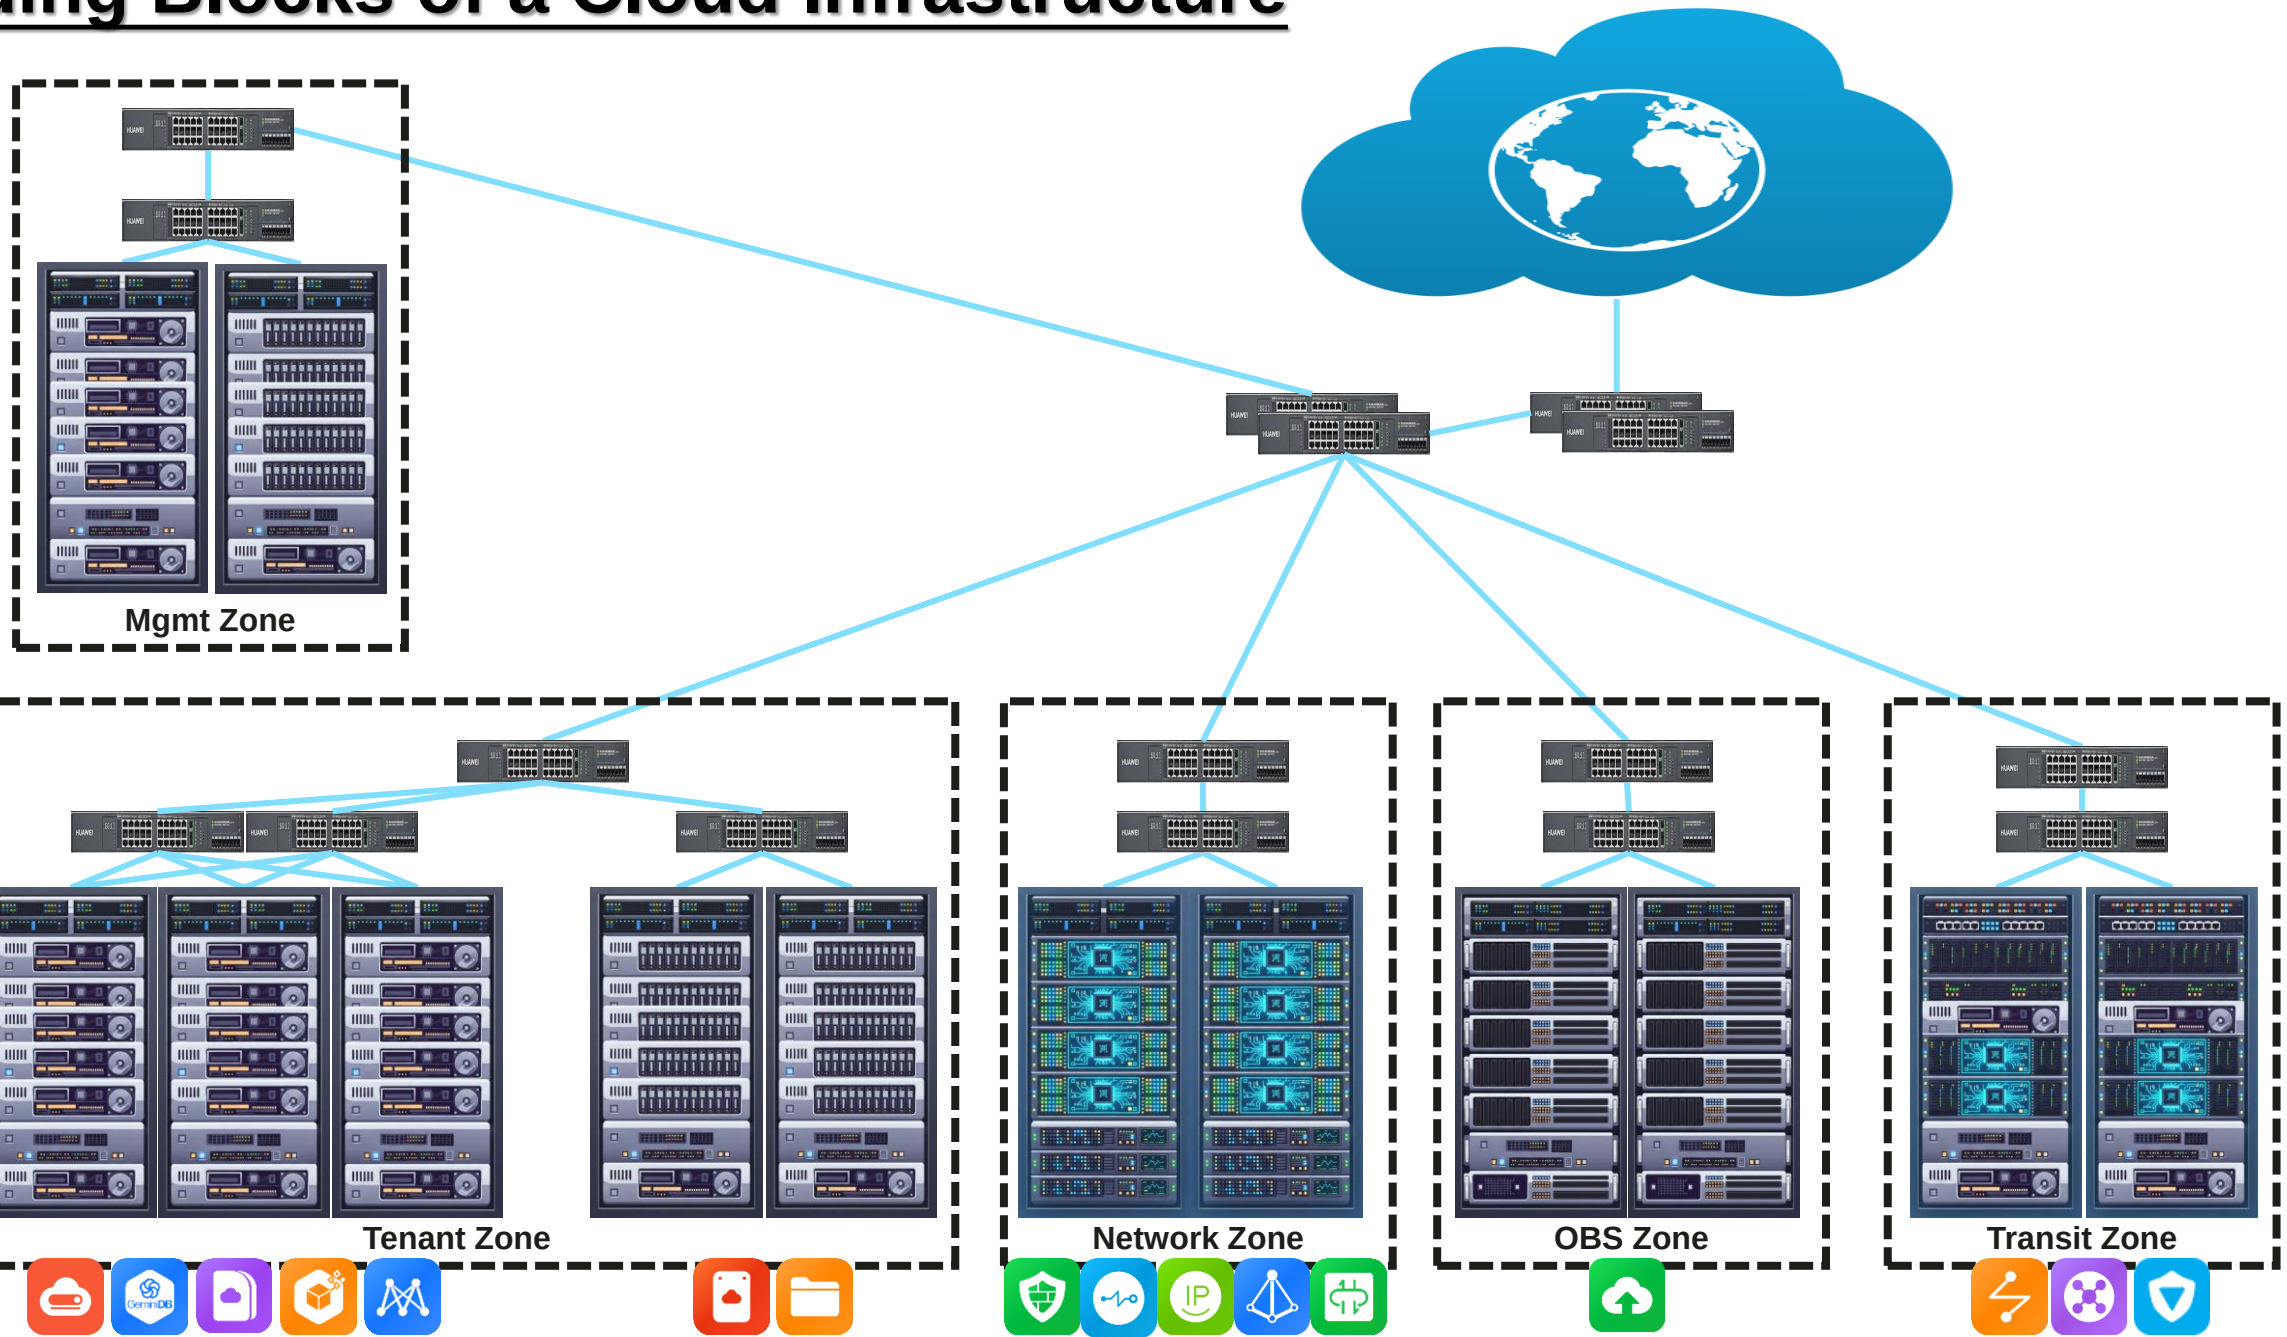

Data Center Network and Wide Area Network

HUAWEI

WAN#1 v2.0

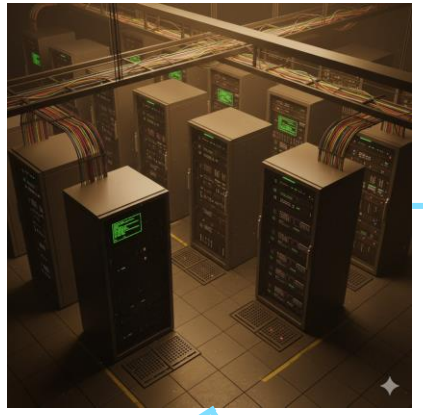

Internet

WAN#2 v3.0

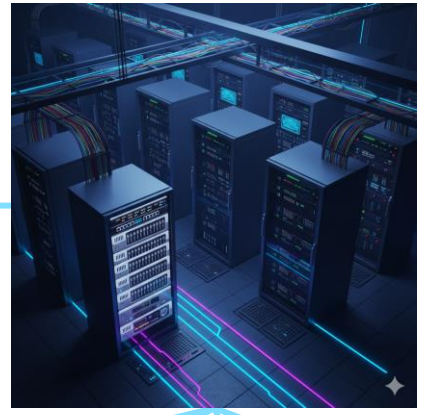

Heterogenous Hardware and Design

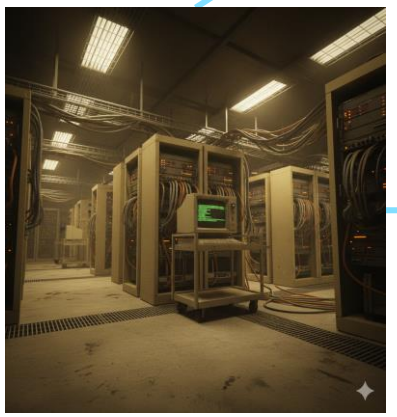

DCN#1 v1.0

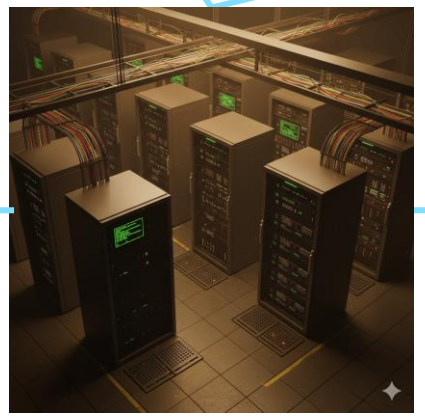

DCN#2 v2.0

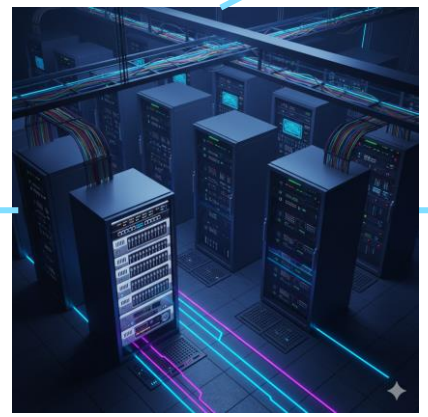

DCN#3 v3.0

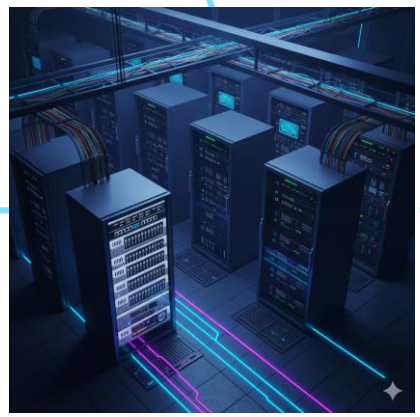

DCN#4 v3.0

DCN#5 AI/ML v1.0

Global Cloud Infrastructure

HUAWEI

Servers **20%** growth

33 Public Regions

0 P1
incident in
last four years

MTTR33.6 min

Fundamental Services SLO
99.999%

Core services SLO
99.99%

Customers
- Corporation: **20000+**
- Premium Users: **3 million+**
- Developers: **10 million+**

Services
- Cloud services: **895**
- Prod environment changes: **9.52 M/year**

Hardware
- Servers: **2 million+**
- Data Center: **1300+**

A global storage and computing network with wide coverage, low latency, and meeting customer experience needs

CloudOcean (Global)
Scale of 1 million servers
China<30ms Overseas<50ms

CloudSea (Regional)
Scale of 100K servers
Latency<10ms

CloudLake (Edge access)
Scale of 5000 servers
Latency<5ms

Network Monitoring

HUAWEI

Multi Customer Impact

Continuous Monitoring and Automation

Node Metrics

End-to-End Metrics

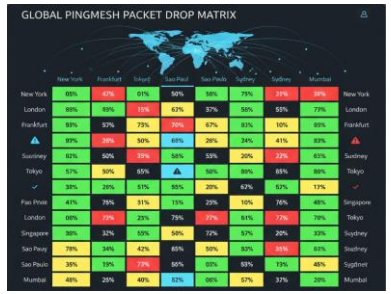

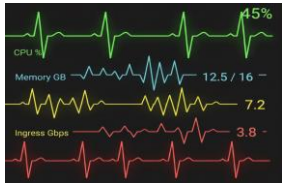

Routers

Hypervisor

Virtual Devices

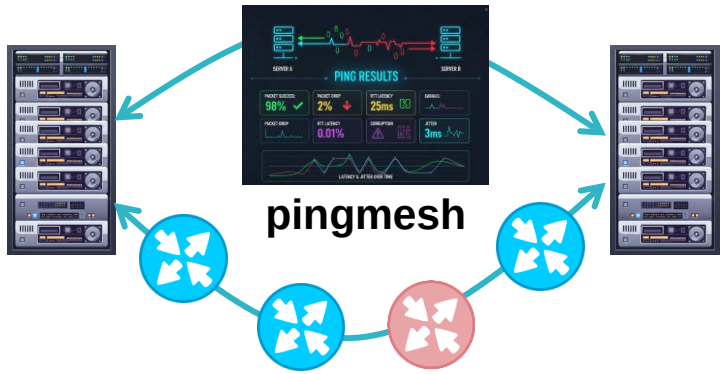

pingmesh

Beyond Conventional Monitoring

Flying under the radar

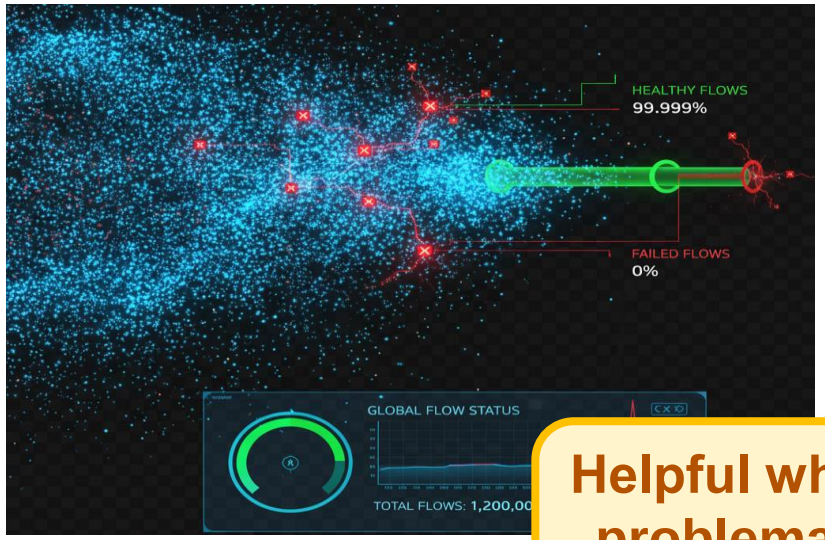

Synthetic vs Real Traffic

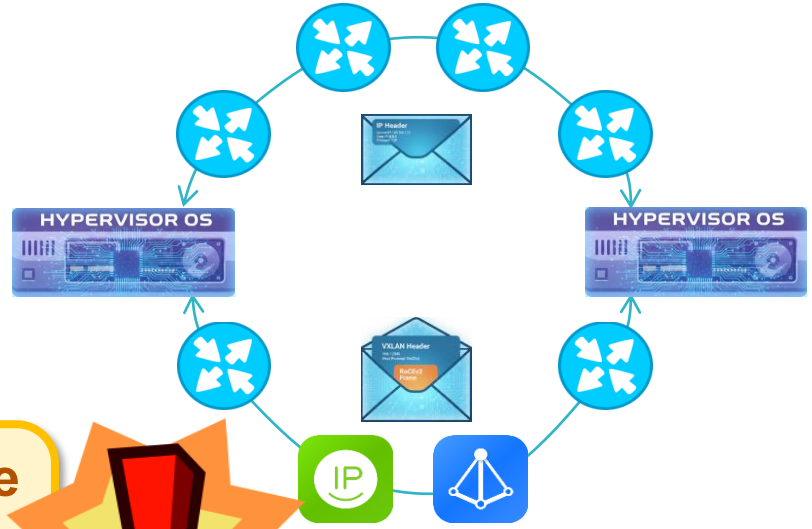

Helpful when "fishing" the problematic flow but not when troubleshooting it.

sFlow + NetStream

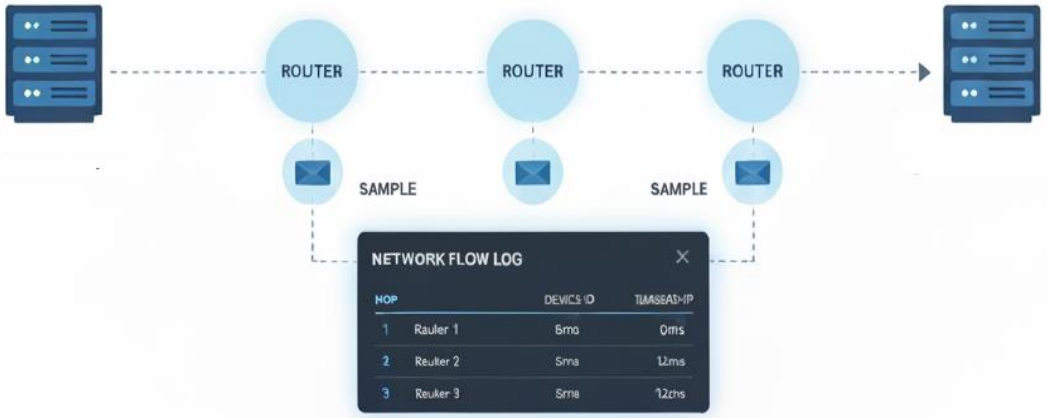

ASIC vs CPU

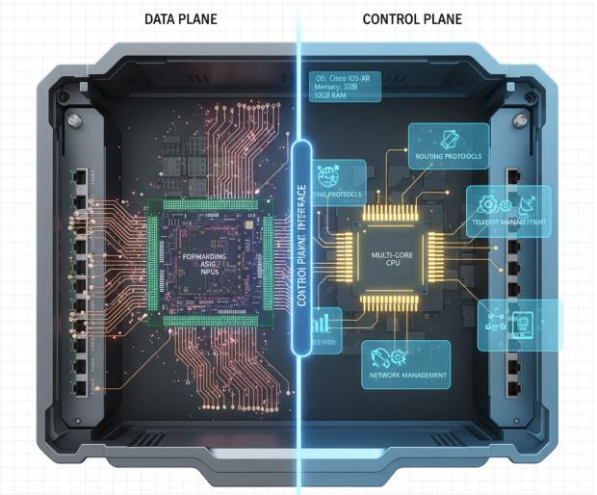

Goal is to track a flow and pinpoint where it is failing!

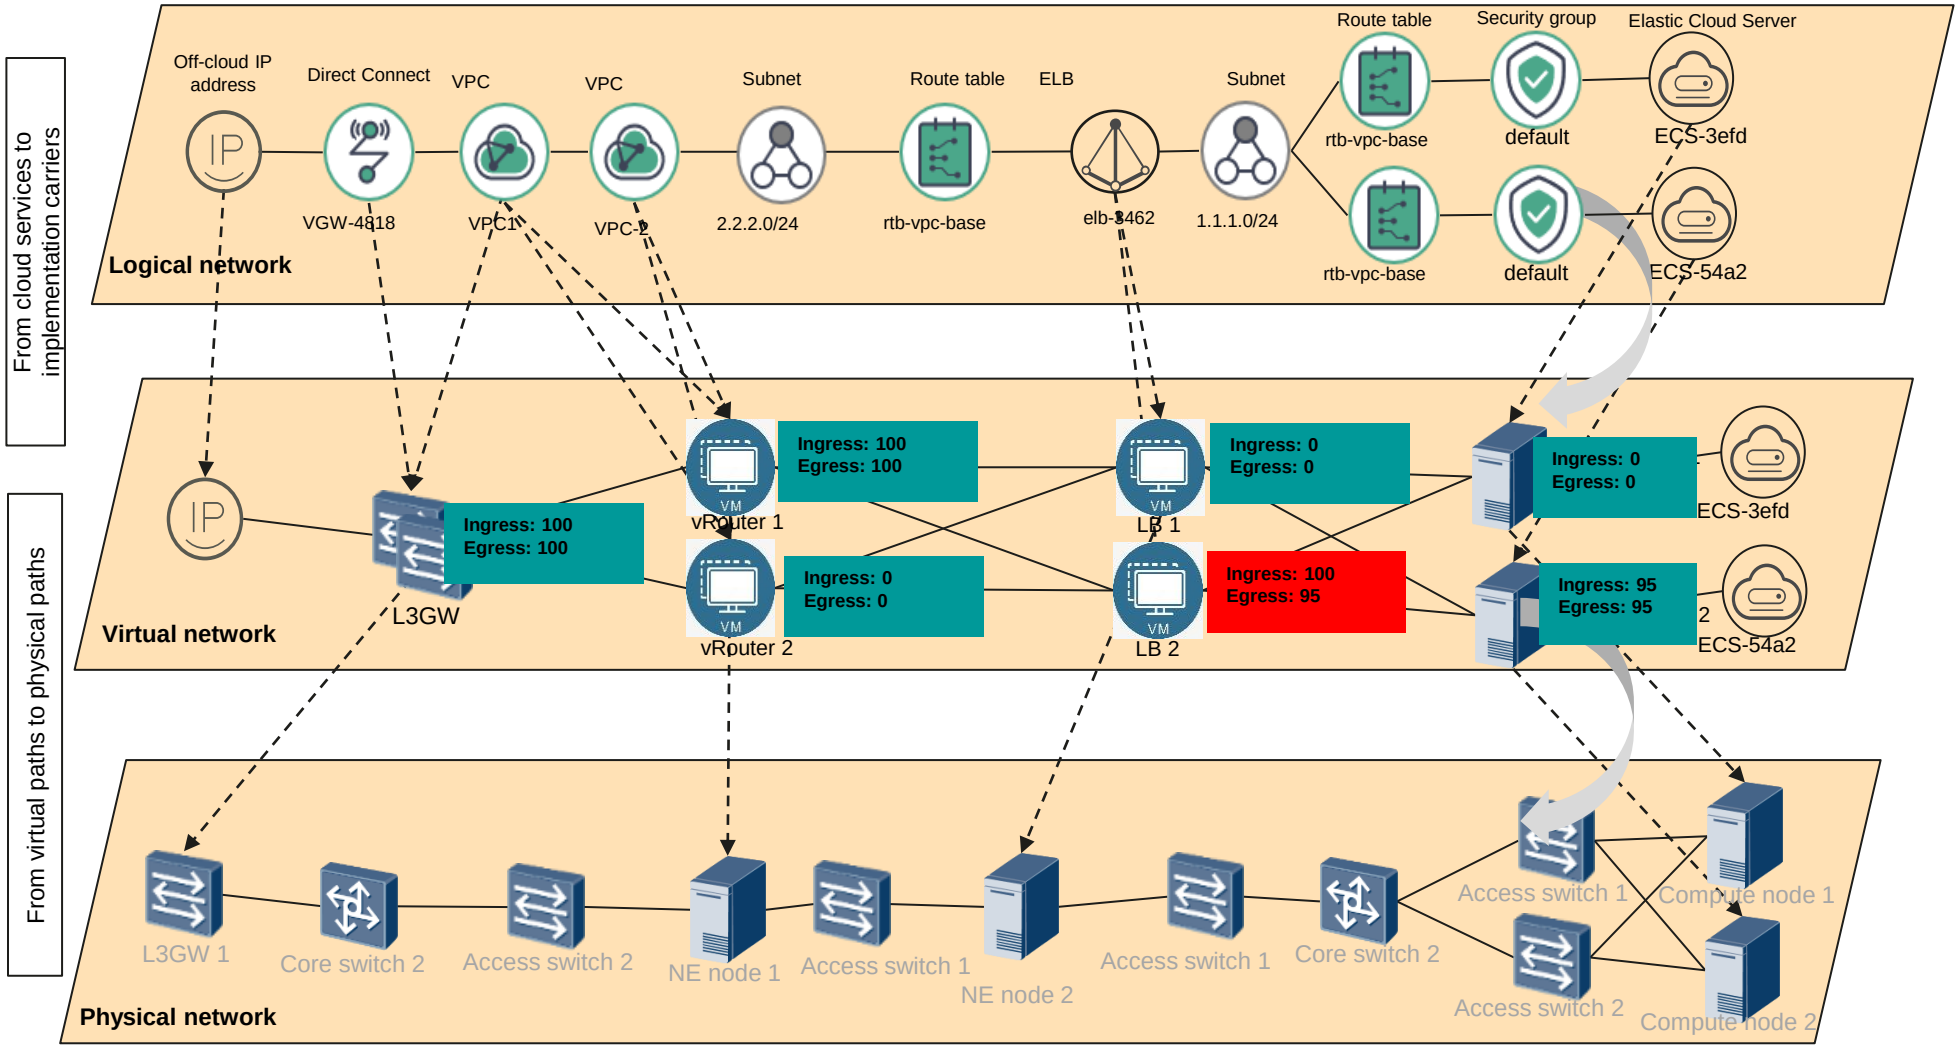

CloudNetDebug Architecture

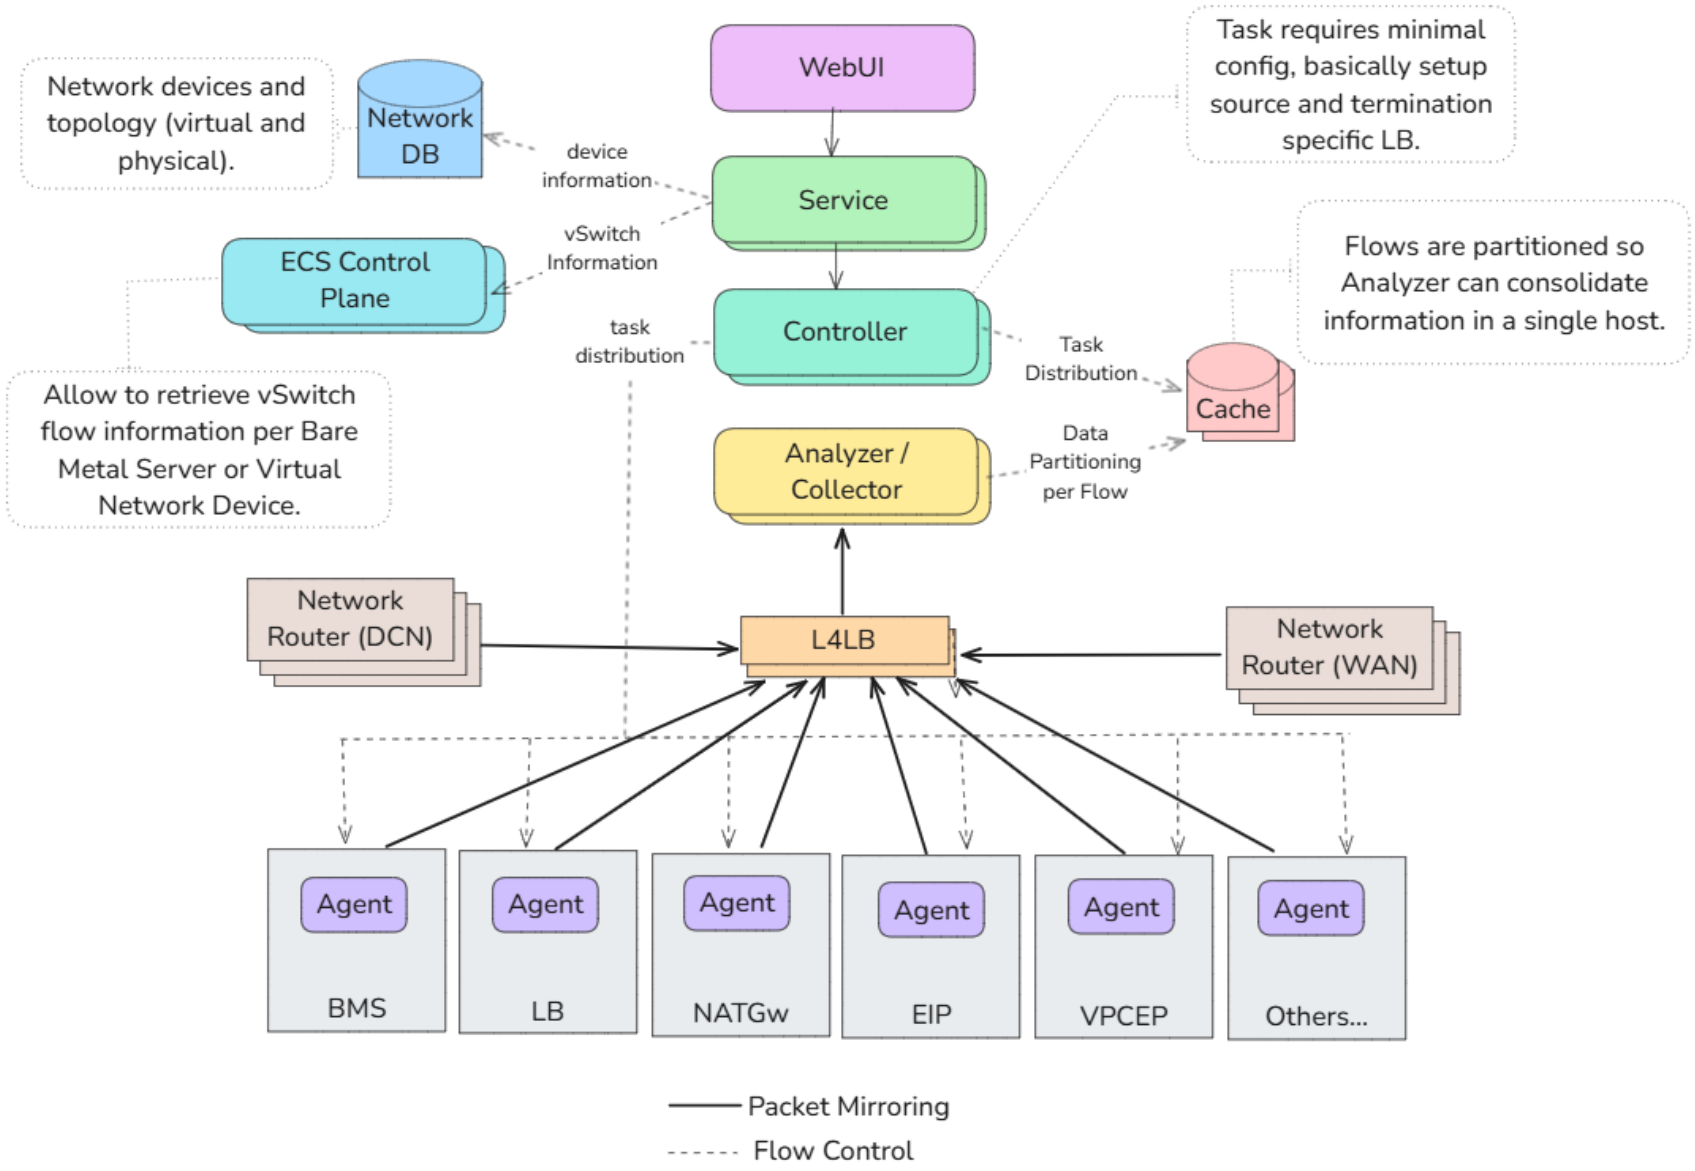

Monitoring Flow at Source and Destination

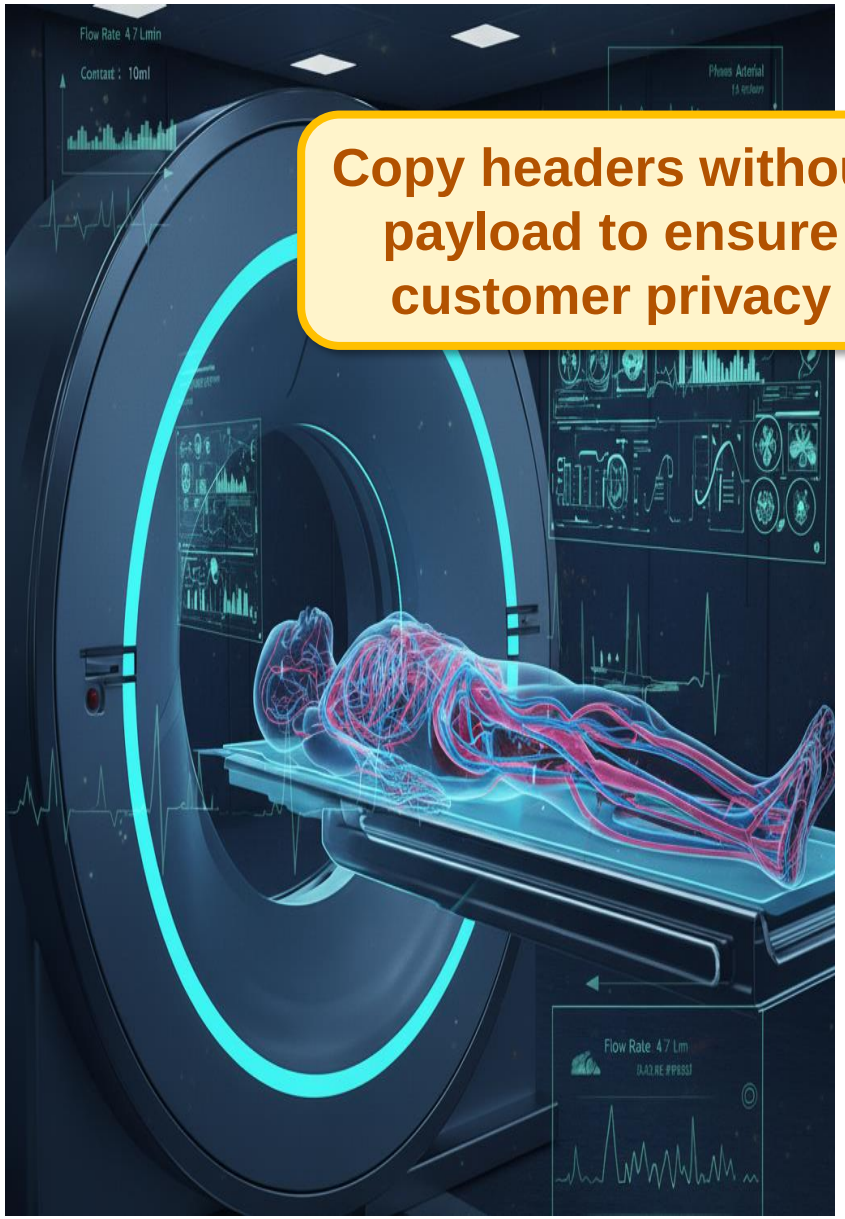

Copy headers without payload to ensure customer privacy

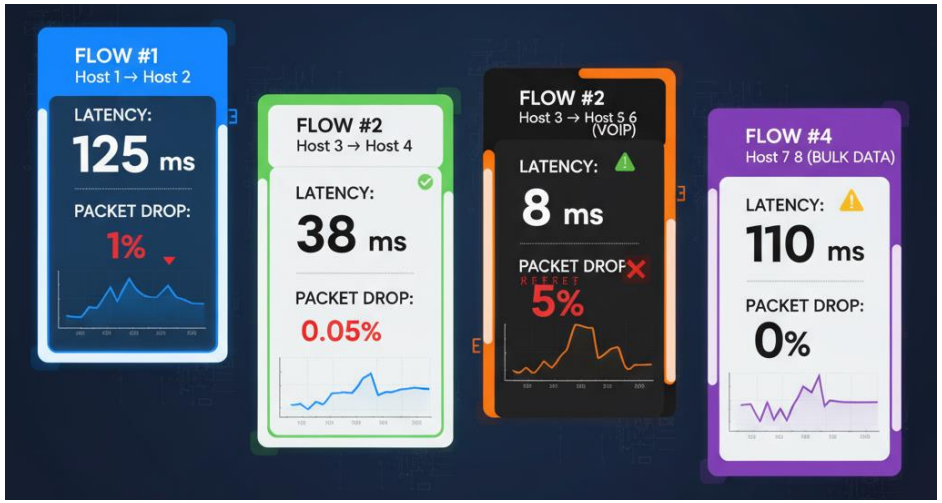

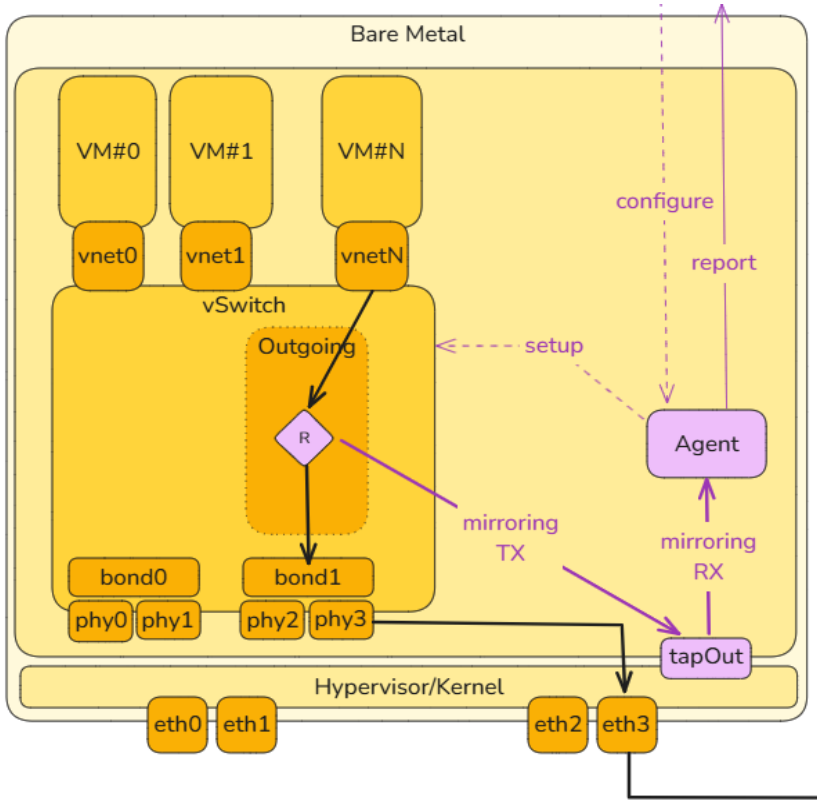

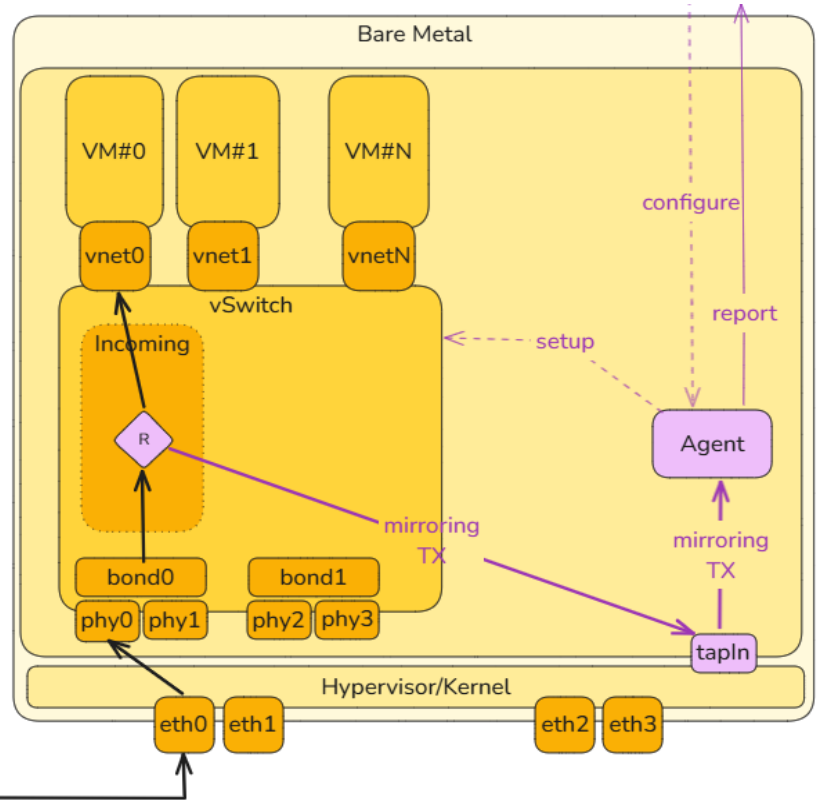

Enabling Flow Monitoring on Conventional Applications

HUAWEI

vSwitch

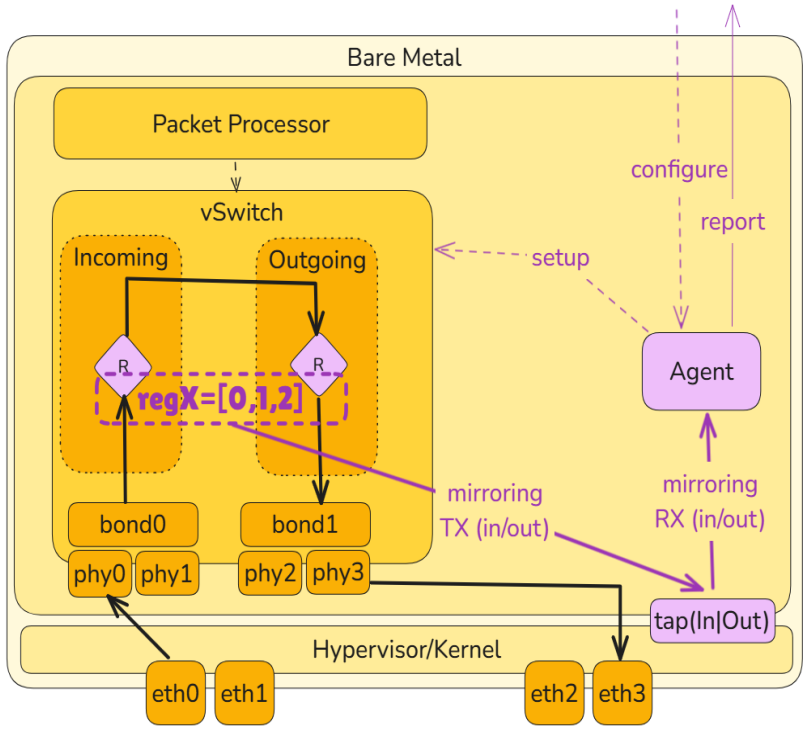

Map incoming and outgoing flows using reg15 config

LVS / Netfilter / Nftables

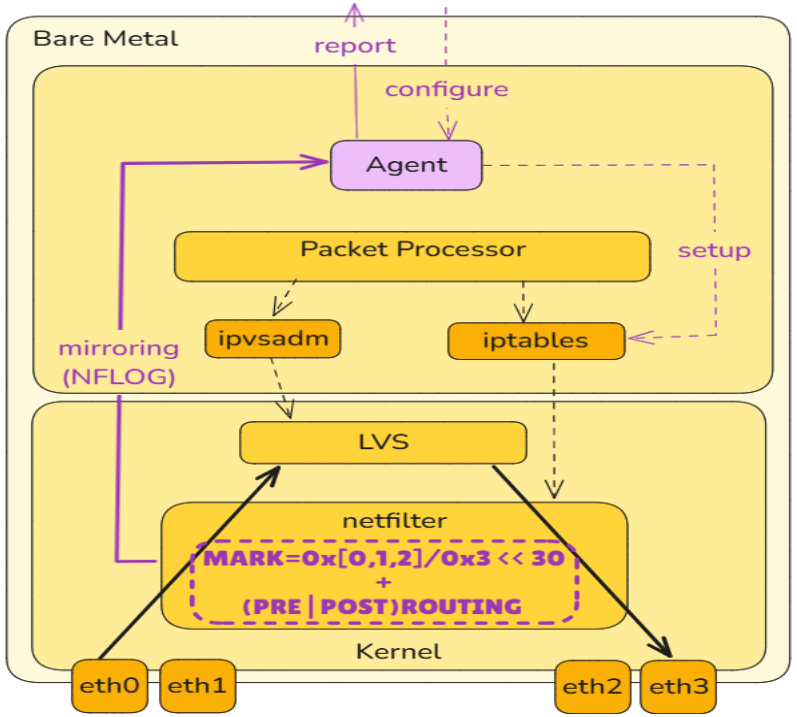

Map incoming and outgoing flows using MARK*

AF_INET (ex: NGINX)

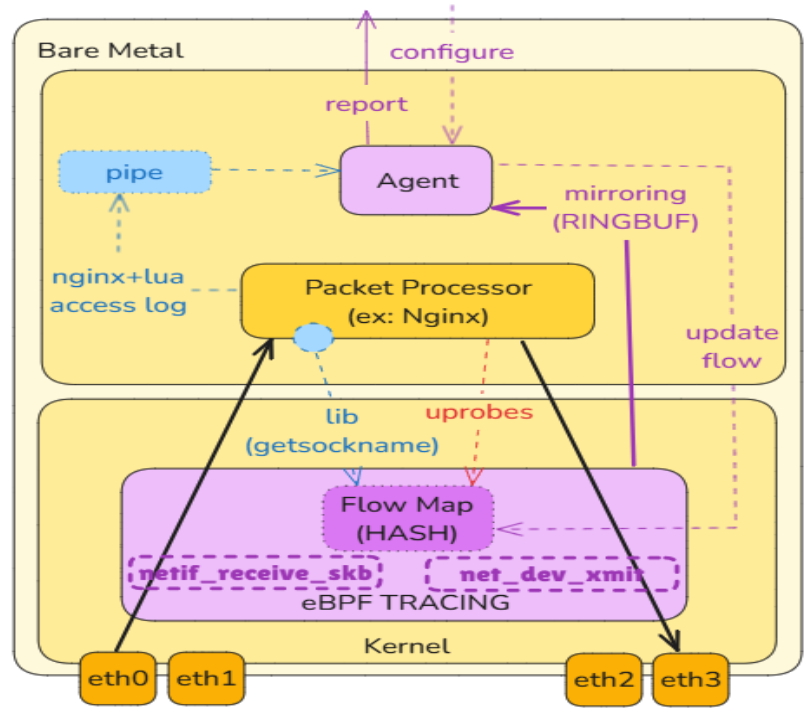

eBPF Tracing to the rescue with flow mapping

Enabling Flow Monitoring for XDP Programs

HUAWEI

eBPF XDP

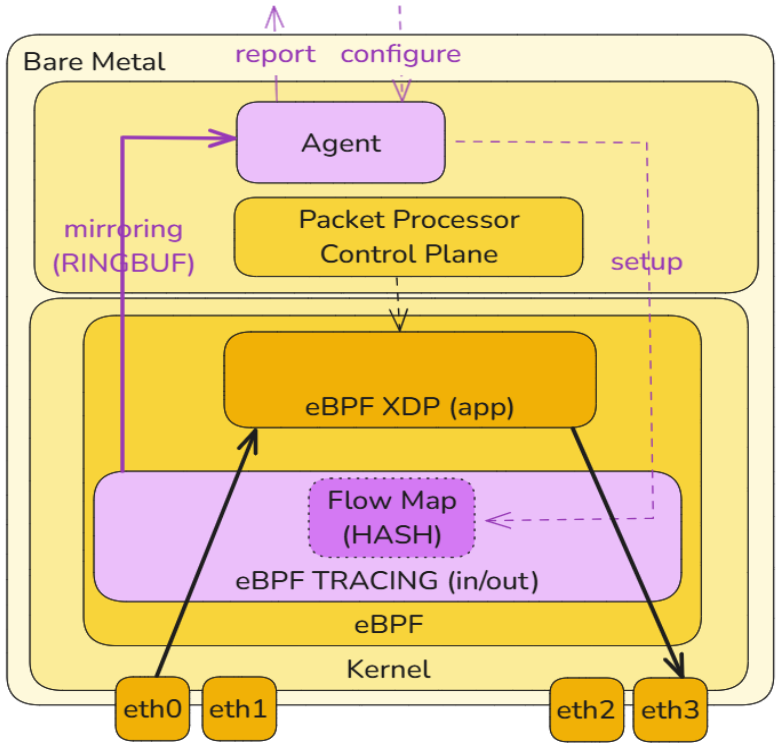

eBPF Tracing to the rescue by using an xdpdump-like tool

AF_XDP

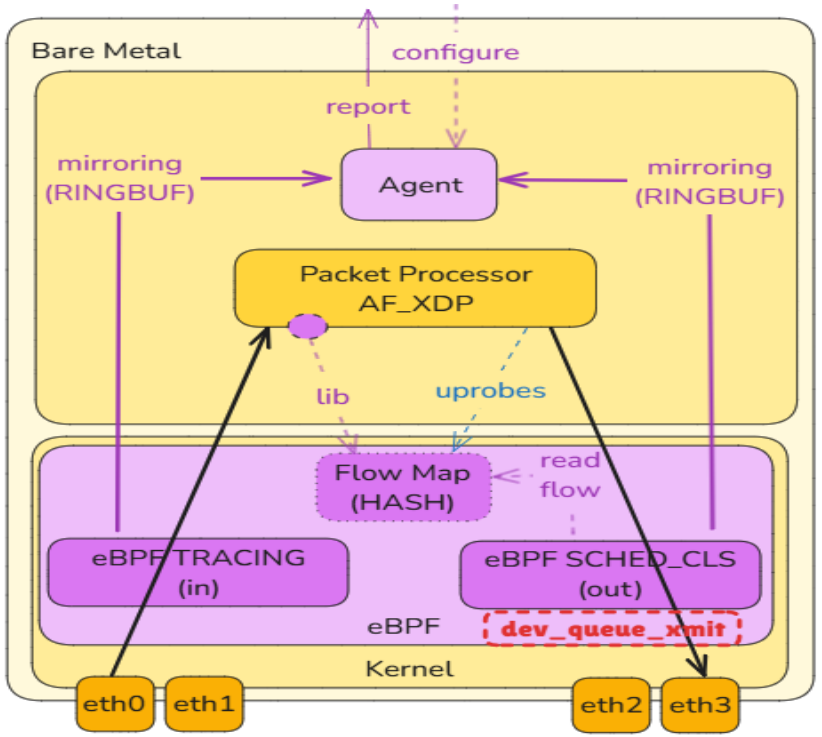

eBPF Tracing for incoming + eBPF TC for outgoing + User Library. This is getting nasty.

XDP Offload

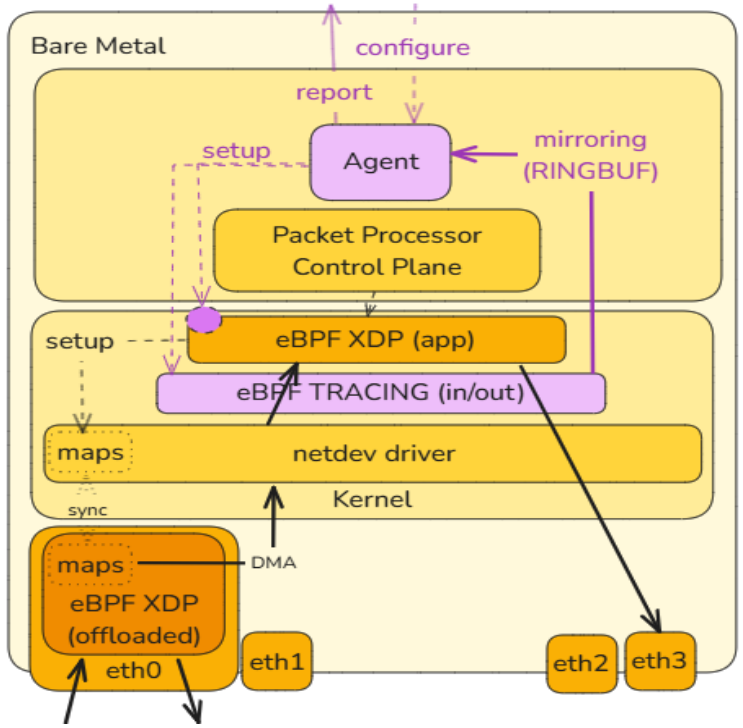

All that and still no guarantee it will work?

Enabling Flow Monitoring for Ad-Hoc Use Cases

HUAWEI

DPDK

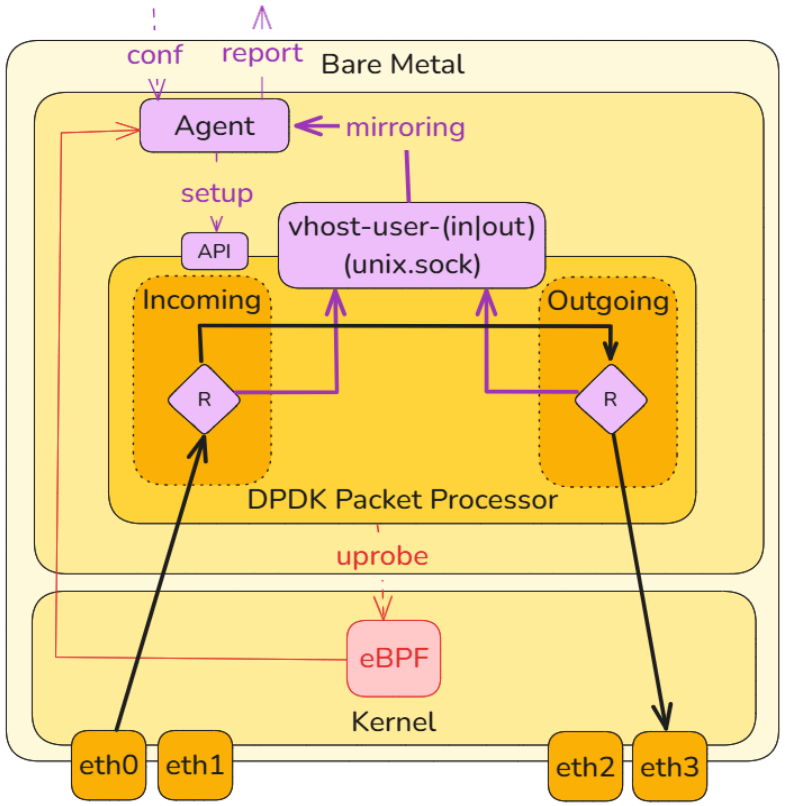

If you have to change, make it in the "right" way.

ENPU

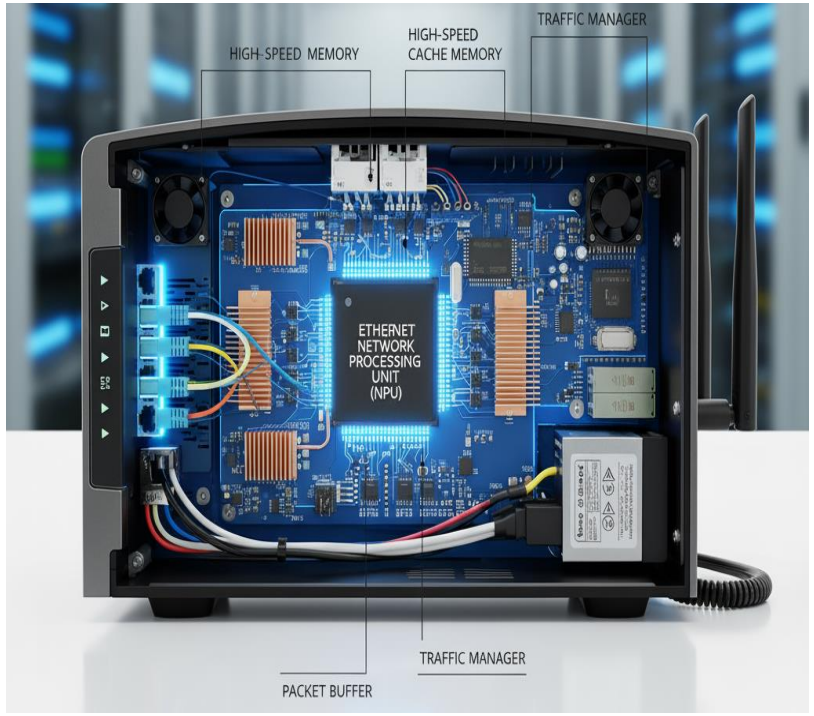

You cannot win every battle.

Others…

Packet Colouring

IPv4 and IPv6

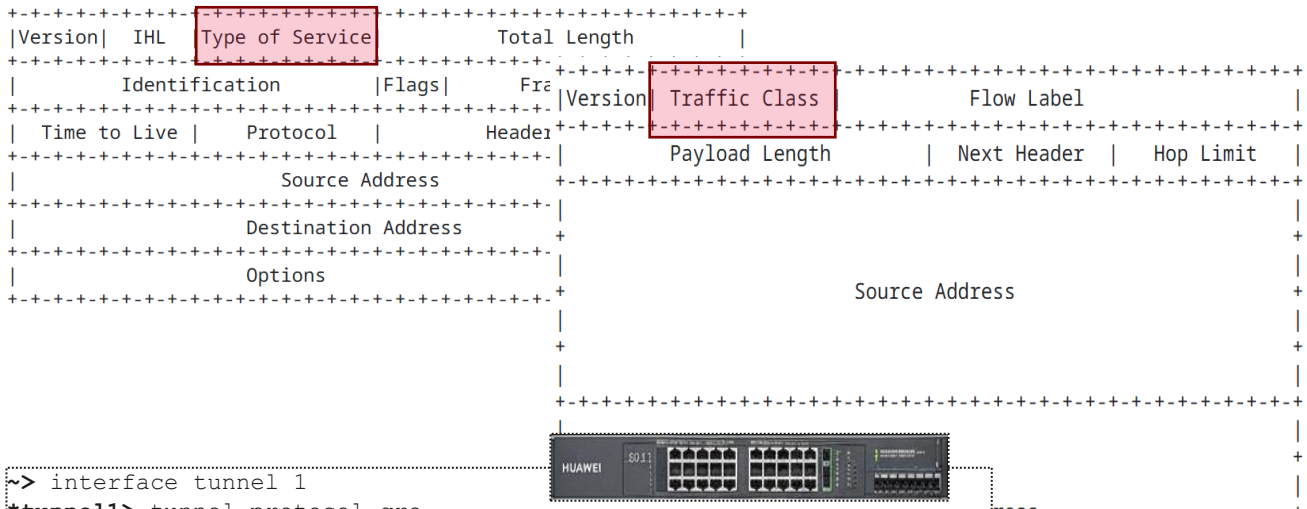


```
~> interface tunnel 1
*tunnel1> tunnel-protocol gre
*tunnel1> source <ip-loopback>
*tunnel1> destination <server-vip>
*tunnel1> quit

~> observe-port 1 interface tunnel 1 truncate packet-length 100

~> traffic classifier colored
*classifier-colored> if-match flag bit-and 4
*classifier-colored> quit

~> traffic behavior mirror
*behavior-mirror> mirroring to observe-port 1
*behavior-mirror> quit

~>traffic policy mirror-colored
*policy-mirror-colored> classifier colored behavior mirror
*policy-mirror-colored> quit

~> interface range GiEth0/0/1 to GiEth0/0/10
*port-group> traffic-policy mirror-colored inbound
*port-group> traffic-policy mirror-colored outbound
```

Inband (IFIT)

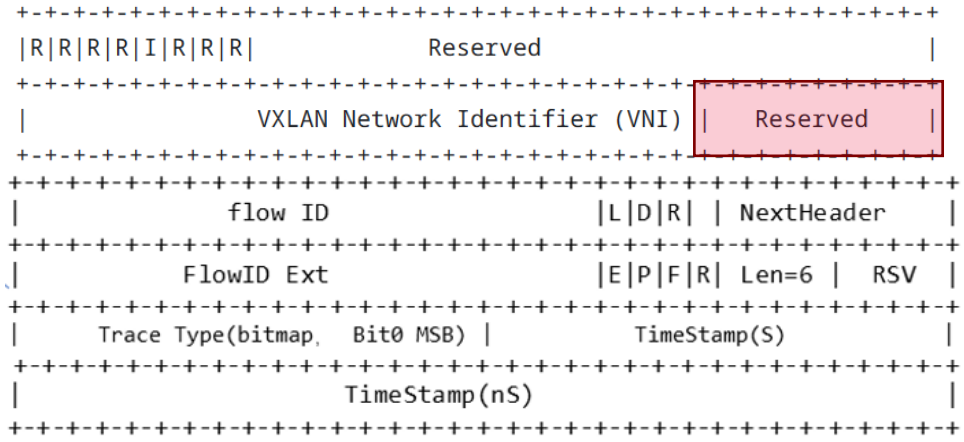

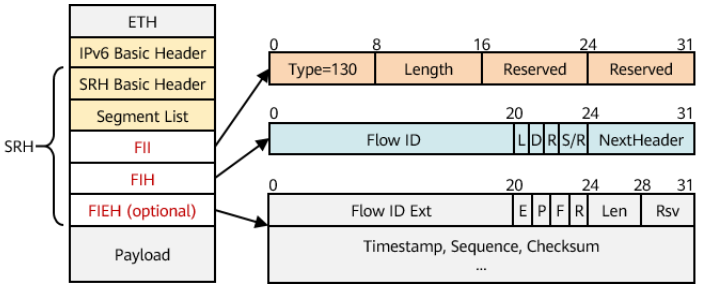

Efficiency and Standardization

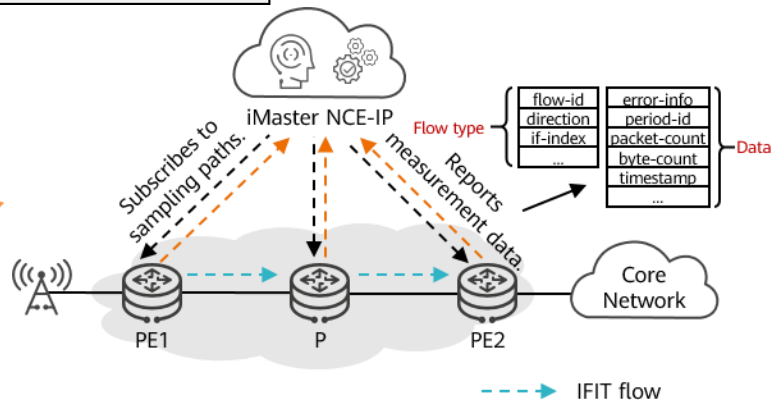

CloudNetDebug UI

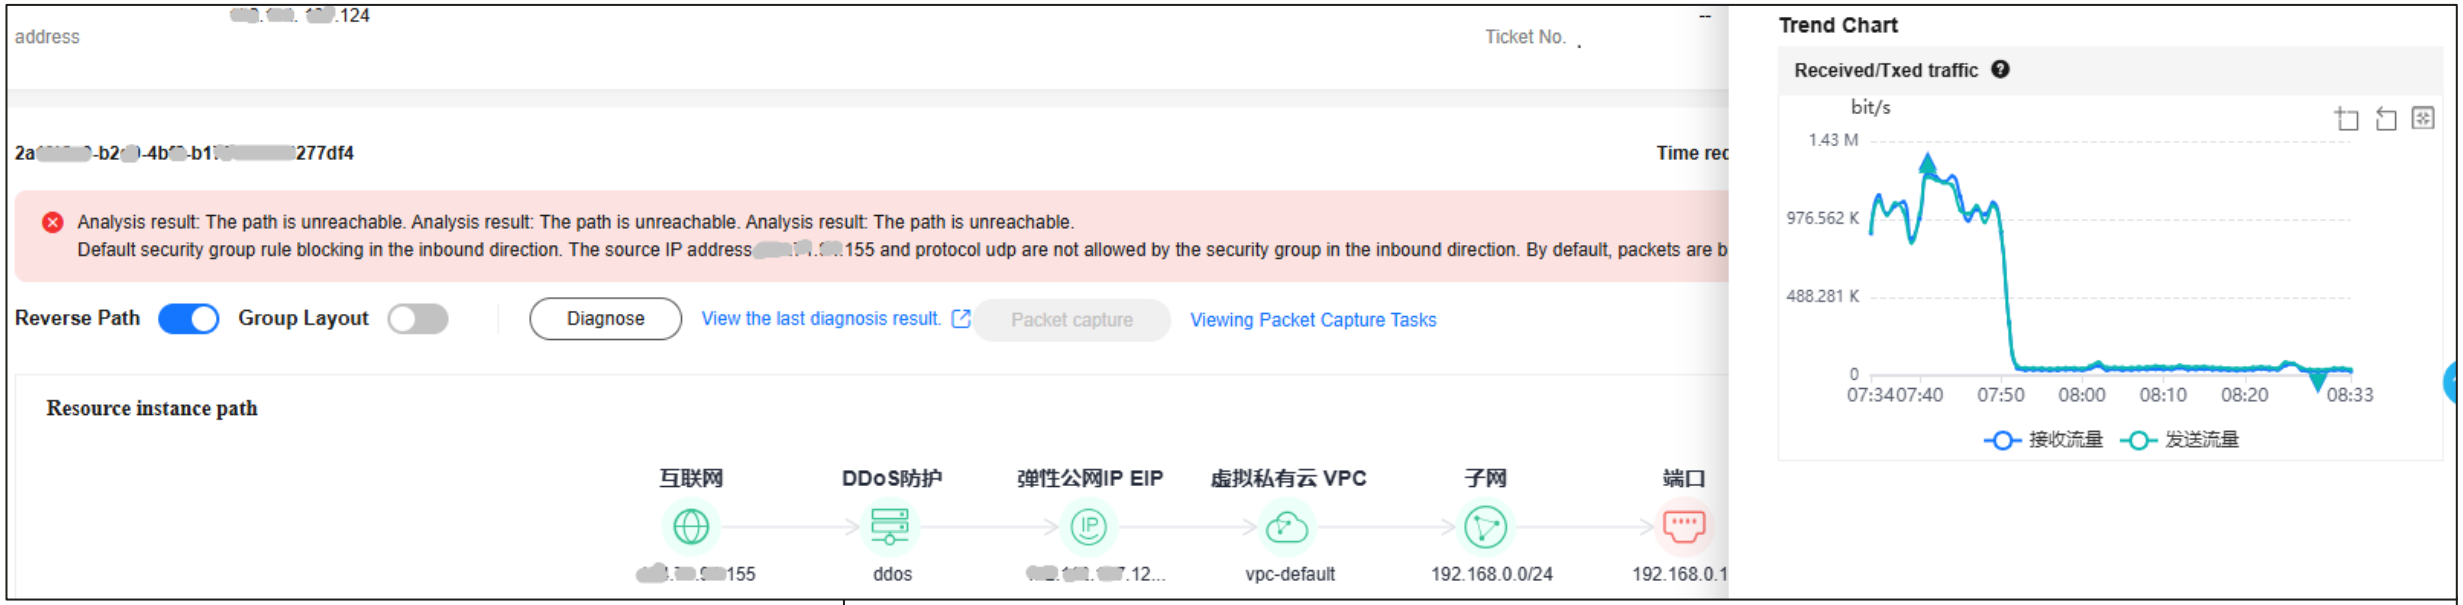

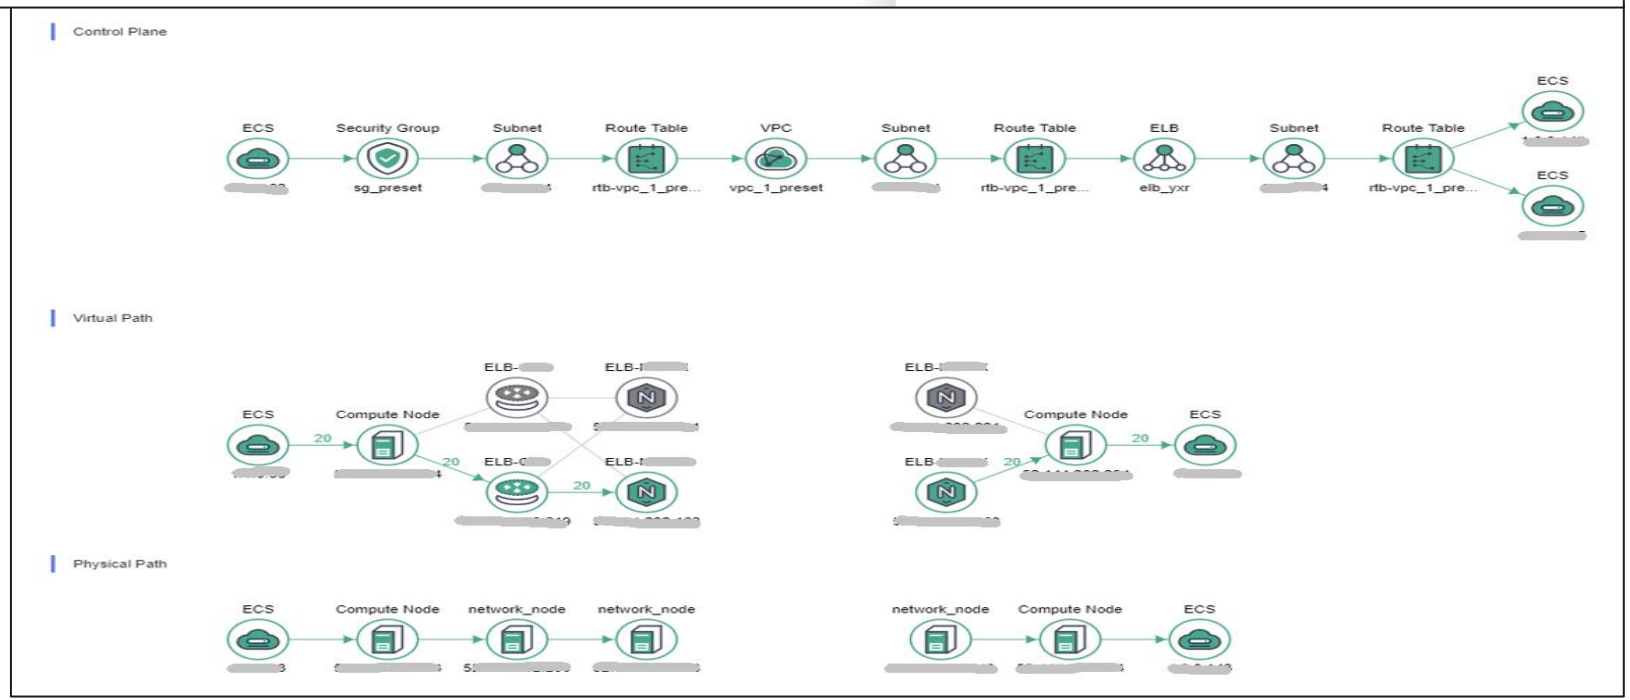

- **Millions of nodes onboarded**
- **+300 users / +2K Investigations**
- **Up to 80% Faster Troubleshooting**

Lessons Learned

HUAWEI

Customer Privacy is Paramount

✓ *A lot can be achieved by only using Metadata*

Safety

✓ *Agent hard limits (CPU, Mem, I/O, …)*
✓ *One Click Stop in case of risk or imminent danger*
✓ *Support Sampling, some flows will have dozens of thousands of packets*

Static Config

✓ *Every change introduces risk. Minimize configuration changes!*

Heterogenous Environment

✓ *Too many versions of hardware and software, you have to support most of them*
✓ *Missing some devices is okay, prioritize the most relevant traffic*
✓ *Understand long term strategy from each team, to influence and correct your system*
✓ *Standardize is the best way to manage complexity*

What's Next?

Close the gap for unsupported Services

✓ *Some network services yet to be supported*

Onboard Physical Devices

✓ *Also having the standardized packet colouring in place everywhere*

Continuous Monitoring

✓ *Instead of being engaged by other alarms or in some cases the customer support*

✓ *Reduce the time to find the "right" flow*

General available for end customers

✓ *If it is useful for SRE, let customer use it will save their time too*

Thank You.