Skyline: A Cloud-Centric Internet Monitoring Engine

Shixian Guo*, **Ziqian Liu***, Yangyang Bai, Yuan Chen, Kefei Liu, Qi Zhang, Songlin Liu, Yang Lv, Jianwei Hu, Gen Li, Zhenyang Zhong, Sisi Wen, Yongbin Dong, Feng Luo, Anjian Chen, Rui Han, Jiale Feng, Lingpei Meng, Siwan Chen, Hang Li, Shuai Xu, Juntao Zhong, Chaoran Hu, Yibo Huang, Yiming Qiu

The Cloud-Internet Ecosystem

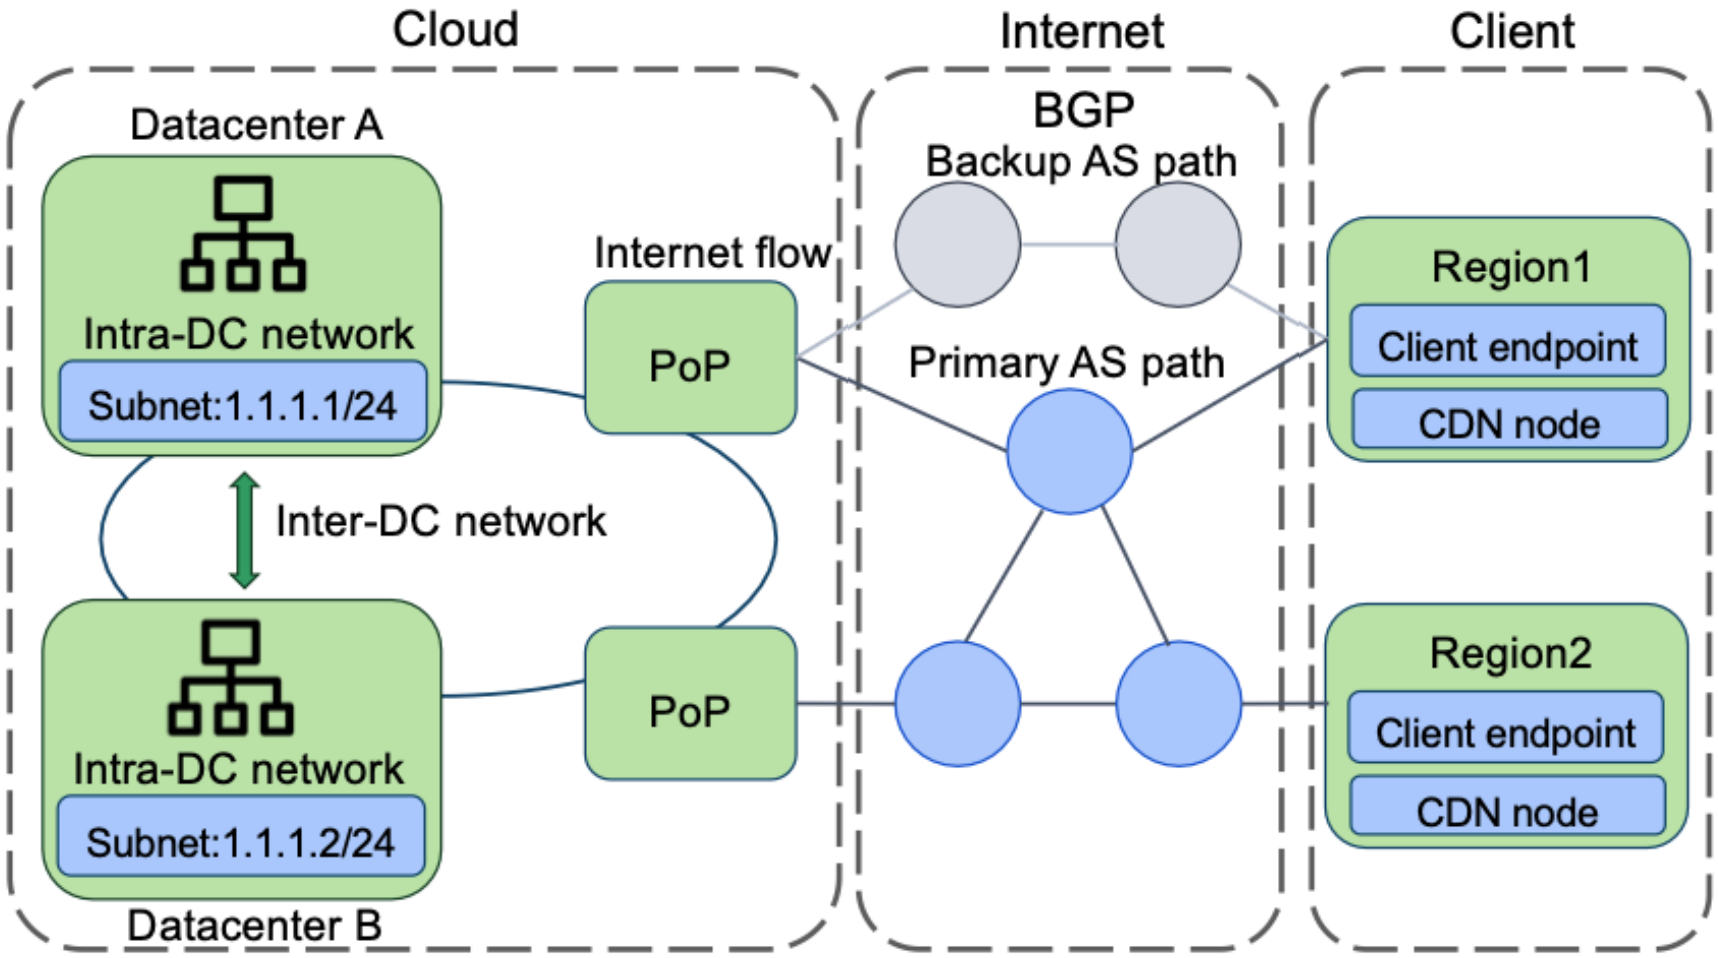

Why It Matters

"Internet faults are a leading cause of cloud outages!"

63%

of all network incidents happen in the Internet

78%

of severe (P0/P1) events are Internet-related

2,000+

incidents detected by Skyline in 2025

Challenges of Internet monitoring

1. **No Visibility:** Cloud providers lack direct ownership, no switch logs or detailed ISP signals.

2. **Vast Scale:** The Internet spans ~200K ASes with constant topology changes.

3. **Diverse Faults:** Incidents range from fiber cuts and BGP hijacks to subtle routing policy errors.

Our question

"How can a cloud provider achieve high-coverage Internet monitoring with limited visibility and control of Internet?"

Key Insight: 3 Coverage Dimensions

1. Traffic Direction

Cover both Cloud → Client (Outbound) and Client → Cloud (Inbound).

2. Incident Lifecycle

Monitor both the incident detection phase and the repair phase.

3. Tenant Granularity

Move beyond simple reachability to L4-L7 per-tenant metrics.

Key Insight: 3 Coverage Dimensions

1. Traffic Direction

Cover both Cloud → Client (Outbound) and Client → Cloud (Inbound).

2. Incident Lifecycle

Monitor both the incident detection phase and the repair phase.

3. Tenant Granularity

Move beyond simple reachability to L4-L7 per-tenant metrics.

Dimension 1: Traffic Direction

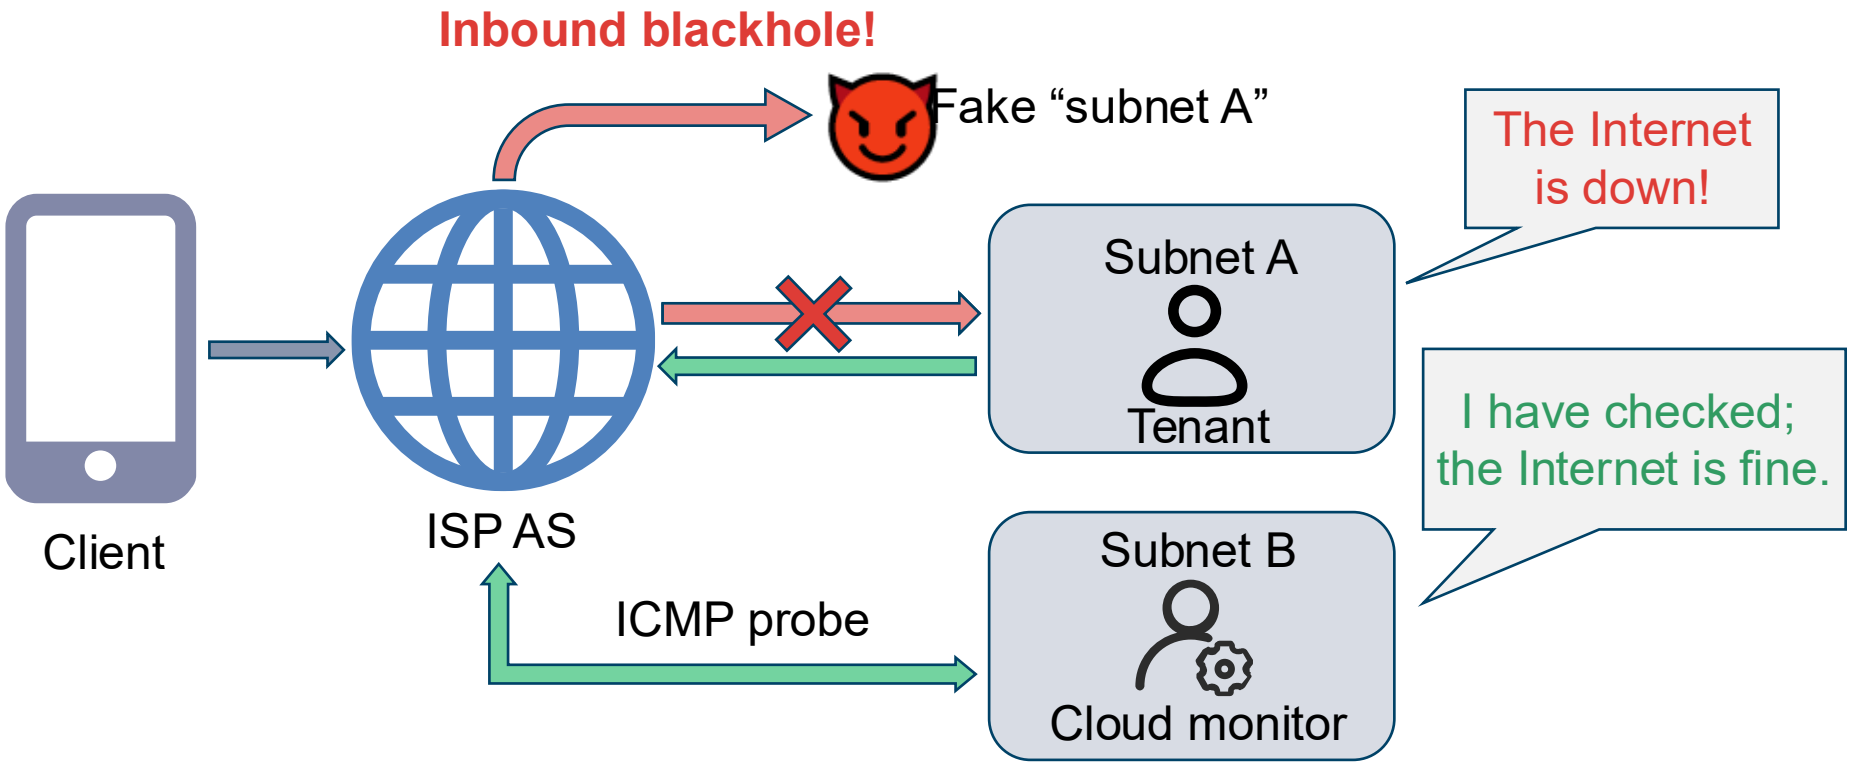

Consequence: Missing inbound (client-to-cloud) direction incidents.

Subsystem 1: BiProbe

Challenge: Cover cloud and client efficiently and non-intrusively?

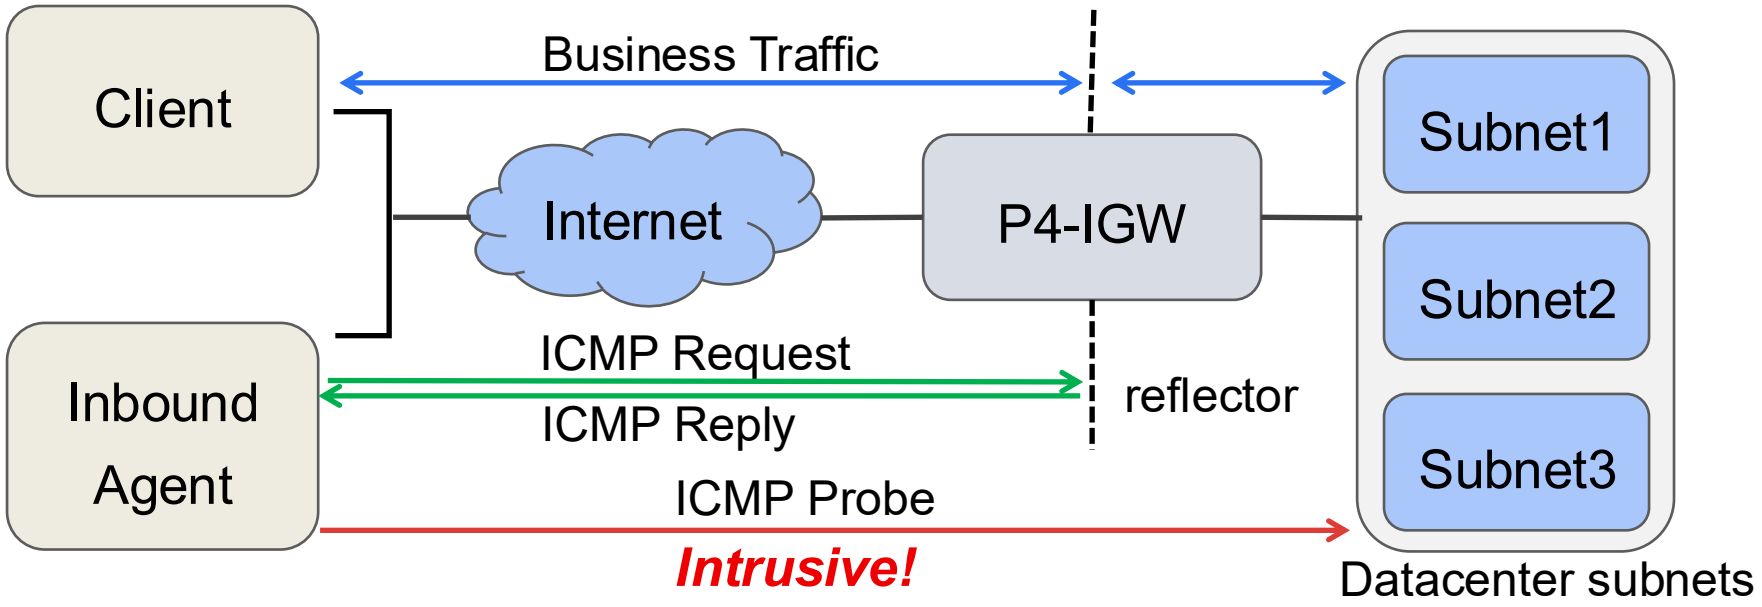

src	dst	IGW	Explanation
inbound agent	subnet1	✅	P4 sends ICMP reply back to inbound agents
inbound agent	subnet2	❌	Inbound probe cannot reach IGW and subnet2
client	subnet1	✅	Regular traffic is forwarded to the tenants

Key Insight: 3 Coverage Dimensions

1. Traffic Direction

Cover both Cloud → Client (Outbound) and Client → Cloud (Inbound).

2. Incident Lifecycle

Monitor both the incident detection phase and the repair phase.

3. Tenant Granularity

Move beyond simple reachability to L4-L7 per-tenant metrics.

Dimension 2: Incident Lifecycle

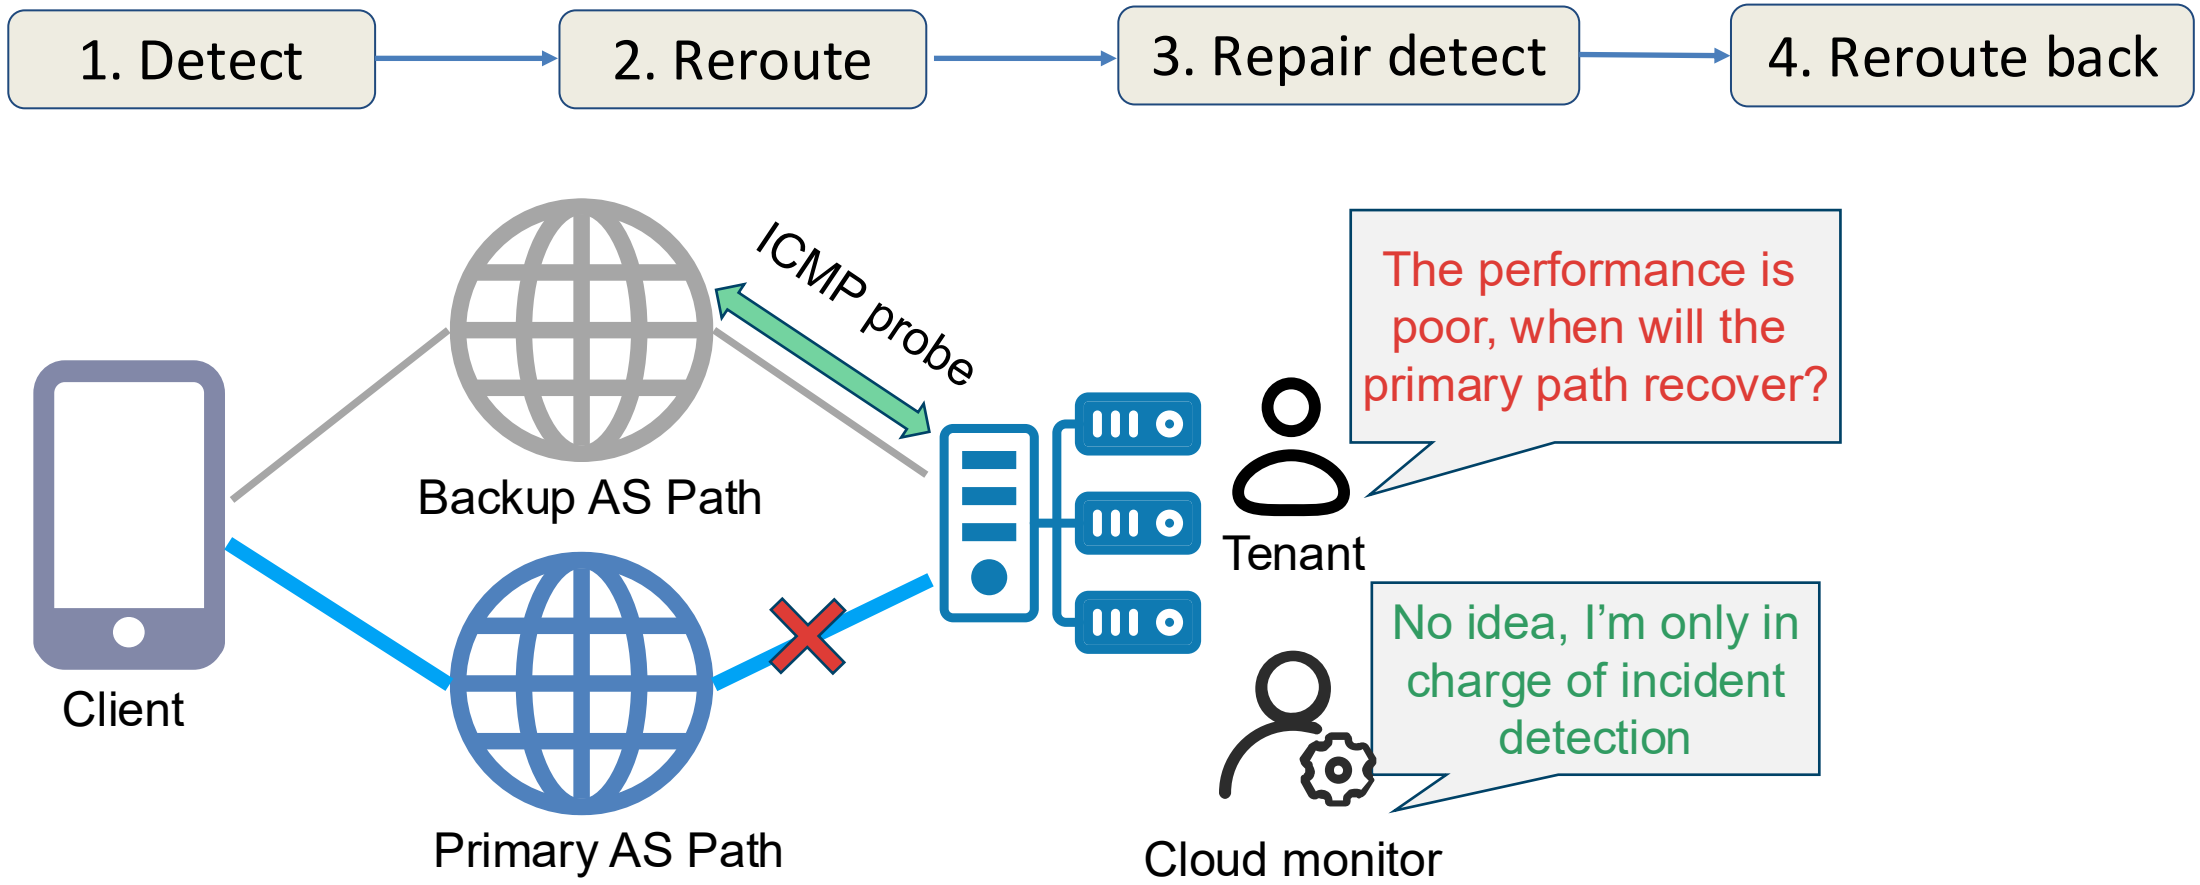

Consequence: Delayed recovery and continued use of expensive/slower backup paths.

Subsystem 2: FlexPath

Challenge: Monitor faulty path after rerouted (path steering)?

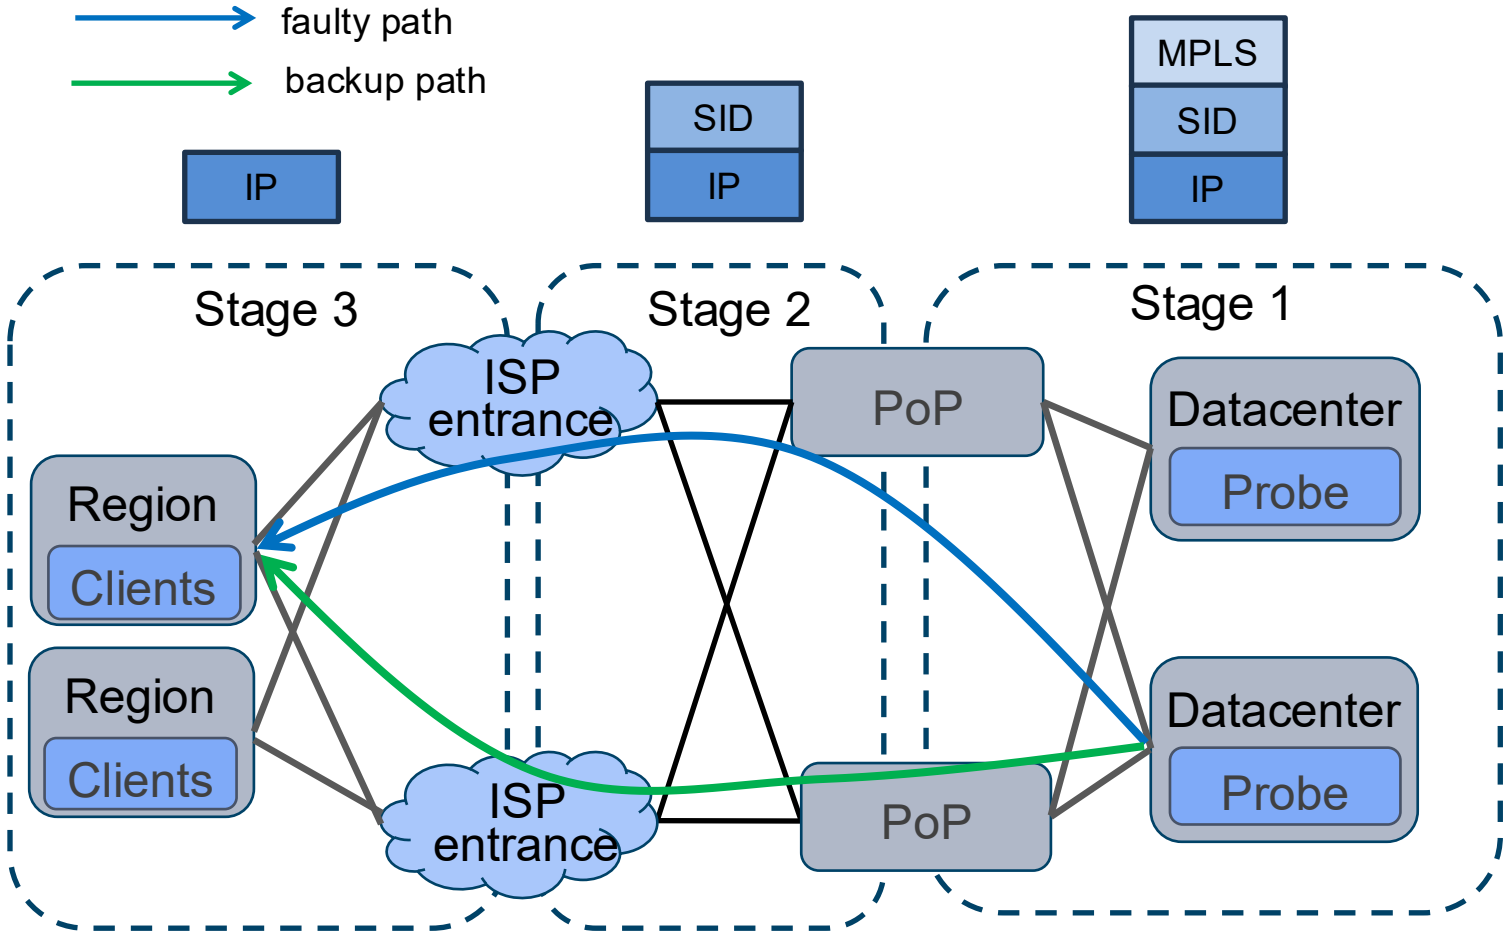

Lifecycle management: identify incident repair (RTT, Drop rate, etc)

Key Insight: 3 Coverage Dimensions

1. Traffic Direction

Cover both Cloud → Client (Outbound) and Client → Cloud (Inbound).

2. Incident Lifecycle

Monitor both the incident detection phase and the repair phase.

3. Tenant Granularity

Move beyond simple reachability to L4-L7 per-tenant metrics.

Dimension 3: Tenant Granularity

Per-tenant upper-layer metrics

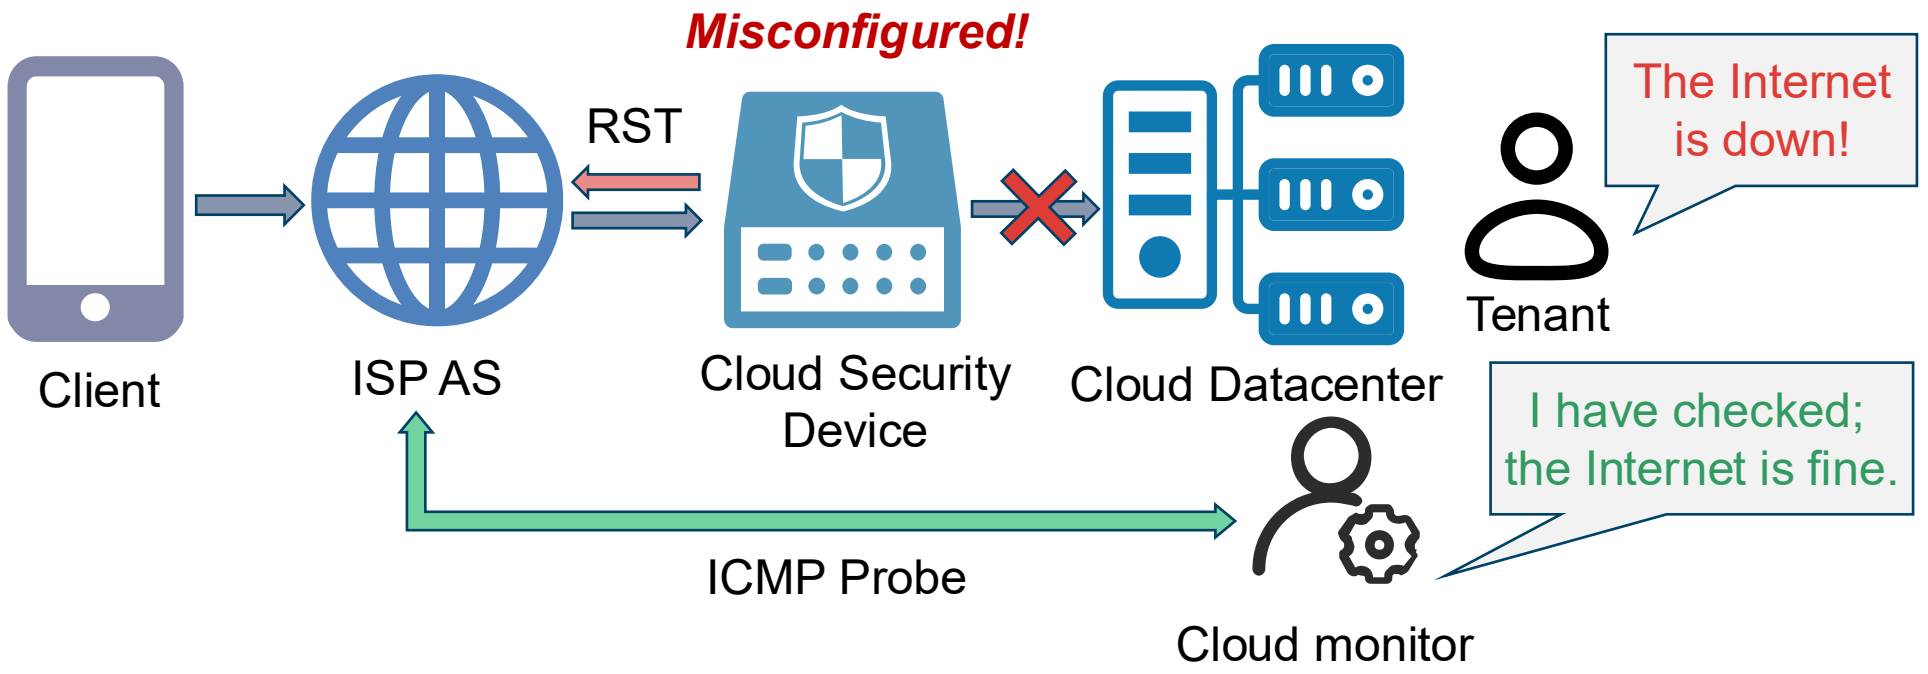

Consequence: Missing per-tenant upper-layer incidents and attestation.

Subsystem 3: FlowHunter

Challenge: Capture upper-layer per-tenant info non-intrusive

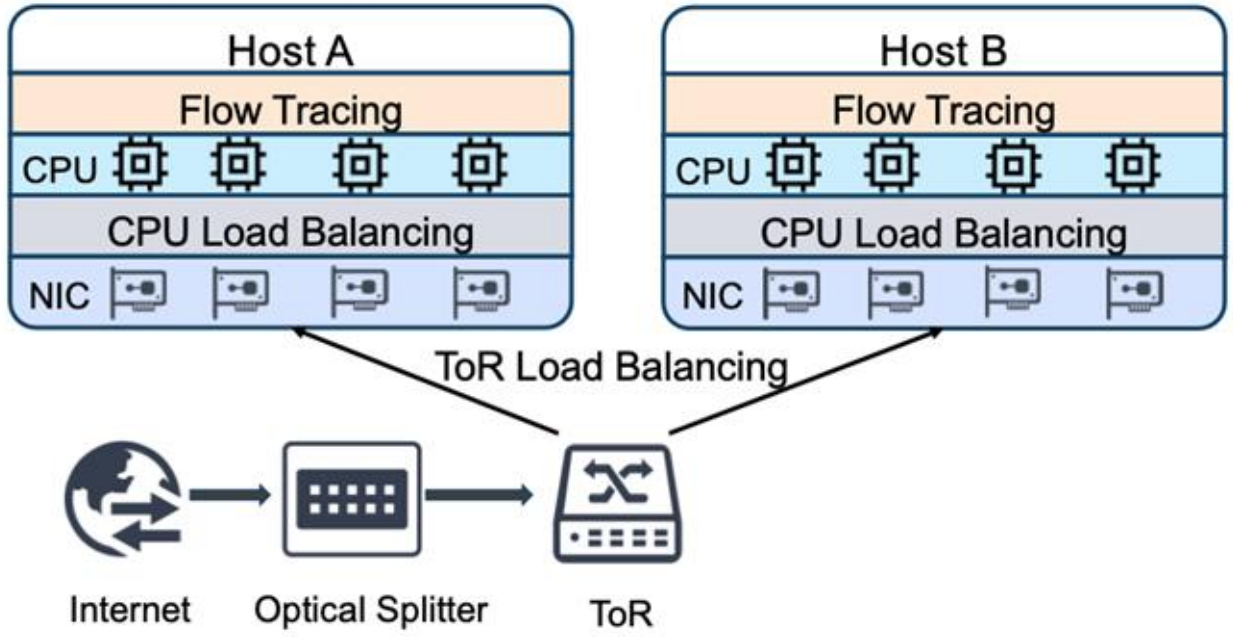

Physical Duplication

Uses **Optical Splitters** to physically duplicate light signals from fiber.

- 100% Passive.
- Zero impact on business traffic.
- Send copy to analysis cluster (DPDK)

Skyline Architecture

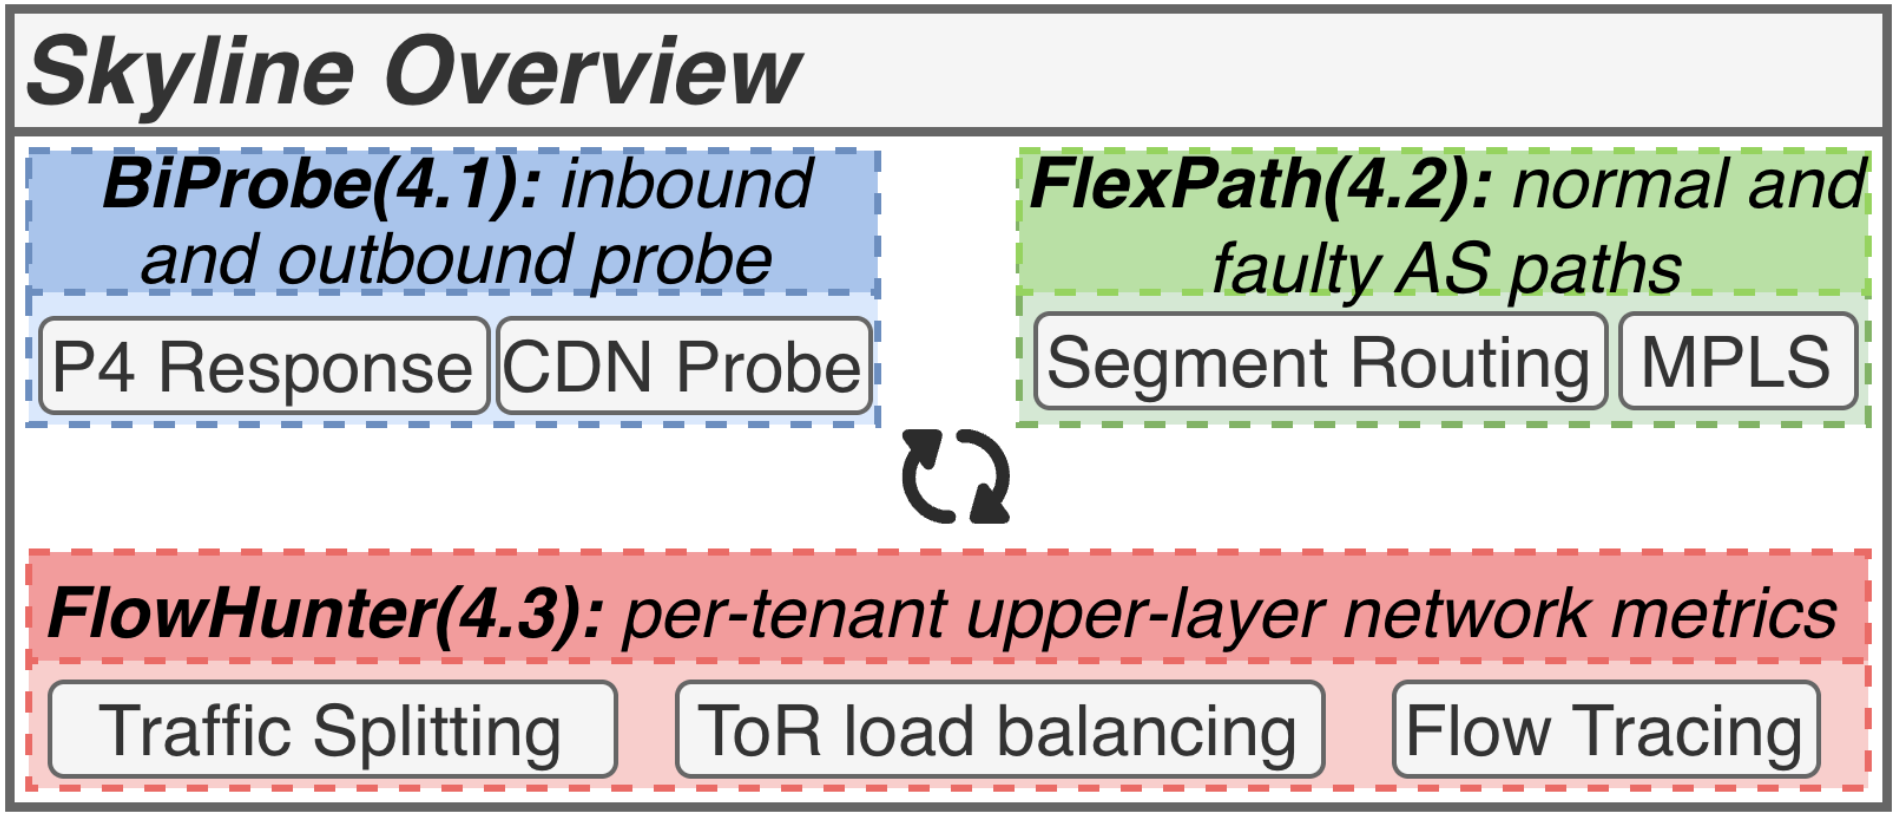

Operational Lessons

High Precision & Recall

Recall: Near-perfect (>99%) for verified incidents.

Precision: ~97% (Low false positives).

Scale: Evaluated over 1 year of data.

Detected 2000+ major issues in 2025 (73% ISP-side, 27% Cloud-side).

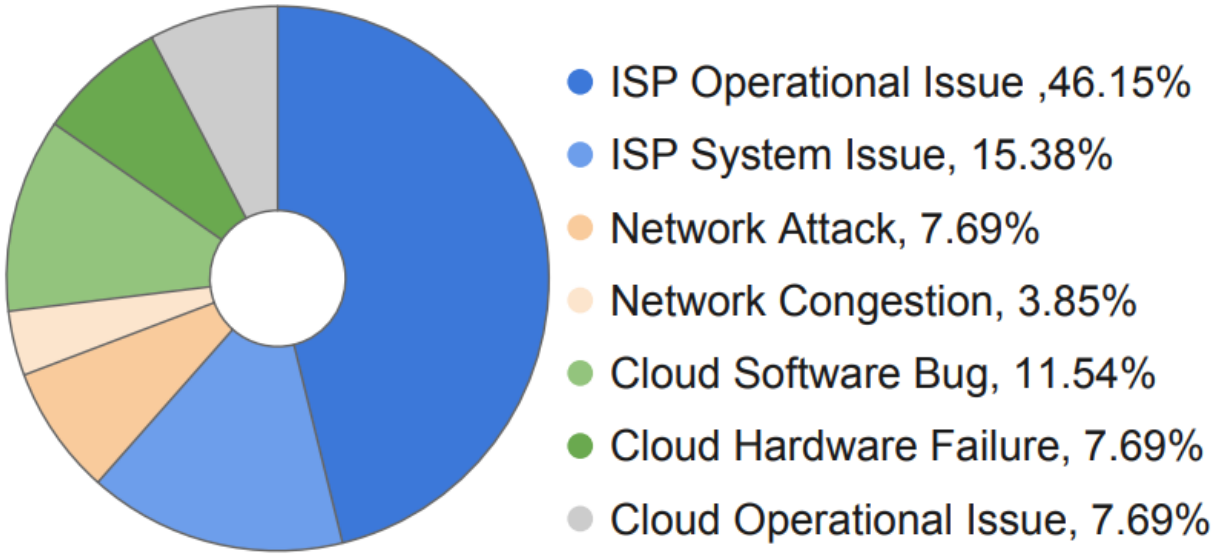

Operational Lessons

Deployment stages

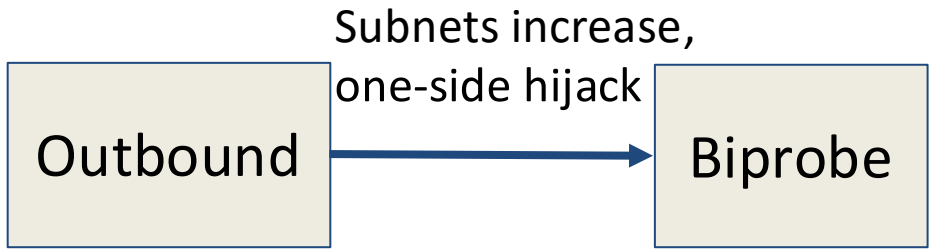

Operational Lessons: Biprobe

Case 1: ISP configuration changes introduced various routing issues.

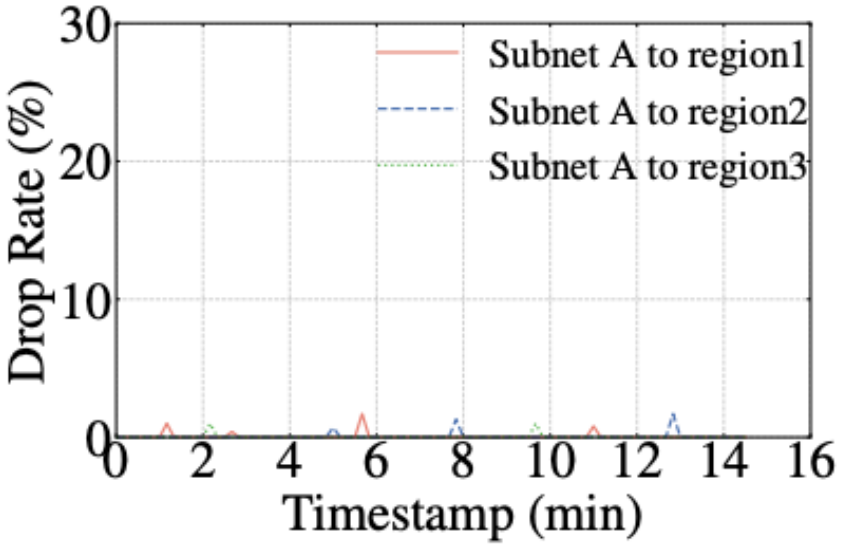

outbound

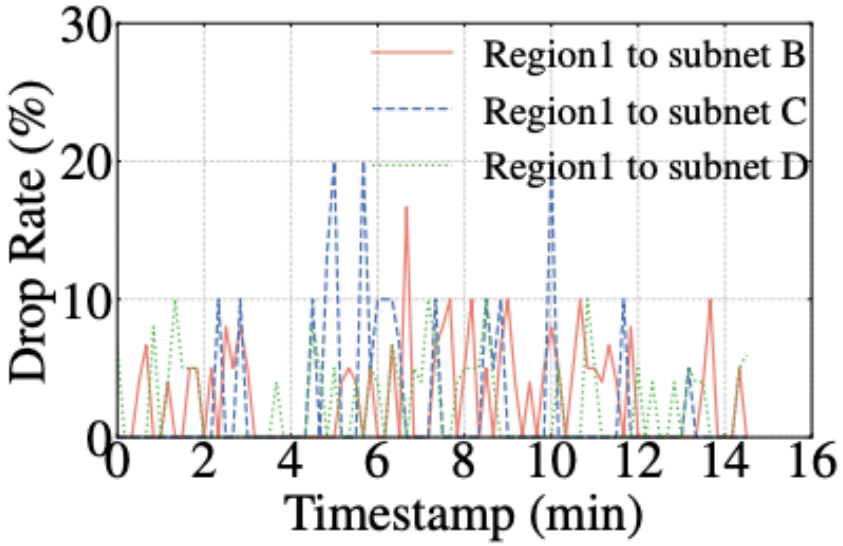

inbound

Operational Lessons

Deployment stages

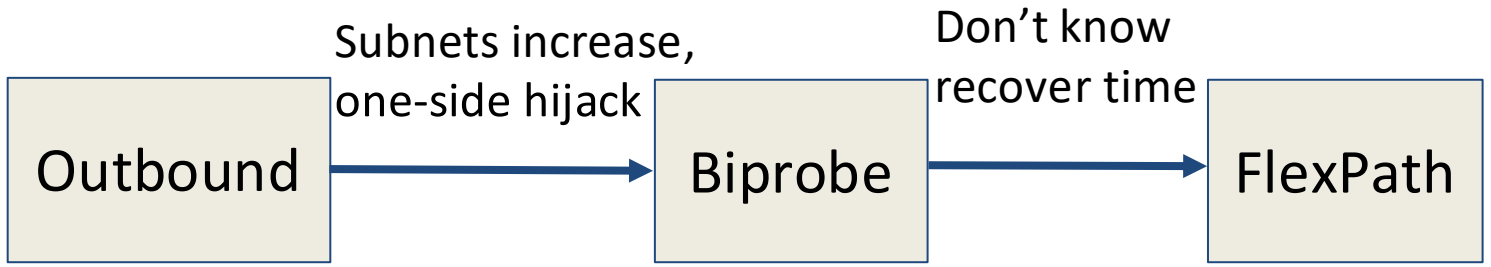

Operational Lessons: FlexPath

Case 2: Internet attack induced large-scale packet loss.

- 15min, detect incident
- 20min, Switch to backup path
- 31min, detect incident repair
- 41min, switch to primary path

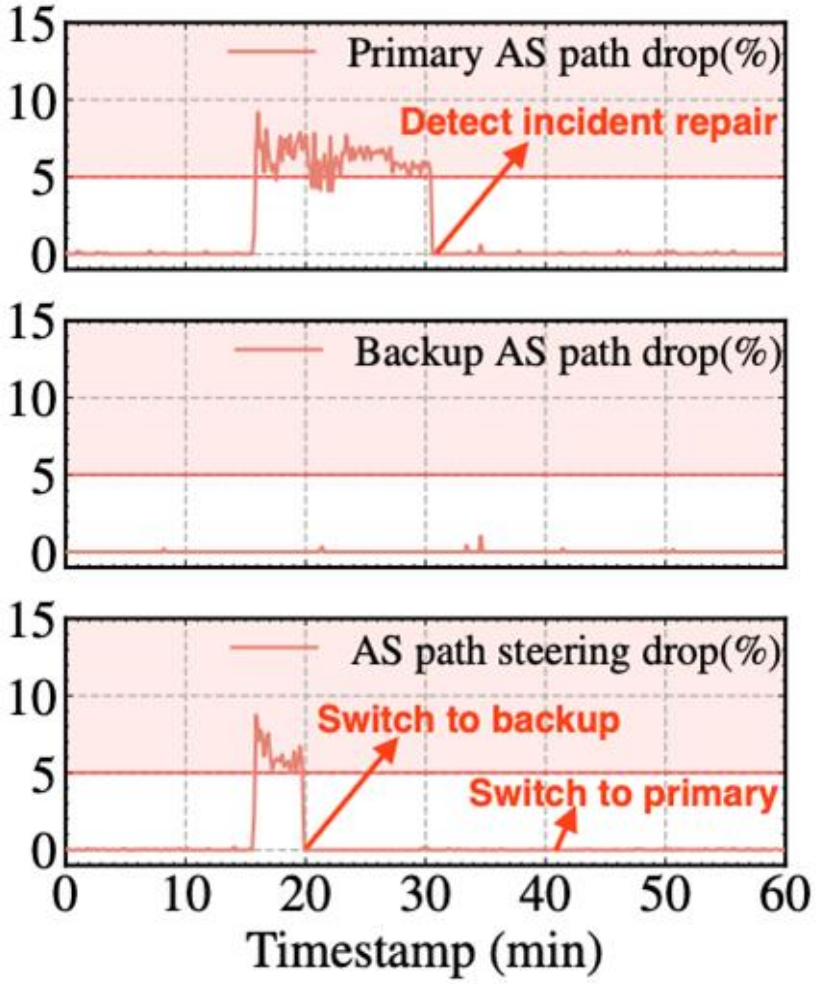

Operational Lessons

Deployment stages

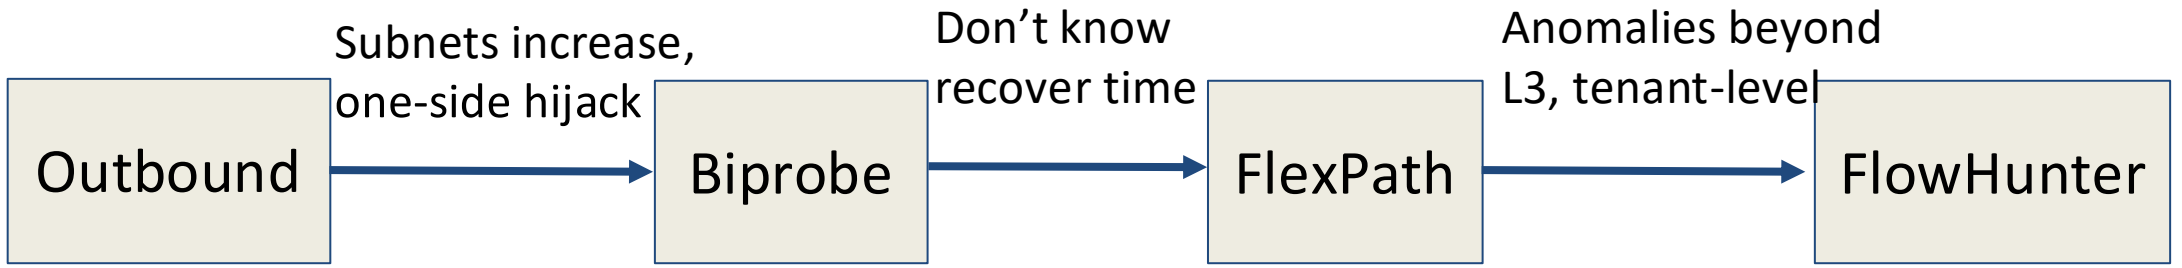

Operational Lessons: FlowHunter

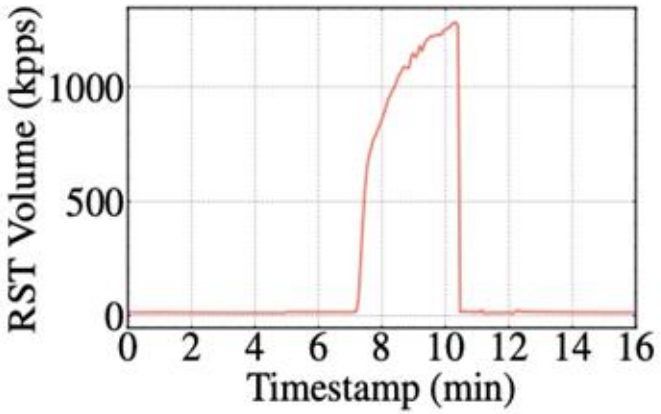

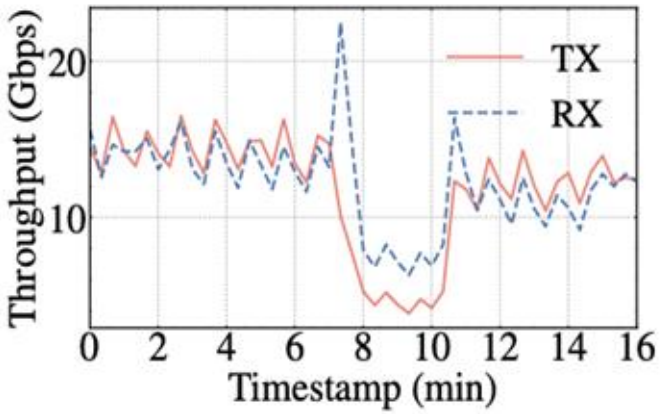

Anomalous RST ***Bussiness traffic tput***

Case3: misconfiguration of security devices

- 7min, FlowHunter detected anomalous RST and TX/RX degrade

Operational Lessons

BiProbe & FlowHunter

	BiProbe	FlowHunter
Mode	Active probing	Passive monitoring
Layer	L3 (ICMP)	L4–L7 (TCP/UDP, HTTP)
Focus	Coarse-grained Internet reachability	Fine-grained per-tenant traffic metrics
Blind spot	L4+ symptoms that preserve ICMP reachability	Incidents where traffic never reaches the PoP

Operational Lessons

Critical Regions of high-priority monitoring:

- High-traffic, business-critical, and latency-sensitive regions
- PoP depending on a single ISP
- ISPs with a frequent failure history
- ISPs with strict authentication mechanisms

Skyline: A Cloud-Centric Internet Monitoring Engine

- Internet is important but the cloud has no control
- We decomposed the Internet monitoring into 3 dimensions
    1. Traffic direction: Biprobe
    2. Incident lifecycle: FlexPath
    3. Tenant granularity: FlowHunter
- We have deployed Skyline over 2 years and proved it is effective at identifying large amount of Internet incidents with high accuracy and low overhead.