IP Measurements (IPM)

SONiC workshop - OCP Global Summit 2024

Clarence Filsfils, Pablo Camarillo, **Ahmed Abdelsalam**, Carmine Scarpitta

Cisco Systems

IP is back and better than ever.

SRv6 uSID

Build anything

Simplified, scalable, and versatile networks that are self-sufficient

Self-sufficiency is standard

End-to-end policy

- From Host to Internet through DC, Access, Metro, Core, Cloud
- No protocol conversion or gateways at domain boundaries

Any service, without any shim

- VPN, Slicing, Traffic Engineering, Green Routing, FRR, NFV

Better scale, reliability, cost, and seamless deployment in Brownfield

Essential embedded assurance

Active probing between Fabric Edges **along all ECMP paths**

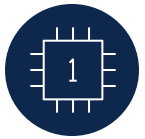

High-capacity probe generation and ingestion powered by Silicon One (14MPPS)

Continuous **routing monitoring**

Advanced analytics and intelligent service optimization driven by AI

Measure everything

IPM

Embedded SLA monitoring and IPM within the network is essential

Problem and Solution Intuition

Network Assurance

- High-investment, low detection
  - Significant investment on external probing appliances (SPs) or complex host-based applications (DCs)
- Blind spots in network monitoring prevent detection of issues
  - Reactive, Not Proactive, Issue Resolution
- Why? The nature of IP is ECMP. The nature of Probing isn’t ECMP.

The nature of IP is ECMP

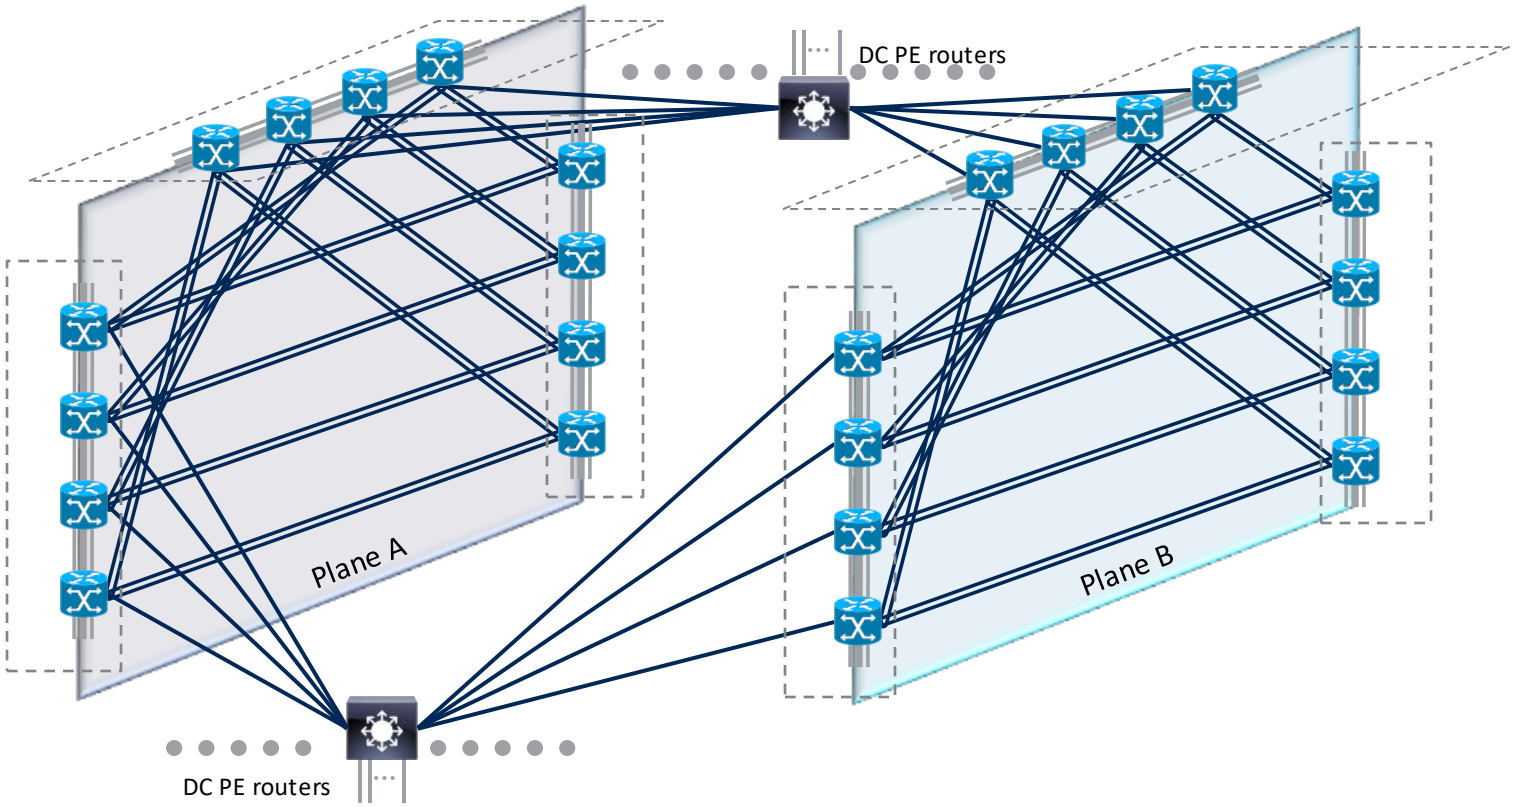

- Simple topologies already exhibit 16 ECMP paths
- The number of ECMP paths grows significantly with larger topologies

Probing neglects the ECMP reality

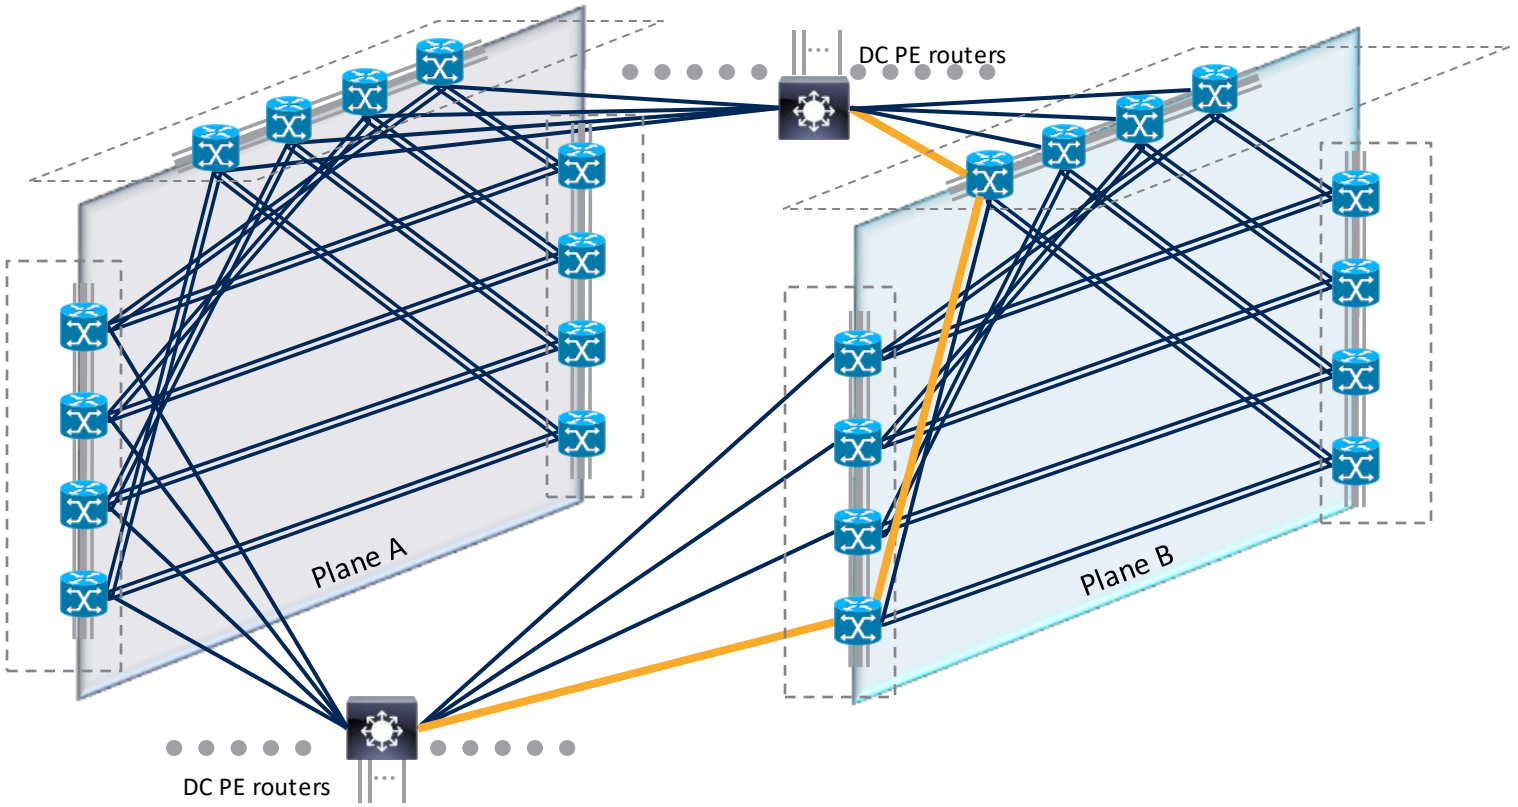

- Legacy solutions do not have the scale to measure all ECMP paths
- Simple probe from Source to Destination.
  - A single 5-tuple. Whatever hashing may be.

A measurement solution MUST:

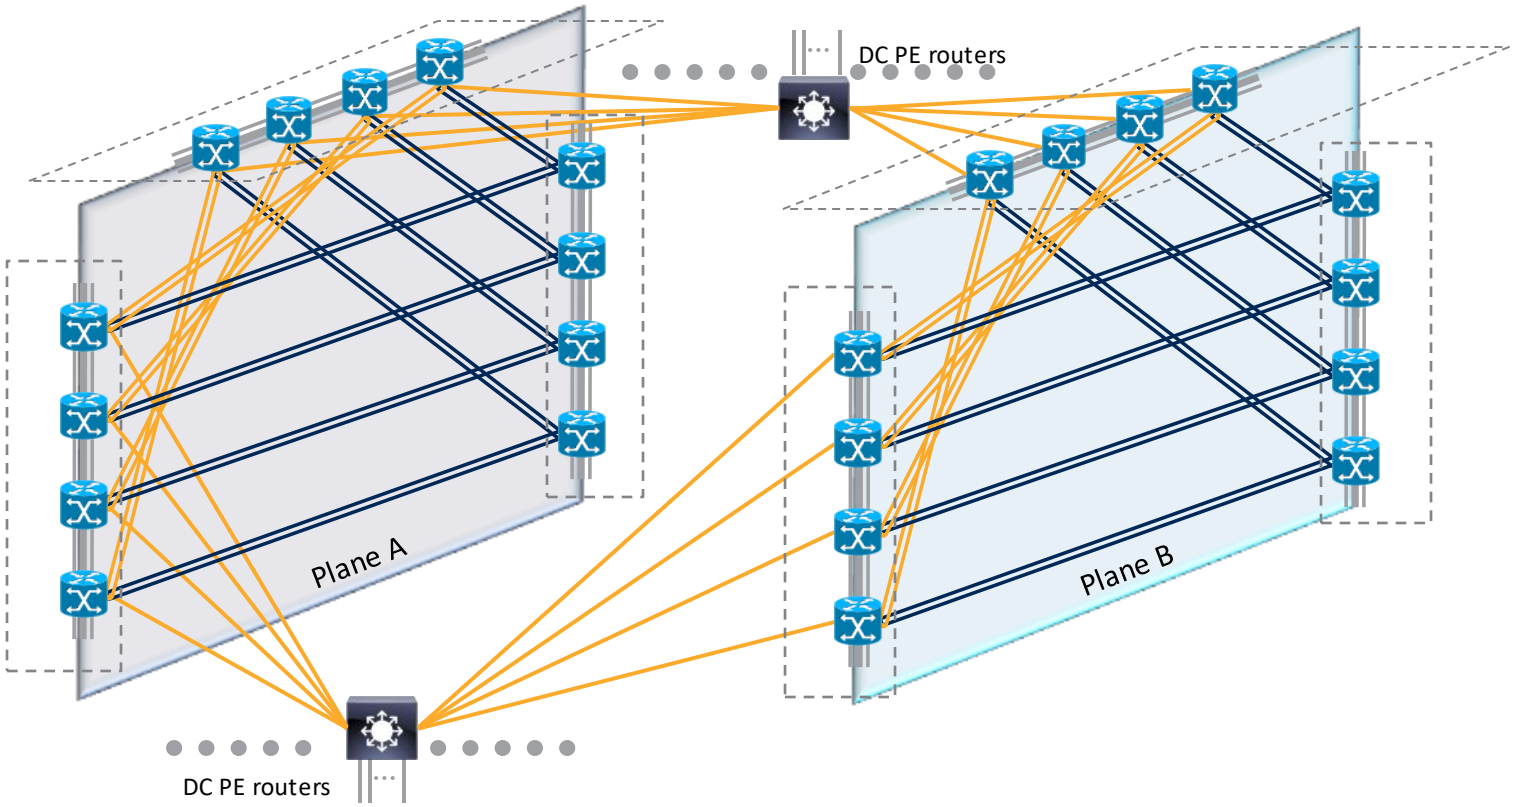

- Discover and monitor all ECMP paths:
- Provide enough PPS to measure all ECMP paths
- Report accurately across ECMP paths

The experience of all clients must be measured

Would a bank accept to monitor < 0.1% of its access?

Legacy solutions are typically 1000 to 10000 times not scalable enough

Legacy coverage is < 0.1%

Operators learn outages from clients

Silicon One provides 14M probes per sec

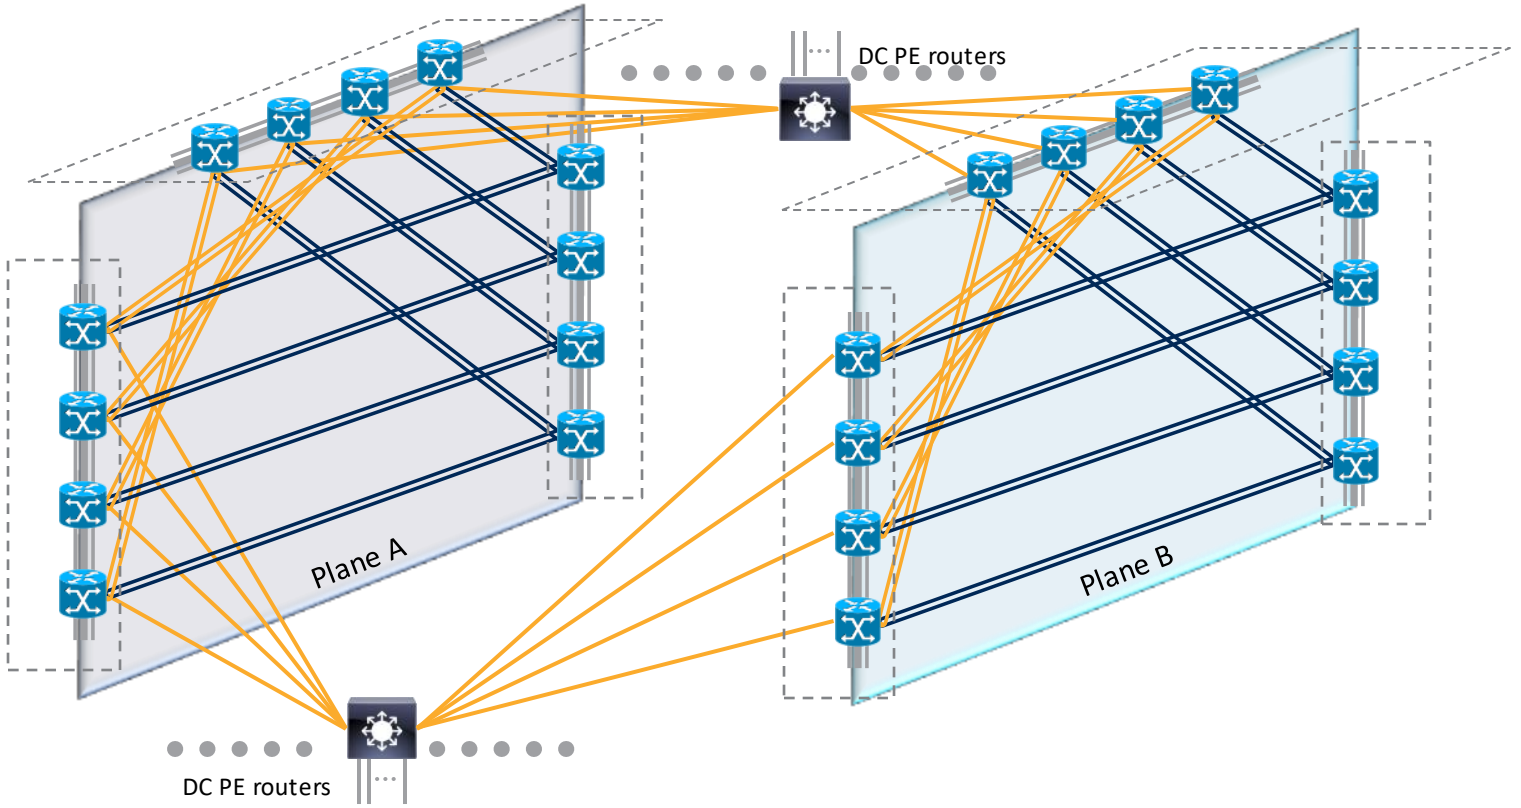

- 1 measurement every msec
- 500 edges
- 16 ECMP paths

8M probes per sec (57% of Silicon One capability)

Accurate Metrics

- 1 bad path out of 8 ECMP
- 12.5% of the clients impacted
- Average hides the issue
- IPM Histograms reports the experience of the whole population

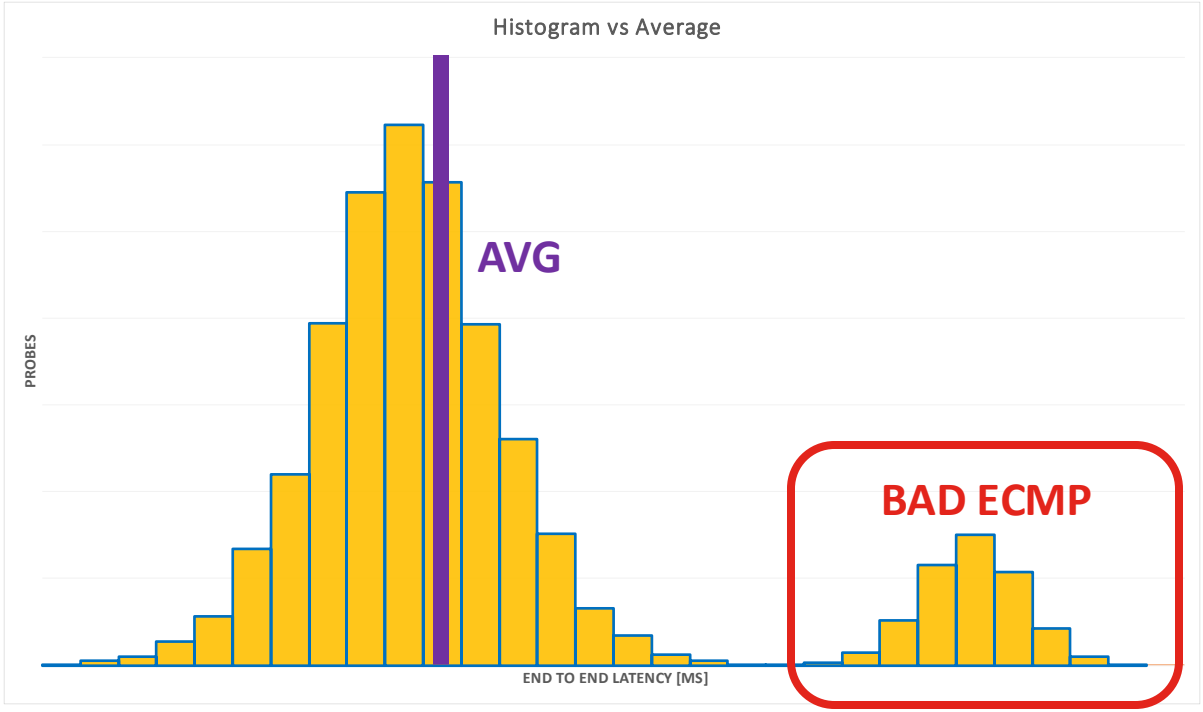

Solution: Accurate and Rich Metrics

- Each probe measures latency, loss and liveness
- One-way measurement instead of 2-way/2
- Per ECMP-path measurement
- Latency histogram instead of min, avg, max
- Absolute loss instead of loss approximations
- Liveness detection

Much Cheaper through Silicon Integration

- Capex Elimination
  - SLA Appliance
  - Router port to appliance
- Opex Elimination
  - Rack Space
  - Power
  - Onsite installation

Standard Based Measurement

- STAMP – RFC8762/RFC8972
- Packet Format:
  - Outer Encapsulating header:
    - > Any IP Encapsulation
    - > SRv6 uSID/GRE/VxLAN/…
  - STAMP measurement packet:
    - > Alternate Marking bit as part of Flow Label
- We monitor the shared transport **AND** the service forwarding path on PEs

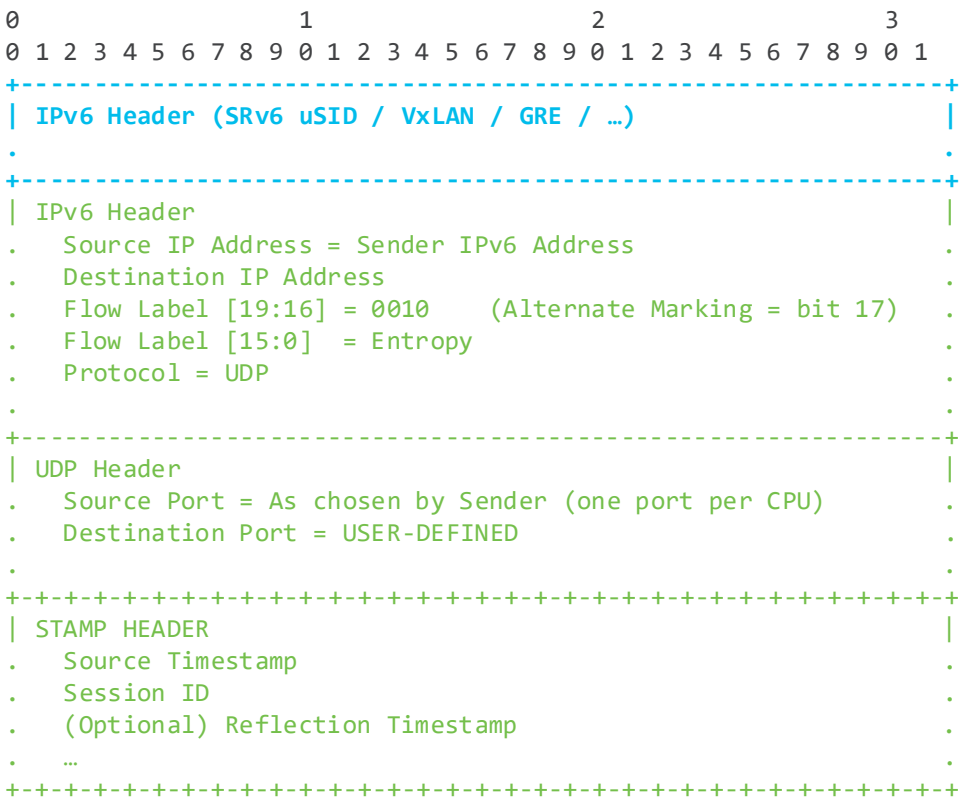

IPM in SONiC

- Supported across the Cisco Silicon One portfolio
- Full stack support in SONiC
  - SAI header
  - S1 SAI Adapter
  - SONiC
- Demo available


CISCO