



IP Measurements

Ahmed Abdelsalam,

Engineering Technical Leader, Cisco Systems

IP is back and better than ever.

SRv6
uSID


Build
anything

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

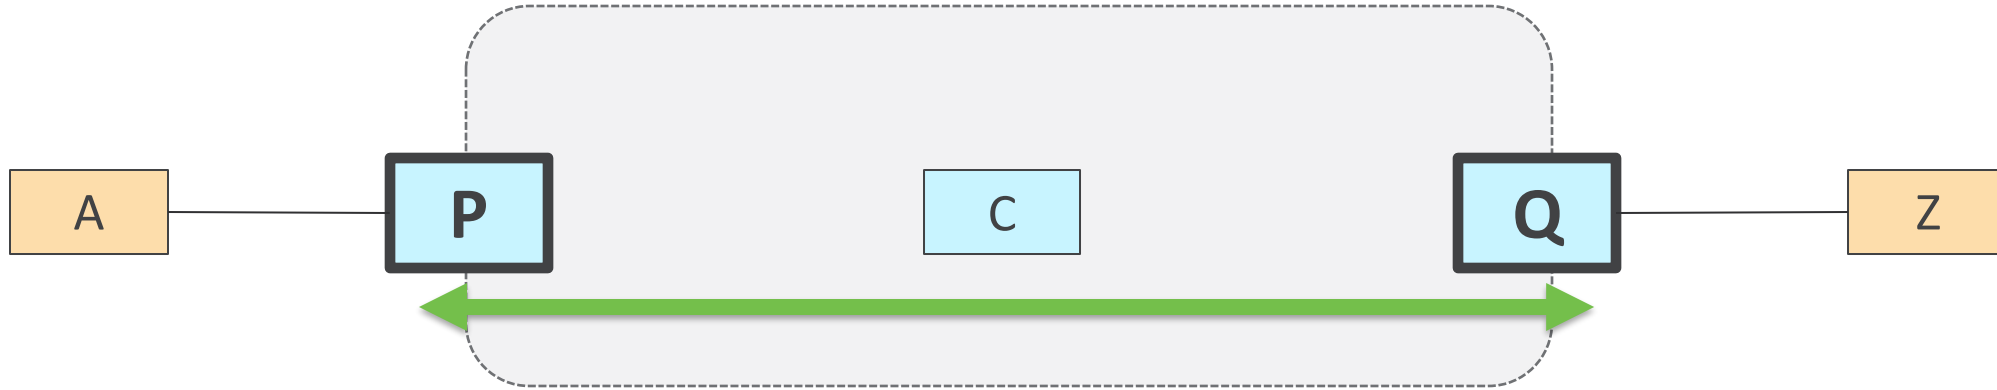
IPM

Embedded SLA
monitoring and IPM
within the network is
essential



IPM

Measurement from Any Edge to Any Edge, across ECMPs



- **Active probing** from any P to any Q via, any ECMP path
- Continuous **routing monitoring**
- **Analytics**
 - Correlation of probe measurement and routing data

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.

Problem statement


- Performance measurement solutions:
 - Lack of scale to measure all ECMP Paths
 - min/max/average doesn't reflect the experience of each user
 - CAPEX and OPEX
- The result:
 - They can't detect the network issues before the customer is impacted
 - Users/Customers detect the network problem then operators troubleshoot.
 - Network Operators repeatedly encountering this issue over time.

The nature of IP is ECMP



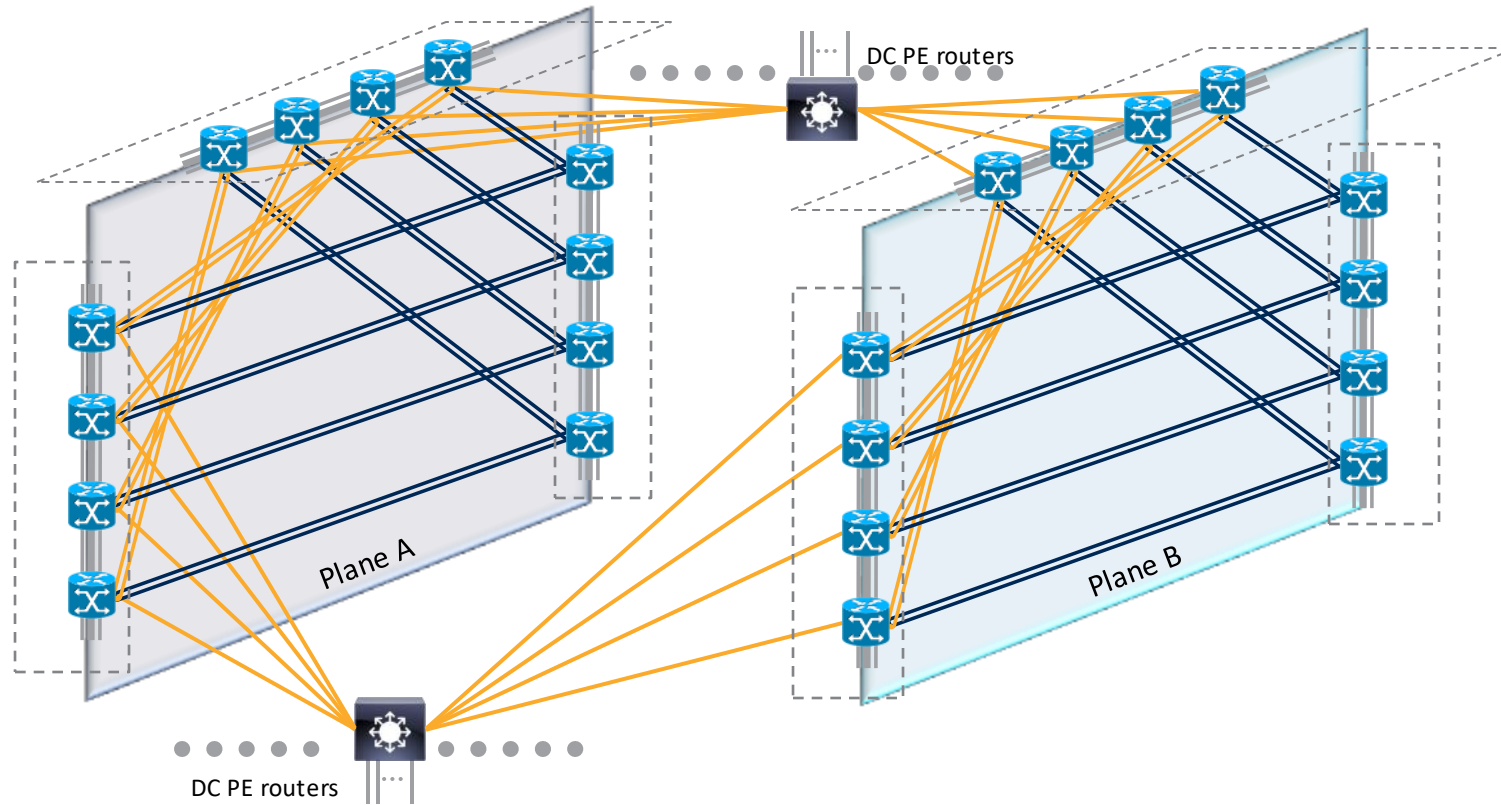
- Simple topologies already exhibit 16 ECMP paths
- The number of ECMP paths grows significantly with larger topologies

Probing neglects the ECMP reality



- Legacy solutions do not have the scale to measure all ECMP paths
- Basic 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



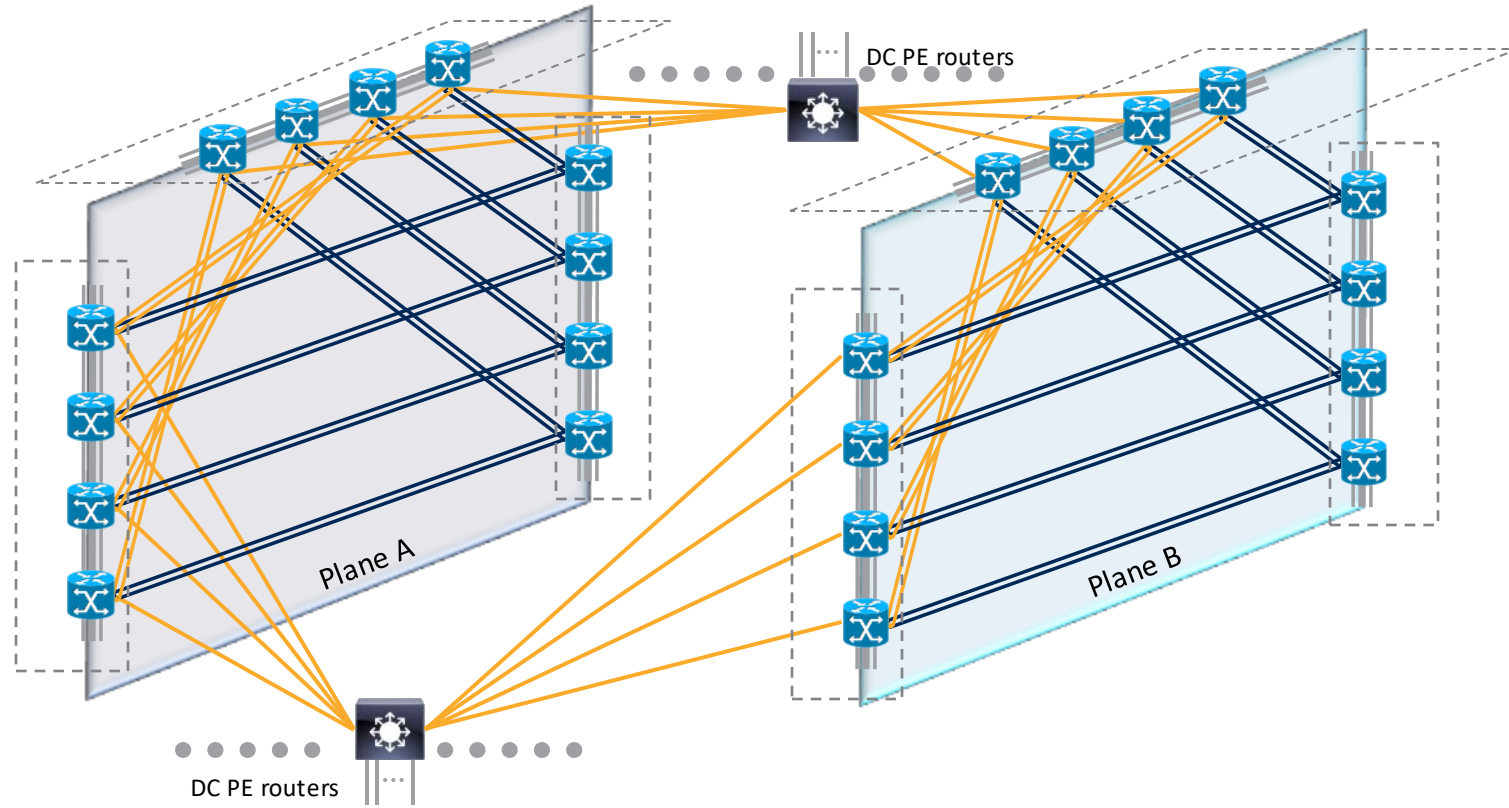
Would a bank accept to monitor $< 0.1\%$ of its access?

Legacy solutions are typically 1000 to 10000 times not scalable enough


Legacy coverage is $< 0.1\%$

Operators learn outages from clients

Probing Scale

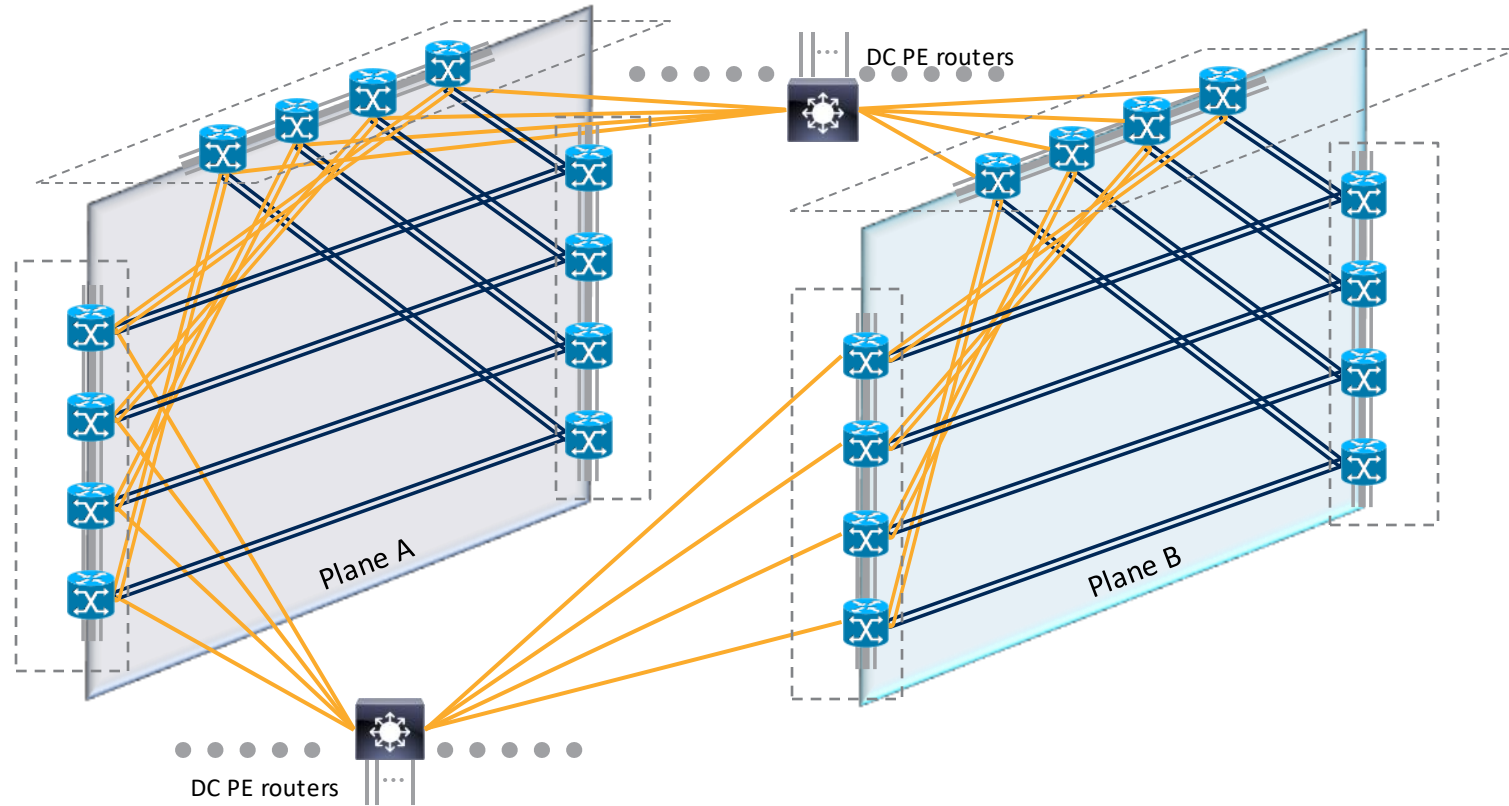


Probing Scale



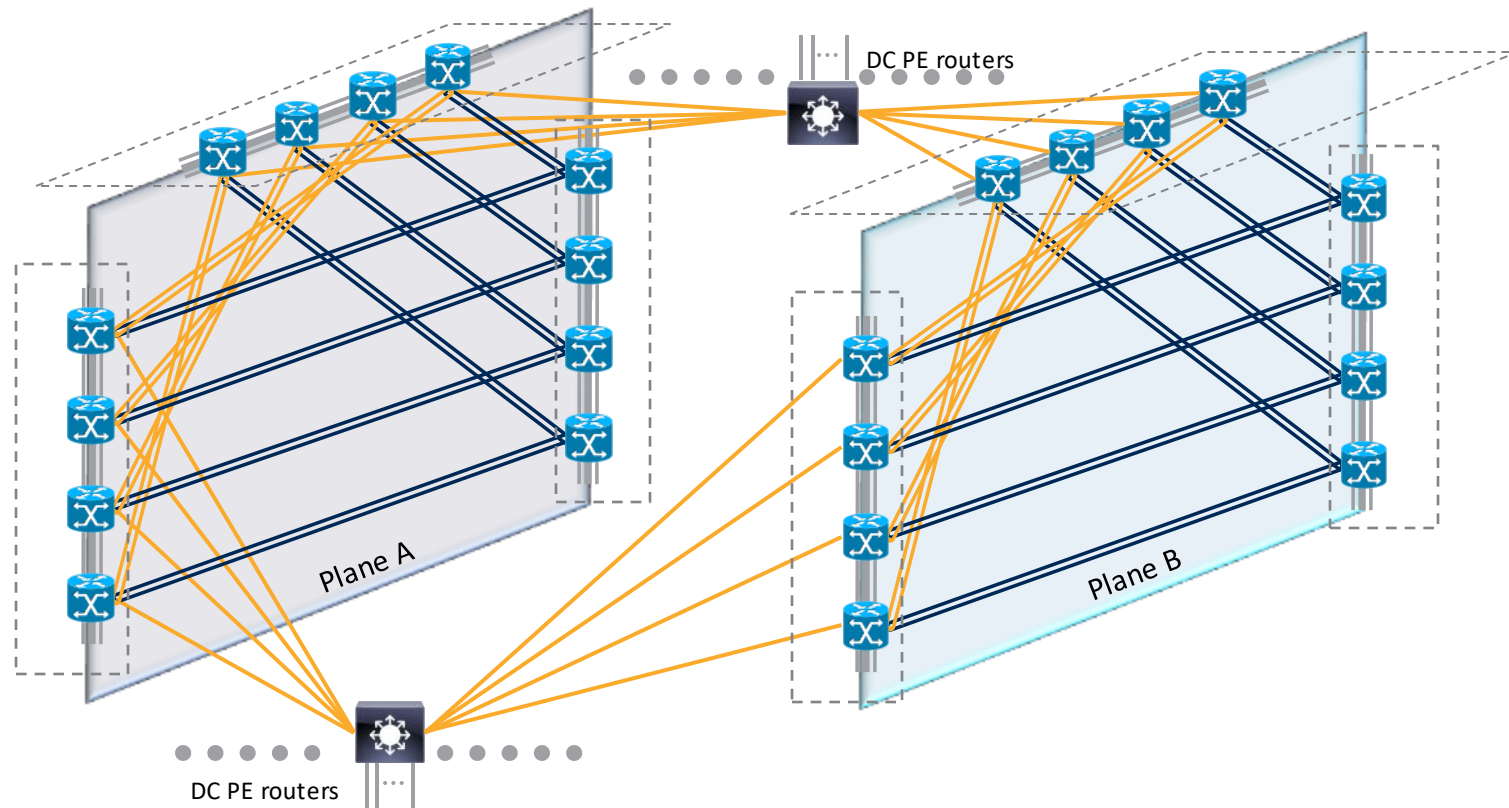
- 1 measurement every msec

Probing Scale




- 1 measurement every msec
- 500 edges

Probing Scale



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

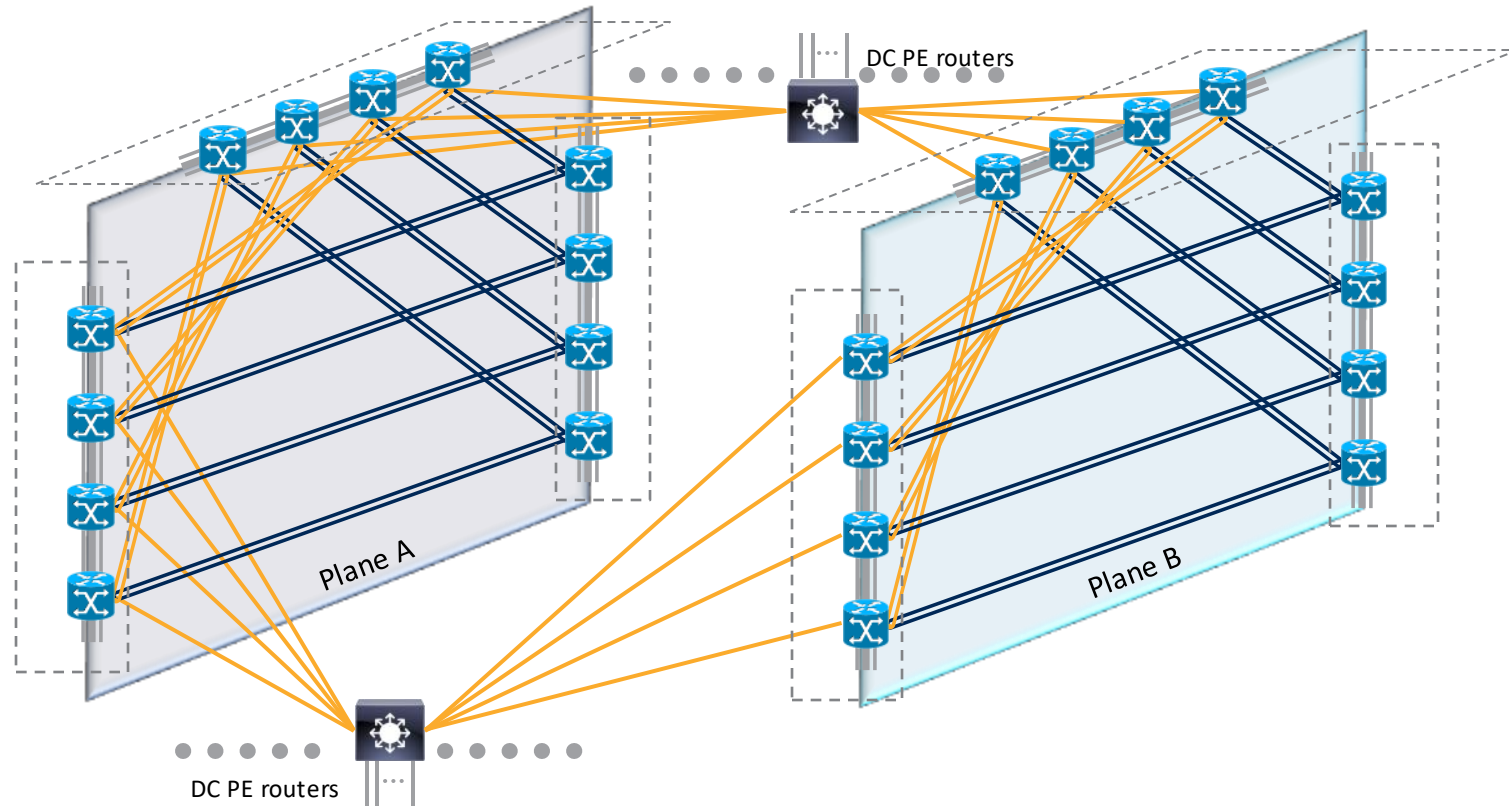
Probing Scale



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

8M probes per sec

Probing Scale




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

8M probes per sec

Legacy solutions are typically 1000 to 10000 times not scalable enough

Legacy coverage is < 0.1%

Probing Scale



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

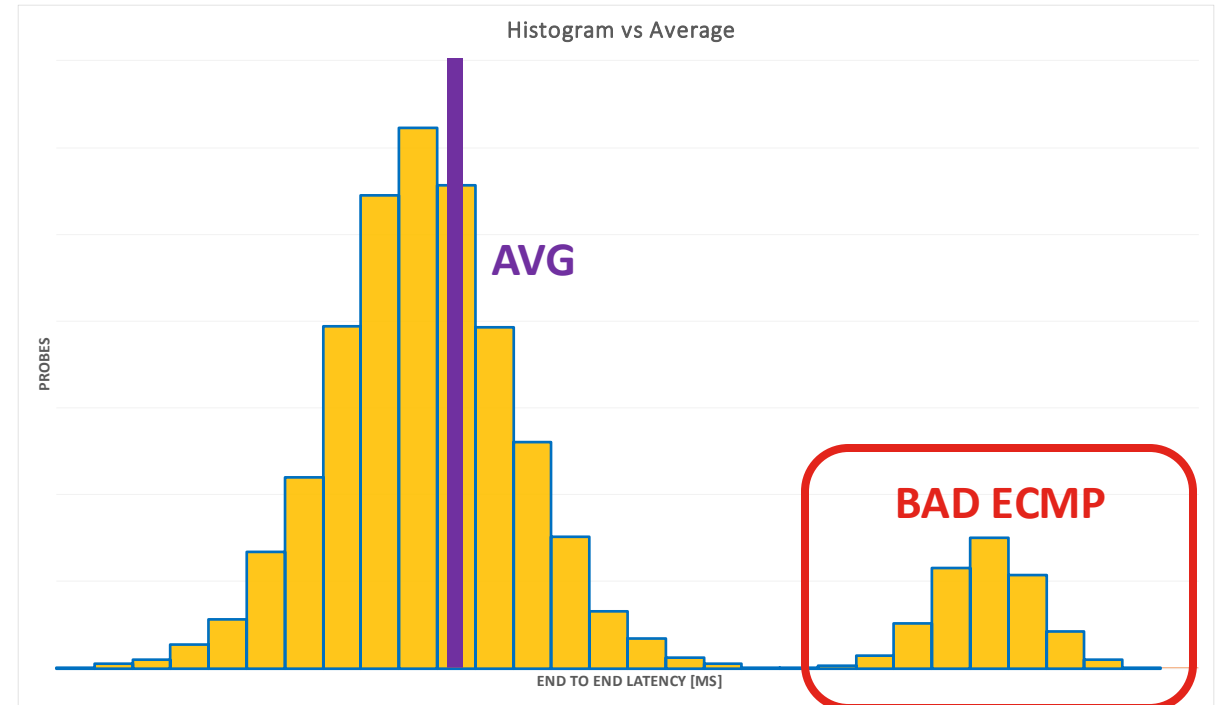
8M probes per sec

Cisco Silicon One provides 14M PPS
Both generation and processing

8 MPPS = 57% of Silicon One capability

Richer Metrics

- 1 bad path out of 16 ECMP
- 6.25% of the clients impacted
- Average hides the issue
- Latency Histograms reports the experience of the whole population



BAD ECMP

Solution: Accurate and Rich Metrics

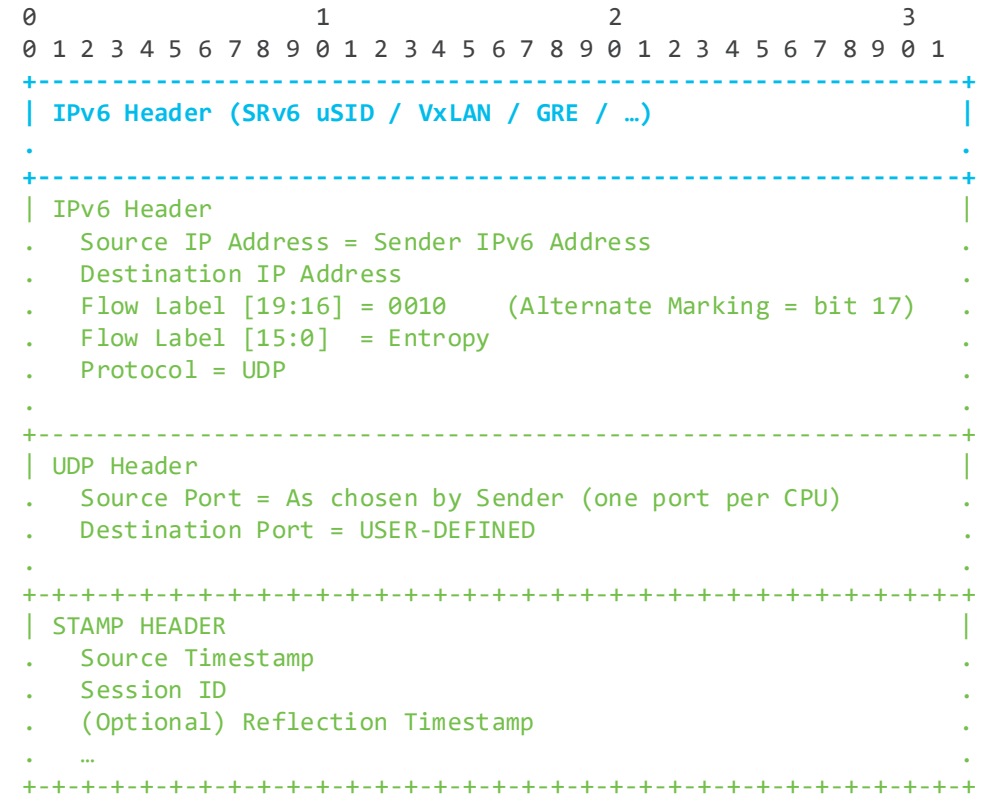
- Each probe measures **latency, loss and liveness**
- **One-way measurement** instead of 2-way/2
- **Per ECMP-path** measurement (leveraging SRv6 uSID)
- **Latency histogram** instead of min, avg, max
- **Absolute loss** instead of loss approximations
 - Alternate marking (RFC9341)
- **Liveness** detection

Much Cheaper through Silicon Integration

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

Standard Based Measurement

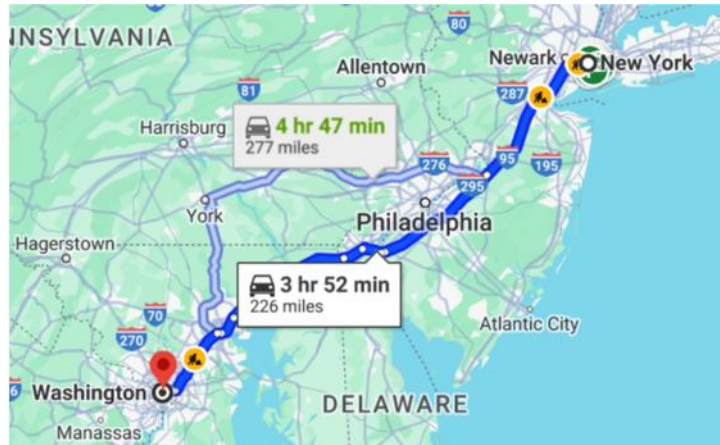
- STAMP – RFC8762/RFC8972
- Packet Format:
 - Outer Encapsulating header:
 - > Any IP Encapsulation
 - > SRv6 uSID/....
 - STAMP measurement packet:
 - > Alternate Marking bit as part of Flow Label
- We monitor the shared transport **AND** the service forwarding path on PEs




Continuous Correlation to Routing

- Time-series of Measurements from any P to any Q along any ECMP path
- Time-series of ECMP routed paths from any P to any Q

Continuous Correlation to Routing



Measured Latency
compared to best topology



Measured Latency
compared to current topology

- Time-series of Measurements from any P to any Q along any ECMP path
- Time-series of ECMP routed paths from any P to any Q

Inference

- Measurement (**PAR, MAD**) report SLA degradation (e.g., loss)
- Without any additional measurement, Routing Correlation allows to infer other (SRC, DST) pairs that are also impacted
 - BRU to MAD/LIS/SEV is impacted
 - LON to MAD/LIS/SEV is impacted



IP is better than ever

IP is better than ever



Build Anything End-to-End

IP is better than ever



Build Anything End-to-End



Measure Everything

