Developer Days
Automation

cisco
The bridge to possible

Integrated Performance Measurement

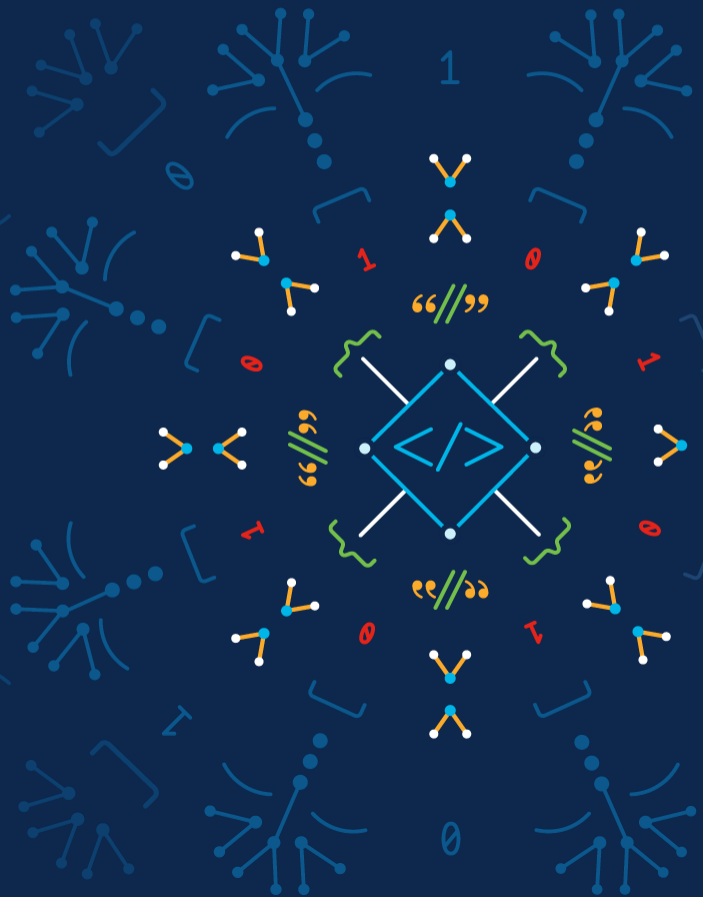

Clarence Filsfils
Cisco Fellow
May 21, 2024

We all love IP

- IP is at the heart of our industry
- We stream music over IP backbones
- Sensors protect our home over IP
- We leverage cloud workloads over IP
- Mobile industry runs over IP
- We got accustomed that IP cannot achieve anything alone
- Plethora of shim layers: MPLS, GTP, NSH, UDP/VxLAN…
- These shim’s cannot be combined and require expansive translations

Segment Routing

- In 2012, a team at Cisco initiated a bold move: let’s strengthen IP and allow IP to deliver any service by itself
- First phase: SR-MPLS: drastic simplification and scale up of MPLS
- Second first: SRv6 uSID: deliver any service over IP, without any shim

IP is back and better than ever.

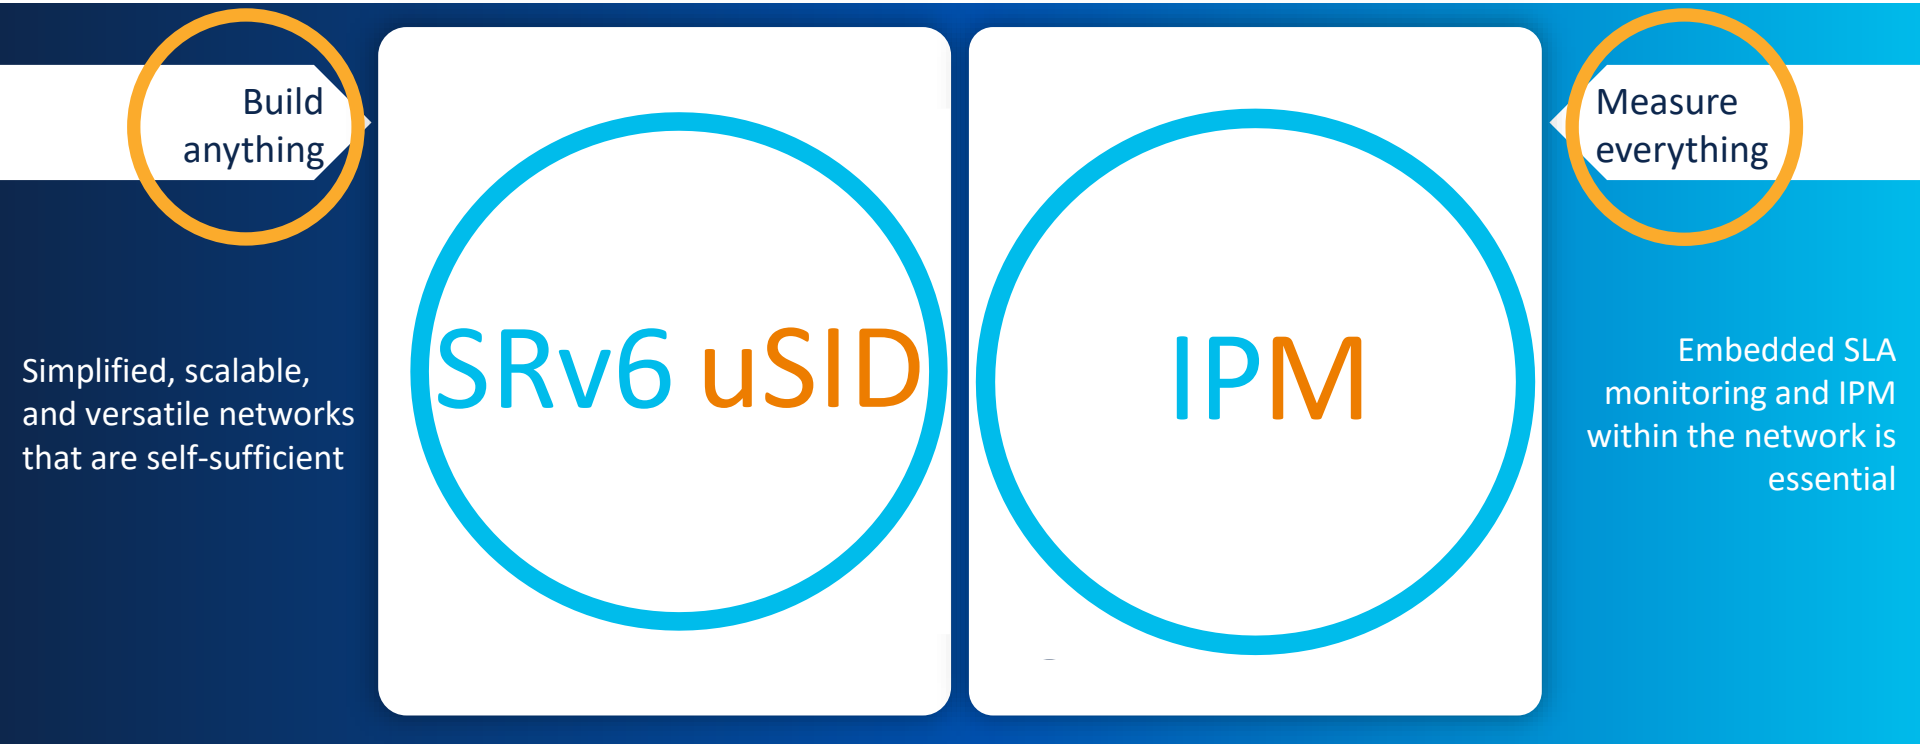

SRv6 uSID

- Build Anything
  - Any combination of underlay, overlay, service chaining, security…
  - VPN, Slicing, Traffic Engineering, Green Routing, FRR, NFV
- Any Domain
  - Access, Metro, Core, DC, Host, Cloud
  - End-to-End Stateless Policy
  - No protocol conversion or gateways at domain boundaries
- Seamless Deployment in Brownfield
- Built day-1 for Automation
- Standardized, Rich Eco-system, Rich Open Source (SONiC)

Outperform MPLS/VxLAN

Outperform MPLS - Daniel Voyer (Bell Canada)

- Native Optimum Slicing
  - SLID is encoded in Flow Label
- HW Linerate Push: 3 times better
  - J2 uSID linerate push: 30 uSIDs >> 10 MPLS Labels
- HW Counter and FIB consumption: 4 times better
  - uSID requires 4 times less counters and FIB entries than MPLS
- Routing scale: 20 times better
  - uSID supports summarization. MPLS requires host routes.
- Lookup efficiency: 2 to 3 times better
  - uSID can process 2 to 3 SIDs in a single lookup (LPM nature)
- Load-balancing: optimum and deterministic
  - uSID provides HW friendly entropy (fixed offset, shallow)

Bell SRv6 uSID Deployment
Paris 2022

Outperforms VxLAN – Gyan Mishra (Verizon)

- Seamless Host support for Network Programming
  - 6 uSID's in outer DA: RFC2460 IPinIP with opaque DA
- TE in the DC
  - elephant flows exist, asymmetric fabrics exist, TE is needed
- TE in the Metro/Core from the host
  - An SRv6 uSID DC allows for the application to control the network program in the metro/core without complex DPI and protocol conversion at the DC boundary,
- uSID DC provides lower MTU overhead (~5%)
  - Lower MTU overhead means lower DC cost
- Vendor, Merchant and SONIC/SAI maturity
  - uSID support across DC vendor (Cisco), Merchant (Cisco, Broadcom, Marvell), Sonic/Sai (Alibaba deployment)

SRv6 uSID DC Use-Case
Paris 2023

Rich SRv6 uSID Ecosystem

Network Equipment Manufacturers

Cisco, Huawei, Nokia, Ciena, Juniper Networks, Ericsson, ZTE, Arrcus, Arista, H3C

Merchant Silicon

Broadcom, Marvell, Barefoot Networks (Intel)

SRv6 uSID

Open-Source Applications

Pyroute2, IPTABLES SERA, GoBGP, exa, BGP, eBPF, cilium, Wireshark, FRRouting, Snort, BIRD

Open-Source Networking Stacks

SONiC, Linux, DPDK.io, P4, OpenStack, OPEN Compute Project, Kubernetes, OvS Open vSwitch, VyOS

Smart NIC / DPU

intel, NVIDIA

Partners

kaloom, Trend Micro, ENEA Qosmos Division, NoviFlow, ixia A Keysight Business, spirent Promise. Assured., Red Hat, f5, 6WIND, Keysight Technologies

SRv6 is Proposed Standard

Architecture
- SR Architecture – RFC 8402
- SRTE Policy Architecture – RFC 9256

Data Plane
- SRv6 Network Programming – RFC 8986
- IPv6 SR header – RFC 8754

Control Plane
- SRv6 BGP Services – RFC 9252
- SRv6 ISIS – RFC 9352
- SR Flex-Algo – RFC 9350

Operation & Management
- SRv6 OAM – RFC 9259
- Performance Management – RFC 5357

Strong Commitment and Leadership

Editor of 96% IETF RFCs
Co-author of 100% IETF RFCs

Over 80000 uSID routers deployed

Inter-DC/Metro Traffic Engineering across all of China
Eddie Ruan

14k+ devices, 70% services on uSID
Akash Agrawal

Simplicity Always Prevails

A Typical Deployment

segment-routing.net

Join our next uSID/IPM event – 9 & 10 October 2024

Stay up-to-date

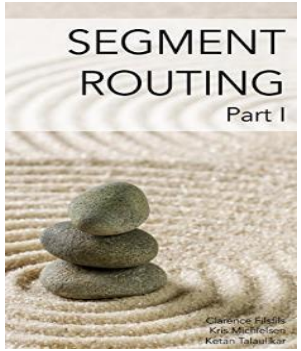

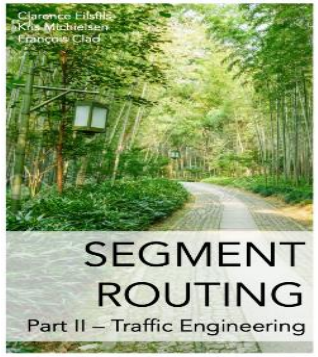

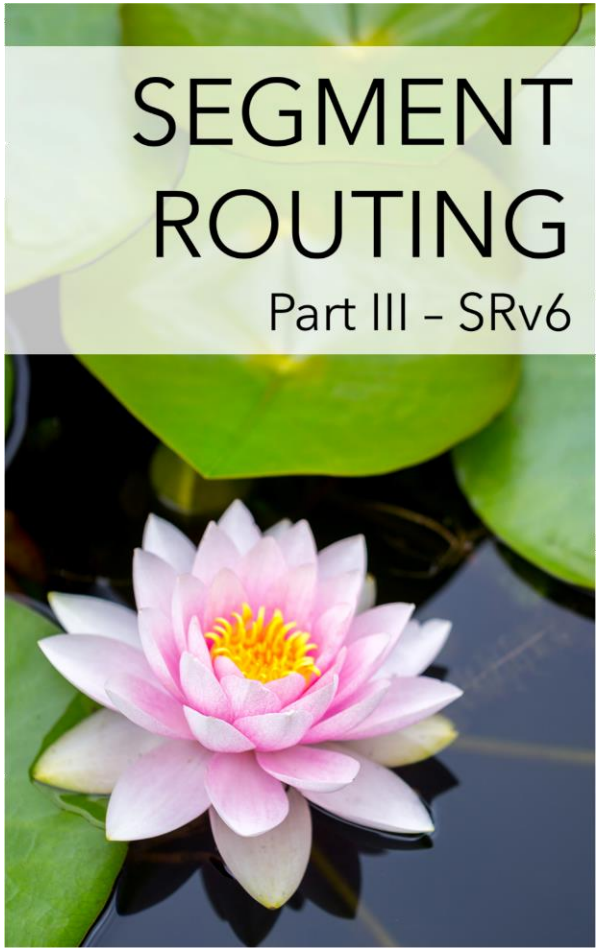

segment-routing.net

The nature of IP is ECMP

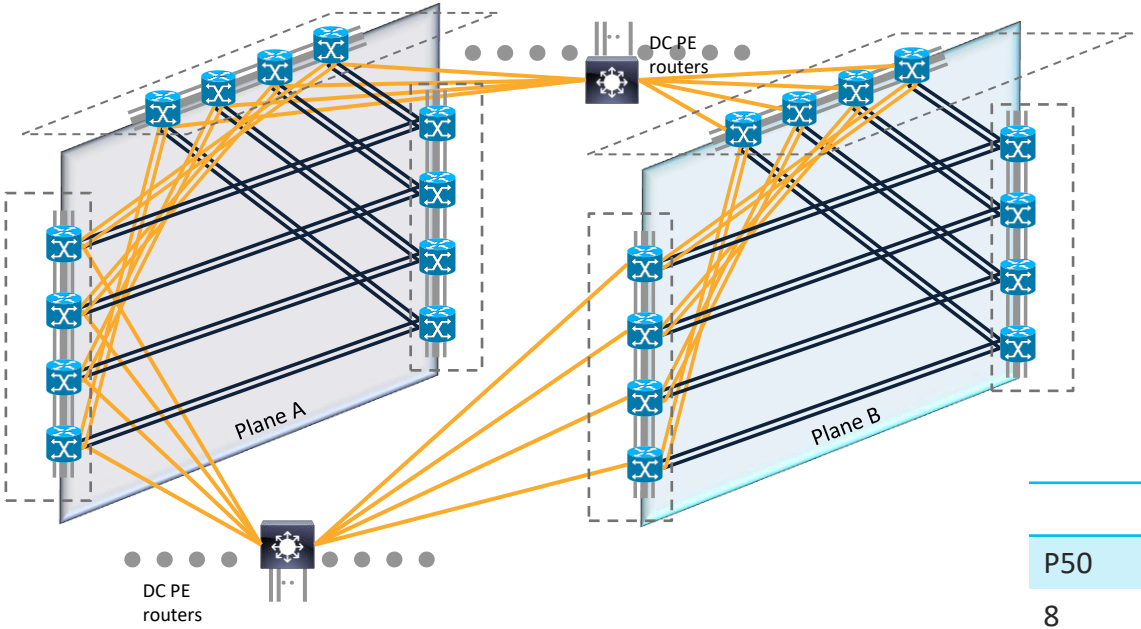

ECMP Paths		
P50	P75	P90
8	16	32

- Legacy solutions do not have the scale to measure all 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

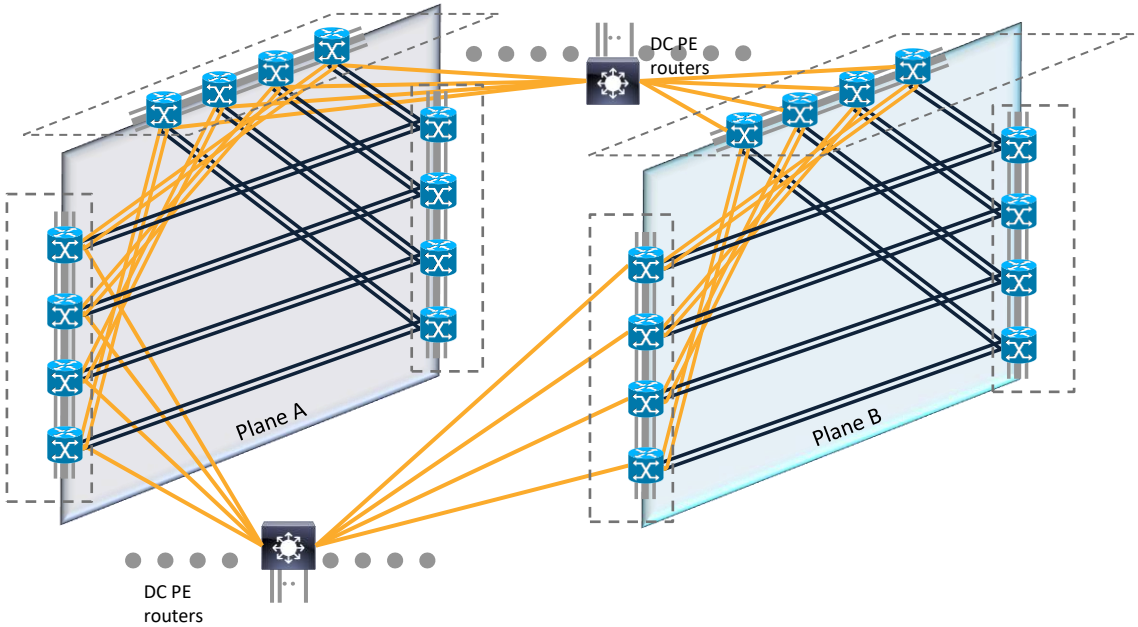

- 1 measurement every msec
- 500 edges
- 16 ECMP paths

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

Richer 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

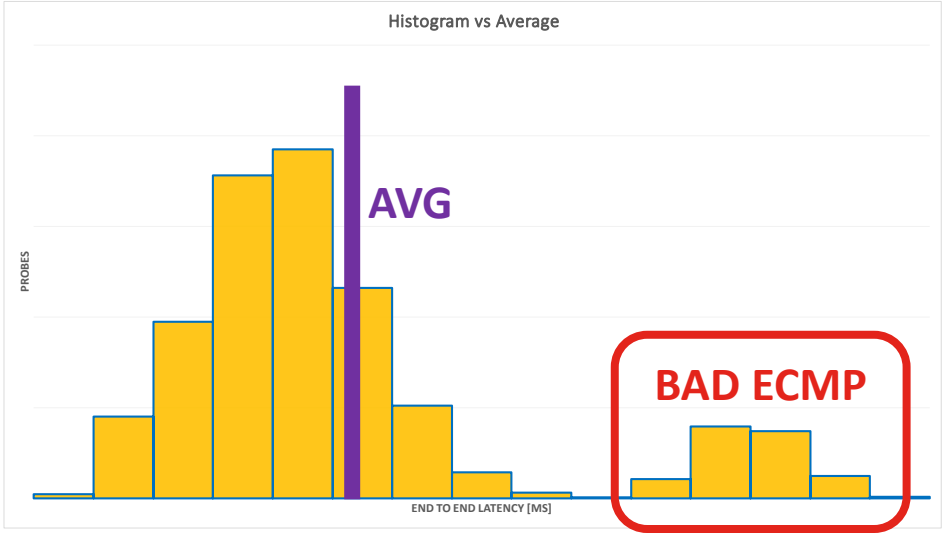

Any IP Fabric, Any Edge to any Edge, Any ECMP Path

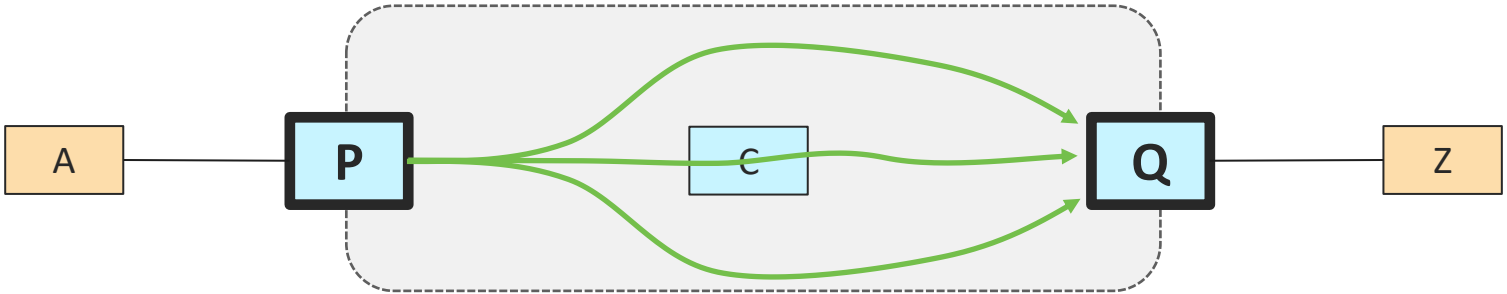

- Absolute Loss
- One-Way Latency (20nSec)
- Liveness (sub-2msec)
- Standard: STAMP (RFC 8762 & RFC 8972)

Much Cheaper through Silicon Integration

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

Continuous Correlation to Routing

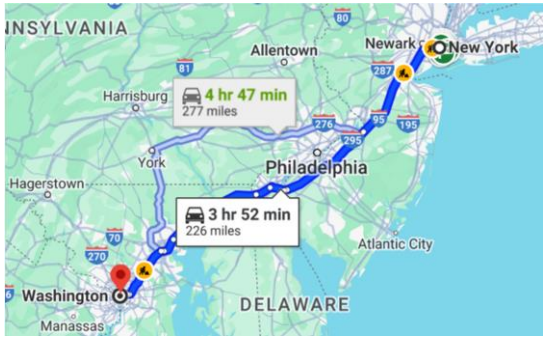

Measured Latency
compared to best topology

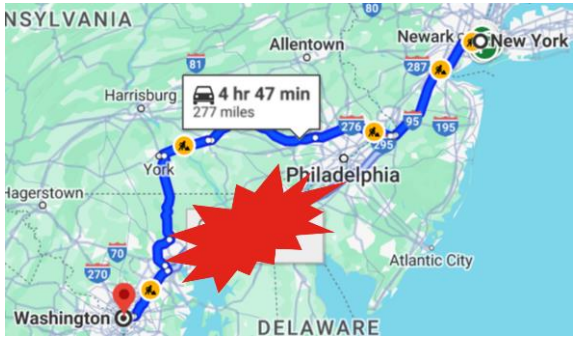

Measured Latency
compared to current topology

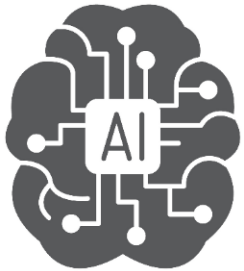

- Time-series of Measurements from any P to any Q along any ECMP path
- Time-series of ECMP routed paths from any P to any Q

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

Michael Valentine, Technology Fellow, Network Architecture,
Goldman Sachs

- uSID and IPM Use-case
- IPM Silicon Integration and Metrics
- Link

Bart Janssens, Senior Specialist Packet Architecture, Colt Technology Services

- Routing Analytics
- Accedian Skylight
- Deployment and Use-Cases
- [Link]

Next phase

Accedian Skylight

End to end service visibility

Insights with micro-second granularity

Faster to launch and monetize SLA-backed services

New on-demand services based on power and SLA based metrics

Silicon One

- **5/9** of service availability with automation
- **Sub-50ms** restoration times
- Up to **97%** power savings
- Accedian Skylight agents to **monitor** **cloud-based 3rd party networks**

\+ Integrated Performance Measurement

Routed Optical Networking

Helping Colt and its customers **lead in performance and sustainability**

Gyan Mishra, Associate Fellow
Verizon

- DC use-case
- uSID and IPM
- Lightweight Host Routing (LHR)
- Link

Eddie Ruan, Senior Staff Engineer
Alibaba

- uSID - Deployment Experience
- SONIC Experience
- Link

IP is better than ever

Build Anything End-to-End

Measure Everything

The bridge to possible