Open. Together.

OCP
SUMMIT

EW:SONiC

In-band Network Telemetry an Open and multi-platform Network Analytics for Disaggregated Data-Centers

Roberto Mari, Product Management, Barefoot Networks

Agenda

- Benefits of a Programmable Data-plane
- Barefoot OCP Ecosystem
- Advanced Data-Plane Telemetry
- Deep Insight Analytics
- Data-Plane Telemetry (SPRINT) multi-platform Demo
- Summary and References

Benefits of Programmable Switches

Reduced Complexity

Strip out **Complexity** that isn't needed by the User, Application and Operator

Data Plane Telemetry

Instrument the Data Plane with Barefoot **SPR**INT™ Scalable, Programmable, Real-Time, In-band Network Telemetry (INT)

Future Proof

Continuously deliver and **Evolve** features on the same Hardware

Scale

Scale data-plane resources to match the needs of Hyper Scale & Service Infrastructure

Speed & Agility

Adapt & Innovate at **Speed** of software

Barefoot OCP Ecosystem – Best of Breed

Network Operating Systems

SONiC

ipinfusion™

ONL Open Network Linux

STRATUM

White Box Hardware (ODMs)

Edge-corE® NETWORKS

WNC Wistron NeWeb Corp.

Supported Devices and Platforms

Lihua Yuan edited this page 2 days ago · 48 revisions

Following is the list of platforms that support SONiC. Last updated Mar 2018.

Switch Vendor	Switch SKU	ASIC Vendor	Swich ASIC	Port Configuration	SONiC Image
WNC	OSW1800	Barefoot	Tofino-T10-018D	48x25G+6x100G	SONiC-ONIE-Barefoot[6]
Edgecore	Wedge 100BF-32X	Barefoot	Tofino-T10-032D	32x100G	SONiC-ONIE-Barefoot[6]
Edgecore	Wedge 100BF-65X	Barefoot	Tofino-T10-064Q	65x100G	SONiC-ONIE-Barefoot[6]

Barefoot ASICs

6.4 Tb/s

3.2 Tb/s

2.4 Tb/s

2.0 Tb/s

1.8 Tb/s

12.8 Tbps

8.0 Tb/s

6.4 Tb/s

INT Data-Plane Telemetry

Data-Plane Telemetry - Different Methods

“*Track every flow*”: Flow Reporting

- Monitor and report every flow’s path and latency

“*Track every drop*”: Drop Reporting

- Mirror every dropped packet along with the drop reason

“*Track every congestion*”: Congestion Reporting

- Produce packet-level snapshots of a congested queue
- Detect, characterize, and analyze microbursts

Flow Reporting: INT End-to-end Mode

Leverages In-Band Network Telemetry (INT)
https://github.com/p4lang/p4-applications/blob/master/docs/INT_v0_5.pdf

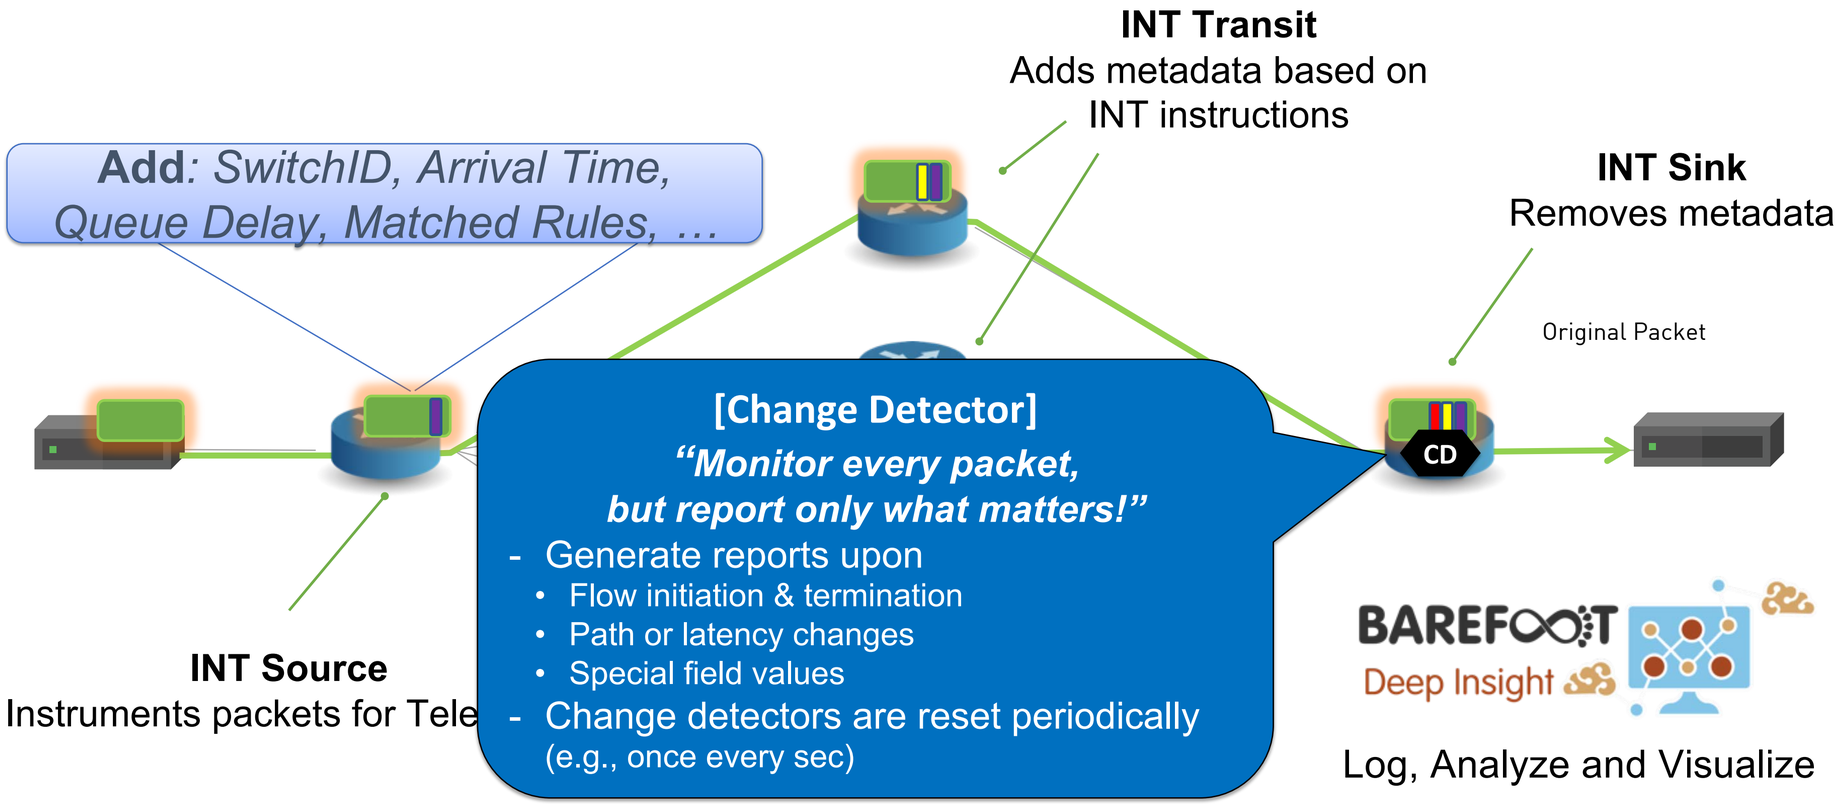

Flow Reporting: INT Hop-by-hop Mode

Report-traffic volume increase is marginal!

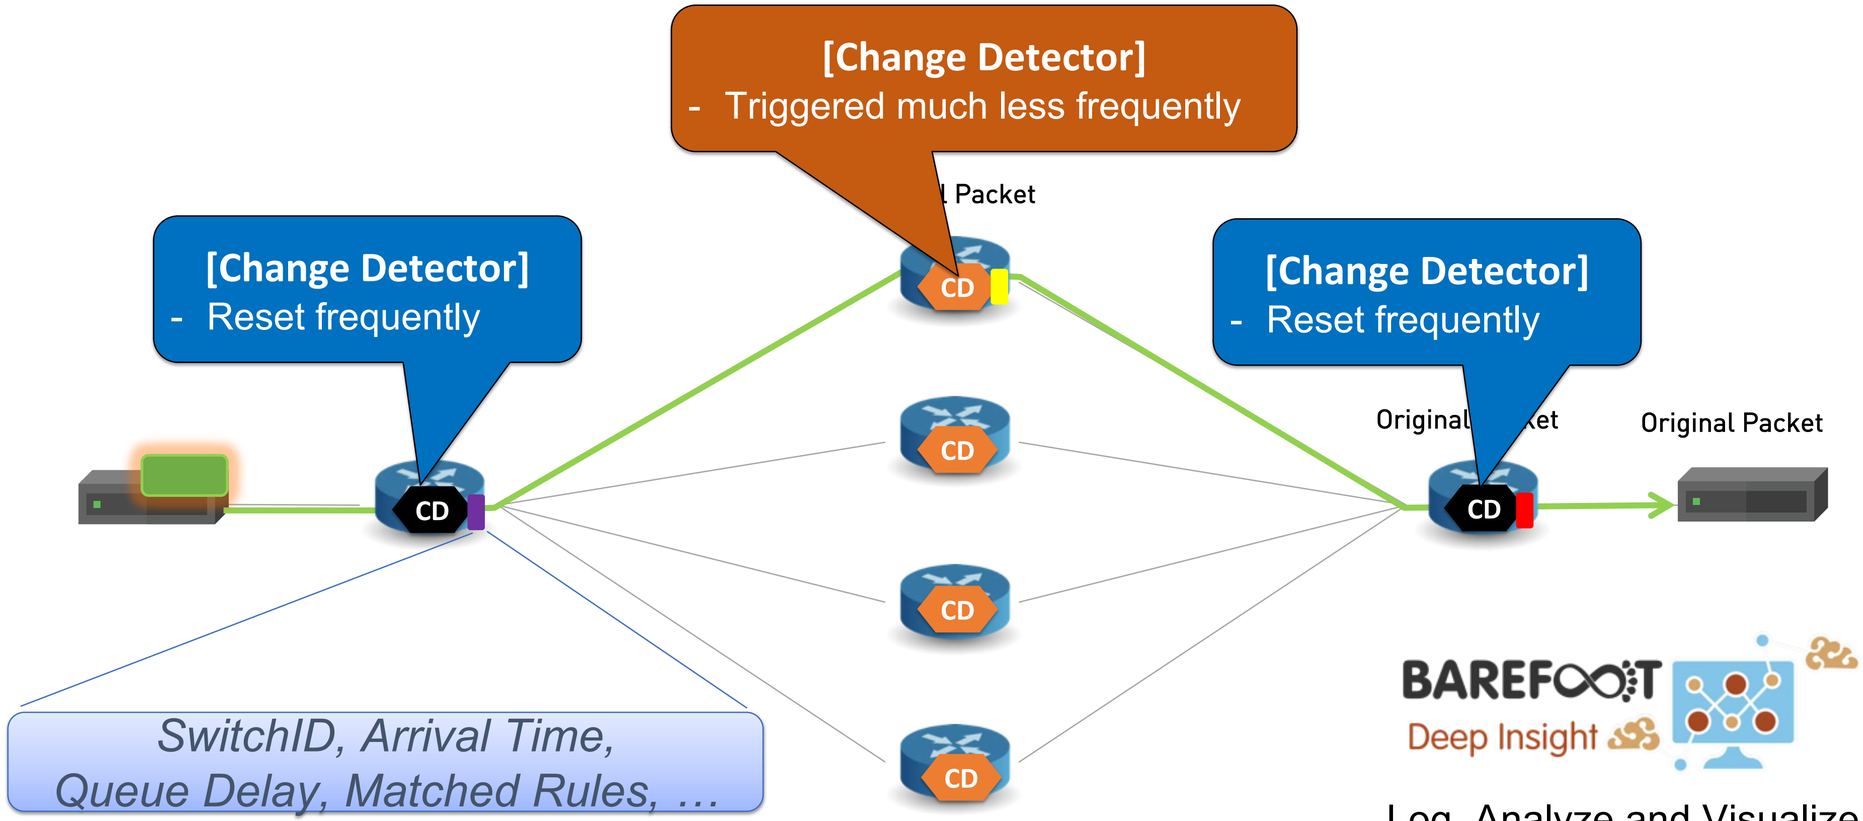

BAREFOOT Deep Insight

Log, Analyze and Visualize

Two Different Modes of Flow Reporting

INT Hop-by-hop Mode

- **Each network device** generates its own flow reports.
- **Deep Insight Analytics** receives reports from different network devices, each describing the telemetry metadata.

- Does not modify original packets
- Functionally equivalent to end-to-end mode
- No dependencies on INT marking (e.g., reserving a DSCP bit)

INT End-to-end Mode

- **Metadata is embedded** in between the original headers of data packets as they traverse the network
- The sink node generates flow reports using data carried in the packet.

- Relies on correct network configuration for packet modification

SPRINT: A Fully Featured, High-Performance INT

FULLY COMPATIBLE SUPERSET OF A VANILLA INT IMPLEMENTATION

S

Smart

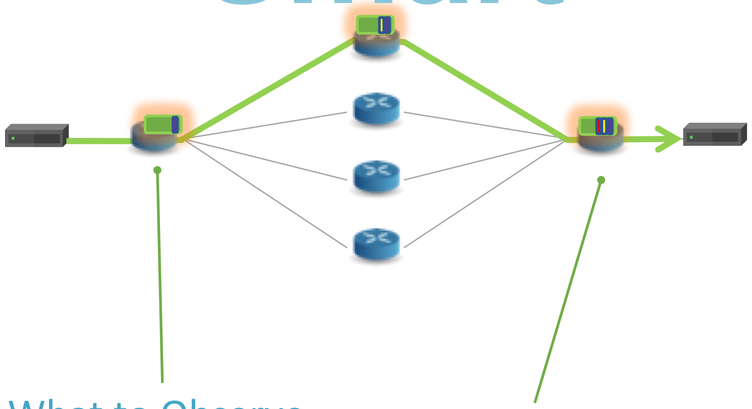

- ✓ What to Observe
- ✓ What to Collect
- ✓ Intelligent Triggers
- ✓ Built-in Load Balancing

P

Programmable

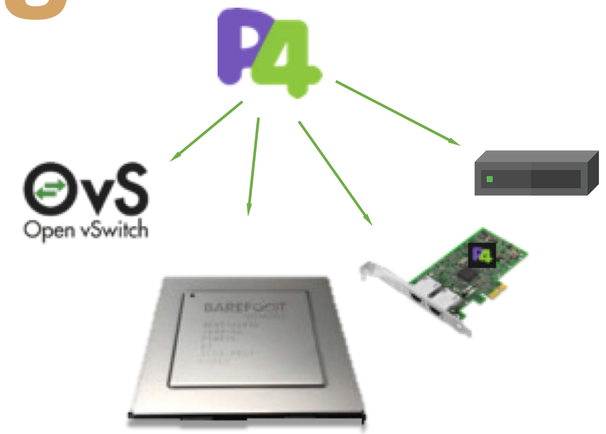

- ✓ Adapt to customers requirements
- ✓ Flexible encapsulation through P4
- ✓ Open specifications and ecosystem

R

Real Time

- ✓ Data-plane Streaming
- ✓ Packet-by-packet Anomaly detection
- ✓ Real time Analytics with Deep Insight

Barefoot Deep Insight Analytics

SPRINT Data-Plane Telemetry

- In-Band Network Telemetry (**INT.P4**)
- Intelligent Deduplication and Triggers
- Hardware Primitives
- Line Rate Monitoring

3rd Party Network Management Solutions

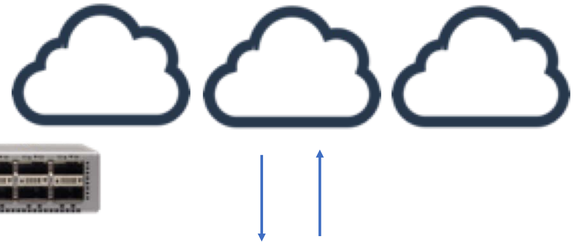

Answer for **Every** Packet...

1. How did it get here?
2. Why is it here?
3. How long was it delayed?
4. Why was it delayed?

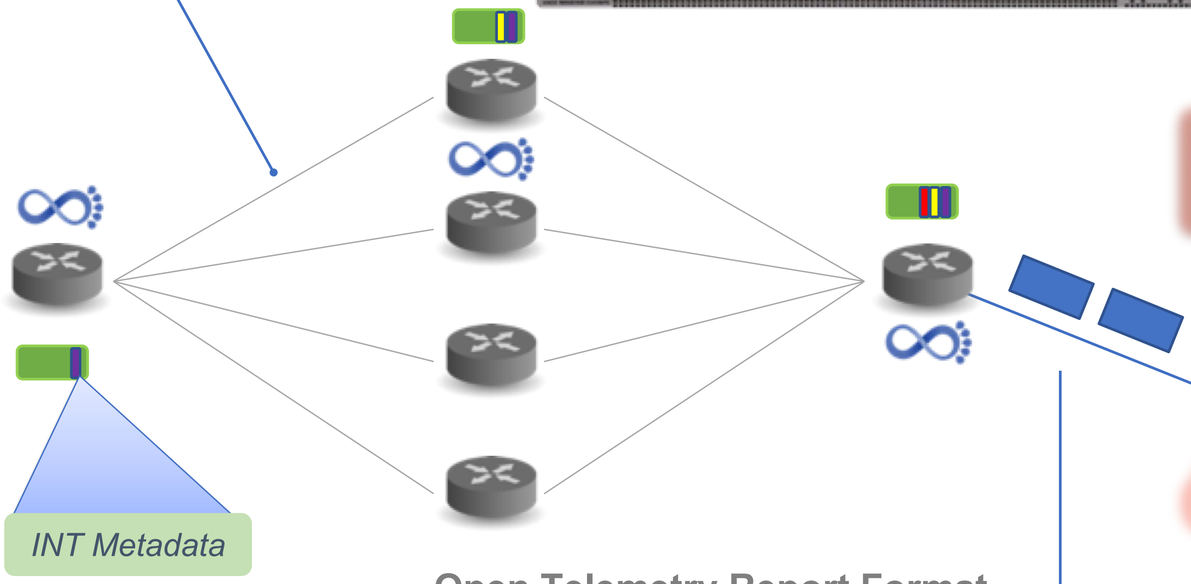

Deep Insight
Open Northbound APIs

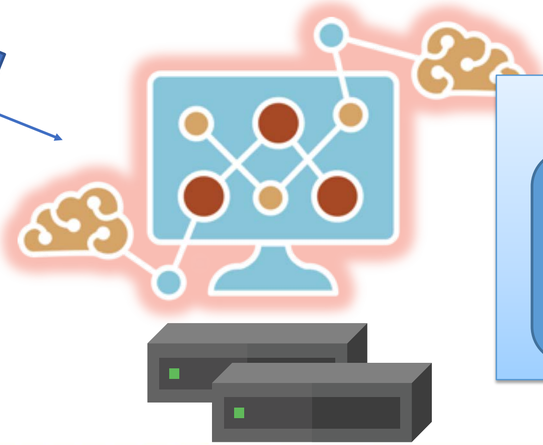

Deep Insight Analytics Software

- Real-time Anomaly Detection
- Rich Analytics
- Dashboards and Drill down workflows
- Modular Architecture for scale-out
- Runs on Commodity Servers

Open Telemetry Report Format
defined by the **P4.org**
Applications Working Group

Deep Insight Real-Time Analytics – Key Capabilities

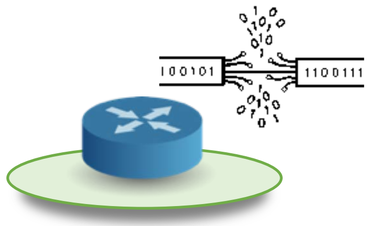

Packet Drop Analysis

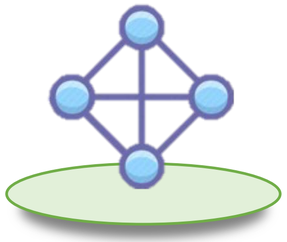

Topology Discovery

UI with anomalies reports and Dashboards

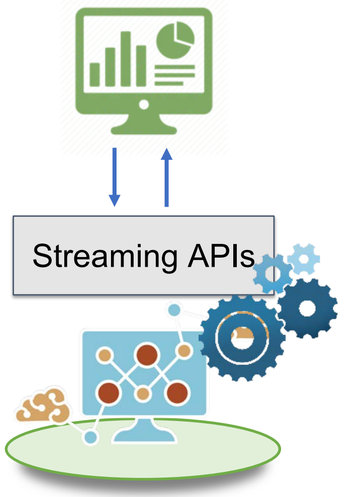

Real-Time notifications to 3rd-party Network Management for Closed Loop

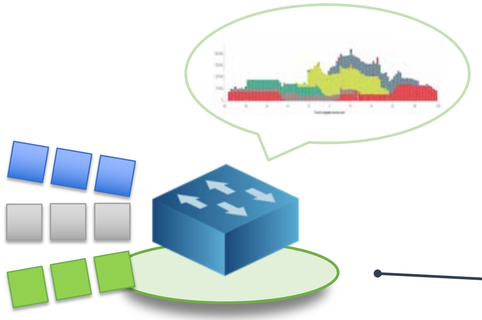

Congestion Analysis

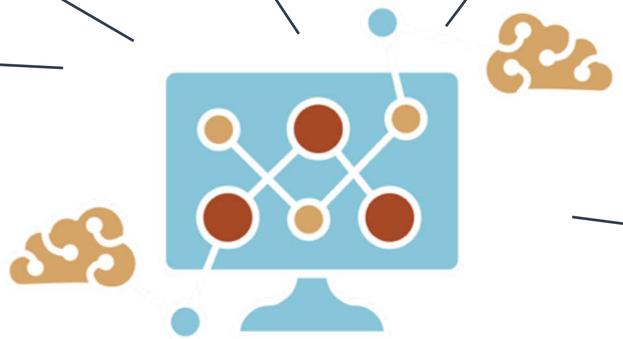

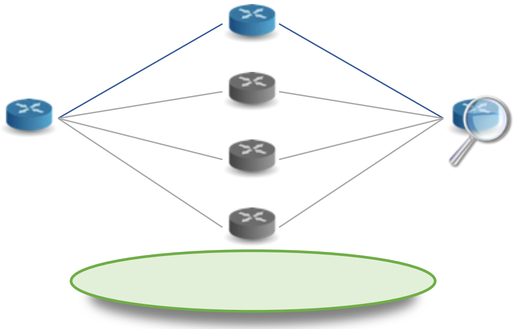

Path and Latency Tracking

Single Server

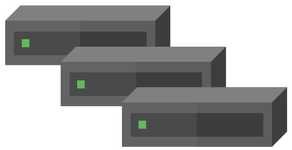

Cluster Mode

Data Retention and Historical Analysis

SPRINT Data-Plane Telemetry Demo

SPRINT Demo with SONiC with OCP Platforms

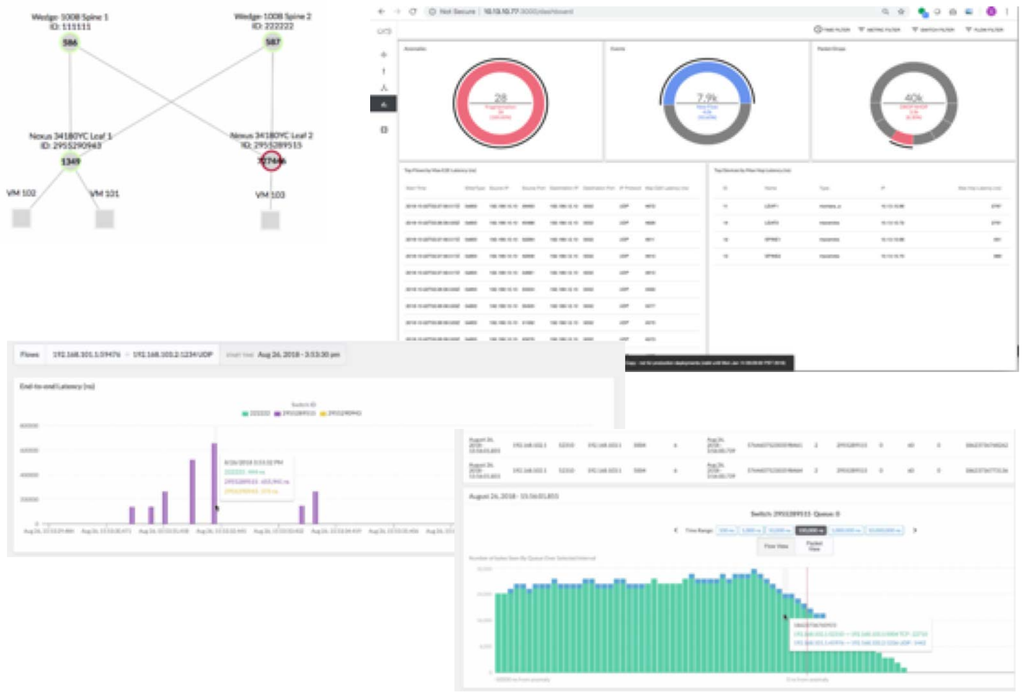

Spine – 2x Wedge 100BF-65X
(Hop-by-hop INT)

Leaf – 2x Wedge 100BF-32x
(Hop-by-hop INT)

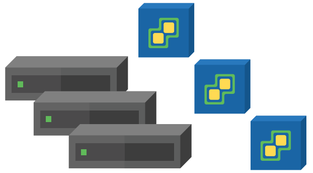

Virtualized Servers
with VMs

(**1x** server)

References - More about SPRINT

- Integrated with SONiC and Commercial Open NOS

  SONiC (via SAI), Stratum, IP Infusion OcNOS

  https://github.com/Azure/SONiC/blob/gh-pages/doc/barefoot_dtel/Dtel-SONiC.md

  https://github.com/opencomputeproject/SAI/blob/master/doc/DTEL/SAI-Proposal-Data-Plane-Telemetry.md

- Open-source implementation of INT and spec in P4

  P4 Apps WG (https://github.com/p4lang/p4-applications)

  IETF In-situ OAM (https://datatracker.ietf.org/wg/ioam/about)

Summary – SPRINT Data Plane Telemetry

OCP Ecosystems

Available Today on OCP Barefoot Platforms, support up-streamed to SONiC

Proven

Deployed in Production at Tier 1 & Tier 2 Hyperscalers and TELCOs

Open

SAI Implementation

P4 open specification

Accurate

Uncompromised Telemetry


Open. Together.
OCP Global Summit | March 14–15, 2019
OCP
SUMMIT