The Trend to Edge Computing

Diverging Switch Form Factors

David Iles – Mellanox Technologies

Workloads Moving Off Premise

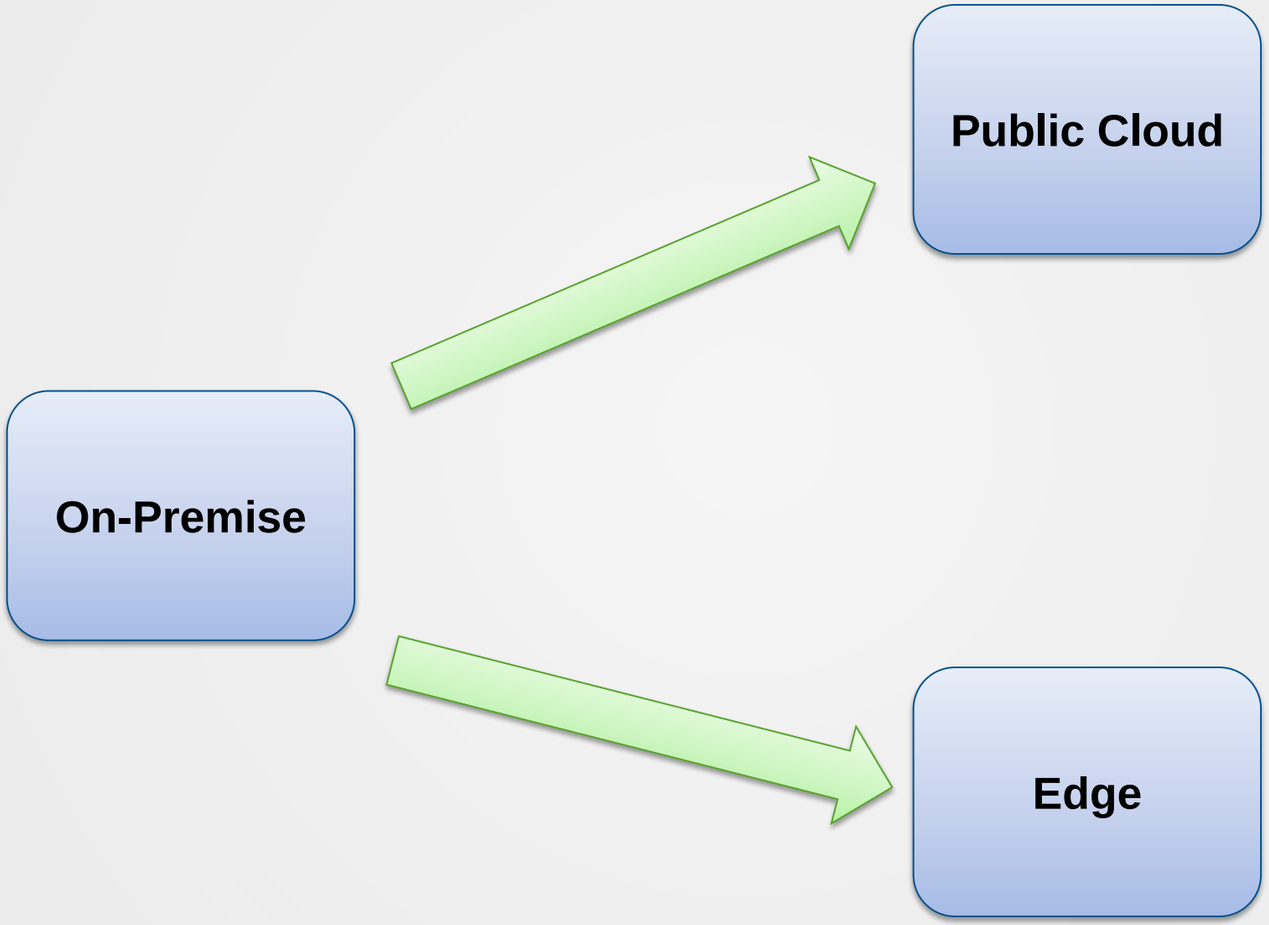

amazon web services

Microsoft Azure

Google Cloud Platform

EQUINIX

DIGITAL REALTY

Data Center Ethernet Switch Ports (10G & Above)

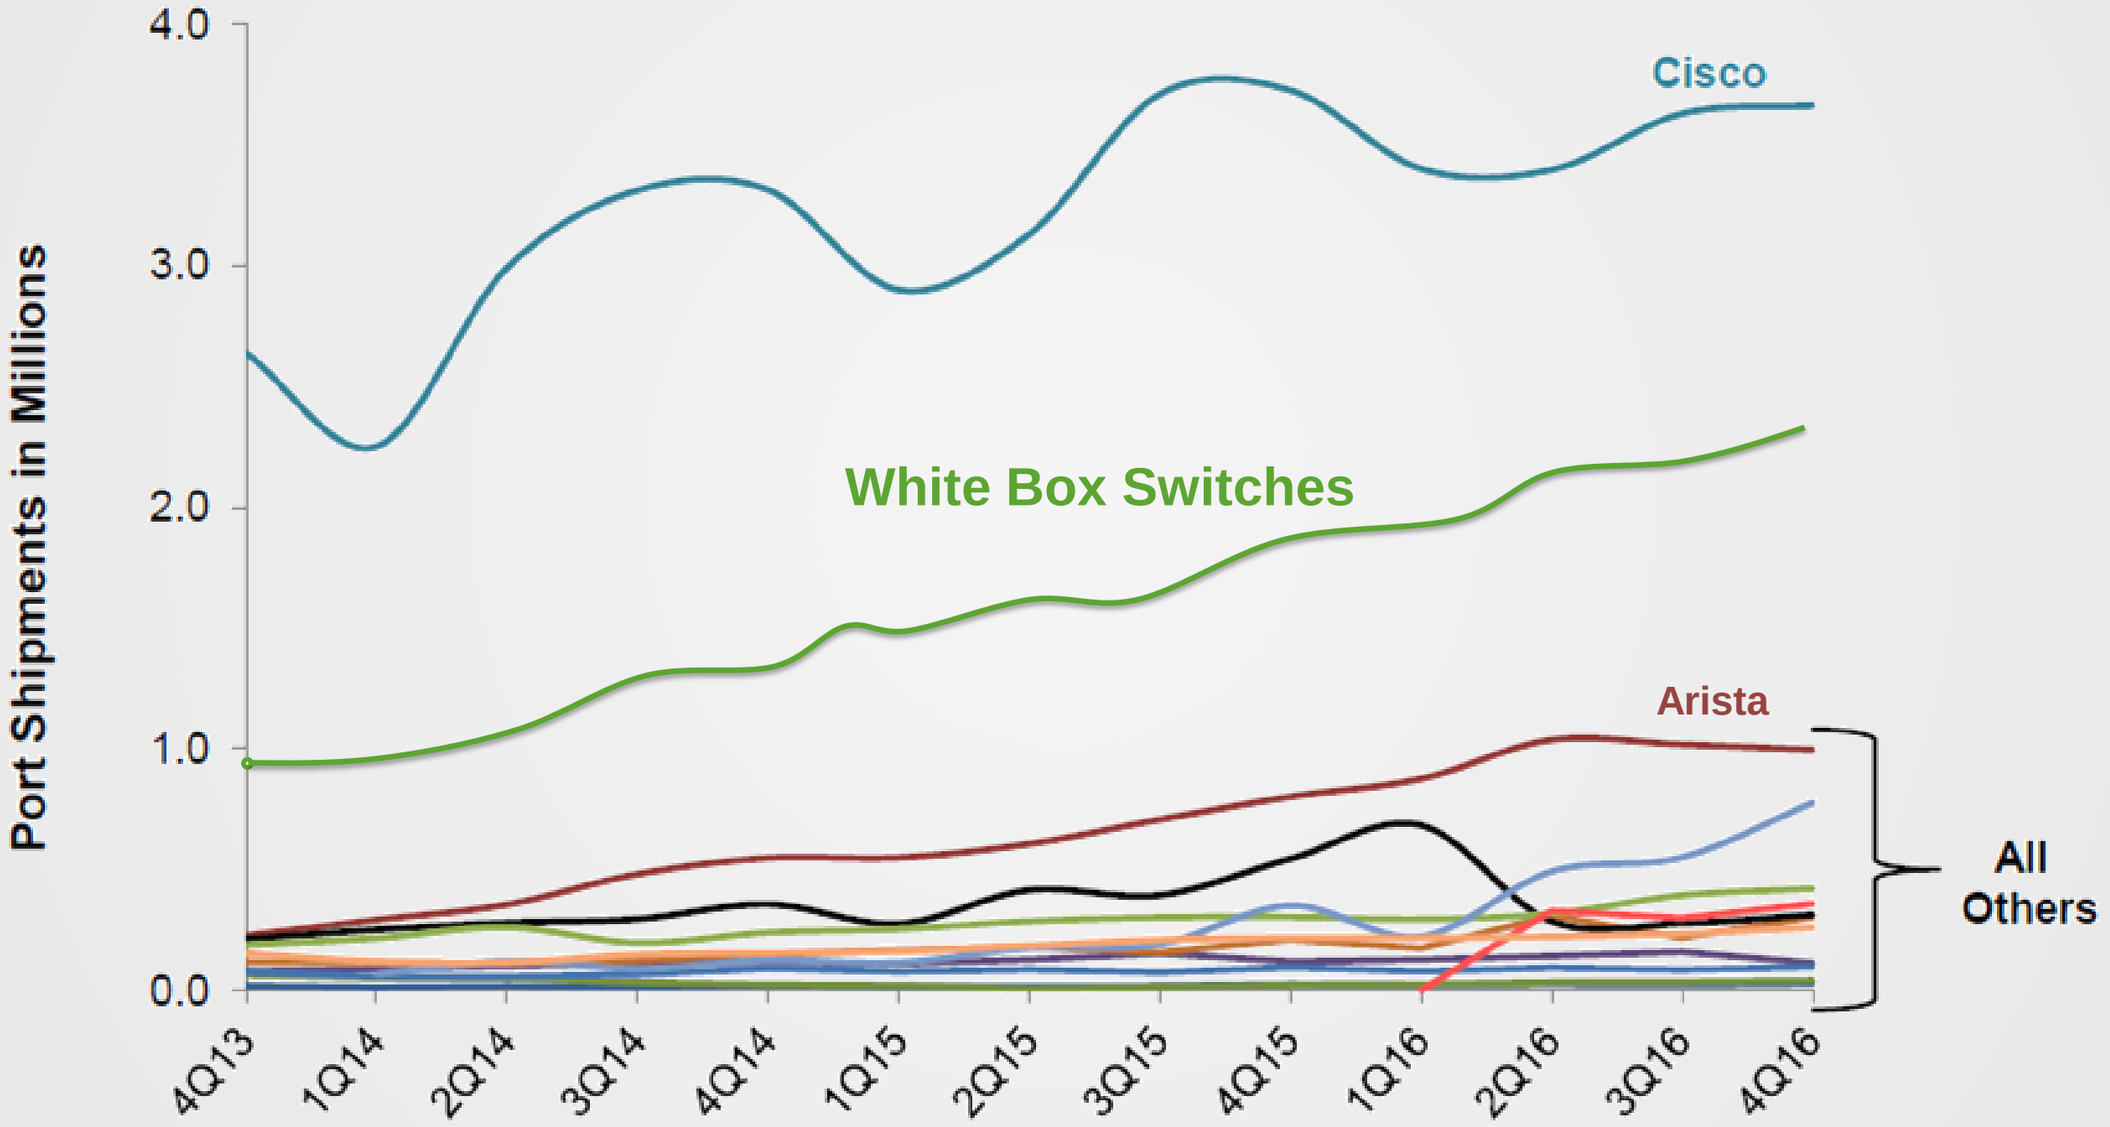

Source: Crehan Research and Mellanox Estimates

White Box for Public Cloud

…a mixed bag

- Whitebox Switches Pros
  - Right price
  - Right port count
  - Right NOS
- Whitebox Switch Cons
  - Limited VXLAN scale
  - VXLAN + 100G + Routing
  - RoCE limitations
- Future for Public Cloud
  - 200/400 Gigabit Ethernet
  - COBO – Onboard Optics
  - 277V Power Supplies
  - Large scale tables/tunnels

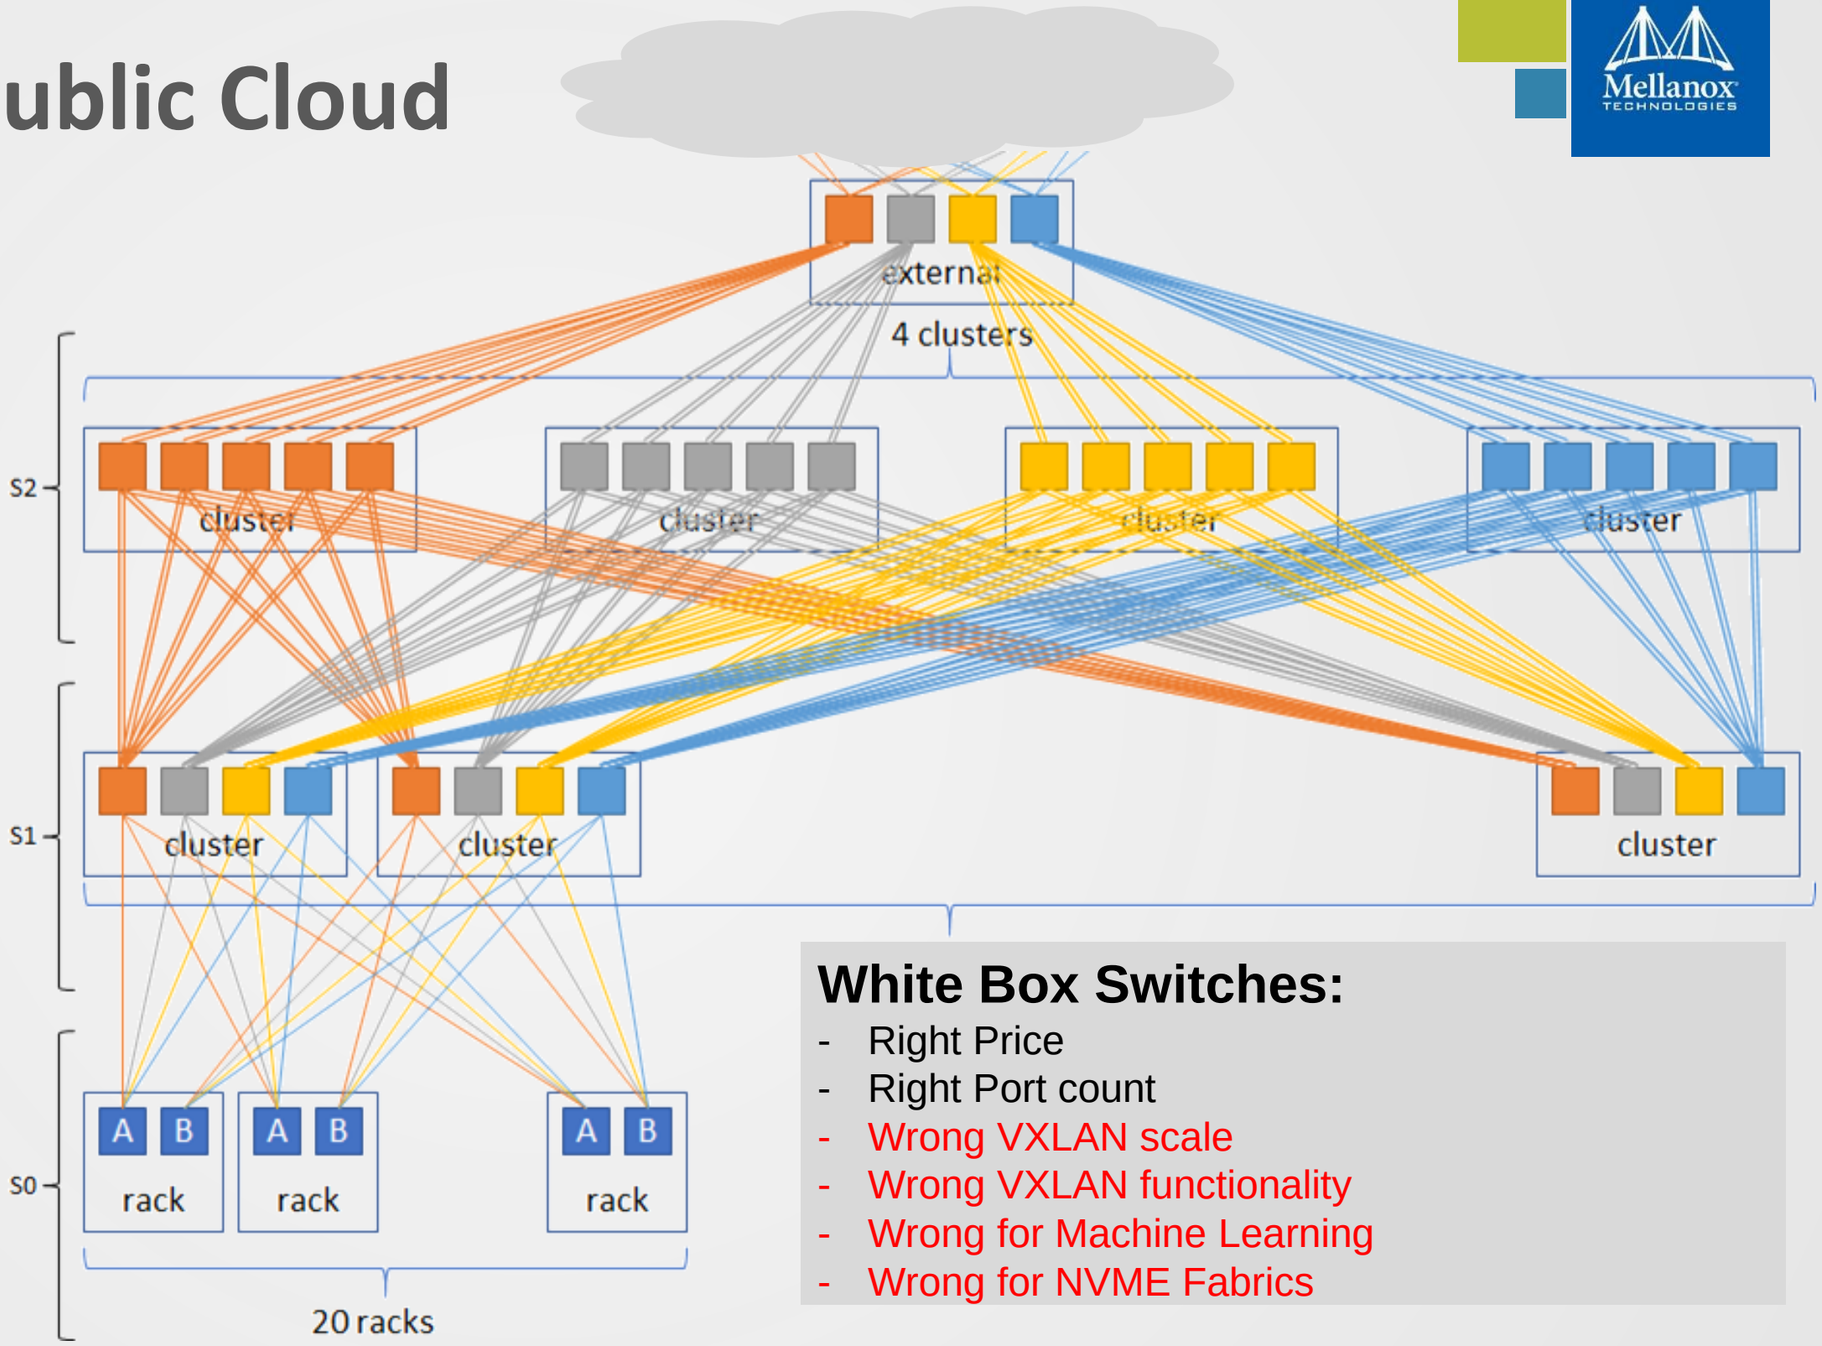

White Box Switches:

- Right Price
- Right Port count
- Wrong VXLAN scale
- Wrong VXLAN functionality
- Wrong for Machine Learning
- Wrong for NVME Fabrics

Workloads are Moving Around

VXLAN is the tunnel of choice

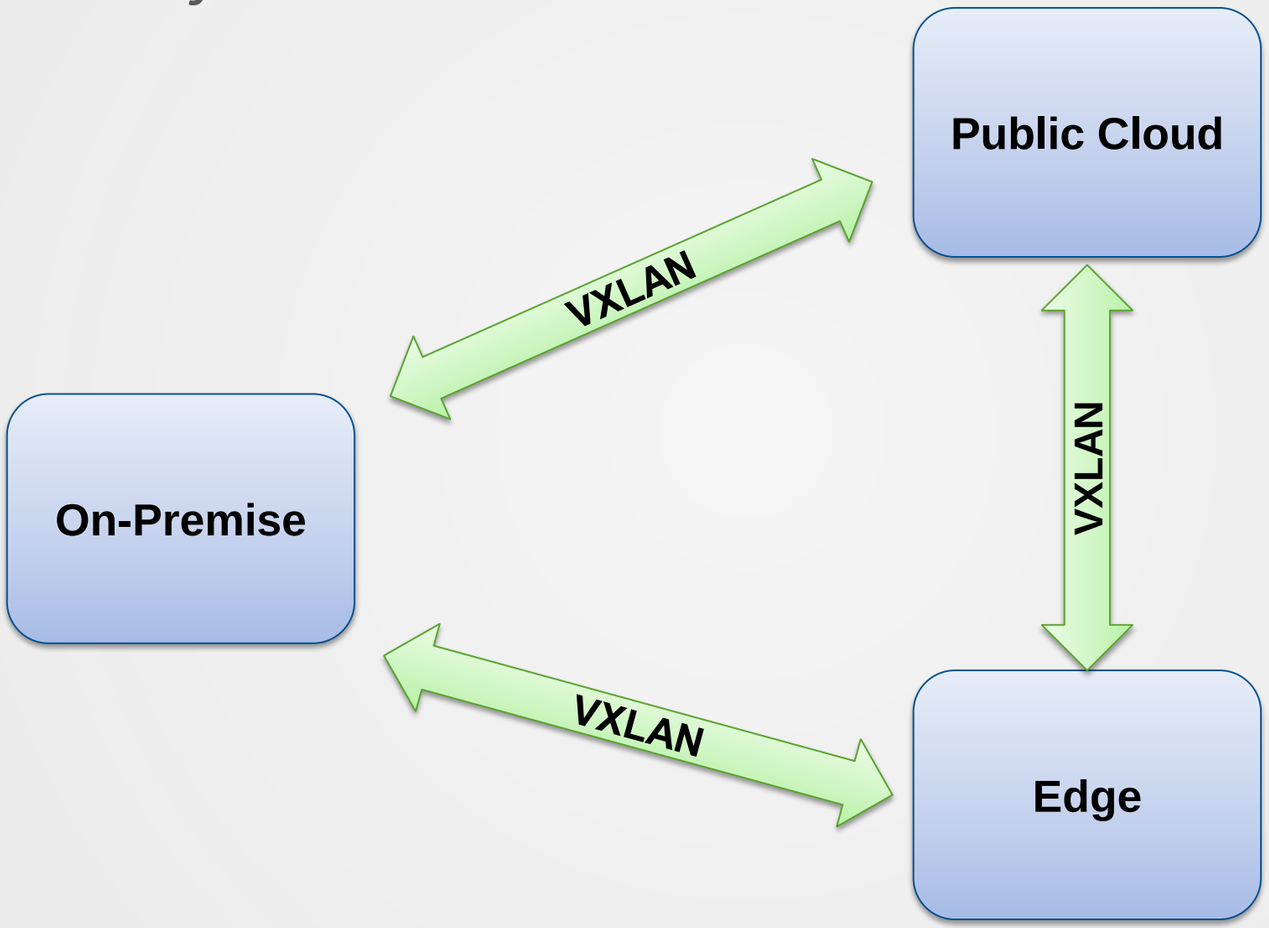

Cloud Connect Acceleration

- Millions of VXLAN tunnels
- 100K Tunnels per switch

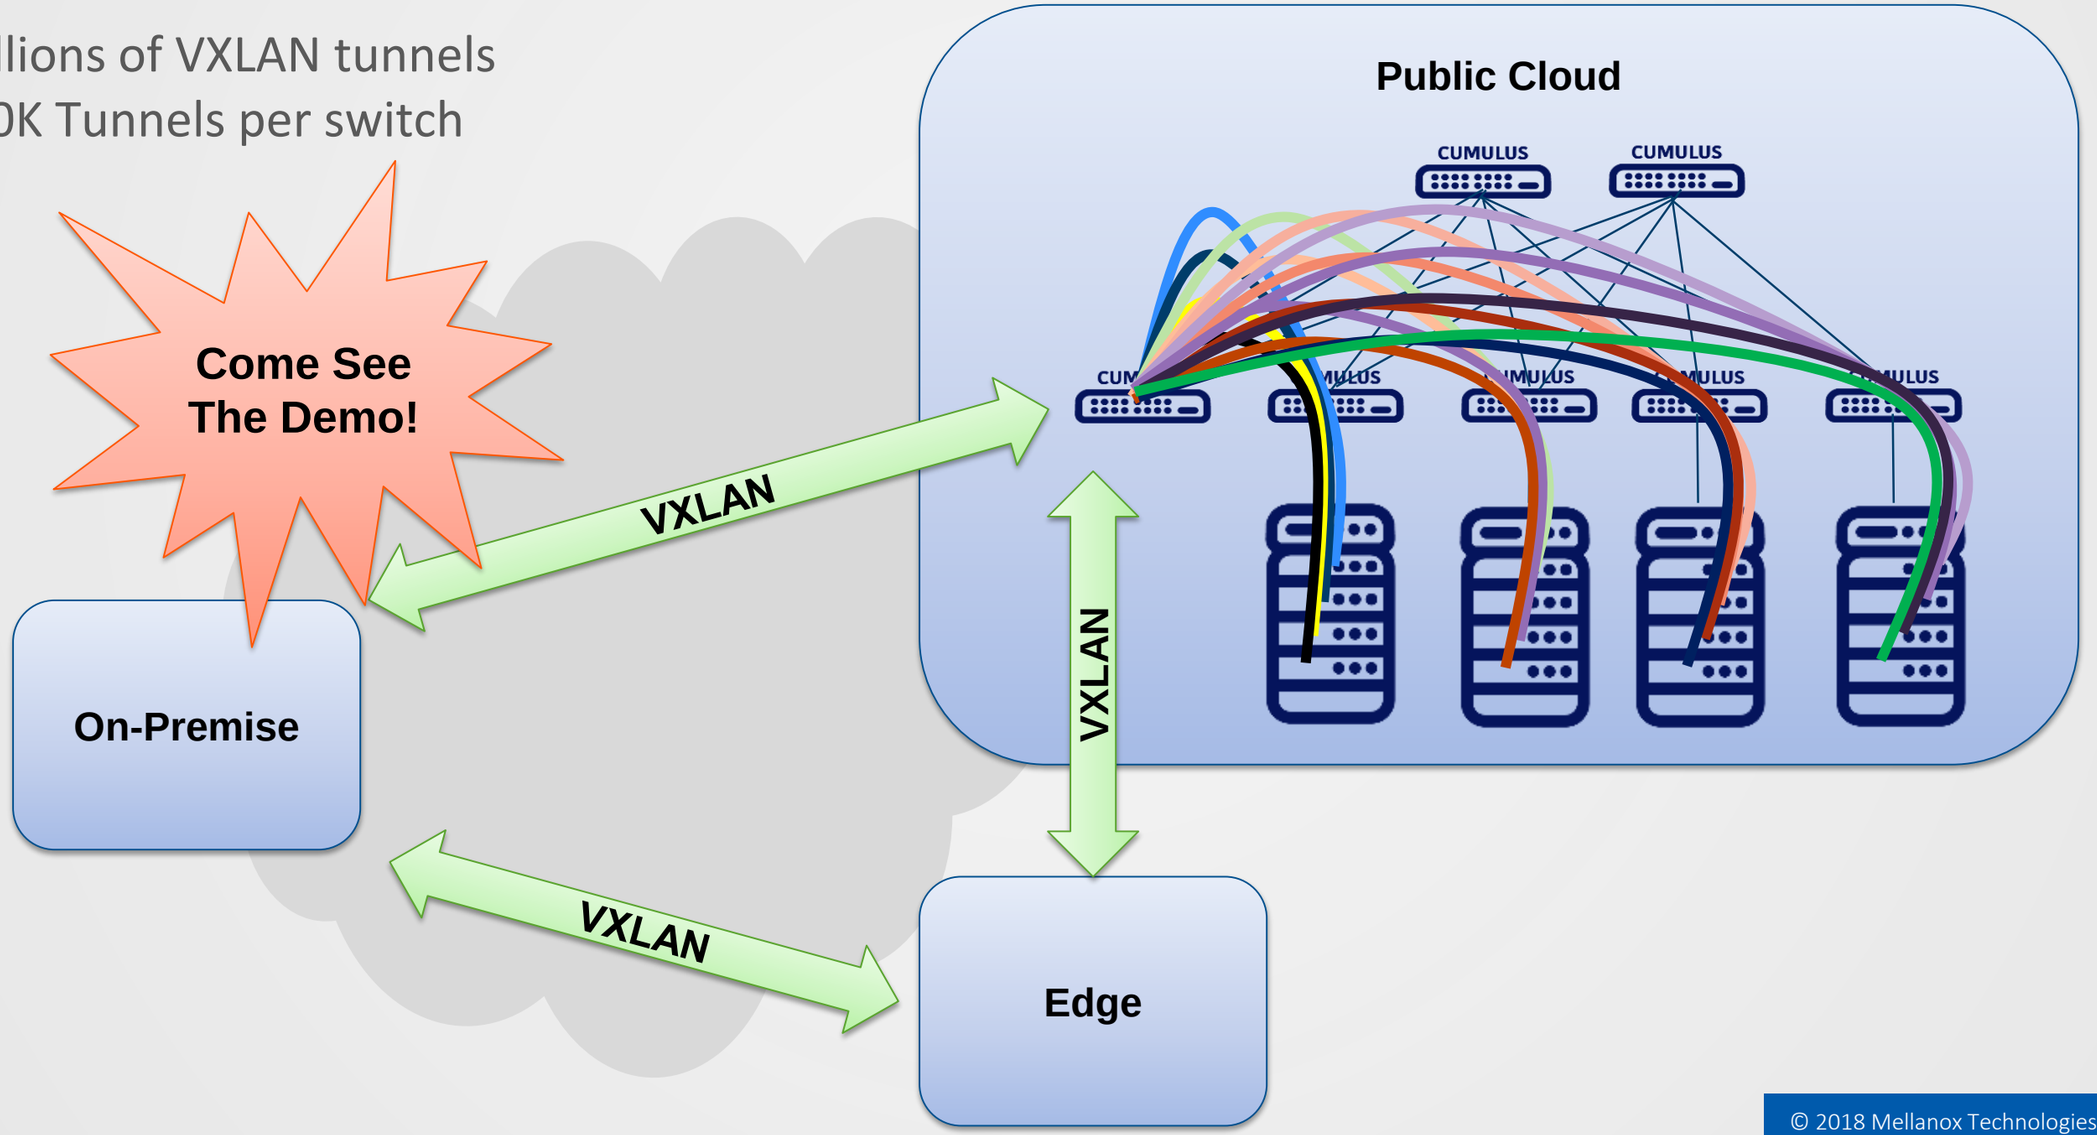

Not All Workloads are Moving Off Premise

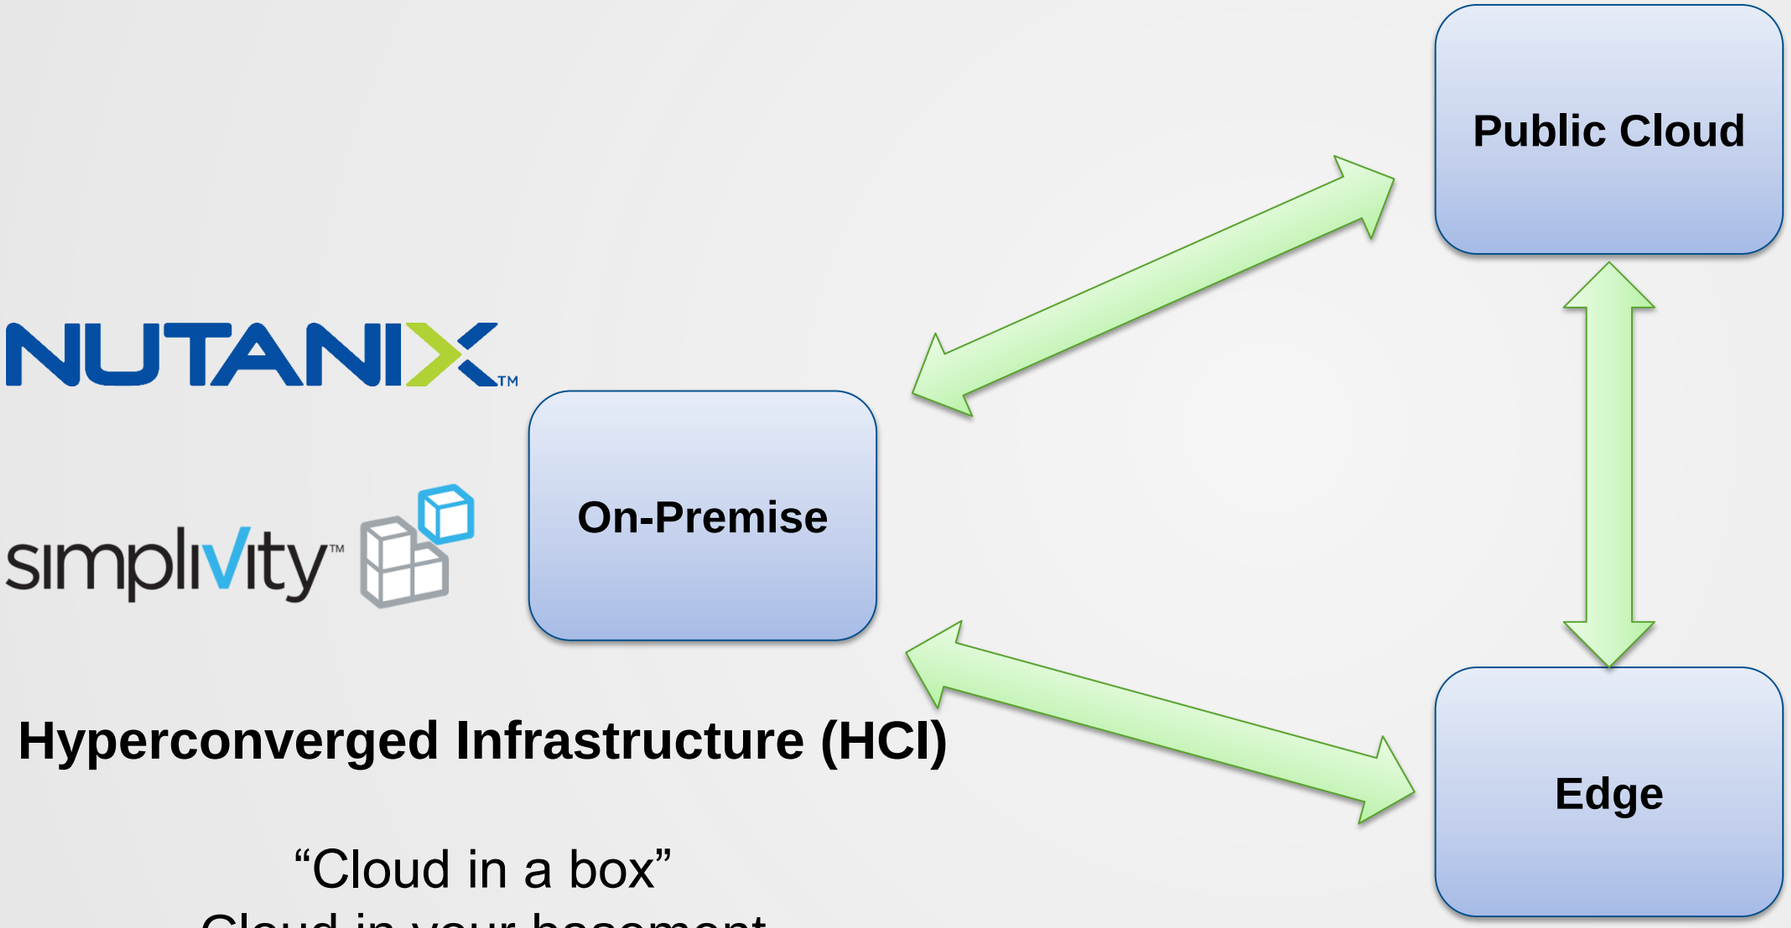

Hyperconverged Infrastructure (HCI)

"Cloud in a box"
Cloud in your basement

Switches Optimized for On Premise Computing

SWaP = Size, Weight & Power

Partial Rack Design

(10/25GbE downlinks, 100GbE uplinks)

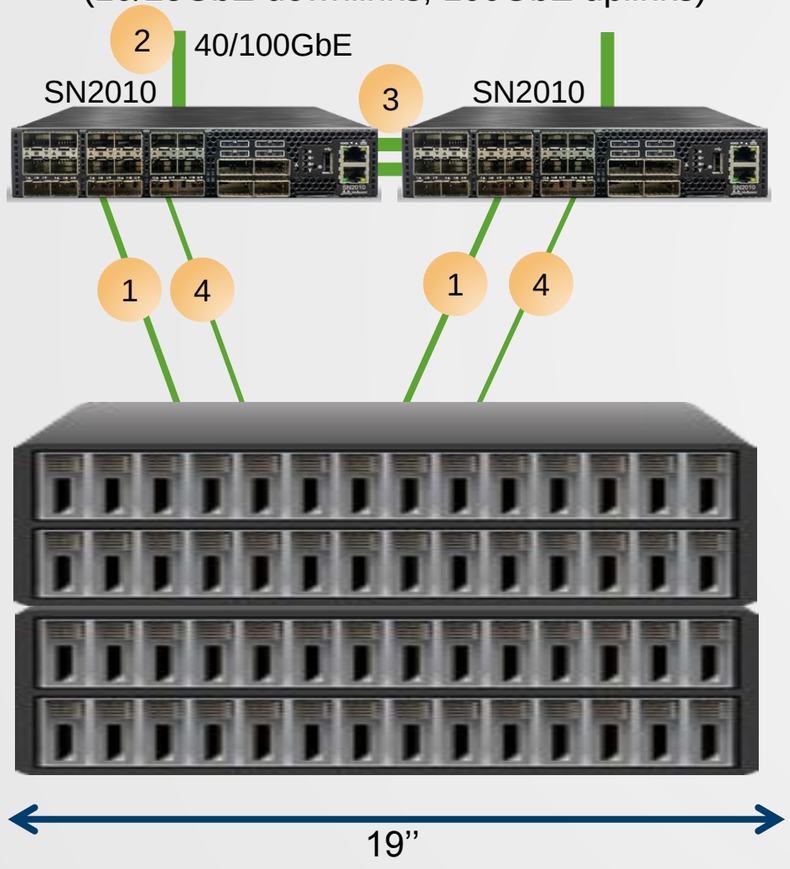

1. 10/25GbE link: SFP28 to SFP28

2. 100GbE Uplink: QSFP28 Transceiver

3. 100GbE MLAG links: QSFP28 to QSFP28

4. 1GbE Transceiver

- ½ 19” width, 1RU height
- 57W typical power
- 18x10/25GbE + 4x40/100GbE
- On-switch Containers

Performance

10G Optimized

25G Ready

Best $/Gb/s

SWaP for On Premise: Small, Fast, Flexible with Management Integration

Edge Computing

What's the opposite of data center consolidation?

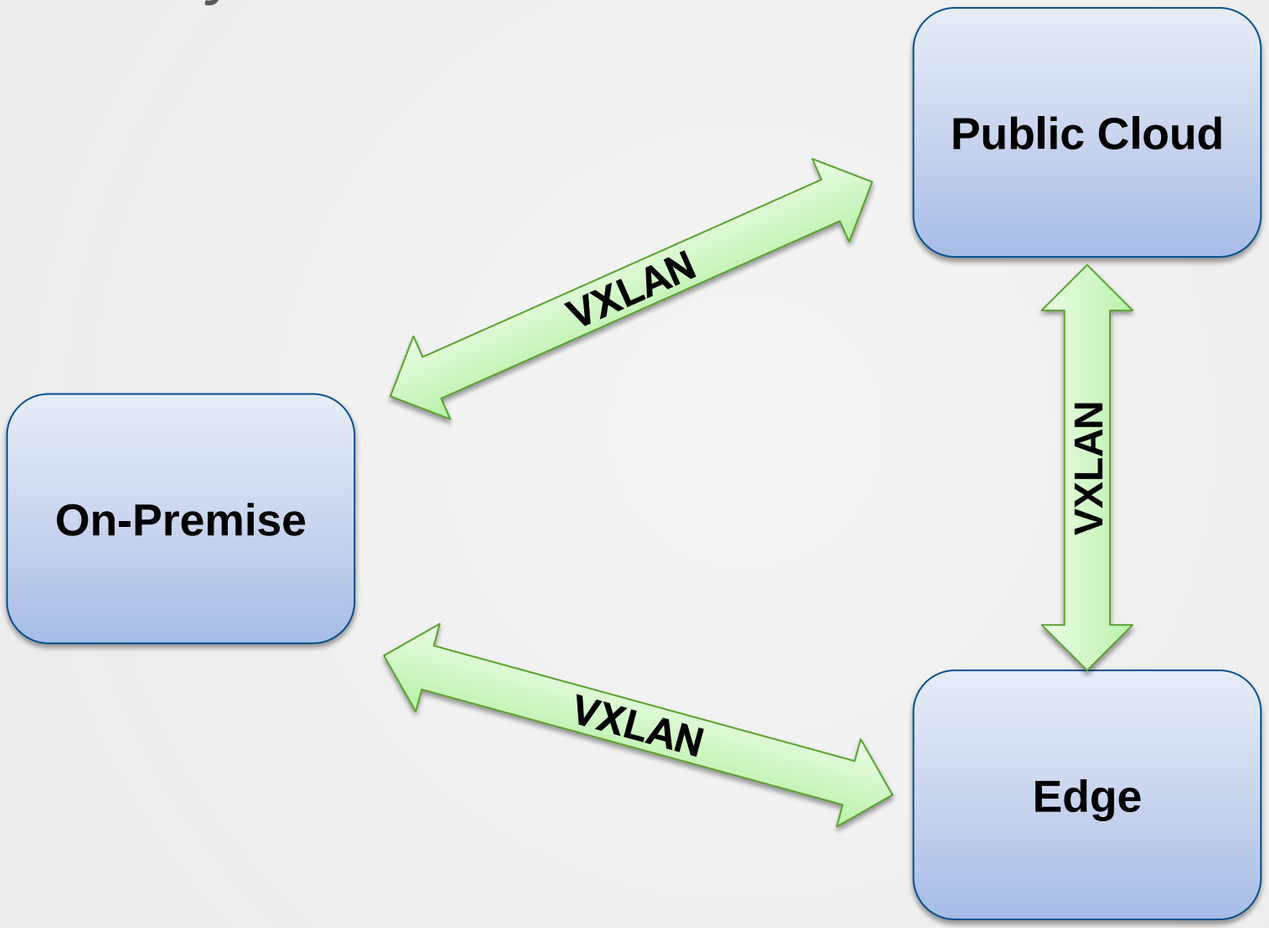

- ***Enterprise Edge***
- ***Media Edge***
- ***IOT Industrial Edge***
- ***Mobile Edge***

Edge Computing Projection

By 2020:

50 Billion

Devices Connected to the Internet

50% of Data

Will be Processed at Network Edge

according to IDC

$1.7 Billion

2015

$6.3 Billion

2020

Micro Data Center Market

(just facilities - does not include compute, network, storage)

according to MarketsAndMarkets

What is Driving the Growth?

Events

Disasters

Transportation

Edge / Micro Data Centers

Colocation

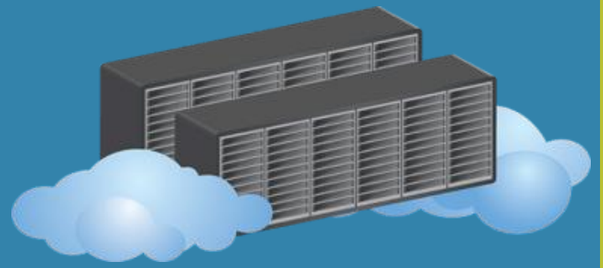

Hybrid Cloud / ROBO

IoT: Industrial

IoT: Smart Cities

Mobile Edge

Virtual Reality

Autonomous Vehicles

What is a Micro Data Center?

- Small facility placing compute & storage close to users

- Under 1 MW Power
  - Sometimes under 100 KW

- Locations
  - Cell towers
  - Central Offices
  - On-premise / business park / sports arena
  - On Transportation – ships, airplanes, submarines

- Challenges
  - Power, Cooling
  - Limited Space
  - Rapid/Remote Deployment, Zero Touch
  - Ongoing operations/monitoring

On Premise

On Transportation

Wherever You Want

Cell Sites

Networks Optimized for Edge Computing

SWaP - Size, Weight, and Power Optimized for Micro Data Centers

- Performance
- High Availability
- Simple
- Automated
- Scalable
- Cost Efficient

Spectrum™

- ✓ 2 Switches in 1RU
- ✓ Ultra Low Power <90 watts
- ✓ Highly Scalable Data Center Interconnect (DCI)
- ✓ Zero Packet Loss & Low Latency
- ✓ RoCE optimized switches for NVMe-oF & Machine Learning
- ✓ Zero Touch Provisioning
- ✓ Network Visibility & Telemetry
- ✓ Cost optimized

Protocols vs Telemetry

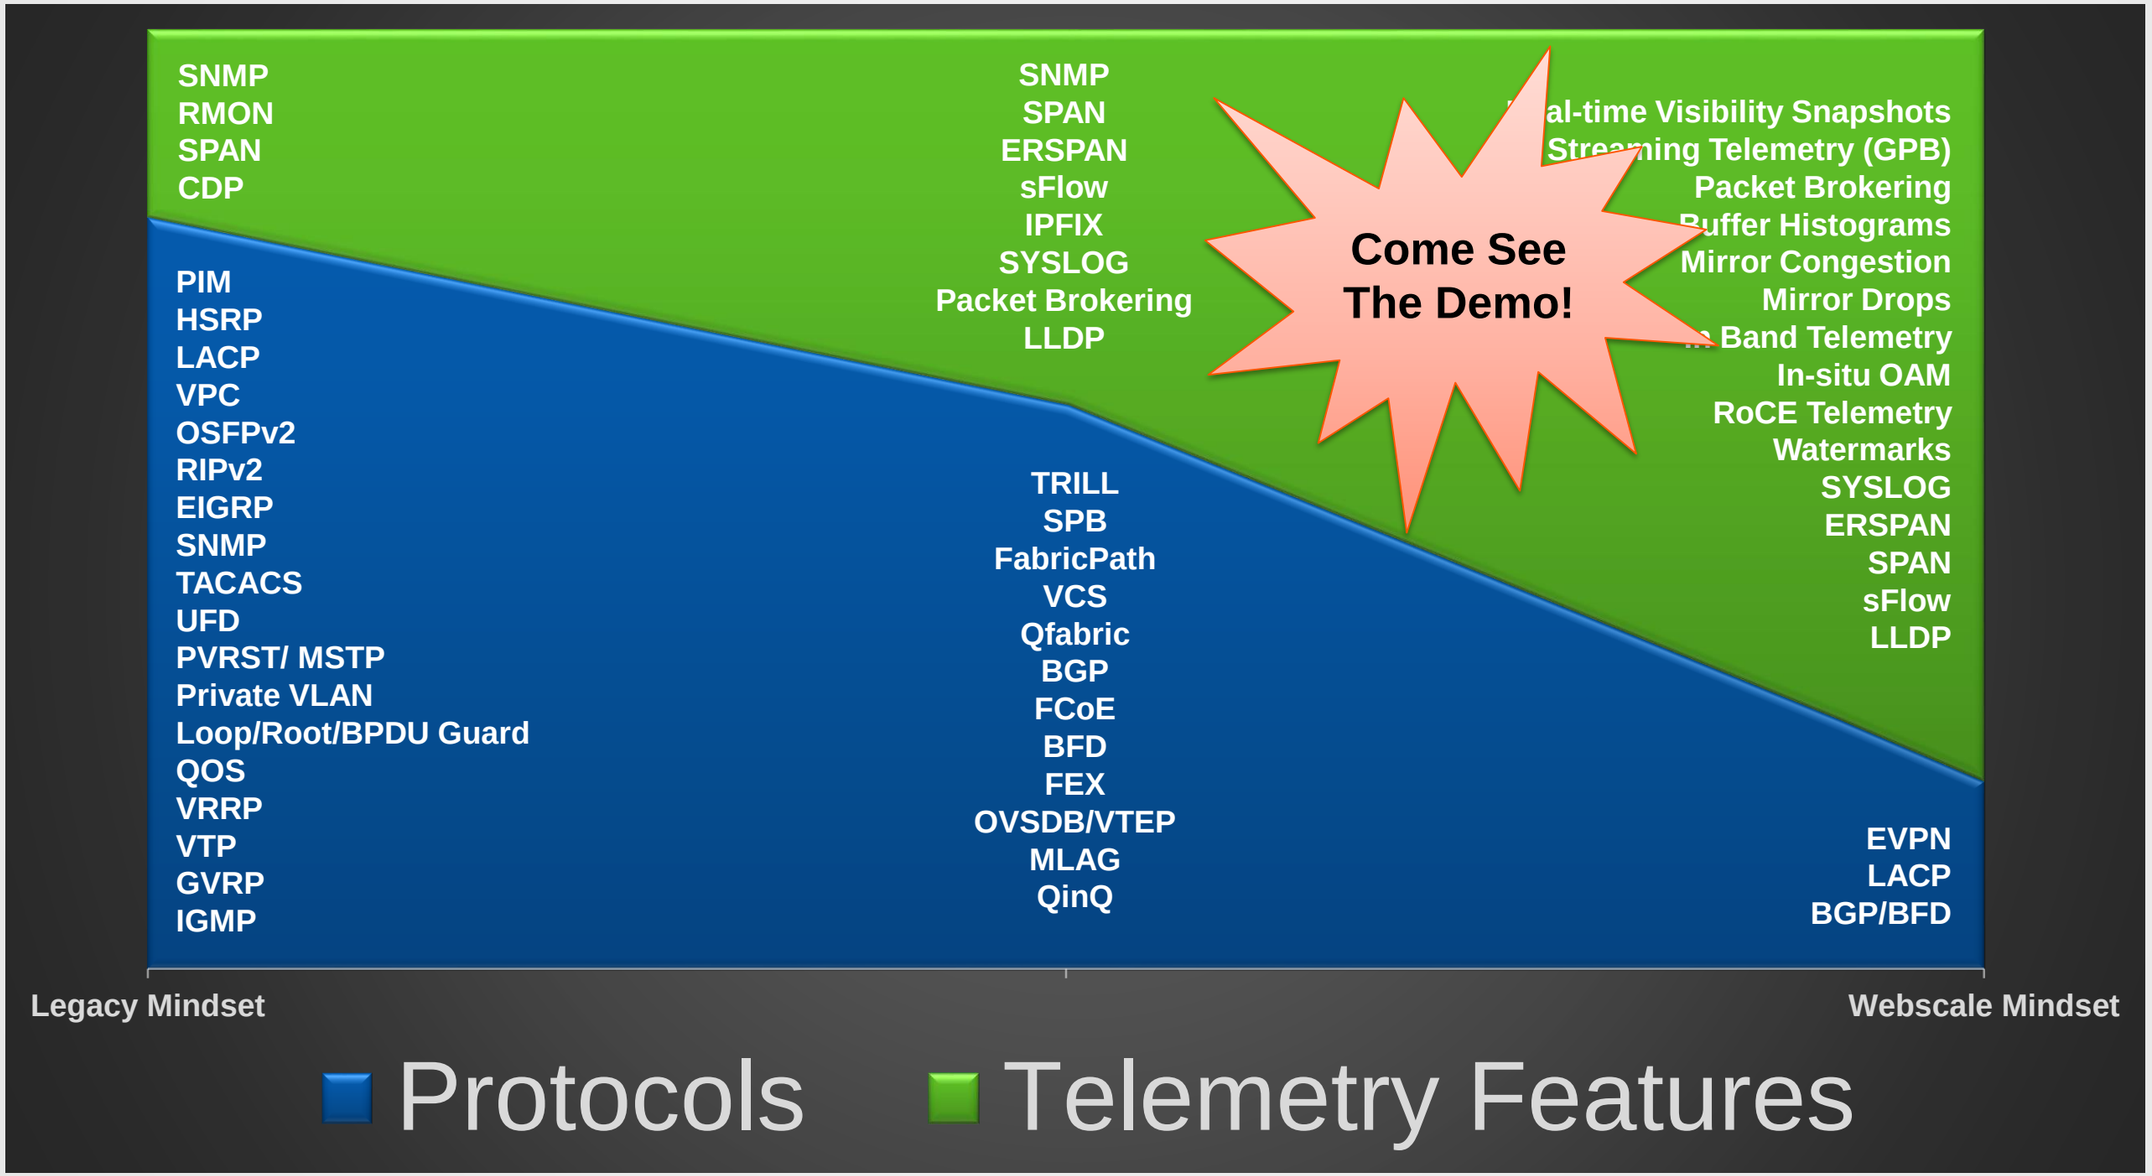

One Unifying Feature: Telemetry

- Exhaustive Packet Drop Counters
  - No mysterious uncounted drops
- In-ASIC Switch Monitoring Tool
  - Bandwidth & queue depth over time
  - Microsecond sampling
  - Hardware generated histograms
    - Per port bandwidth (tx/rx)
    - Per port delays from flow control
    - Per queue monitoring

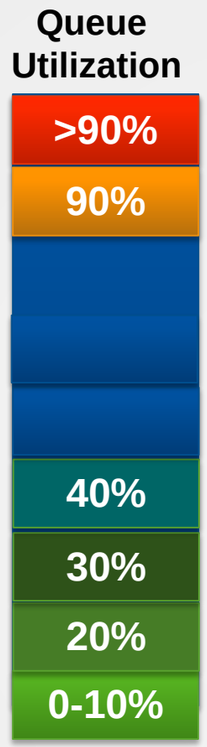

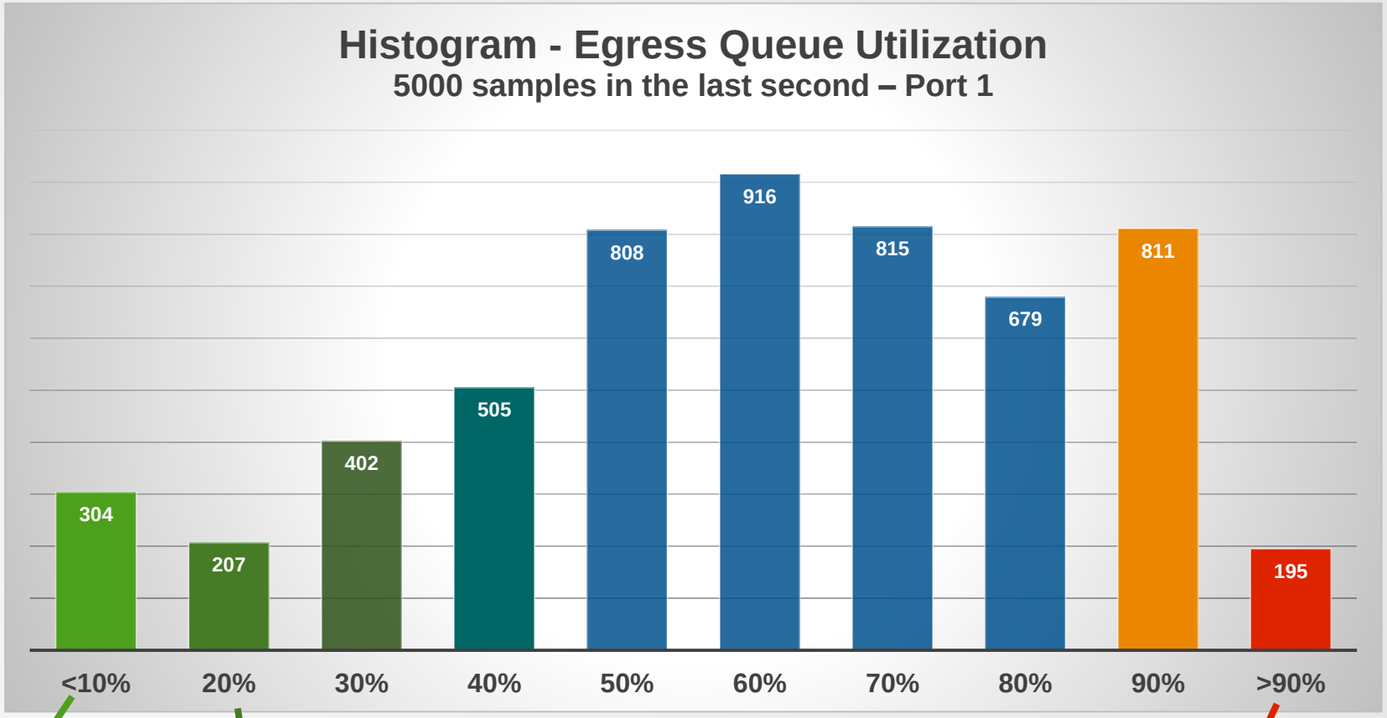

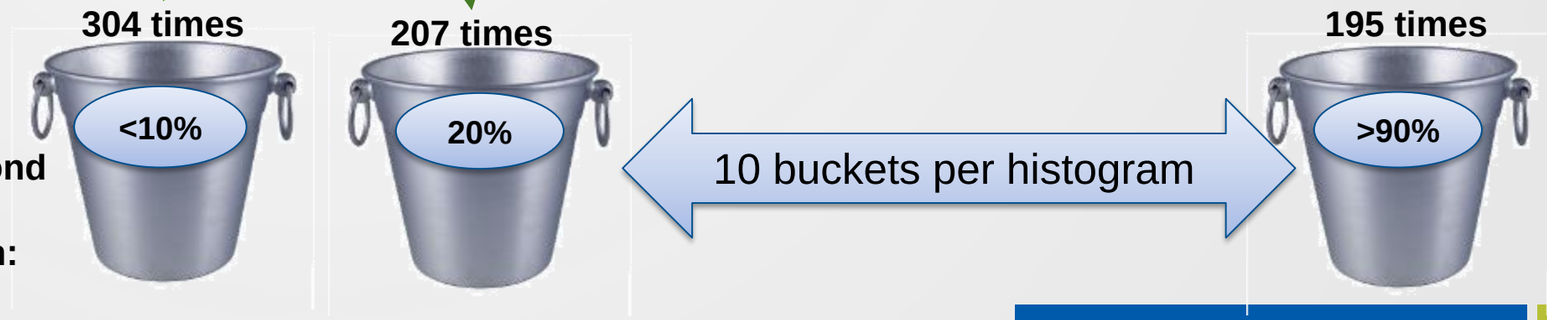

Hardware checked the queue 5000 times in the last second

Here's how many times the queue was at 10% utilization:

Histograms – a Tool for Analysis

They are not just for shopping online

In Band Telemetry

- In-band advanced notifications
  - Monitor network, Detect anomalies
  - IP/TCP/GENEVE/VXLAN options
  - Notify about
    - Switch ID
    - Egress port
    - Egress queue
    - Queuing bytes
    - Queuing latency
    - Time stamp
- Passive in-band telemetry
  - DSCP coloring according to ingress time
  - Count drops on specific flows across fabric
  - Monitor end-to-end high-latency

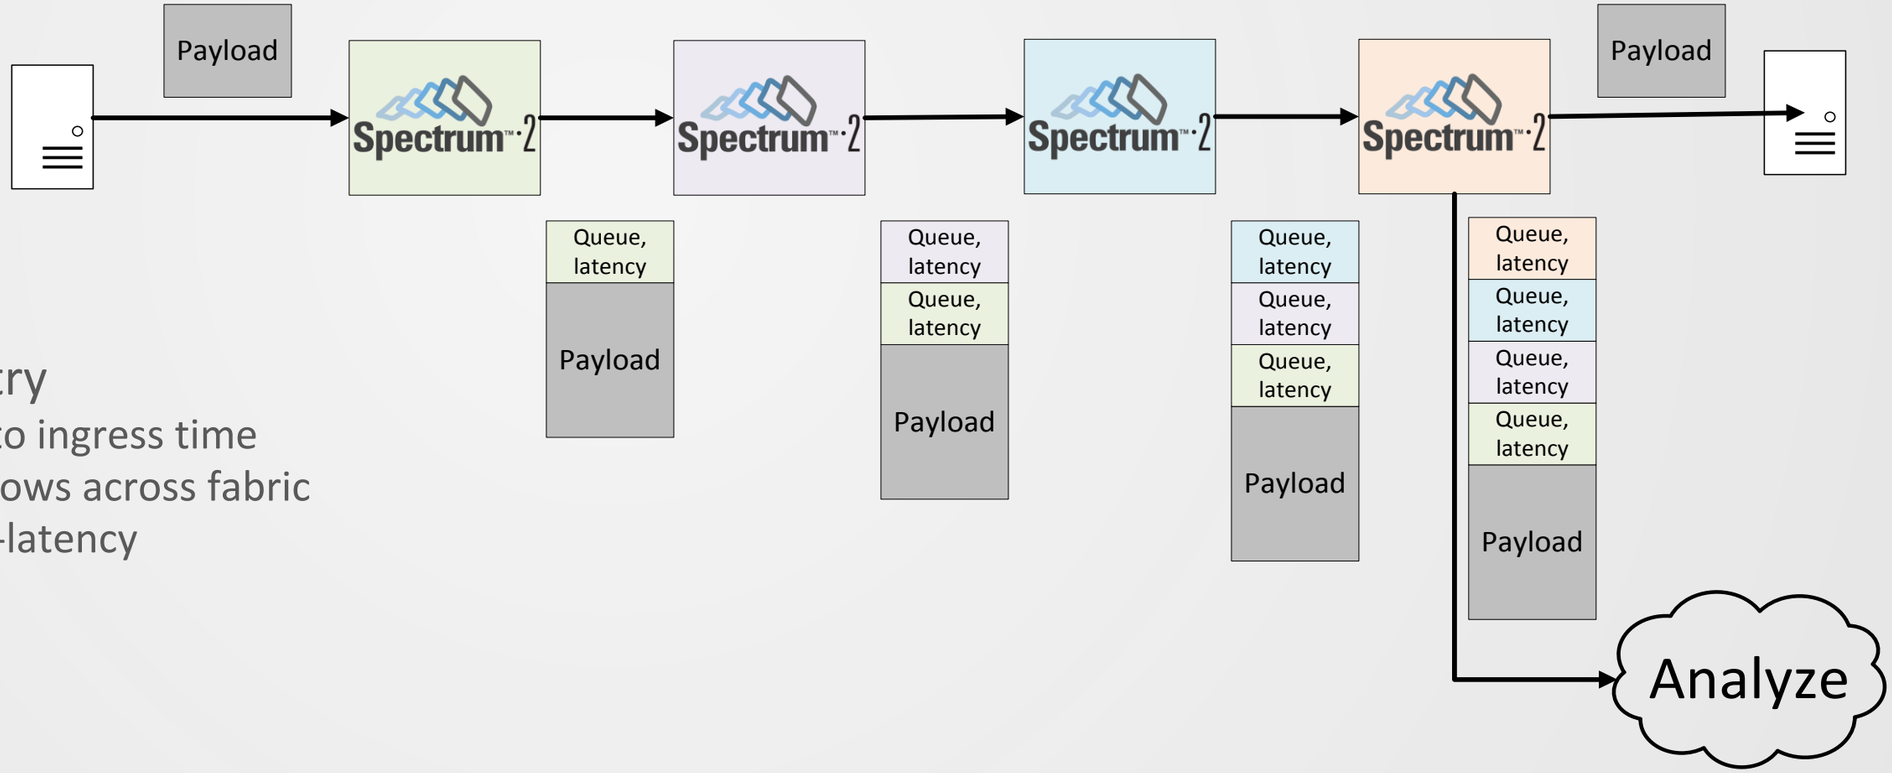

Thank You