



OCP SUMMIT

March 20-21
2018
San Jose, CA

OPEN. FOR BUSINESS.



Switchgear Innovations w/ Cavium-XPliant®

Albert Fishman, Solutions Product Marketing Manager

OPEN. FOR BUSINESS.



Cavium-XPliant Programmable Switches

3 Switch Silicon Generations

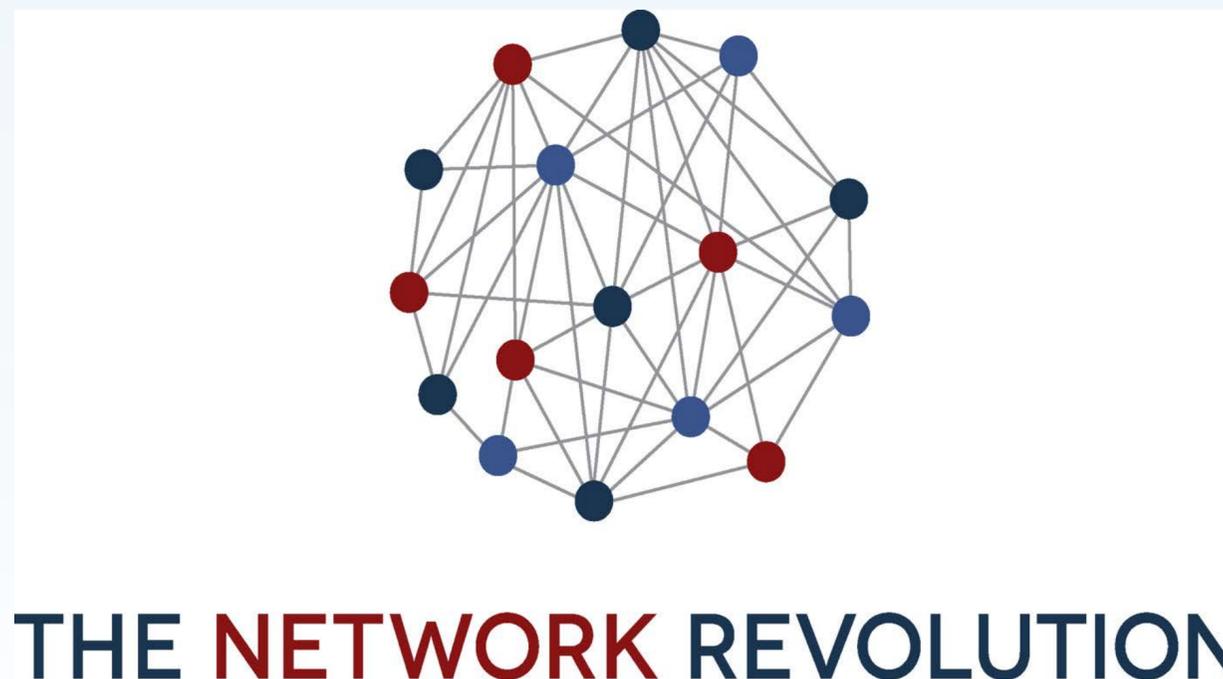
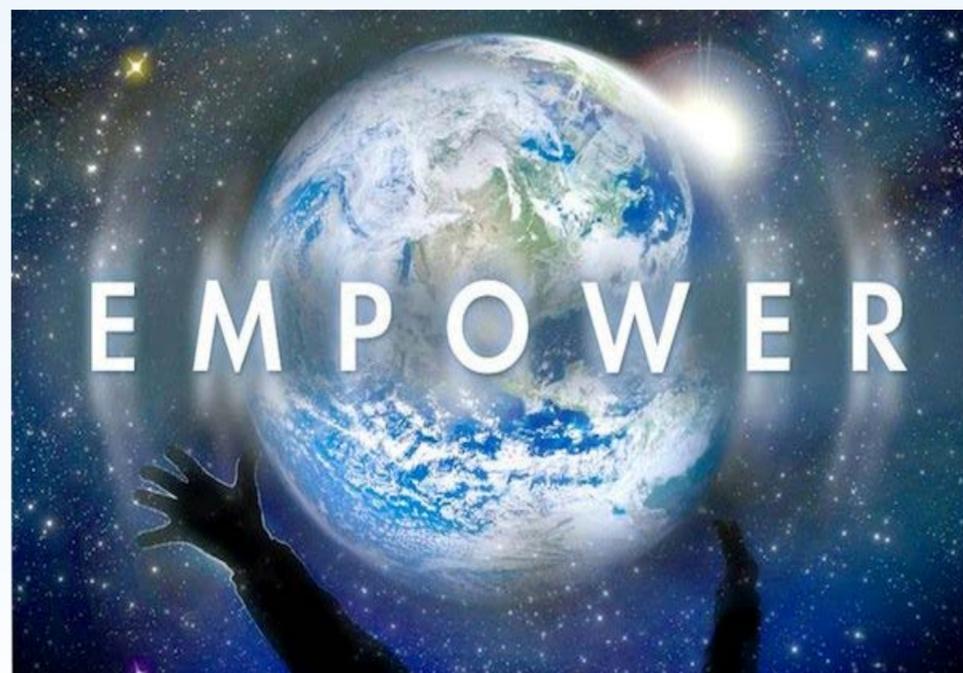
Field **Proven** Solutions

Legacy **Interoperability**

Programmable Data Plane

Leading **OEM** Products

White-Box Options



Cavium XPliant & the OCP Stack



OCP Approved SW
Available w/ XPliant
Programmable Switch



OCP HW Spec/Design
Available w/ XPliant
Programmable Switch



Introducing Packet Trakker™: XPliant Telemetry

Albert Fishman, Solutions Product Marketing Manager

Albert.Fishman@cavium.com

OPEN. FOR BUSINESS.





Packet Trakker: XPliant Programmable Telemetry

ENTER

Packet Trakker: Cavium Programmable Telemetry Suite

Physical Networks
Scale Out

Data Center
Networks

Hyper
Converged
Networks



Telecom 5G
Networks

Enterprise
Networks

Packet Trakker: Cavium Programmable Telemetry Suite

More Services with
Different Demands

Data Center
Networks

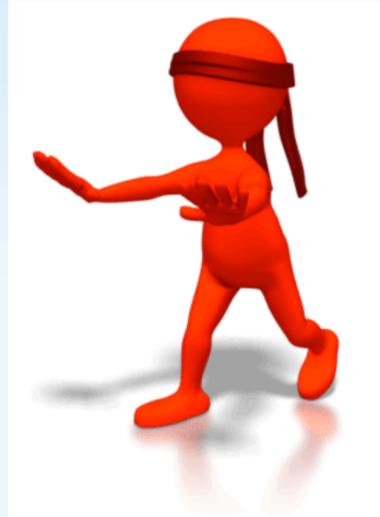
Hyper
Converged
Networks



Telecom 5G
Networks

Enterprise
Networks

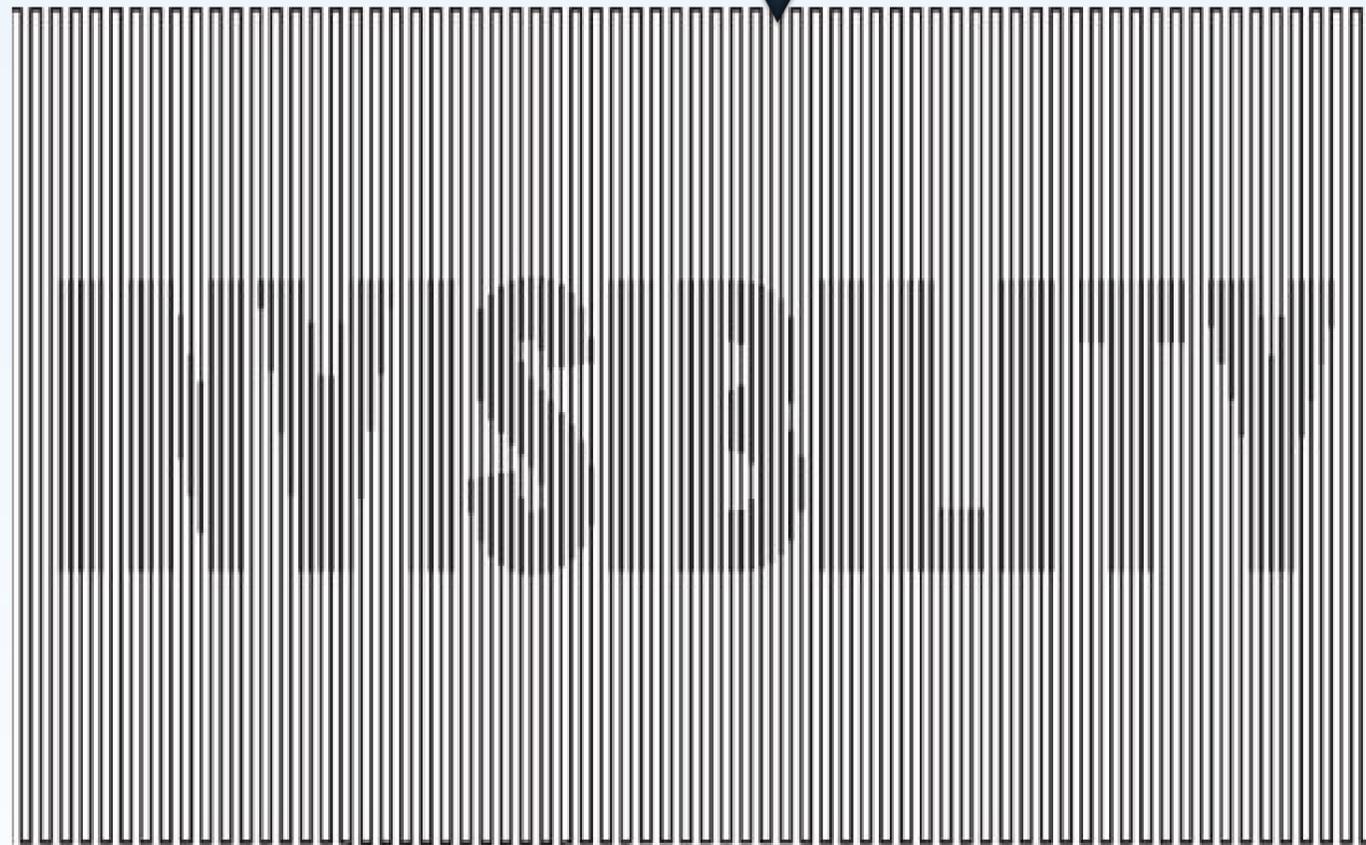
Packet Trakker: Cavium Programmable Telemetry Suite



Management Challenges...

Quality of Experience

Applications Performance



Networks Efficiency

Networks Reliability

Packet Trakker: Cavium Programmable Telemetry Suite

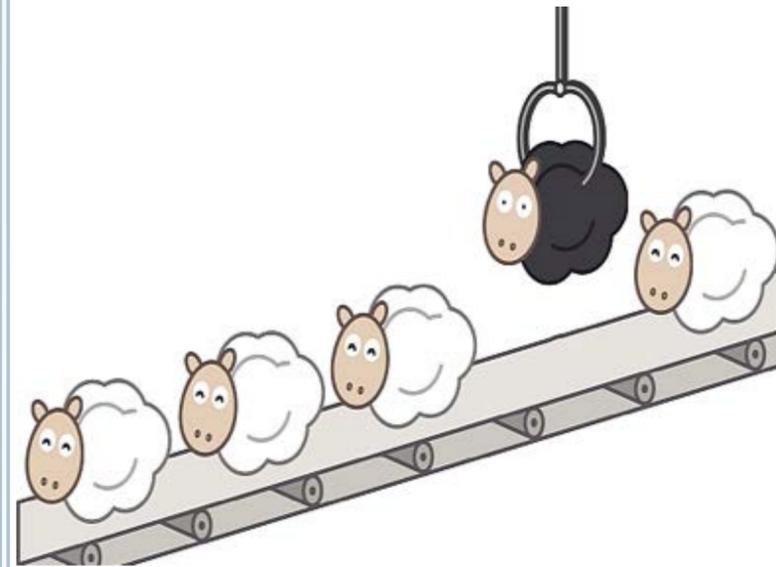


Comprehensive Suite for
Network Observability,
Analytics & Telemetry on
XPliant Switches

Monitor



Detect



React



Packet Trakker: Cavium Programmable Telemetry Suite

Applications Quality of Experience

Microburst Detection

Latency Fluctuations

Early Congestion Alarming

Network Reliability and Efficiency

Paths Tracking Metrics

Exceptions Detection

Troubleshooting
Observability

New SDN Use Cases Will Emerge ...

Packet Trakker: Cavium Programmable Telemetry Suite



Legacy Solutions Pull Model Based on Polling & Probing

- **NOT Event-Triggered**
 - Events occur between the sample or transient in nature can be missed
- **NOT Real-Time**
 - polling done every few seconds is eternity (~150M packets can pass 100G single port in 1 sec)
- **NOT Granular**
 - Unavailability of complete telemetry data, in per flow resolution, makes difficult to measure network performance
- **Limited Hardware Assistance**
 - Software resources are overwhelmed with data extraction which leads to inadequate data resolution and scale

Packet Trakker: Cavium Programmable Telemetry Suite

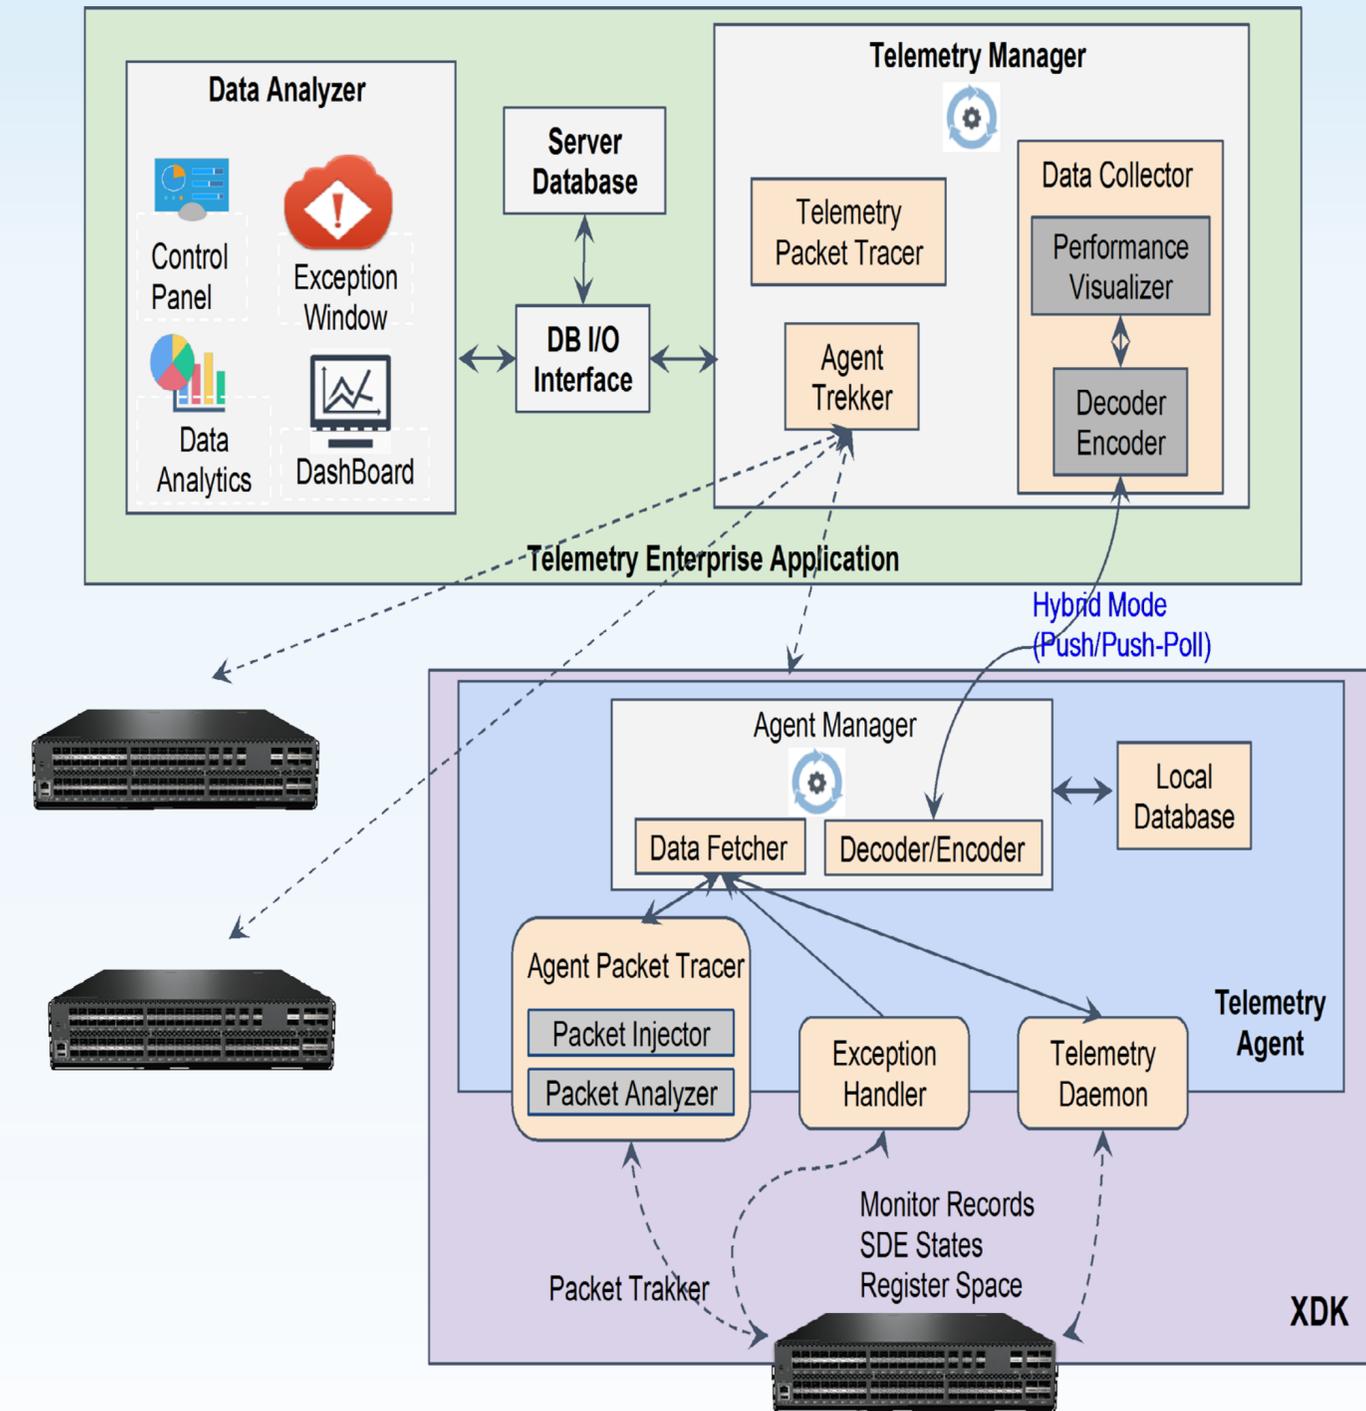
Cavium-Xpliant Push Model
Event-Triggered in Real-Time



- **Event-Triggered**
 - No important or transient event is missed
- **Real-Time**
 - Real-time monitoring and alarming
- **Granular**
 - Complete telemetry data availability, per flow, in ns resolution
- **Hardware Assistance**
 - Telemetry data extraction to Analytics tools with no Host CPU Software Intervention

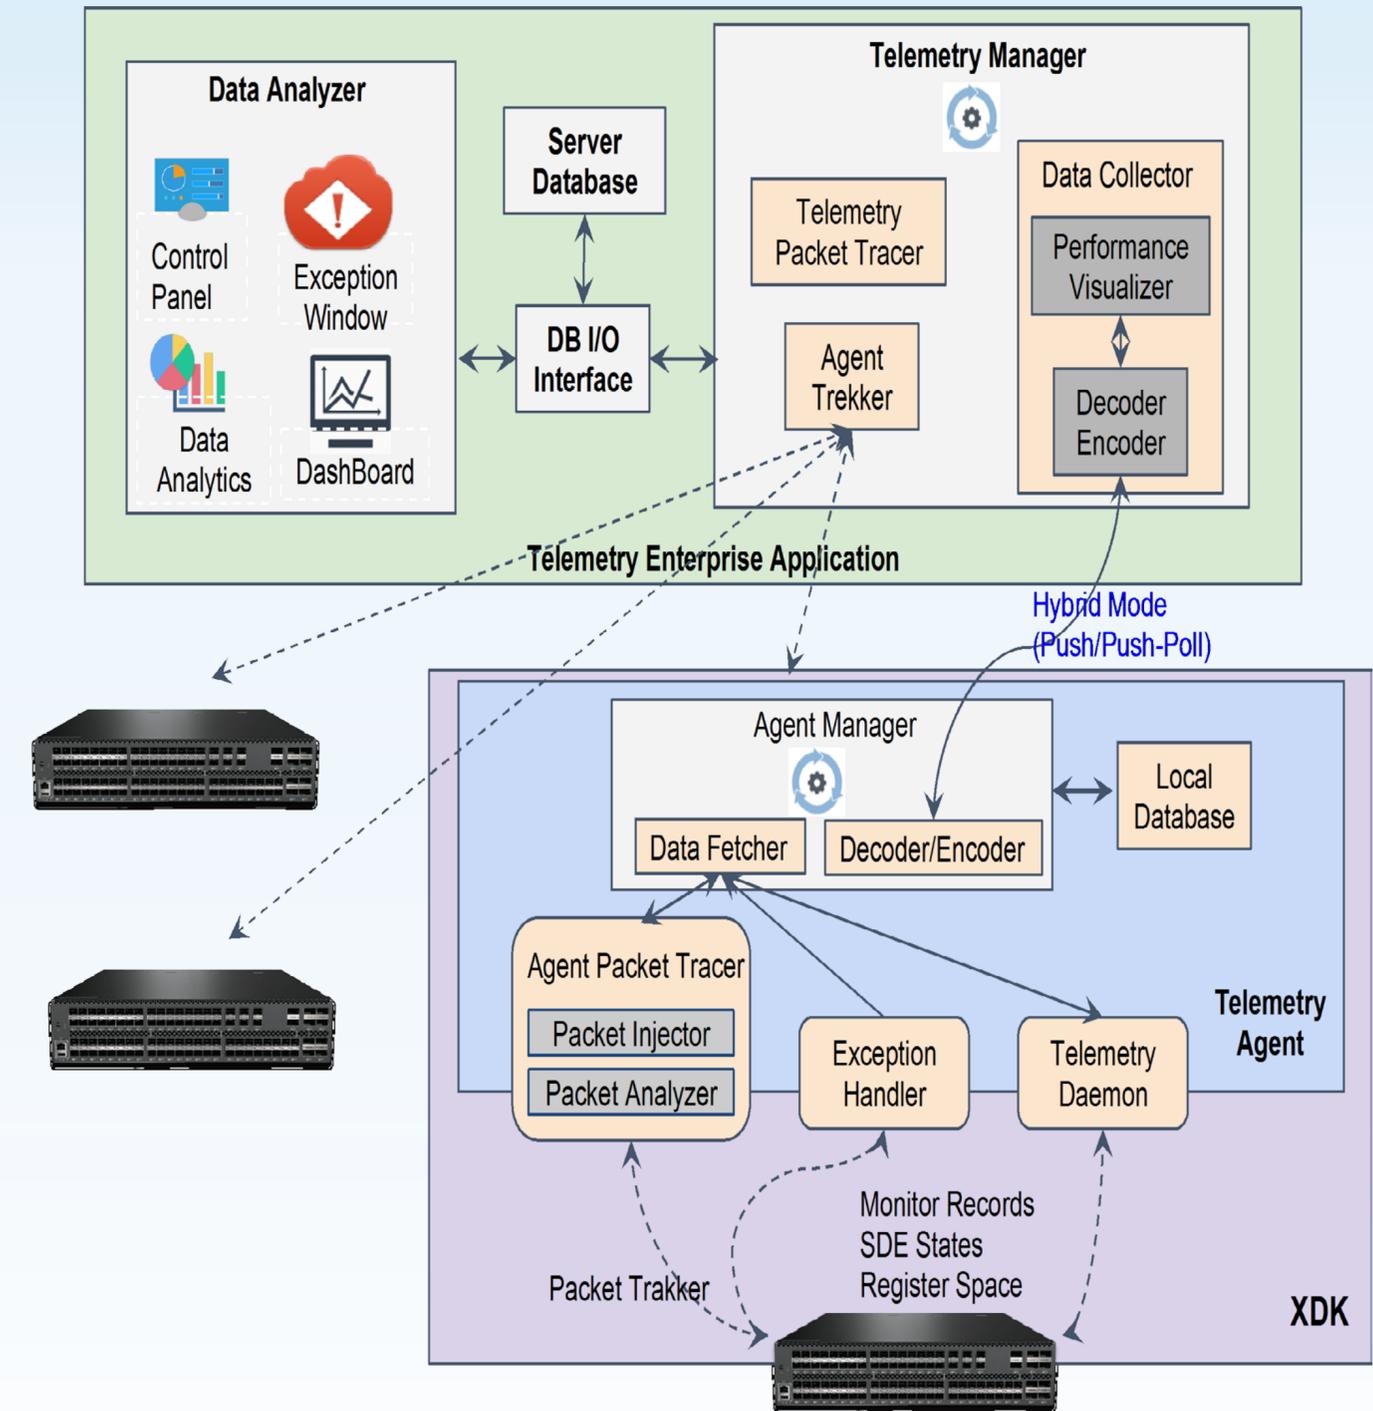
Packet Trakker: Cavium Programmable Telemetry Suite

- **Packet Trakker is**
 - A Comprehensive Suite for Network
 - Observability, Analytics & Telemetry
 - Supports In-Band & Out-Of-Band Telemetry
- **Software Programmable Telemetry on XPliant**
 - HW Flexibility is Leveraged via Programmable SW
 - Hardware Support In Real-Time With ns Resolution
 - Software-Definable States, Counters, Data Export
- **Front-End Software**
 - Data Representation, Visualization & Analytics
 - Configurations & Scriptable Events Response



Packet Trakker: Cavium Programmable Telemetry Suite

- **Supported on ALL XPliant Switches Generations**
 - Portable to XPliant-based White-Box Hardware
 - Including OCP Hardware Platforms w/ XPliant
- **NOS Integration-Ready**
 - By Porting Directly Over XDK or via SAI
 - Including Integrations w/ SONiC & OPX
- **Designed for Extensibility**
 - Modular Design Based on Kafka DB & gRPC
 - Can be Integrated w/External Analytics Tools

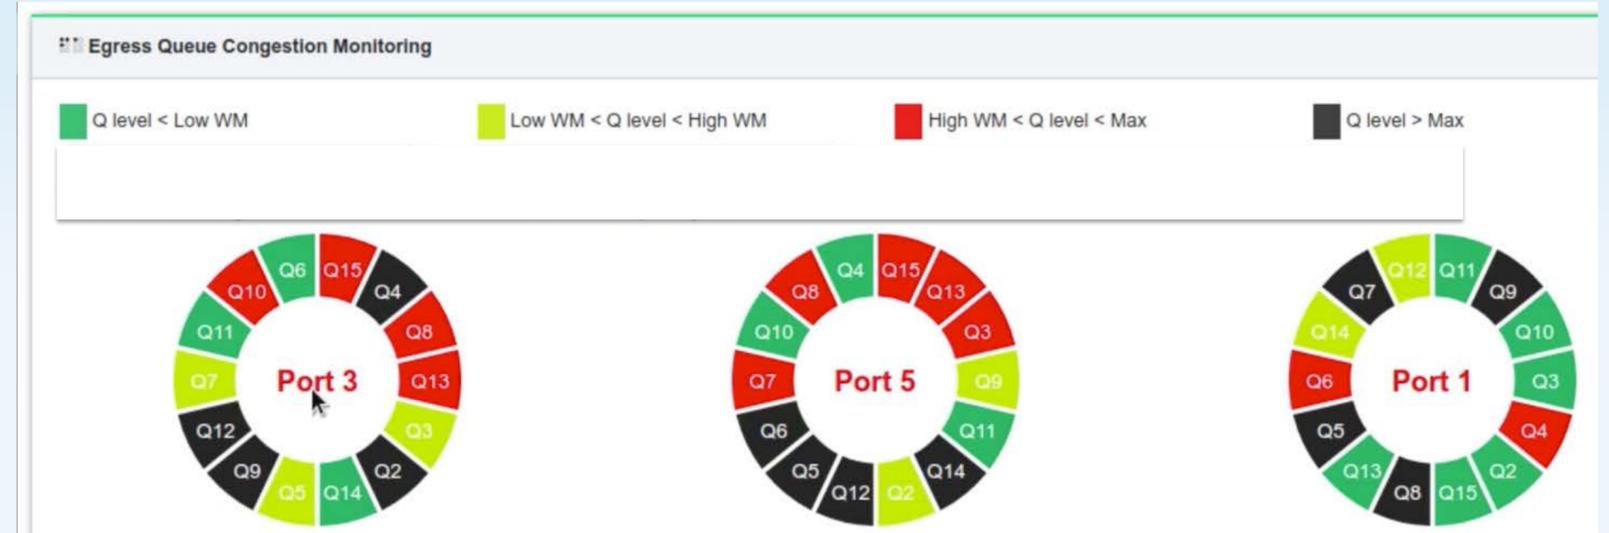
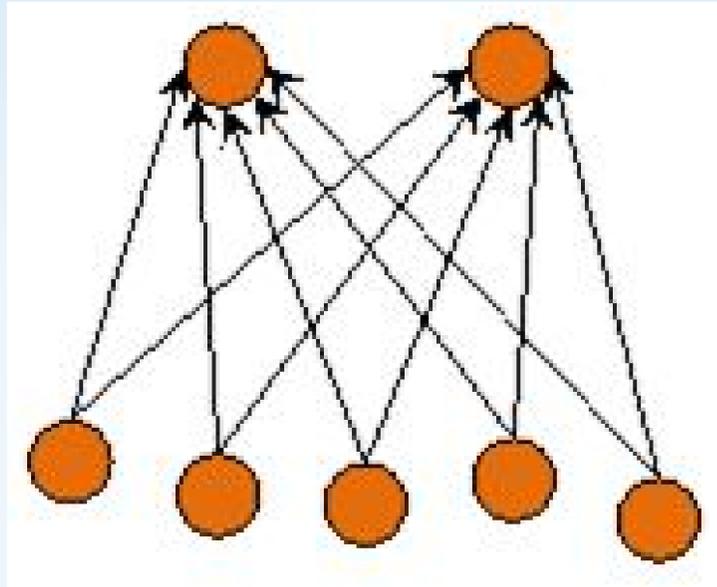




Packet Trakker: XPliant Programmable Telemetry

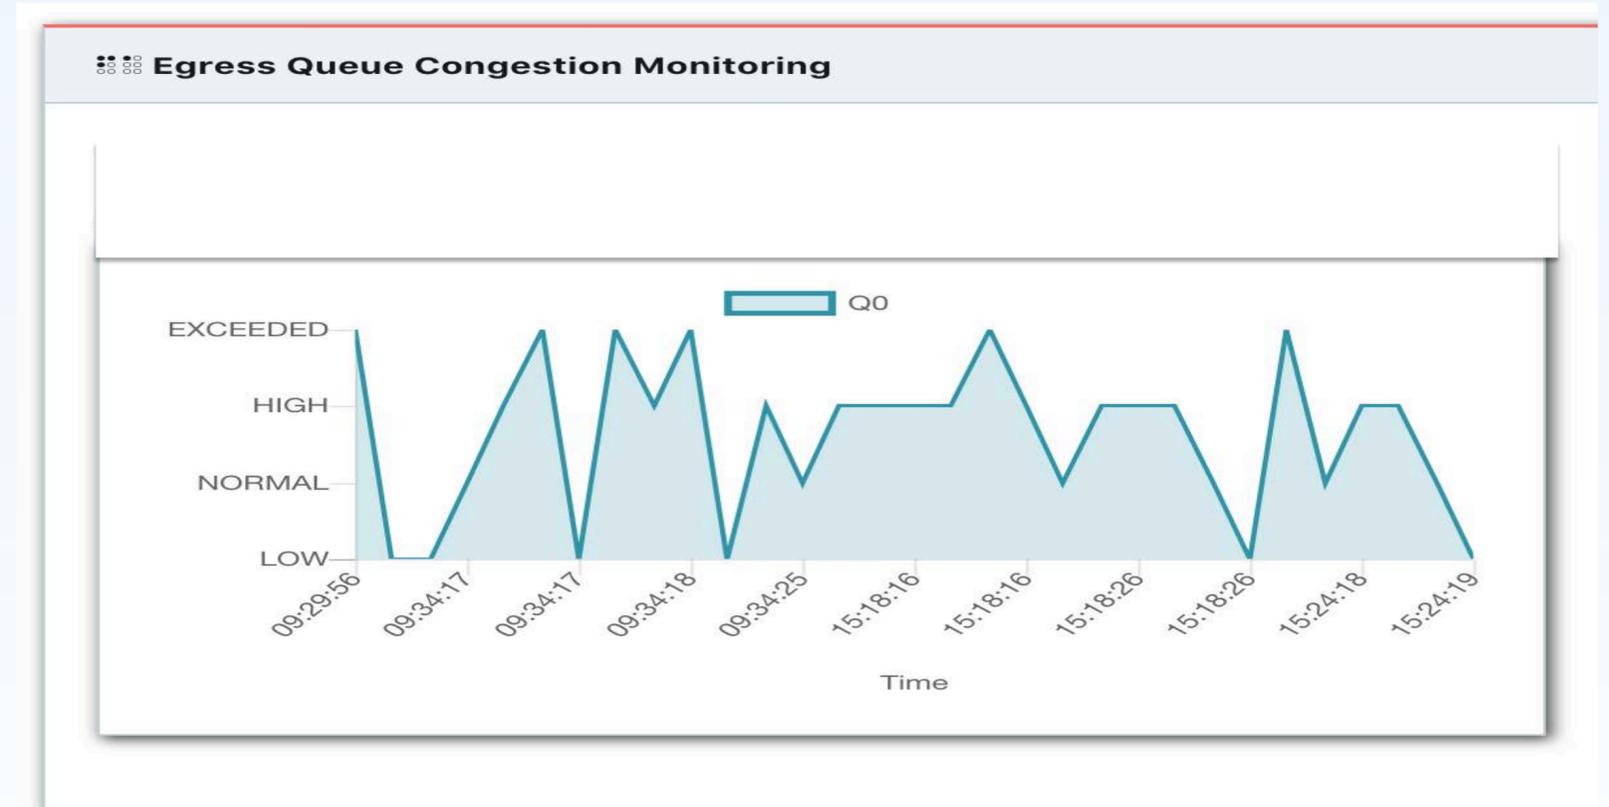
ENTER

Quality of Experience (QoE): Early Congestion Alarming

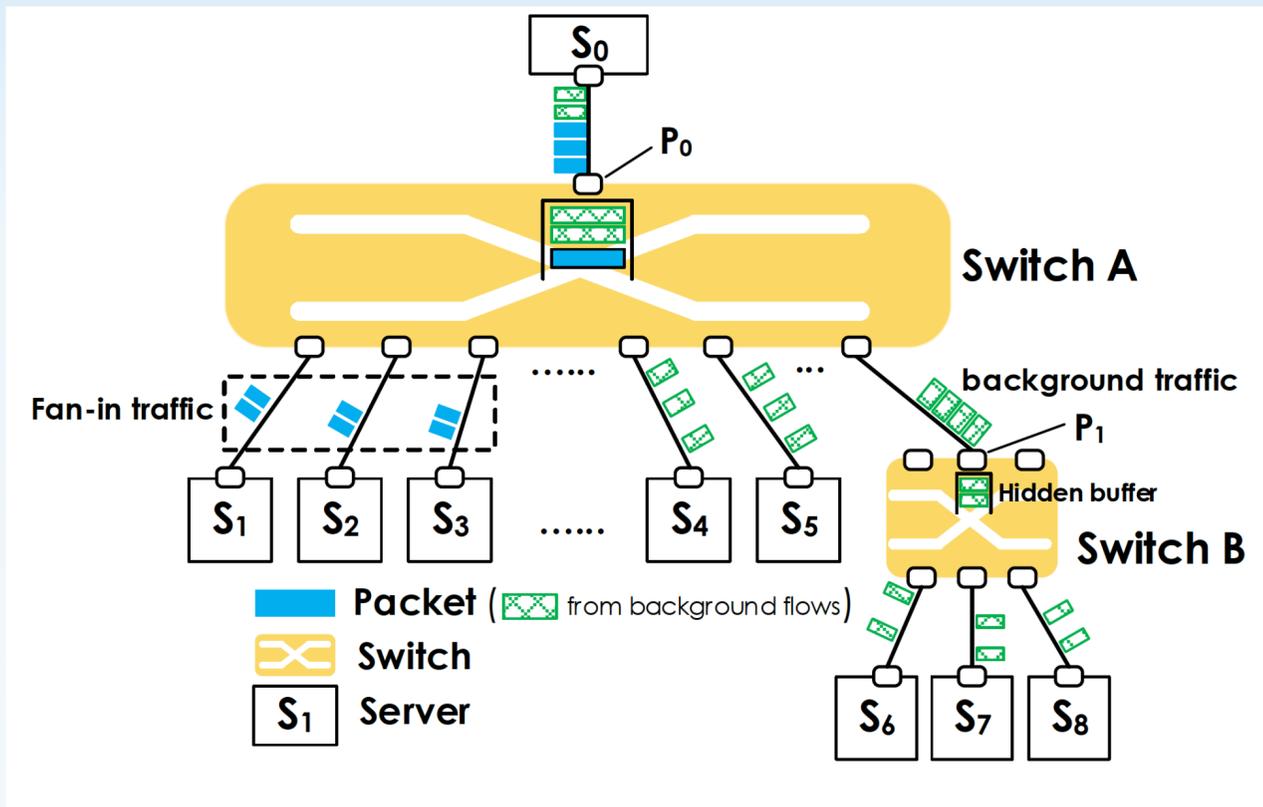


Transmit Queues Build Up → Latencies Fluctuate
→ Leads to Applications Performance Degradation
Can the impact be prevented/responded promptly ?

- **Egress Transmit Queues State Monitoring**
 - Queues States Periodic Reports
 - Queues States Change Alerting
- **Latency Fluctuation Prevention / Handling**
 - Alert on Queues Build-Up in Real-Time
 - Assistance in Aggressors/Elephant/Mice Flows Identification



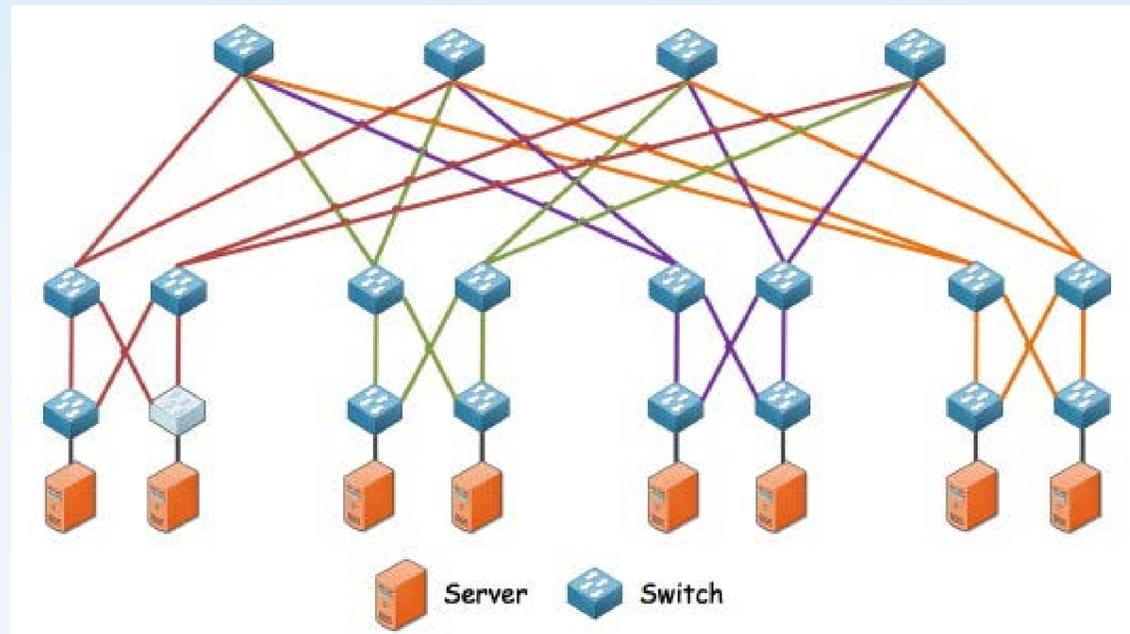
QoE: Microburst Detection & Flows Identification Assistance



- **70% bursts sustained for at most tens of ms**
<https://homes.cs.washington.edu/~qiao/paper/imc.pdf>
- **10% chance being affected by a slow response**
<https://engineering.linkedin.com/performance/who-moved-my-99th-percentile-latency>

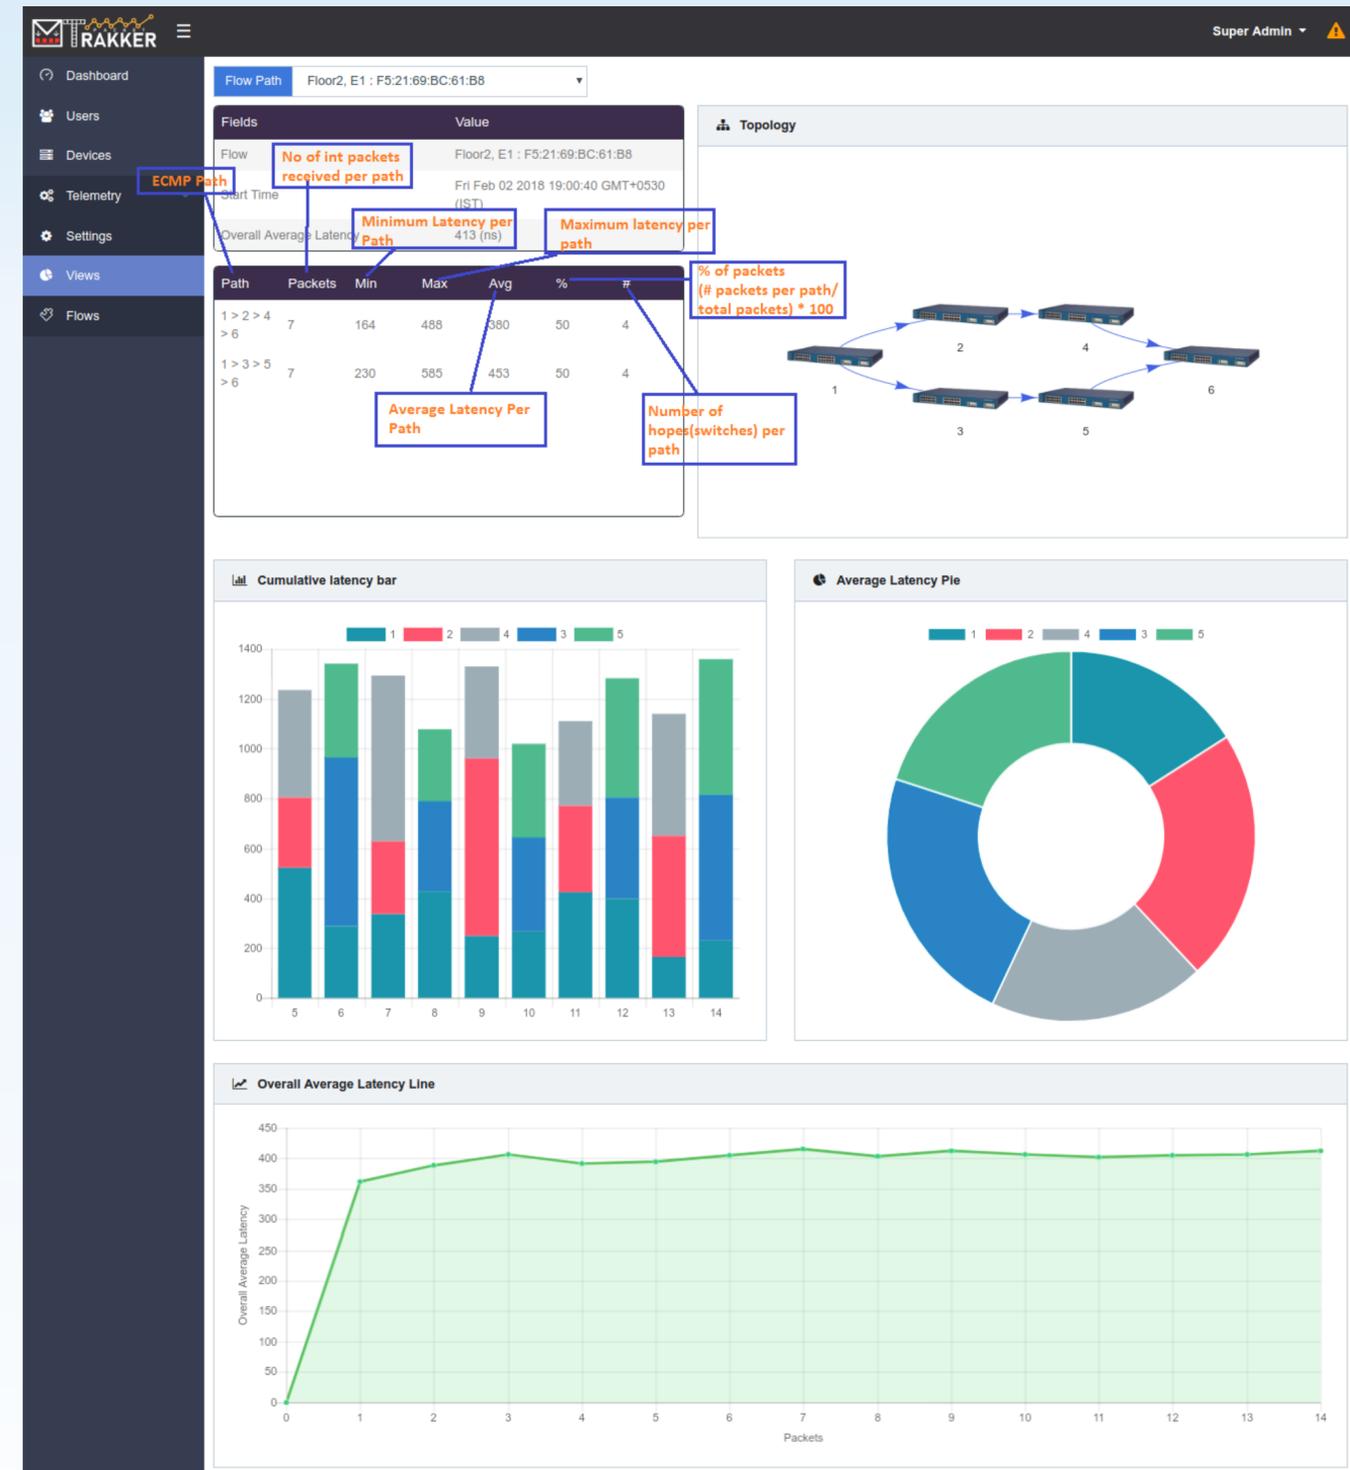
- **Microbursts Detection**
 - In real-time
 - With ns granularity
- **Flows Identification Assistance**
 - Elephant/Mice/Aggressor Flows

Network Efficiency: Latency/Distributions Metrics w/Embedded Telemetry



What ECMP Path My Flow Is Taking ? What Are the Latency / Distribution ?

- In-band (data packets) Switch States Report
 - iOAM, INT, UDP standards & practices
 - User-Defined Transport & Metadata
- Traffic Flow Path Traceability
 - Per Flow, per Path, per Hops Data
 - ECMP Distribution Anatomy
 - ECMP Latency Analytics

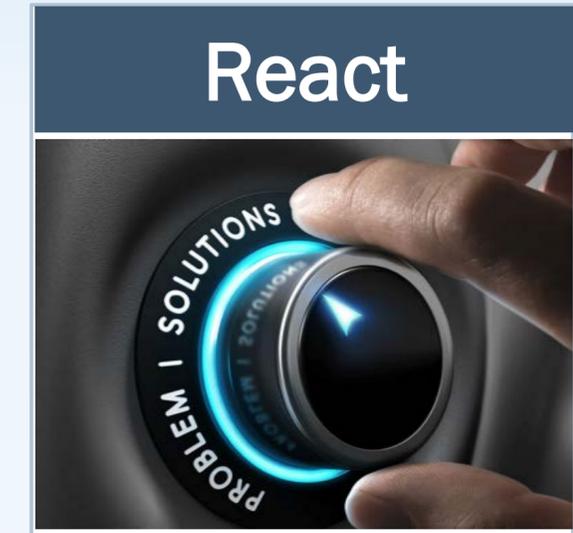
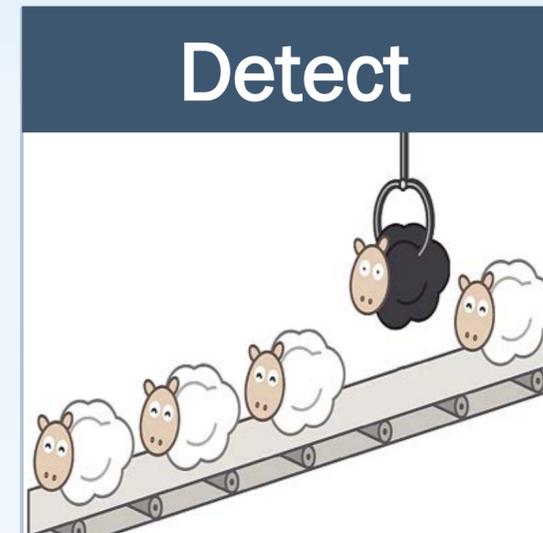


Packet Trakker : Programmable Telemetry Use Cases w/XPliant

- **Applications Quality of Experience**
 - Egress Queuing Monitoring, Microbursts Detection, Early Congestion Alarming, Latency Fluctuations
- **Network Reliability & Efficiency**
 - Misconfigurations & Malfunctioning Exceptions Detection, Black Holing,
 - Metrics For Traffic Engineering Analytics Flows Path, Distribution and Latency
 - In-Band Telemetry: iOAM/INT/Any Transport
 - Custom-Condition Mirroring, Based on Packet Headers, Metadata, Queues States



Packet Trakker : Programmable SDN Telemetry Suite



Contact us for more details:

- Albert Fishman, Solutions Product Marketing Manager
- Albert.Fishman@cavium.com



OCP SUMMIT