

ARISTA

EOS Precision Data Analyzer

Nick Ciarleglio
Nick@aristanetworks.com

Precision Instrumentation – The *Next* Pillar of SDCN

Integrated monitoring and analysis tools have not kept pace with network technology

Customers with high end environments must fly blind or deploy expensive overlay infrastructures

Tool deployments become artificially limited because aggregation networks are disproportionately expensive

The monitoring market is prime for disruption

EOS + Data Analysis Enables Scalable Precision Instrumentation

Would You Drive Blindfold?



Precision Data Analyzer Core Capabilities



Visualization Tools

Home Grown

Open

Enterprise

Core Components

Advanced
Filtering
Mirroring and
Flow Analysis

Precision
Capture and
Time Stamping

EOS Control
Plane
Event
Management
APIs

Advanced Hardware

Hardware Precision Timing

Programmable Data Plane

Precision Data Analyzer Components

Troubleshooting / Application and Network Performance Management / Security

Advanced Filtering
Mirroring and Flow
Analysis

Precision Capture and
Timing

EOS Control Plane

TAP Aggregation

Filtered Multiport Mirroring

Packet Truncation Hardware
Time stamping

LANZ+
Monitor and Capture

Hardware Time Stamping

PTP 1588
Precision Timing

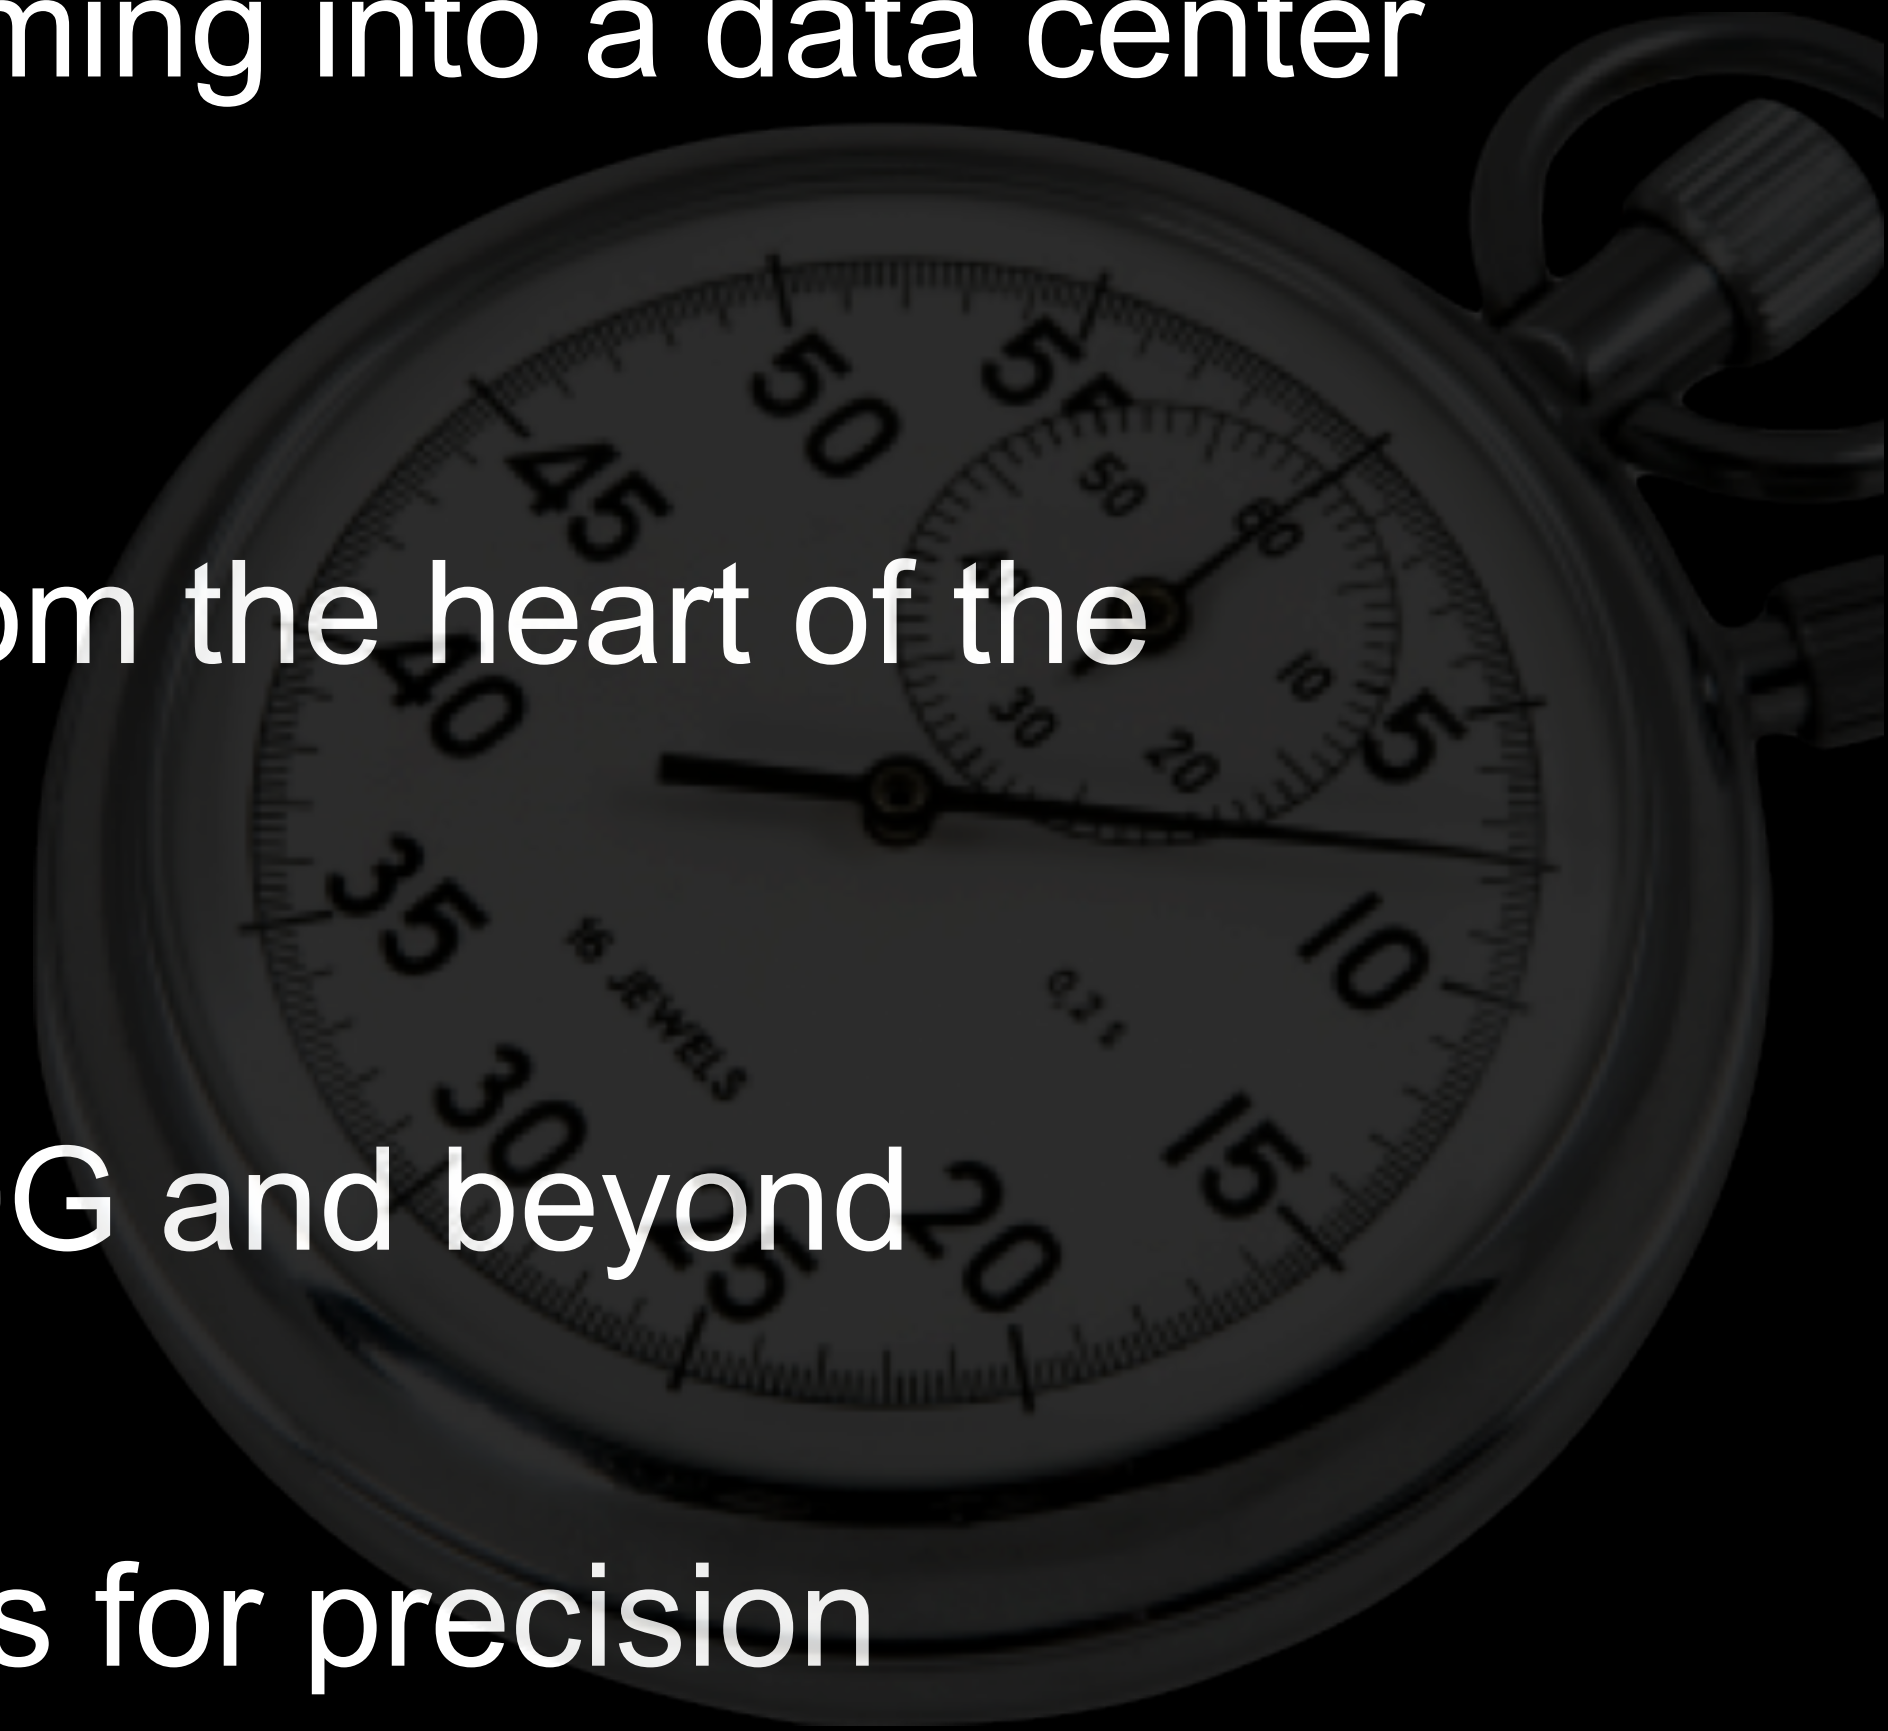
EOS Software Defined
Control Plane

Advanced Event
Management

Open API
Linux Extensibility

Hardware Precision Timing

- Arista first to integrate precision timing into a data center platform
- Enables advanced functionality from the heart of the network
- Supports high densities of 1/10/40G and beyond
- Dramatically lowers the entry costs for precision instrumentation



Advanced Mirroring

- Enhancing standard mirroring features to support flexible monitoring without expensive overlay hardware
- Enable on-box capture to leverage EOS & SSD
- Leverage hardware timing to accurately track packets without the need for 3rd party apps

Advanced Mirroring Feature Overview

- Multi-destination mirroring
 - Replicate traffic to multiple capture/tools
- Advanced load sharing
 - Symmetric traffic distribution
 - Intelligent Load sharing
- Mirror direct to EOS
 - Troubleshoot dataplane activity using TCPdump
 - Capture locally to flash or SSD
 - Integrate with Event Manager for automated captures
- Flow manipulation
 - Add hardware timestamps
 - Truncate packets
 - Filter traffic by ACL

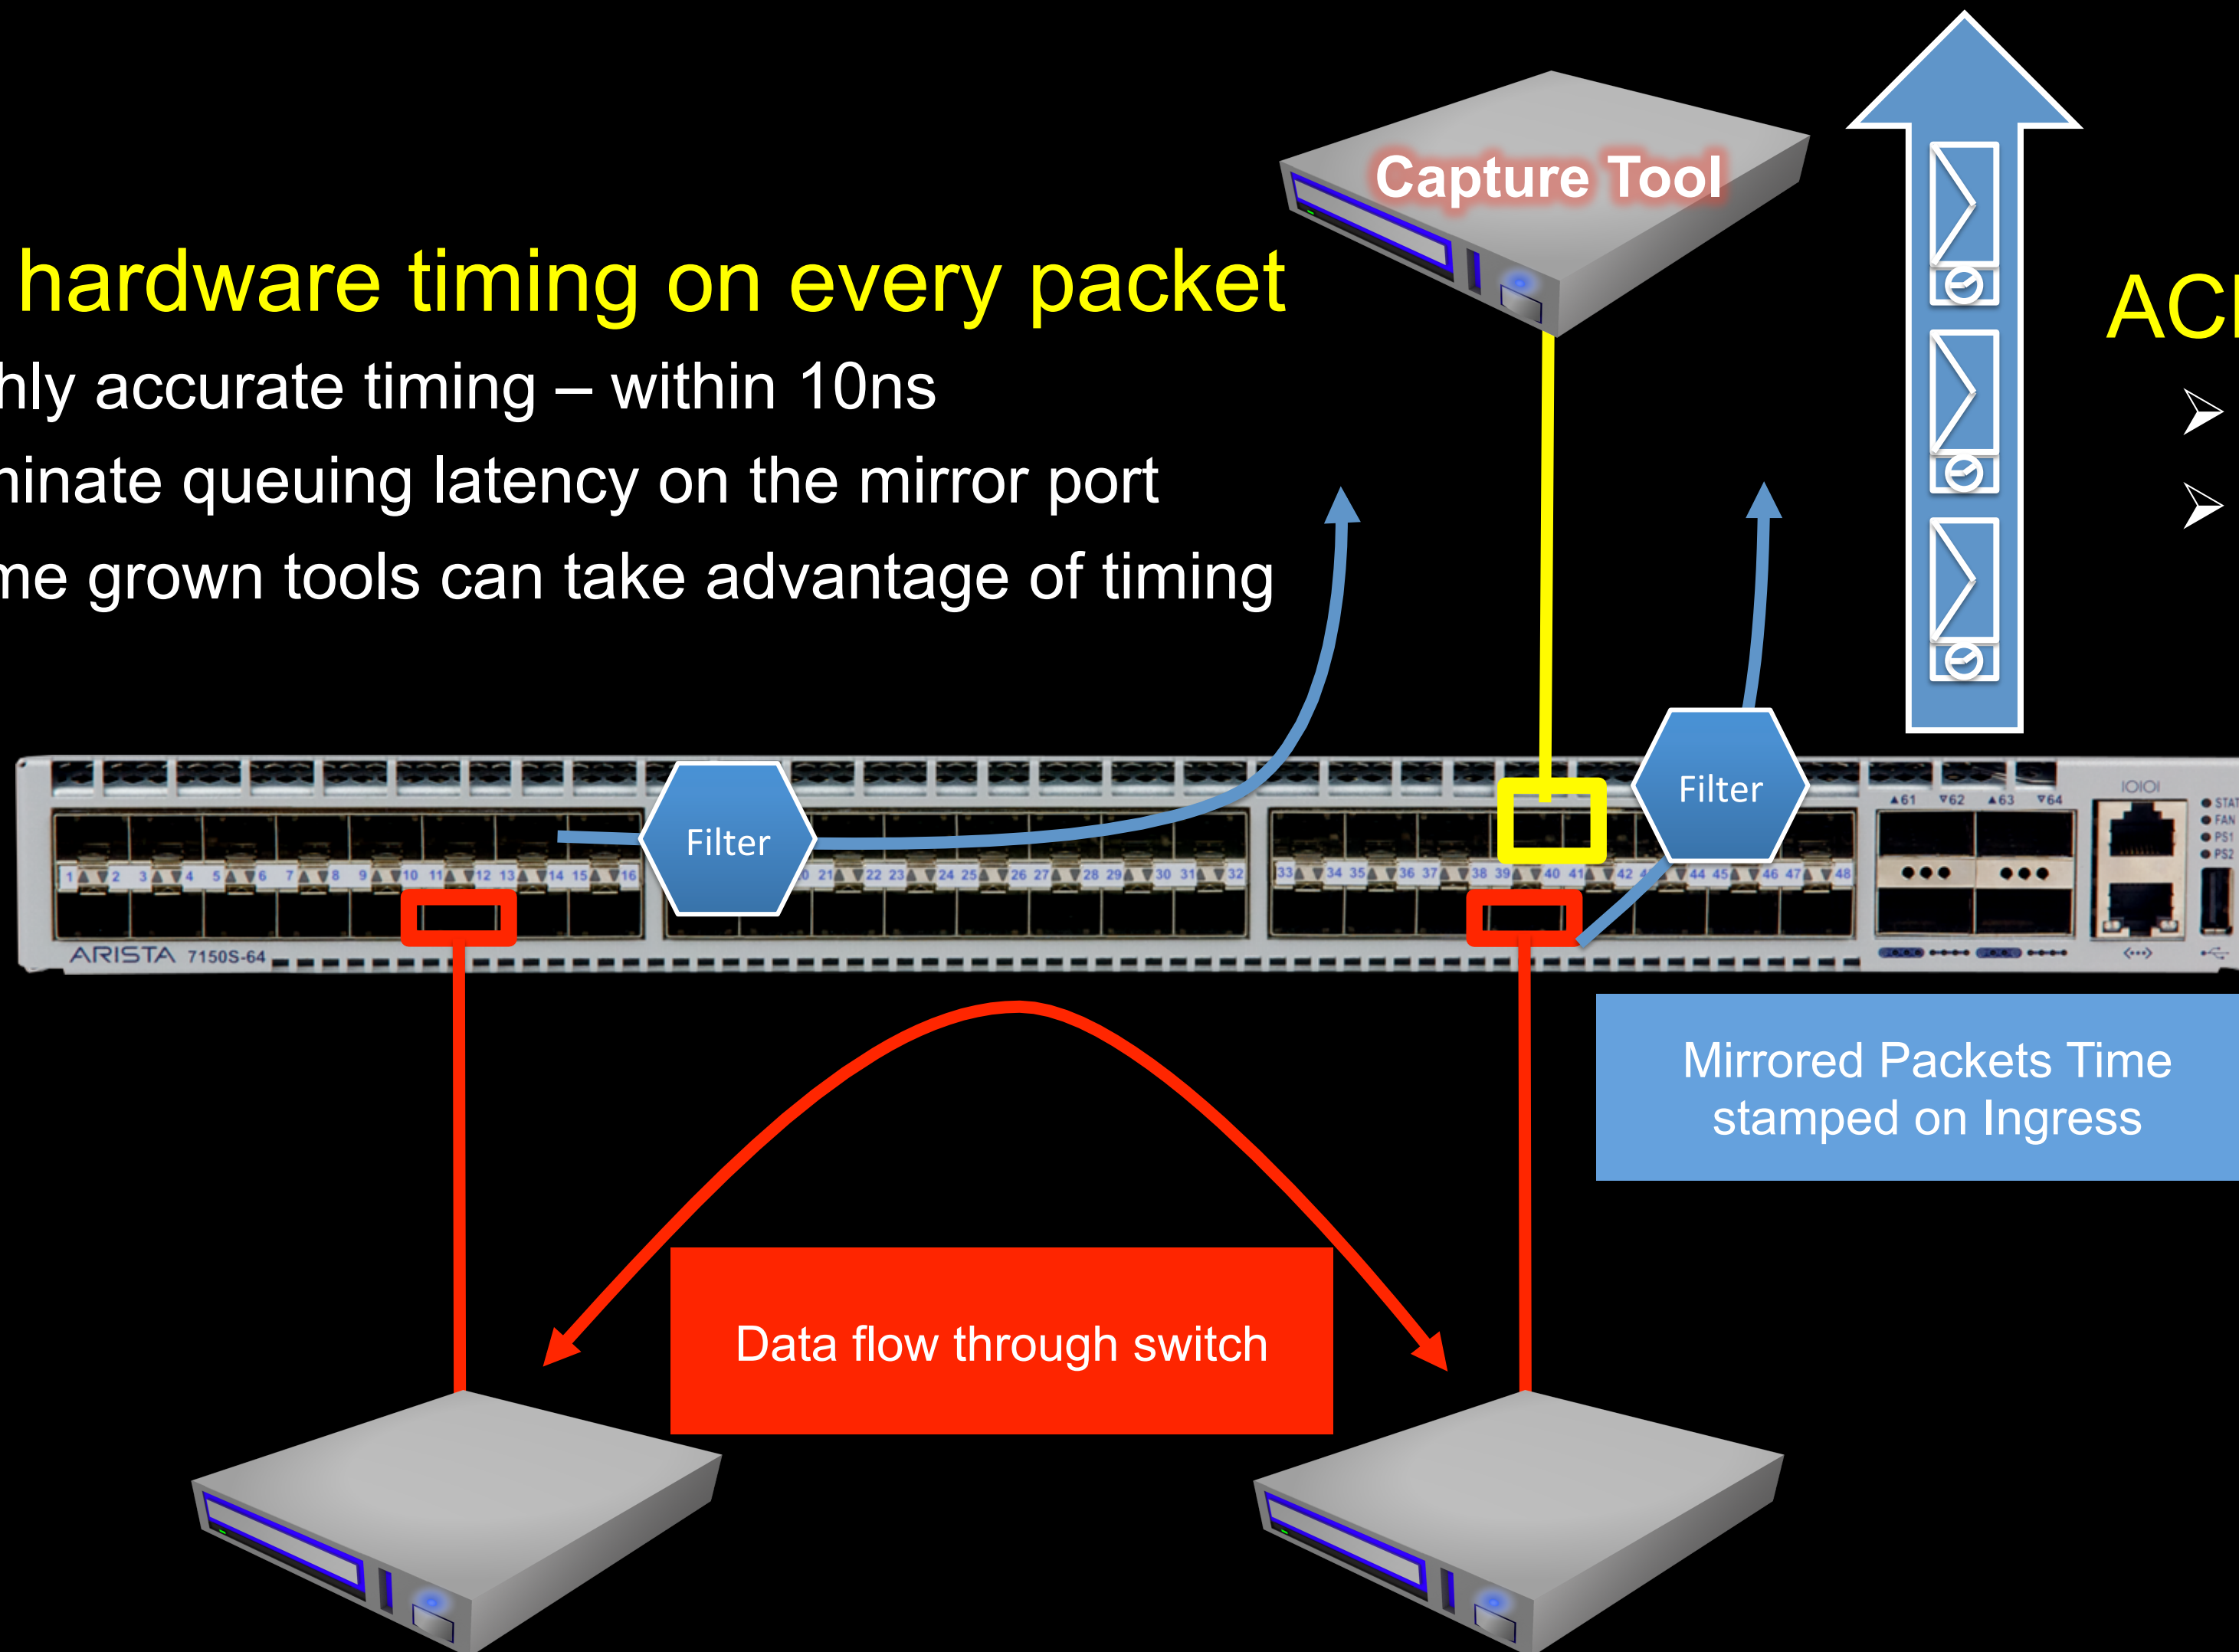
Hardware Timing & Packet Filtering

Ingress hardware timing on every packet

- Highly accurate timing – within 10ns
- Eliminate queuing latency on the mirror port
- Home grown tools can take advantage of timing

ACL Based filtering

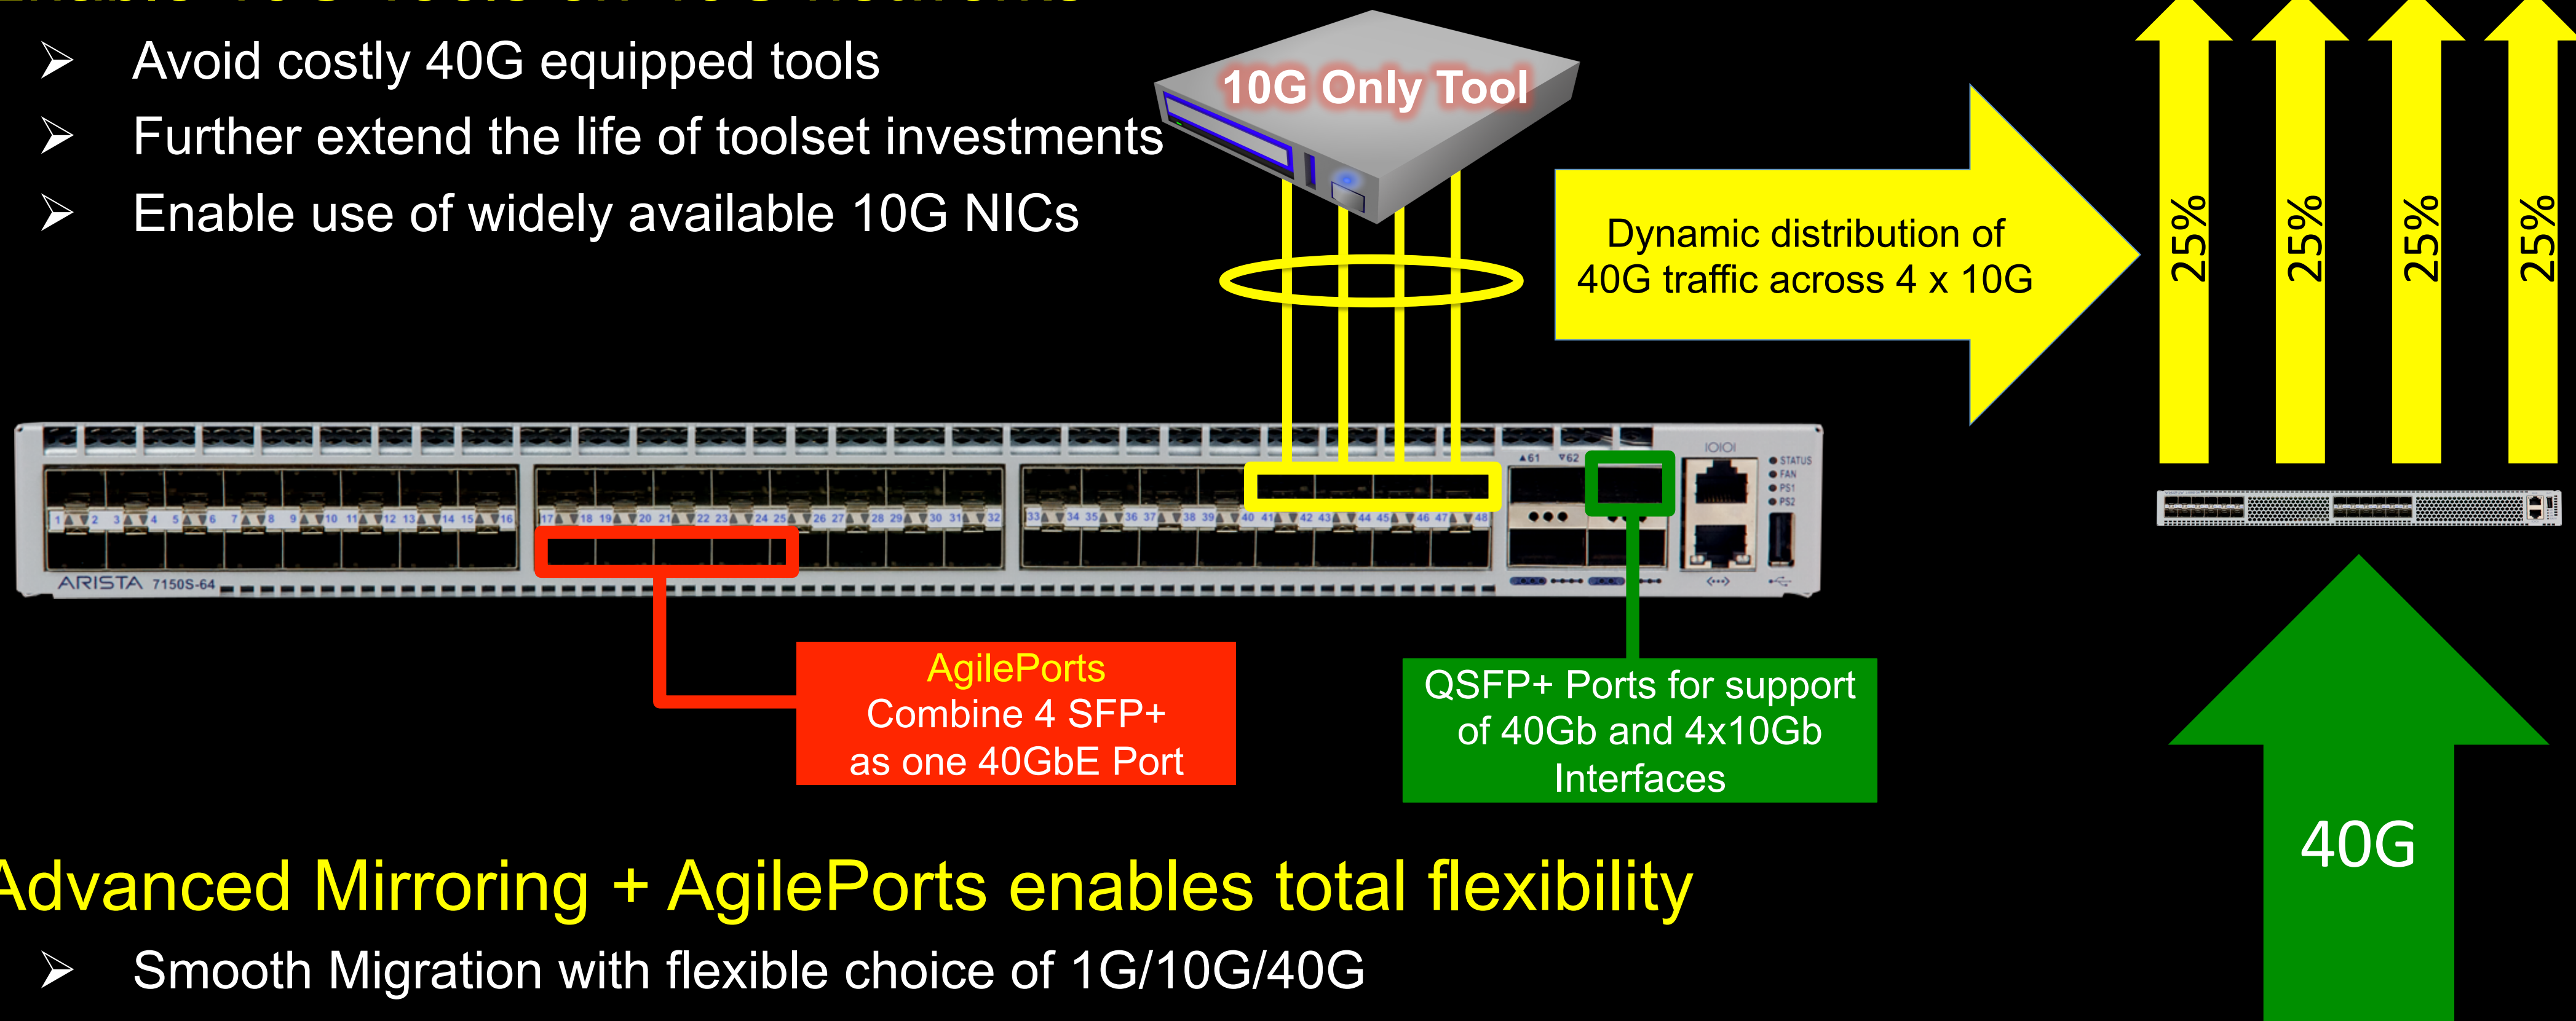
- Reduces load on tools
- Increasing longevity



Intelligent Loadsharing

Enable 10G Tools on 40G networks

- Avoid costly 40G equipped tools
- Further extend the life of toolset investments
- Enable use of widely available 10G NICs



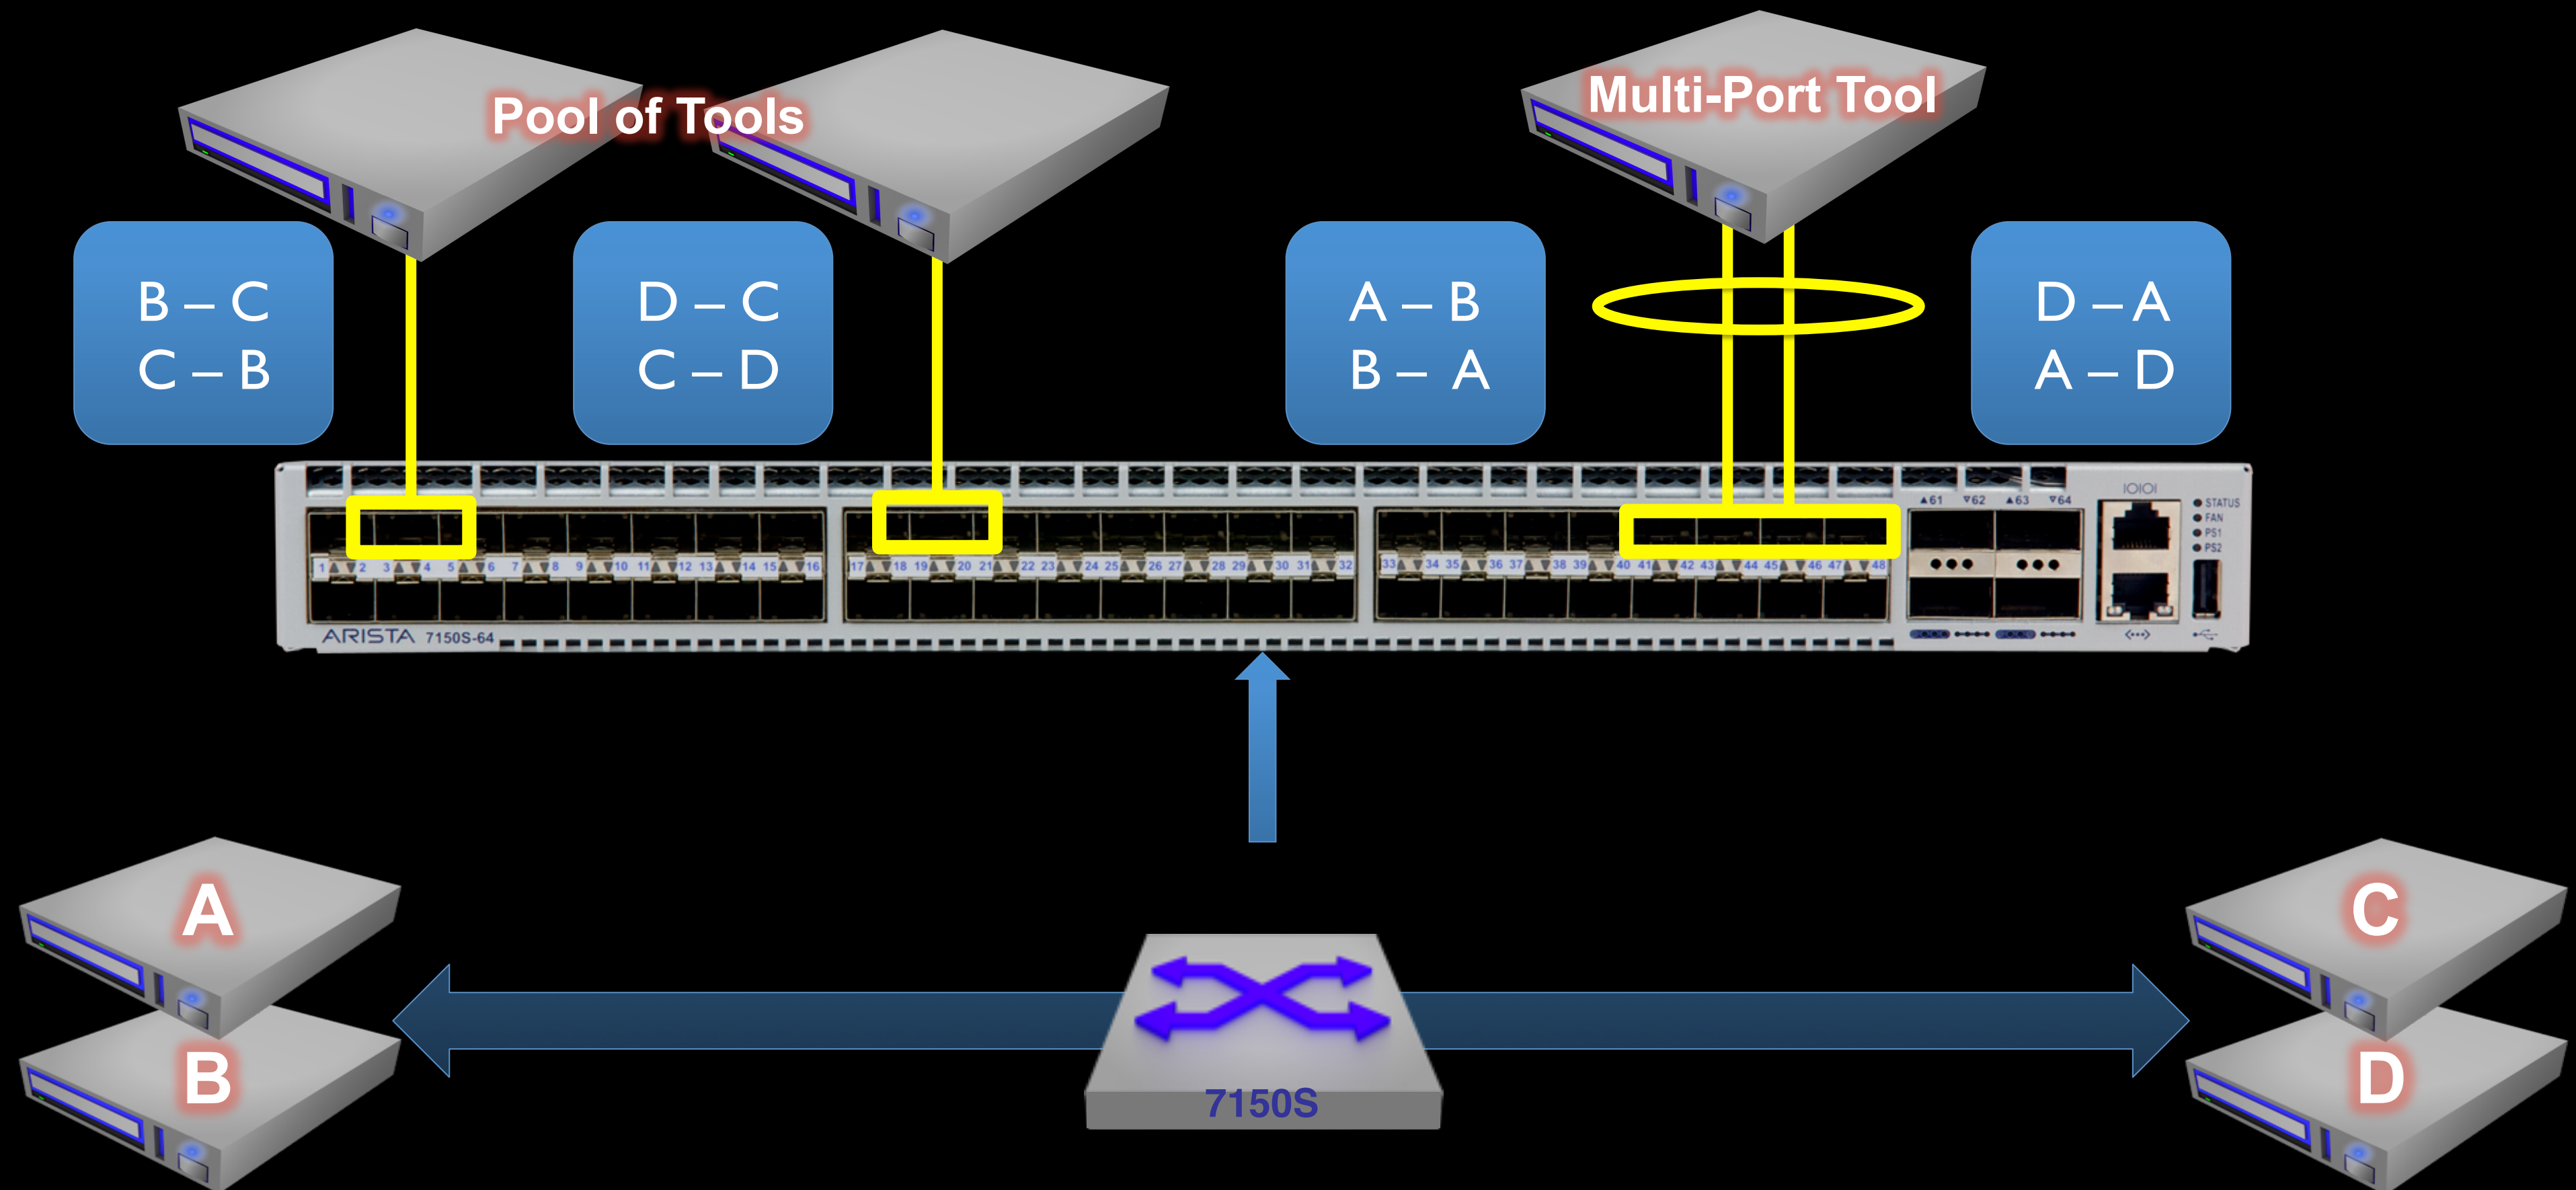
Advanced Mirroring + AgilePorts enables total flexibility

- Smooth Migration with flexible choice of 1G/10G/40G
- Cost Effective migration with 10G transceivers reuse
- Support 10G tools on 40G networks with intelligent loadsharing
- Avoid issues with multi-port LAG hashing

Symmetric Loadbalancing

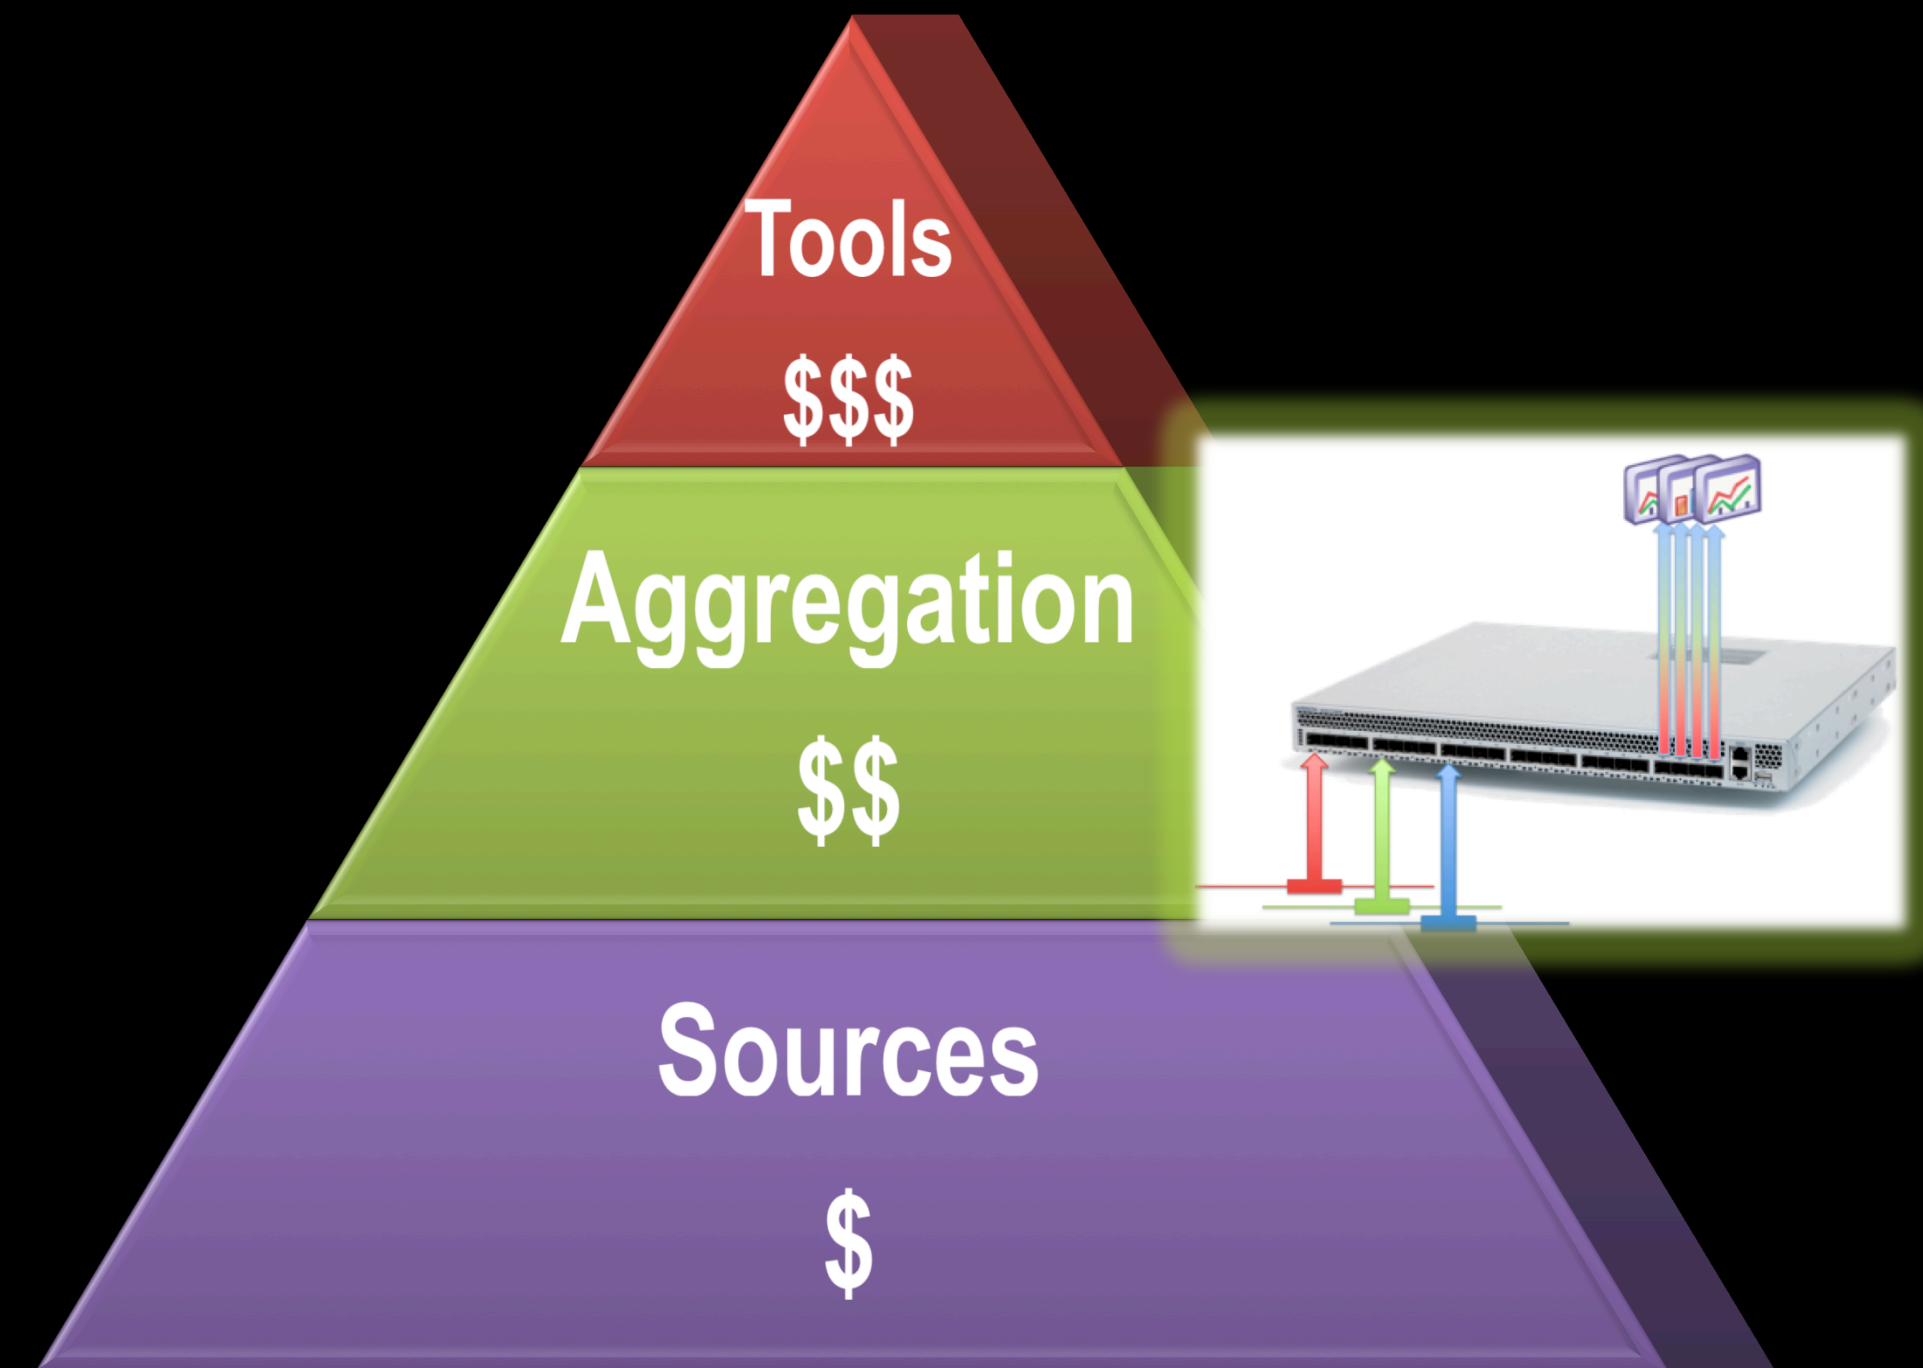
Ensures all traffic in one flow is sent to the same Analyzer

- Enables distribution of aggregated over multiple lower speed tools



What Is TAP Aggregation?

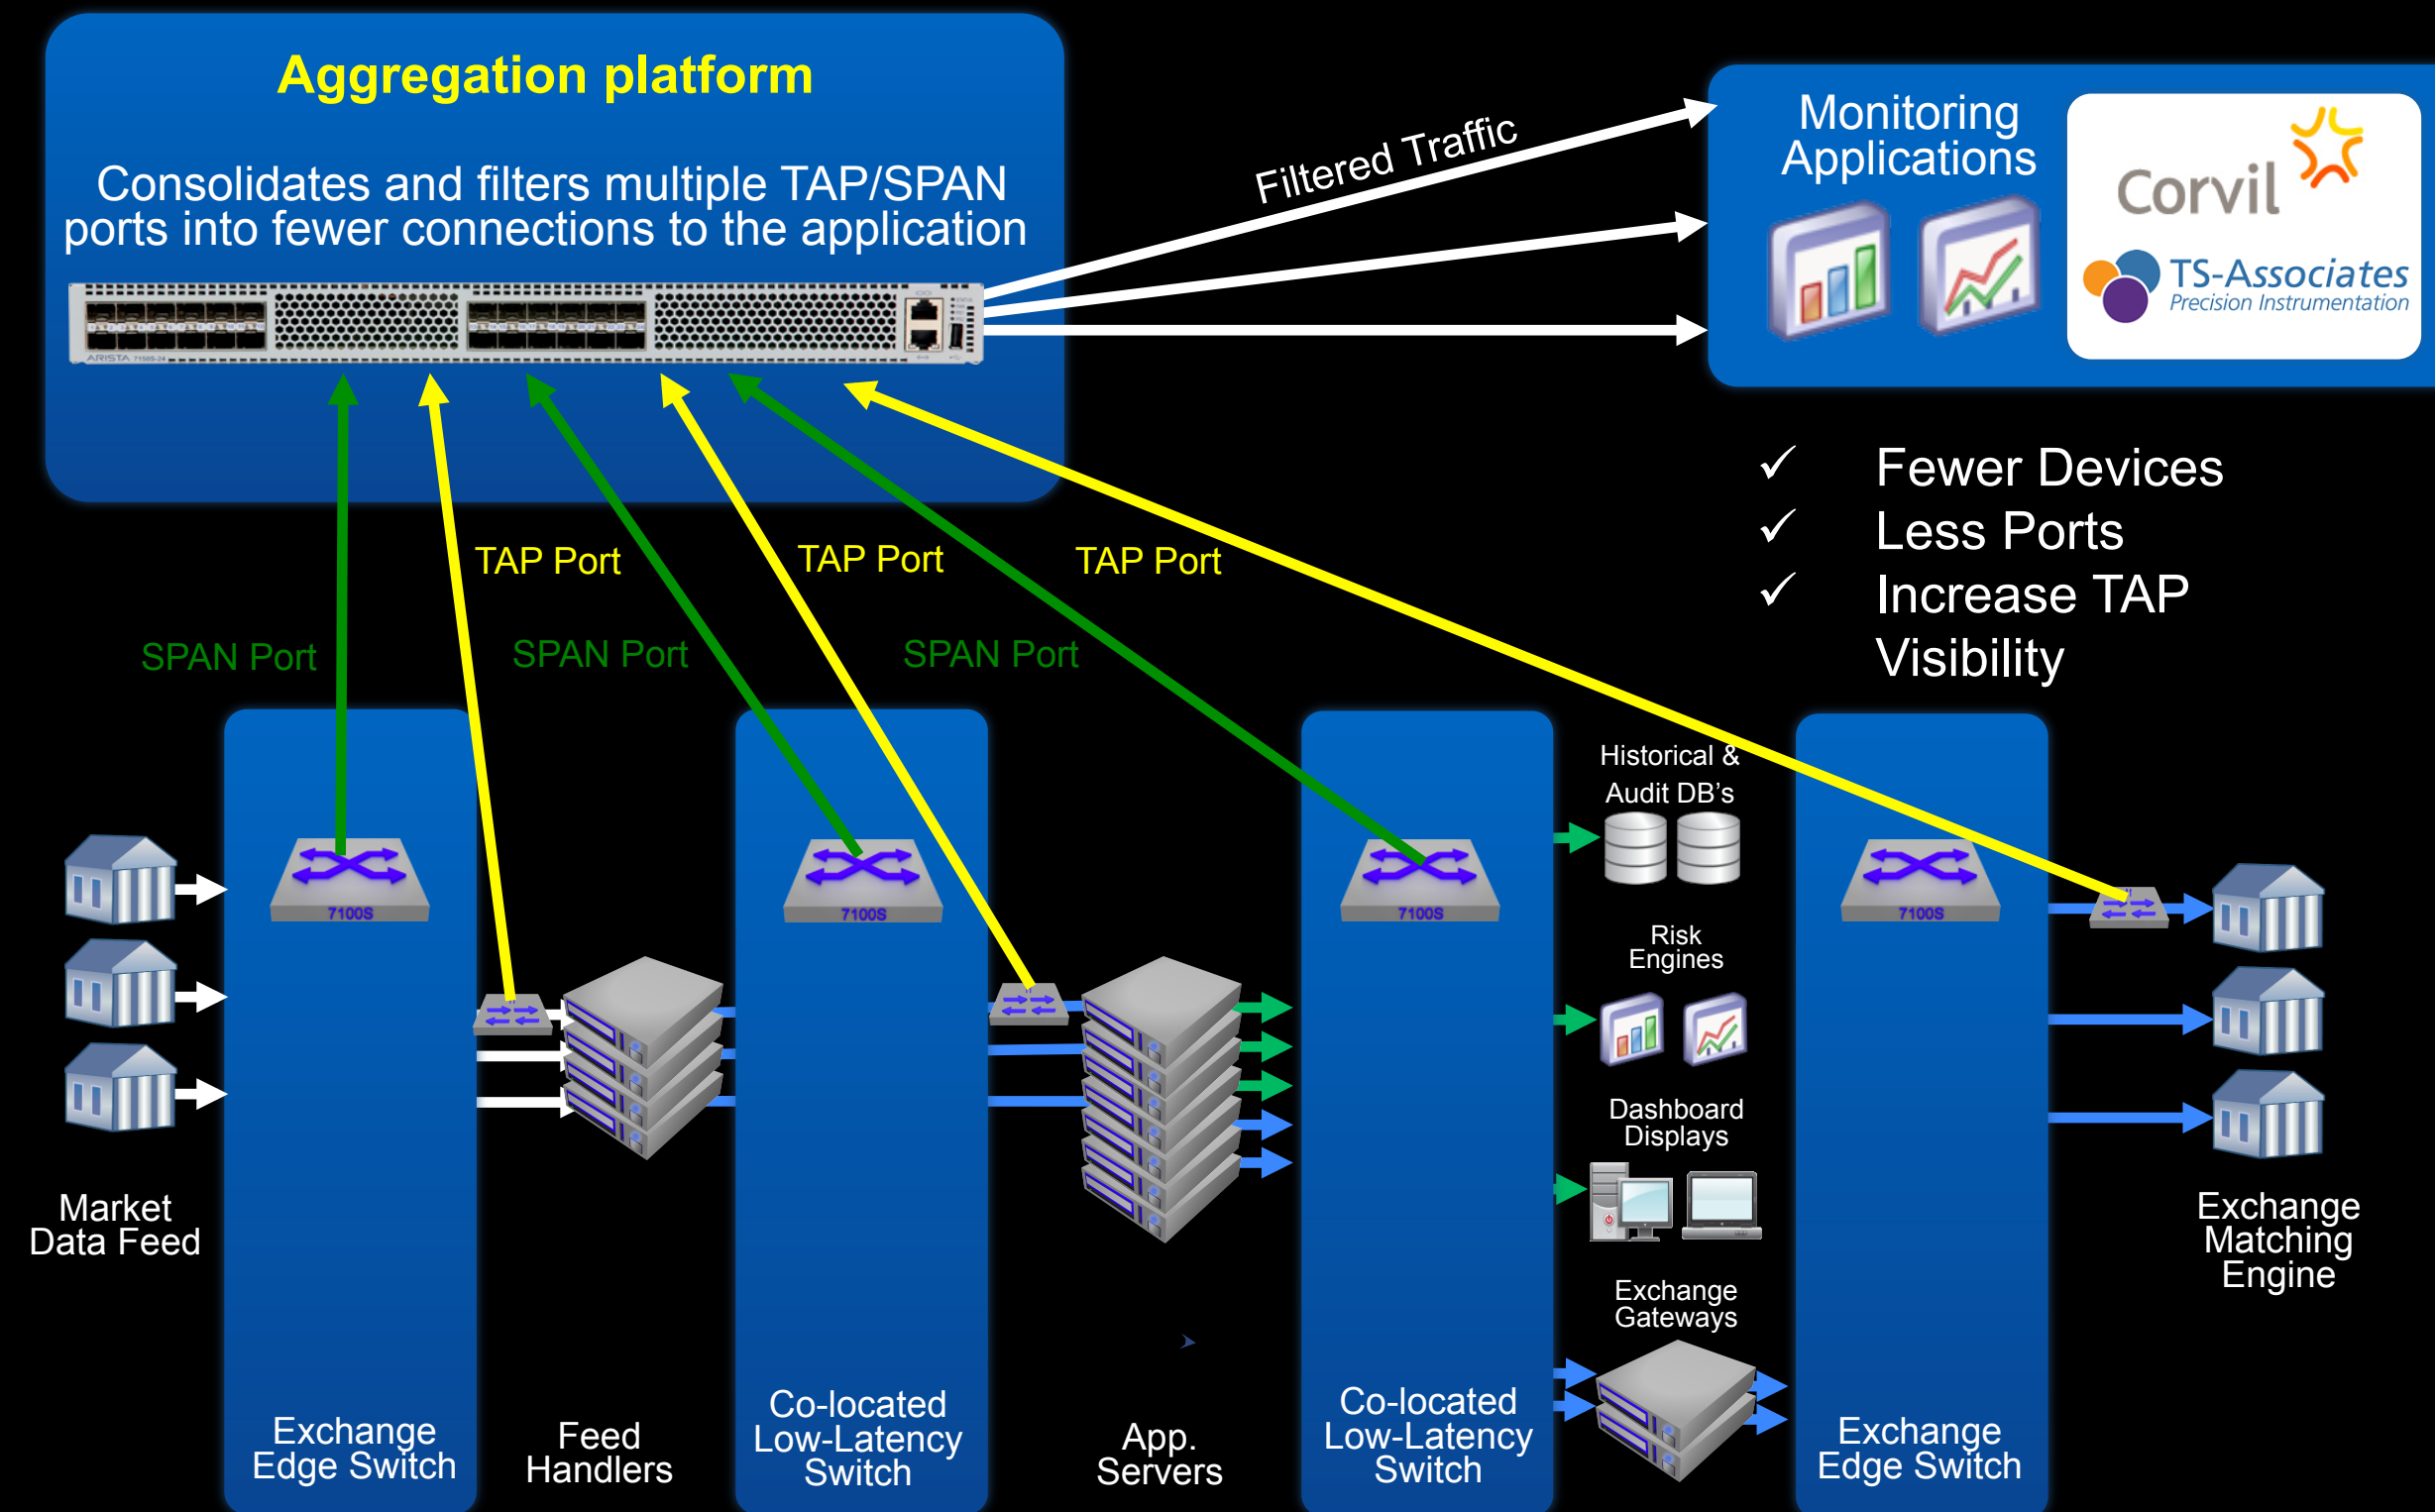
- Customers who need deep network visibility deploy parallel instrumentation networks
- 3 major components:
 - Sources: TAPs (Optical, Electrical), SPAN ports
 - Aggregators: / Matrix Switches
 - Tools: Analysis tools / Packet Capture
- Analyzers are expensive, low density and limited in performance
 - Usually 1-2 x 10G max
- Aggregation layer
 - Aggregates many sources to fewer tool ports
 - Enables duplication or load balancing of traffic to multiple tools
 - Filters uninteresting traffic to unload analyzer



Who needs TAP Agg?

- Financial Services
 - Heavy users in trading and datacenter networks
 - For compliance, performance monitoring and troubleshooting
- High Performance Compute / Clustering
 - Fine grained monitoring of performance between nodes
- Cloud/Web Scale and Service Provider
 - DC wide monitoring infrastructure to allow rapid visibility into hotspots
 - Lawful intercept, capacity planning and customer sat. applications
- Government and Law Enforcement
 - lawful intercept applications requiring transparent traffic replication and filtering

Introducing Integrated SPAN/TAP Aggregation

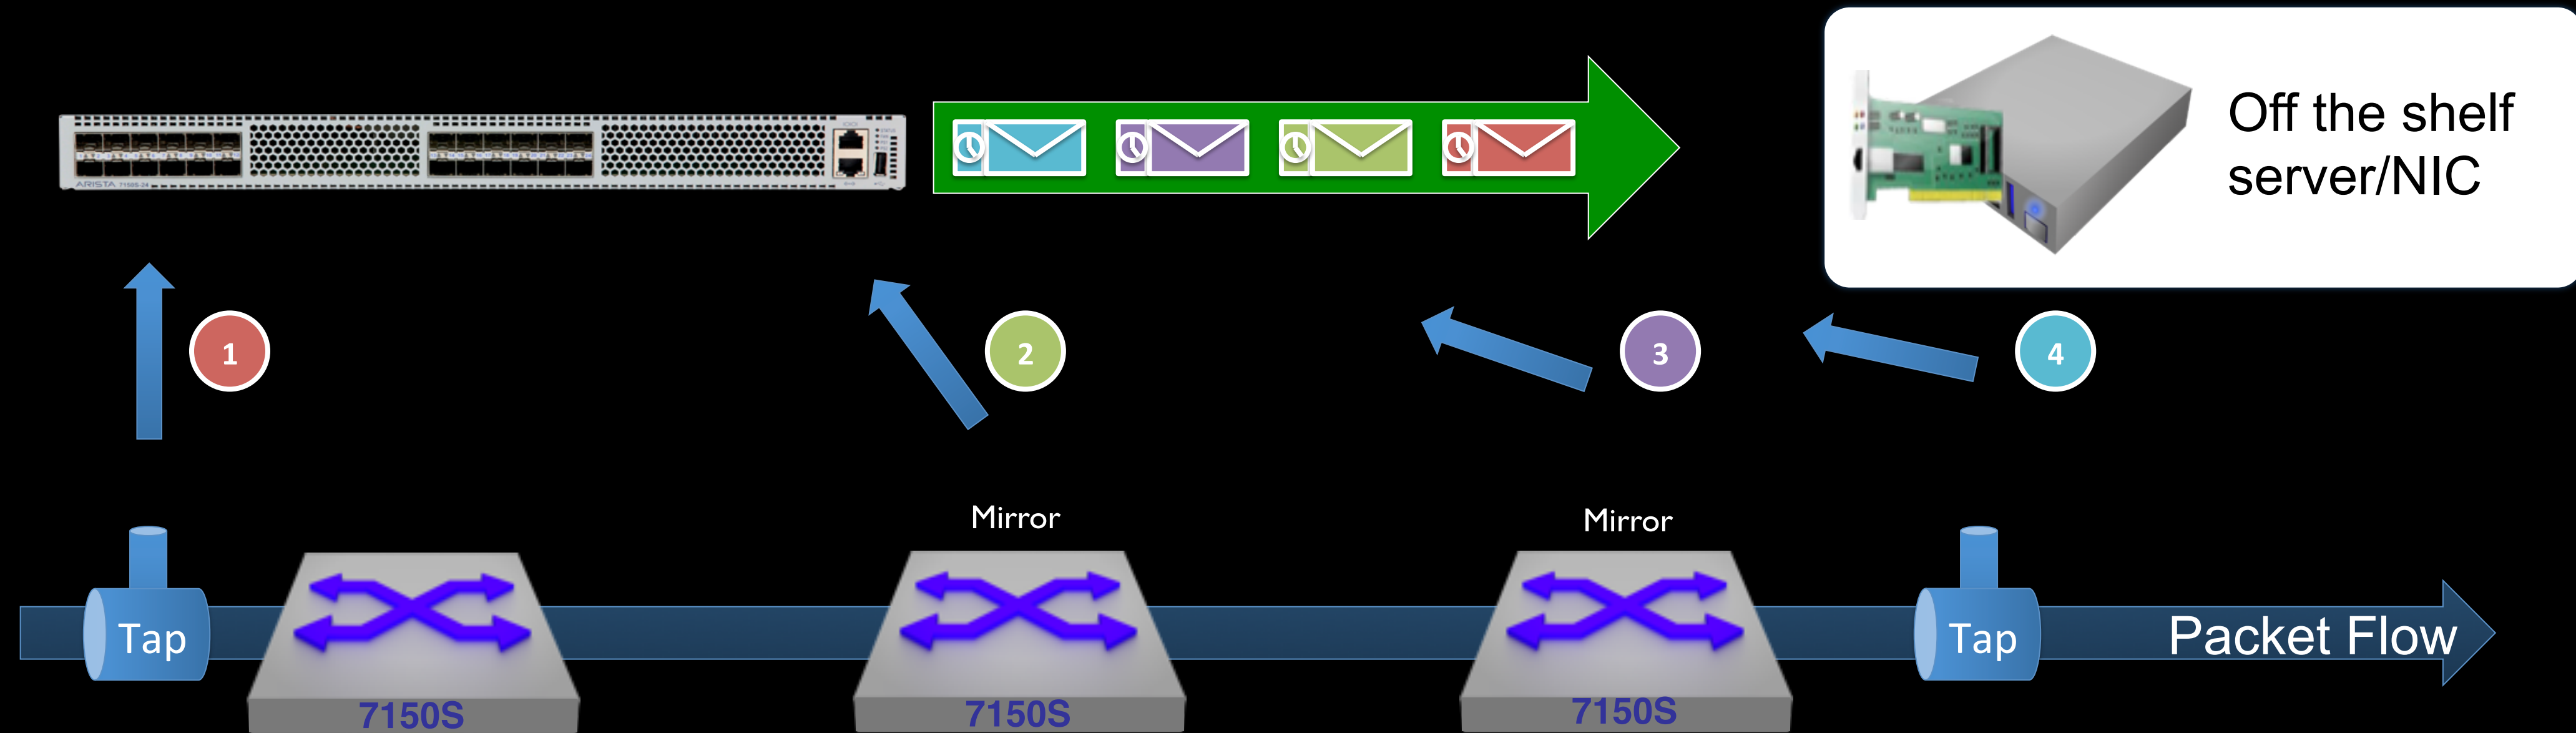


- Integrated SPAN/TAP aggregation functionality
- Advanced features: N:M Replication, Filtering, HW Time-stamping and more
- Substantial Capex and Opex saving vs. proprietary SPAN/TAP aggregators

Unique Feature Set

- ✓ Simple CLI Interface
- ✓ Any:Any Replication
- ✓ Symmetric Loadbalancing
- ✓ HW Time-stamping
- ✓ Intelligent Loadsharing
- ✓ PTP 1588
- ✓ Enhanced L2/3/4 Filtering
- ✓ Source Port Identification
- ✓ Packet Truncation
- ✓ AgilePorts + 40G
- ✓ Buffer Tuning
- ✓ sFlow and LANZ+

Example – Unidirectional Flow Latency Measurement

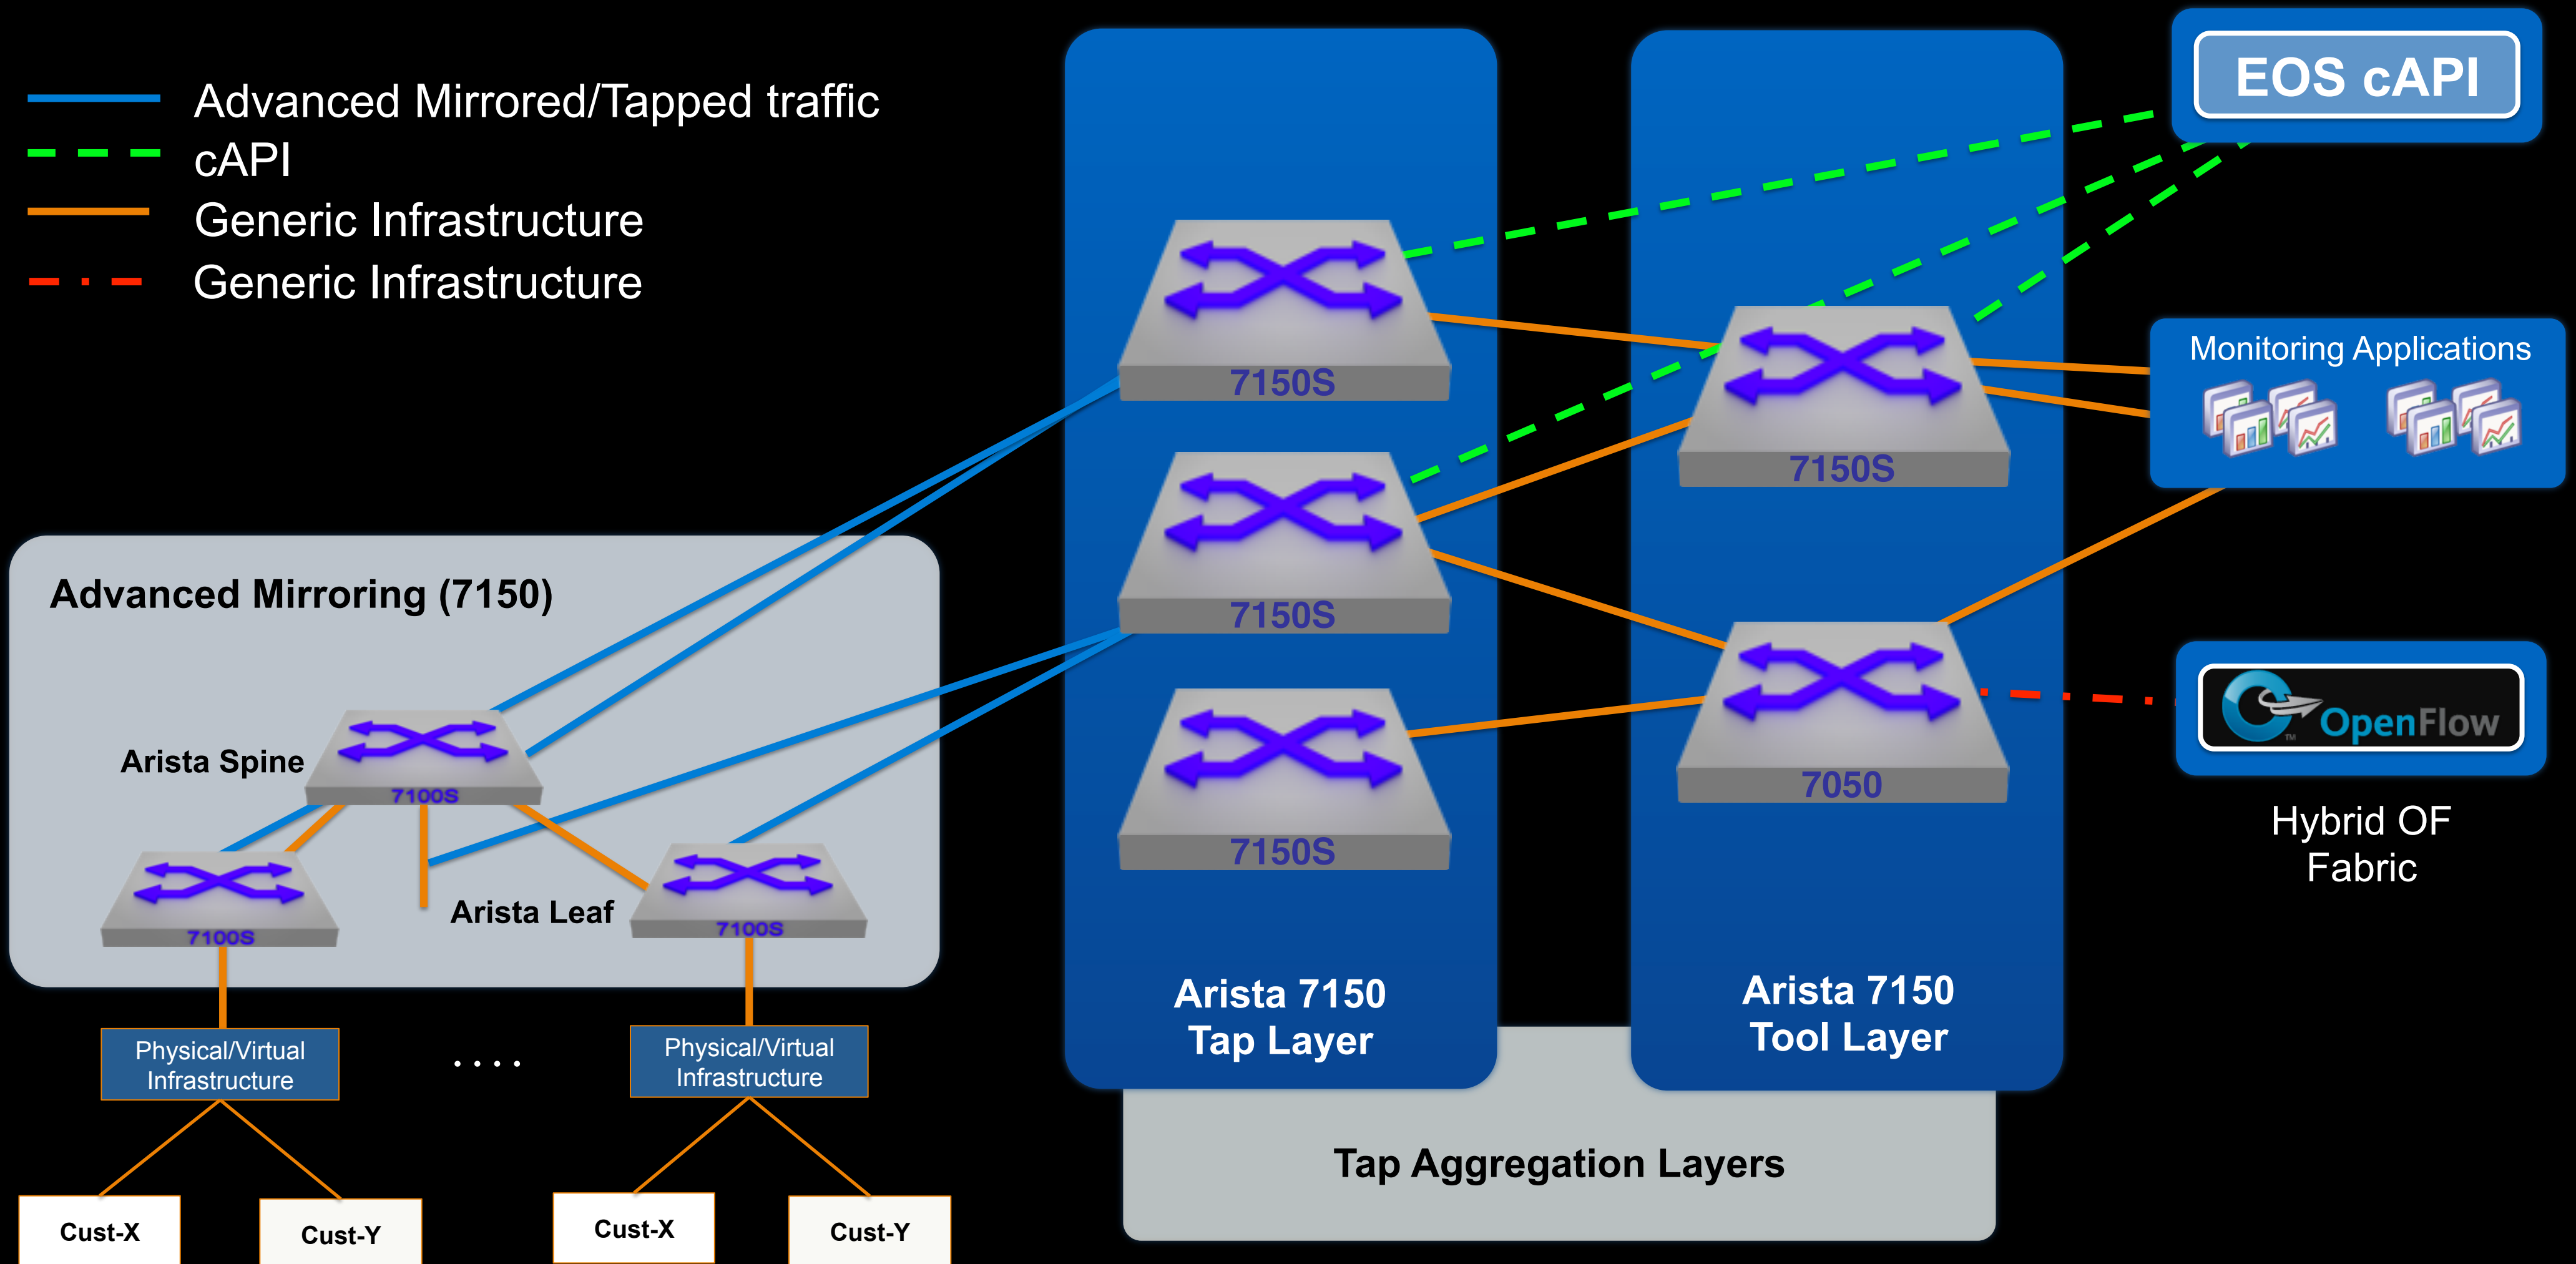


- Packets marked with a per-hop ID
- 10ns precision timing with off the shelf hardware
- Relative timing can be recovered with simple scripts
- 95% reduction in cost vs traditional capture tools

~ `capture$./Decoder.py timestamp.pcap`

HW Timestamp(ns)	Delta(ns)
895645872.61	
895671057.07	25184
895711657.89	40600
895736919.49	25261

Example - High Capacity TAP Fabric Design



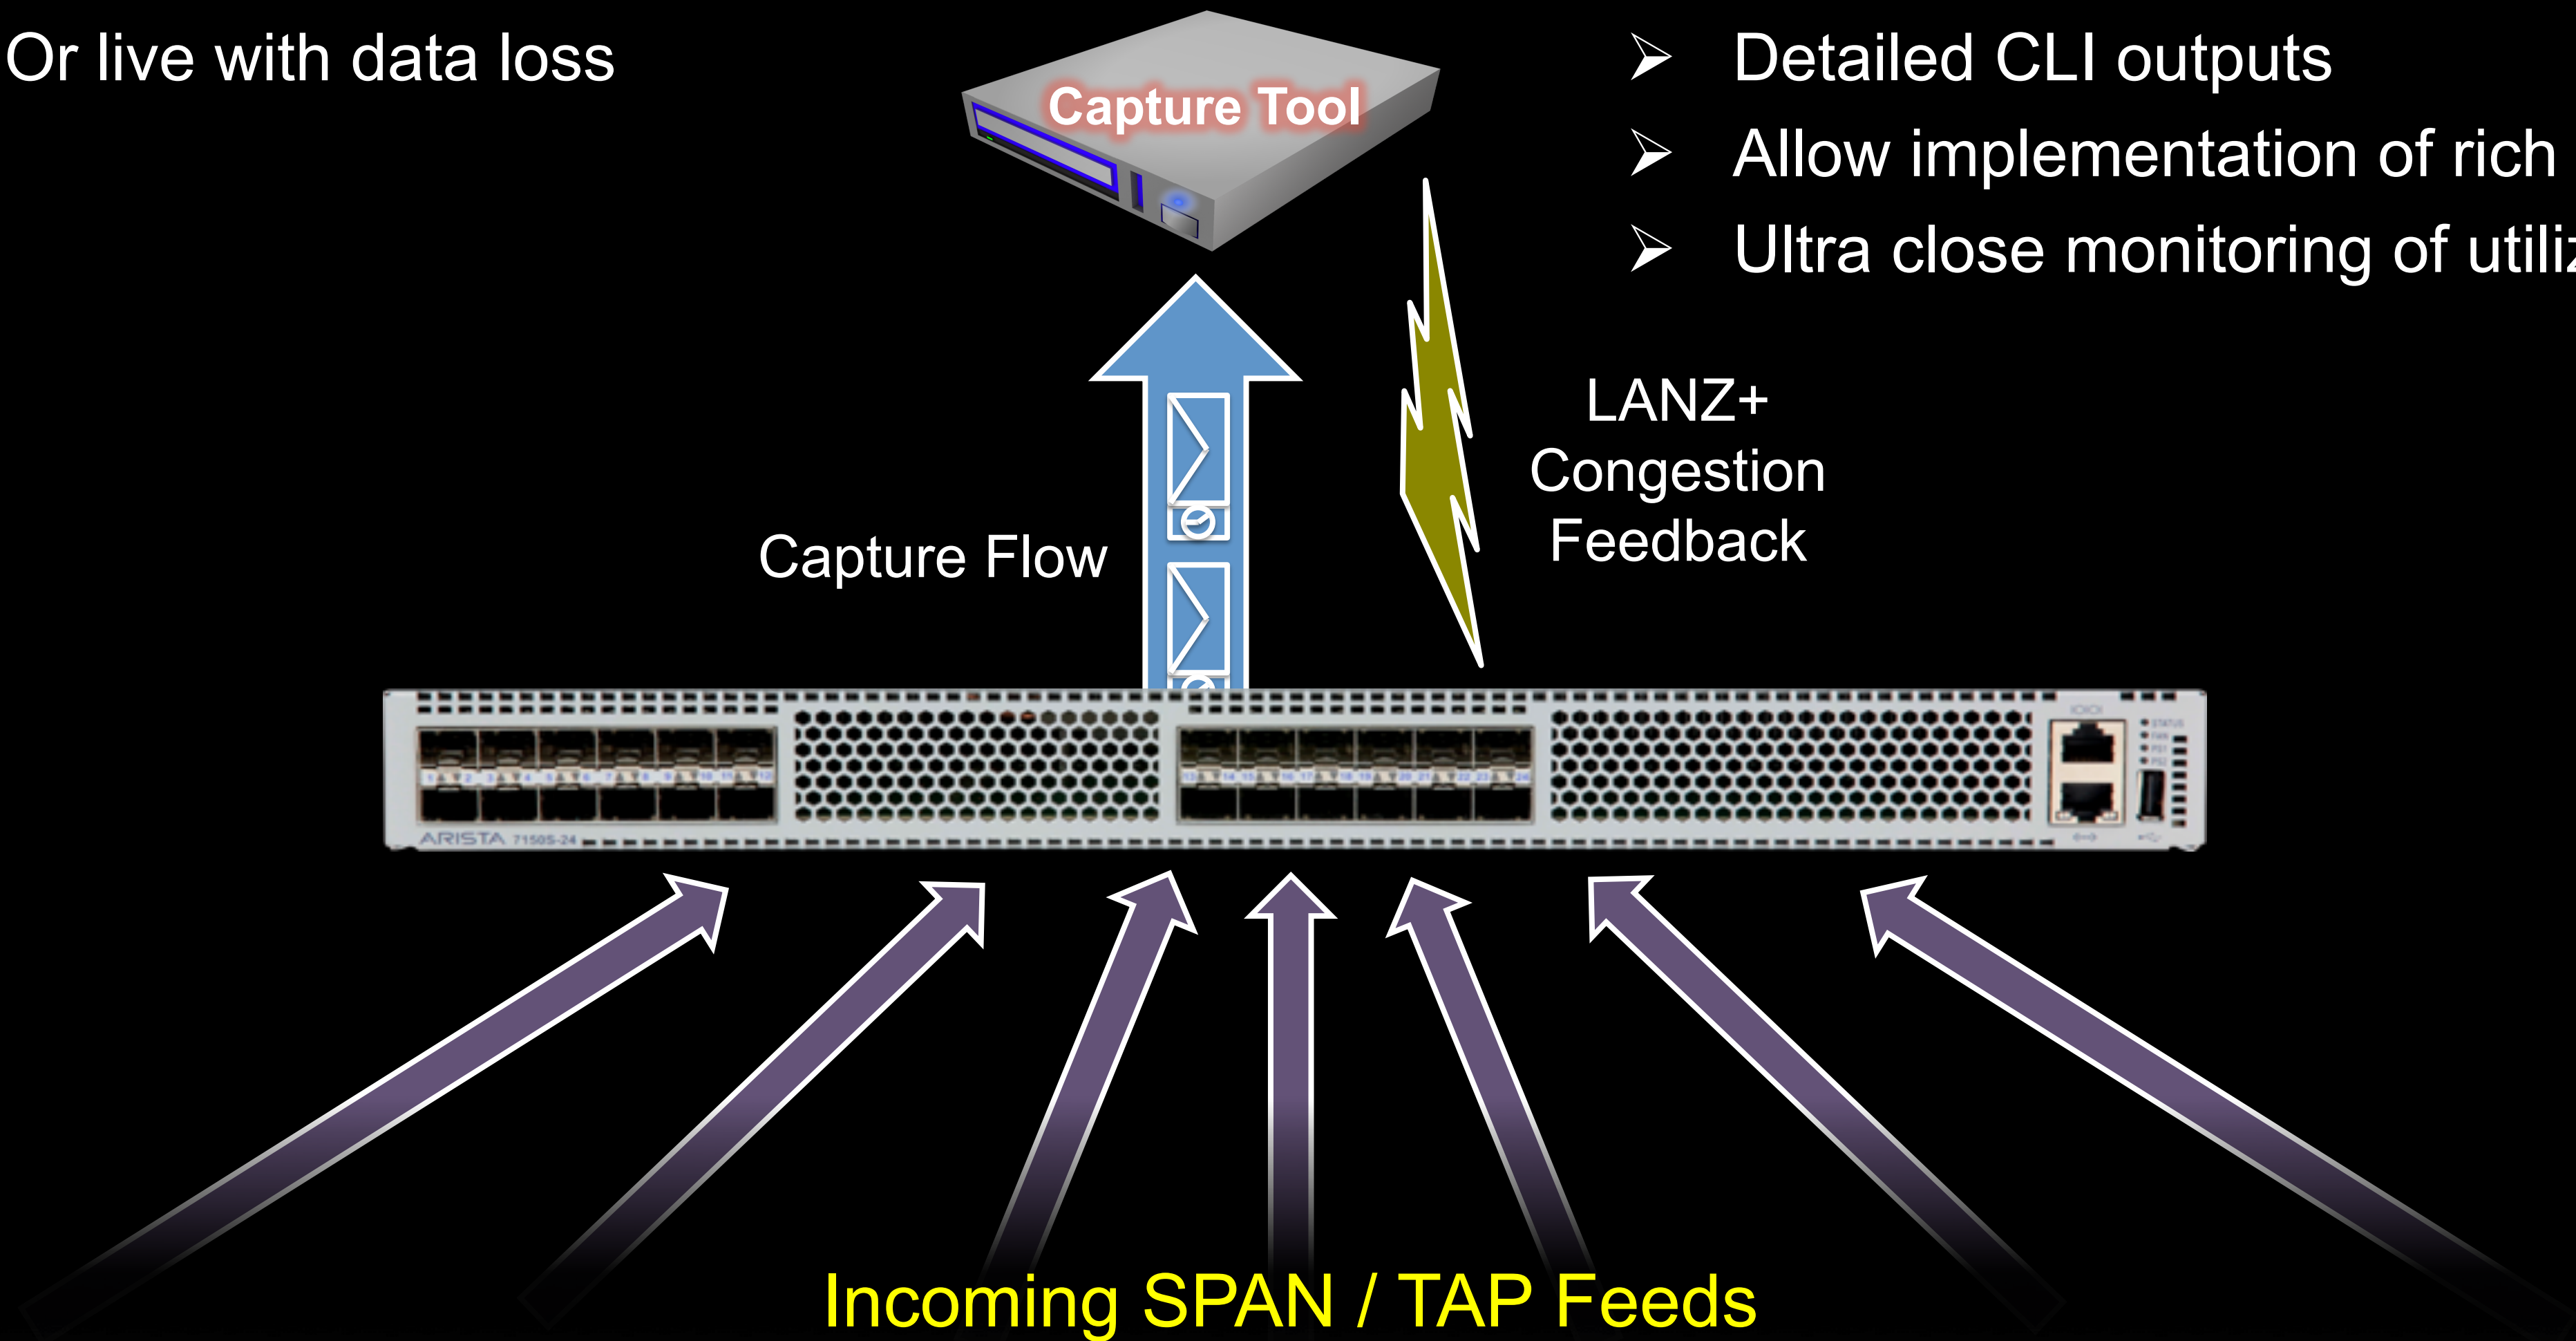
What has LANZ+ got to do with TAP Aggregation?

TAP Aggregation is oversubscribed

- Filter aggressively to avoid traffic loss
- Or live with data loss

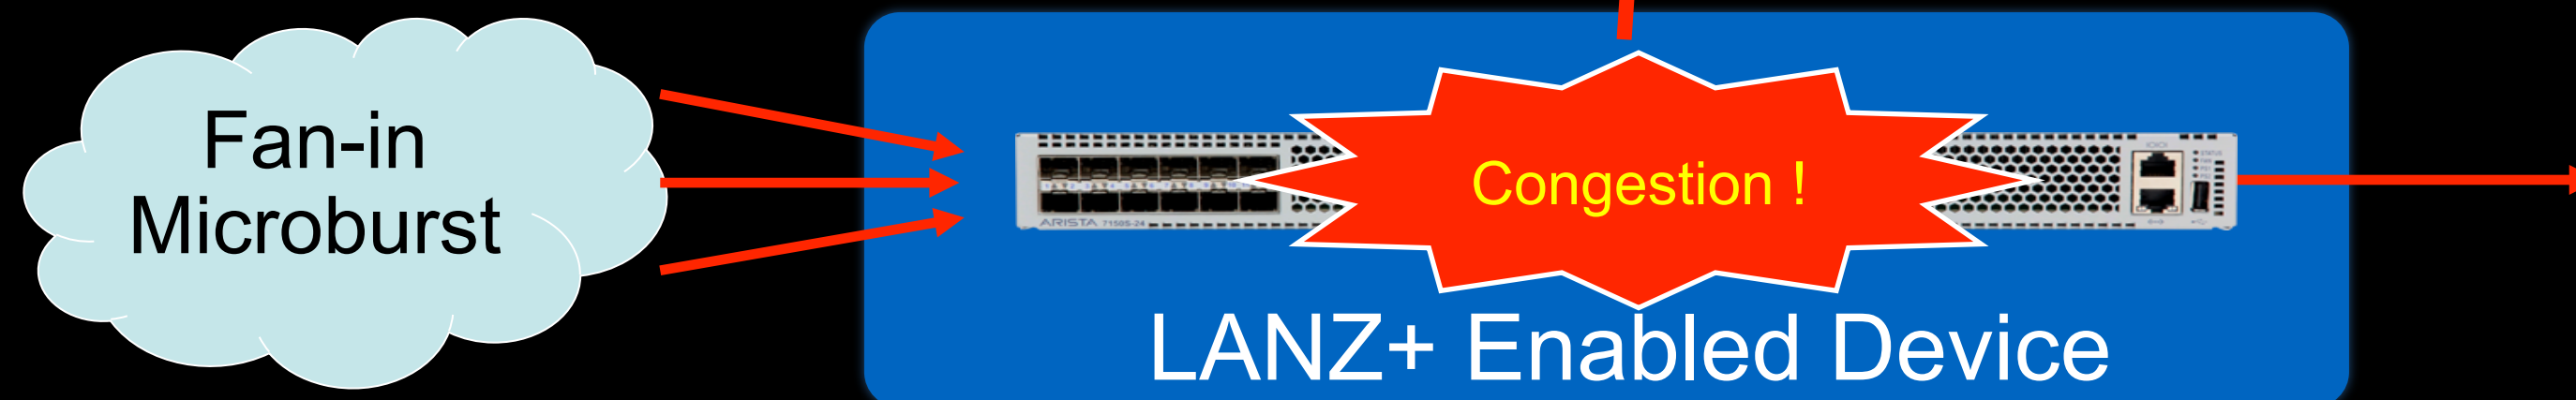
LANZ+ Provides the answer

- Real-time feedback to tools
- Detailed CLI outputs
- Allow implementation of rich filters
- Ultra close monitoring of utilization



LANZ+ Enhanced Streaming

- ✓ Real-time streaming of congestion events
 - ✓ Full Precision timestamps correlate to external events
- ✓ Direct to applications via **open** protocols
- ✓ Analyze network performance in real-time
- ✓ Hotspot early warning system for
 - ✓ Hadoop, Cloud, Cluster, Financial



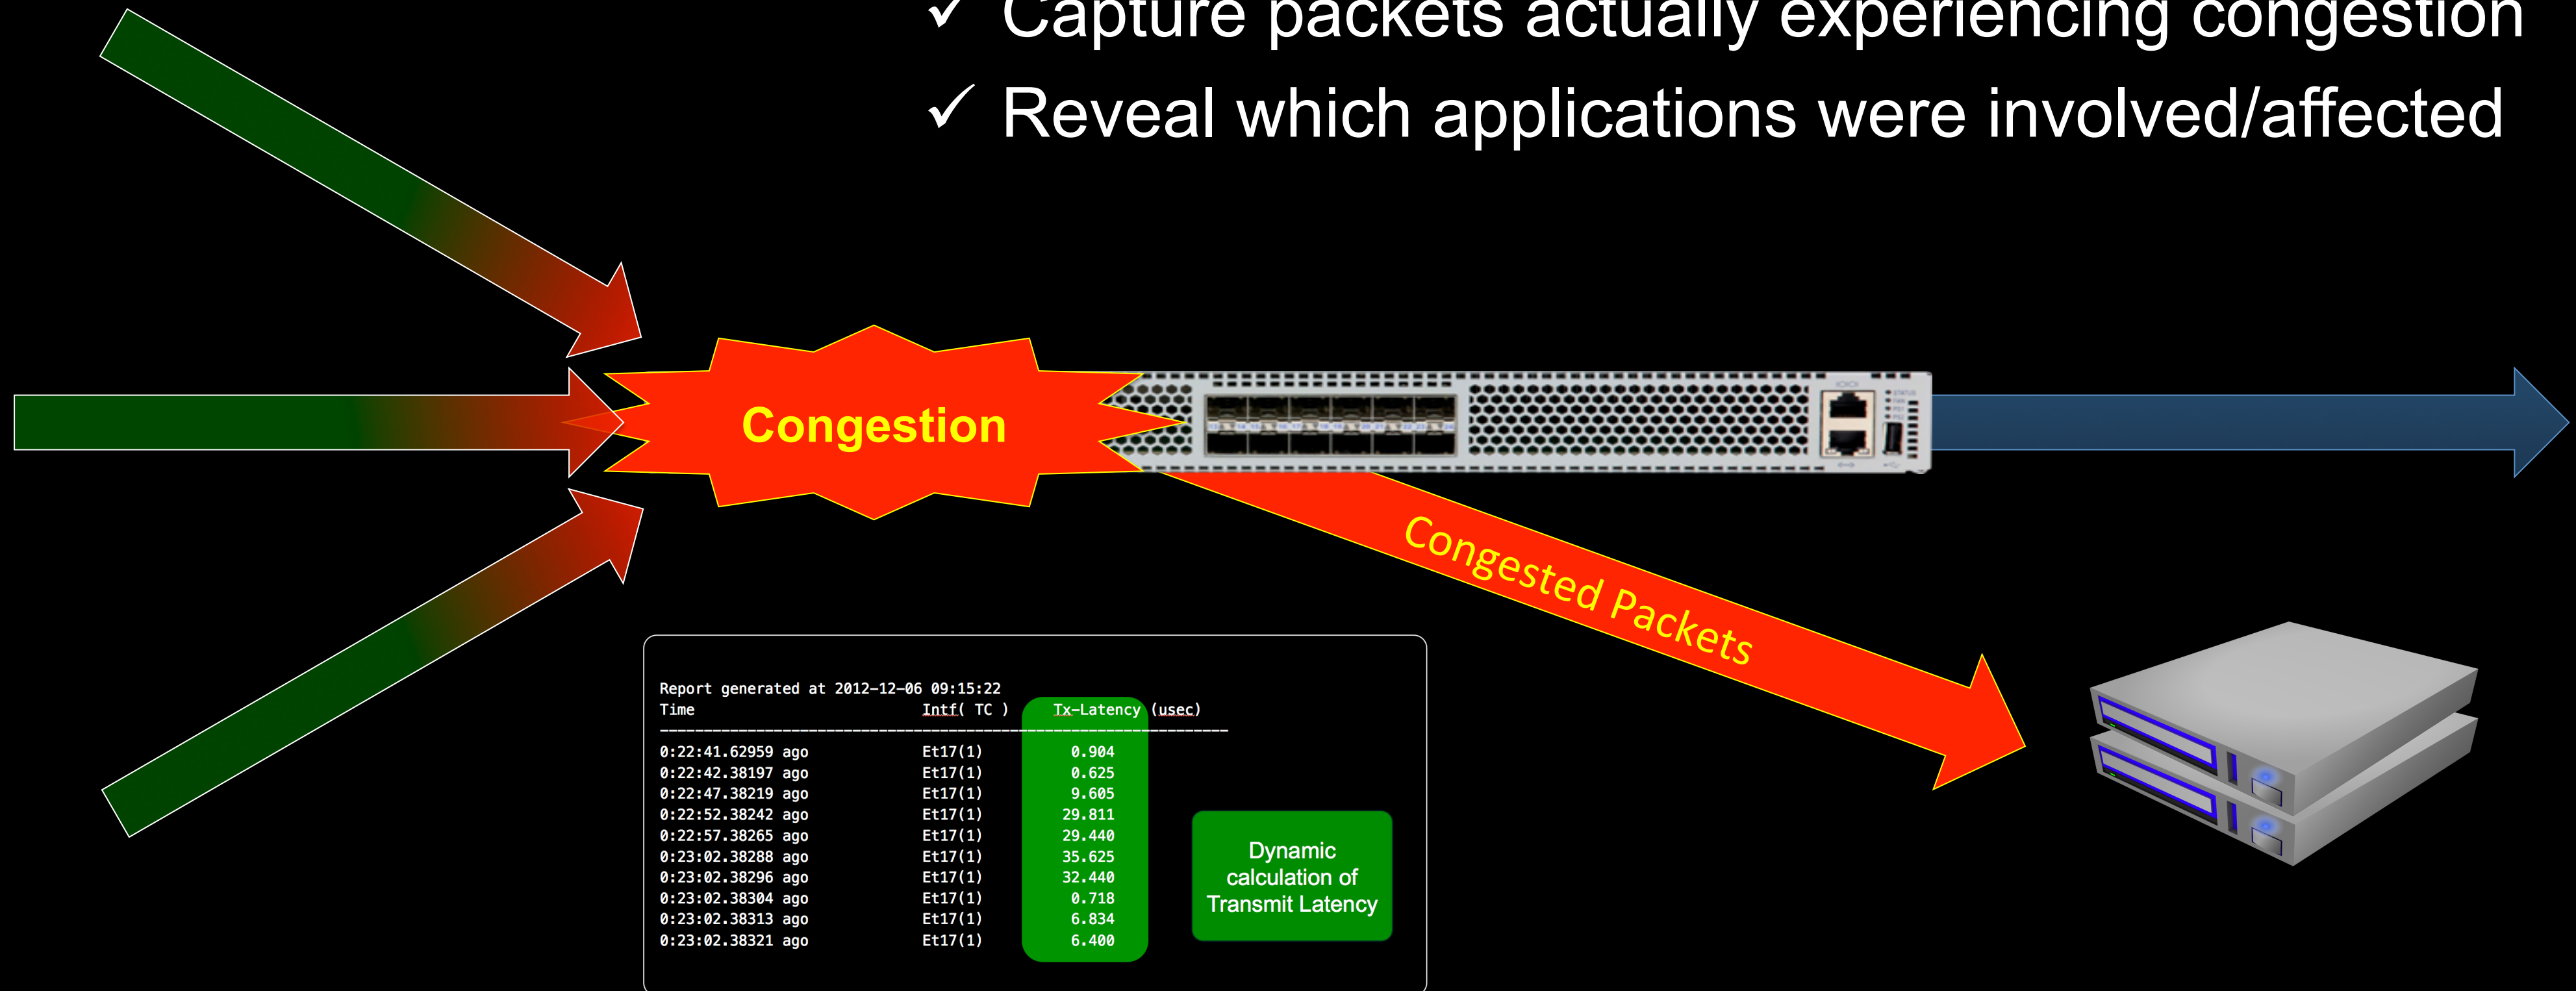
```
CongestionRecord:
  timestamp: Fri Feb 10 19:02:00.37462 2012
  intfName: Ethernet3
  switchId: 0
  portId: 5
  queueSize: 3
CongestionRecord:
  timestamp: Fri Feb 10 19:02:00.37474 2012
  intfName: Ethernet3
  switchId: 0
  portId: 5
  queueSize: 3
```



Correlate Performance with Congestion Events and En-queued Data

Unique - LANZ+ Congestion Capture

- ✓ Congestion threshold triggers data capture
- ✓ Capture packets actually experiencing congestion
- ✓ Reveal which applications were involved/affected



Correlate performance with traffic captured at the time of congestion