

Unlocking the Future of Data Center with AI

Derek Caruso, Solutions Engineer, CAI

December 3, 2024



AI presents challenges for IT teams

30.8%

Of Gen AI testing and development managed by Central IT¹

13%

IT organizations with infrastructure prepared for AI today²

36.2%

IT Staff/ Skill Recruitment/ Training is highest priority to support gen AI¹

Rapid growth in data volume and variety

Unfamiliar application stacks and new, complex infrastructure patterns

Insufficient IT automation and observability

Greater cyber-security threats



New operational silos

Shortage of technical expertise

Disorienting AI hype

High entry cost and lock-in issues

INFRASTRUCTURE DEMANDS



IT Infra and Operations Considerations for AI



Infrastructure and operational consistency— avoid new islands



Optimize for utilization and efficiency in many dimensions—support multiple projects, leverages GPUs wisely, power and cooling needs, lifecycle management



Comprehensive security protocols and measures



Support rapidly-evolving software ecosystem



Manage cloud vs. on-prem vs. hosted model



Straddle the training → fine tuning → inferencing → repeat model

Striving for a Competitive Advantage with AI/ML

Quantitative Modeling and FSI



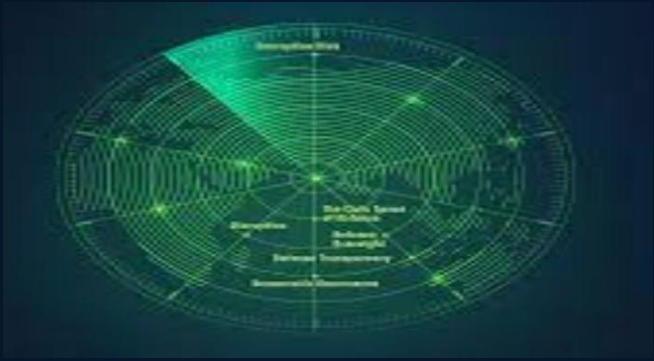
Life Sciences



Engineering



Public Sector and Defense



Striving for a Competitive Advantage with AI/ML

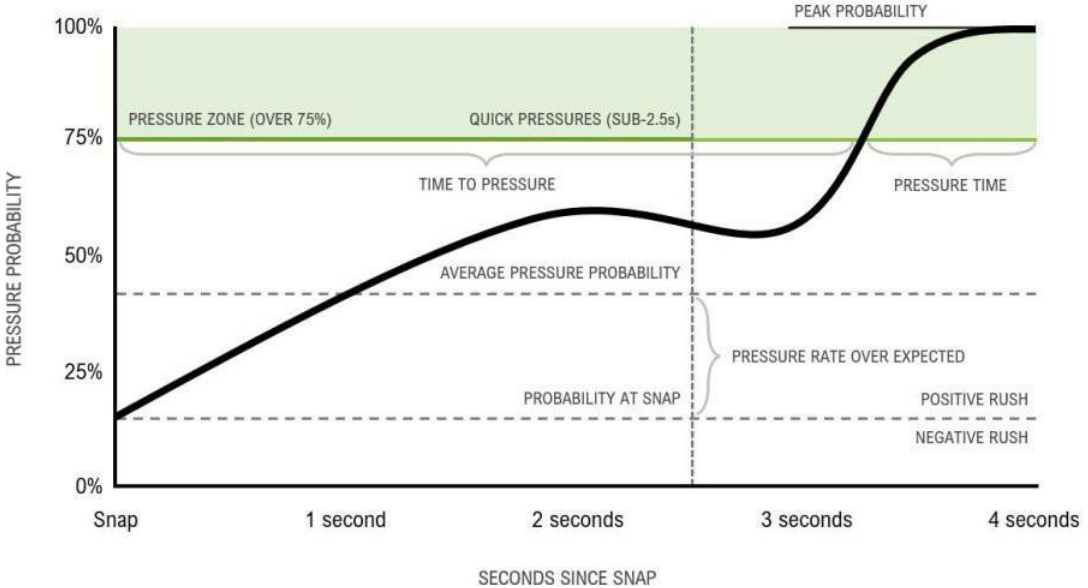
Some Real-World Examples




© 2024 Cisco and/or its affiliates. All rights reserved.

The Anatomy of a Pressure

A breakdown of metrics derived from new Next Gen Stats Pressure Probability Model



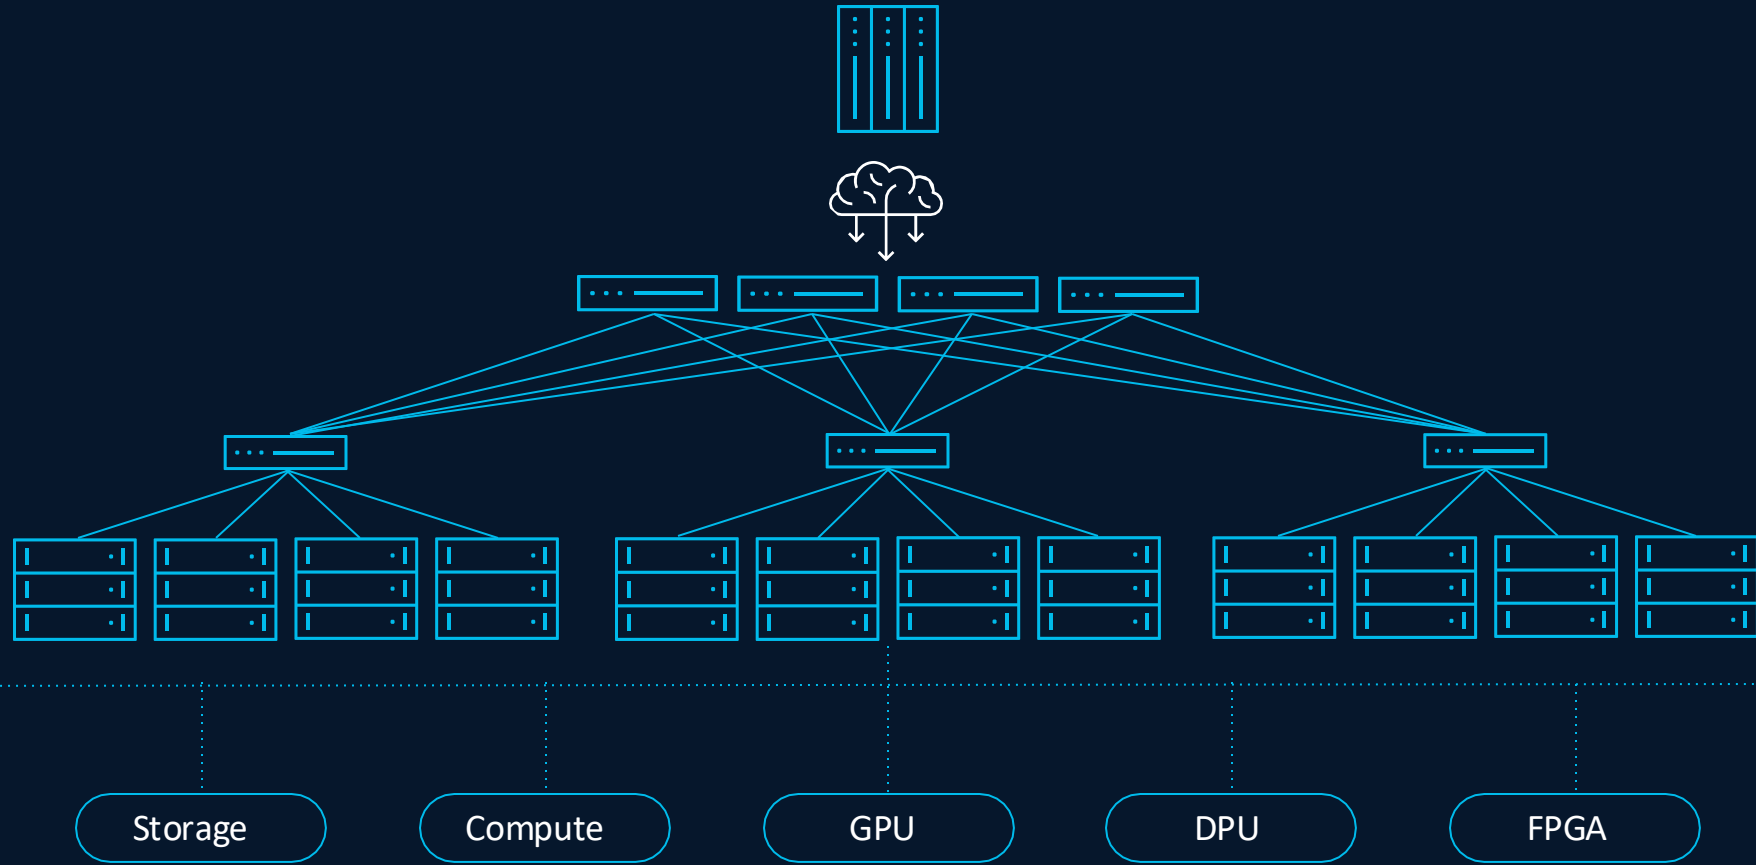
CORNER ANALYSIS - TURN 1



AI/ML Architecture

Lossless, High-Throughput, Low Jitter, Low-Latency Connectivity Fabric

25G/50G/100G/400G/800G



AI Cluster Types – Fabric Requirements

Training

Sometimes called “back-end” Network

Primary component responsible for the AI system's main tasks, such as making predictions, generating outputs, or performing analysis

Tightly Coupled Processes

RDMA (High Bandwidth Flows)

Low Jitter Tolerance

Nonblocking Topologies

Inferencing

Sometimes called “front-end” Network

Receiving and processing input data to be fed into the AI system. Also interacts with external environment or user


Loosely Coupled Applications

TCP (Low Bandwidth Flows)

High Jitter Tolerance

Oversubscribed Technologies

Deployment Options for Ethernet Fabrics



Size of AI clusters (small to large)

CPU

intel. AMD


CPU + GPU

nvidia. intel. AMD

Network is key to AI/ML training job success

Considerations:

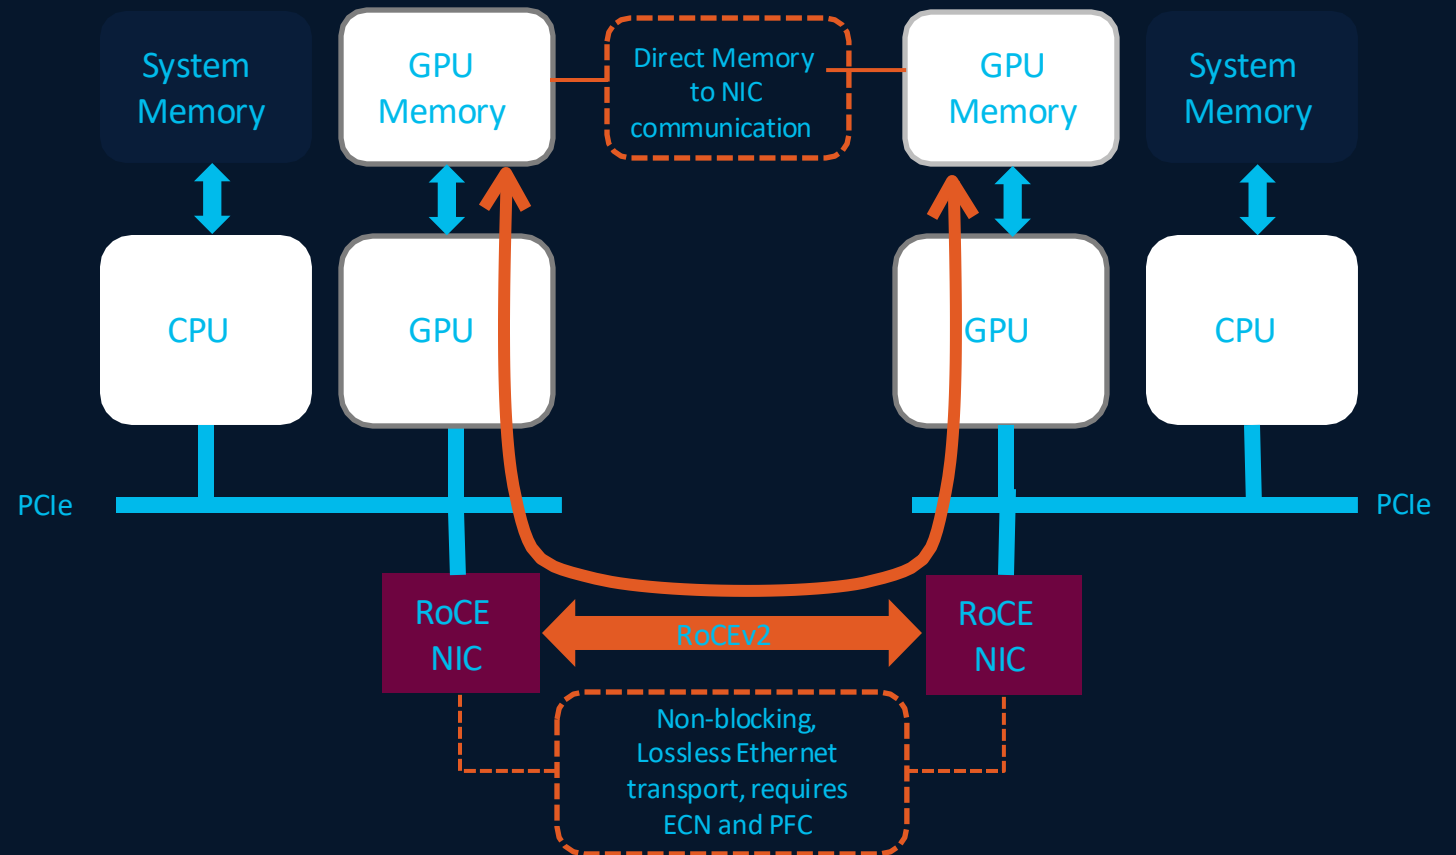
- Jitter / Congestion
- Tail-end Latency / Idling
- Poor performance
- On average 25% of jobs fail
- Cost of a job
- Expensive, wasted resources/time



Bringing RDMA to Ethernet


RDMA over Converged Ethernet (RoCE)

- RoCE enables the RDMA data exchange over lossless Ethernet fabrics
- Provides reliable alternative to RDMA over InfiniBand fabrics
- Manages congestion using DCQCN (ECN and PFC)
- RoCE Version 2 introduced support for L3 routing



Explicit Congestion Notification (ECN)


- IP Explicit Congestion Notification (ECN) is used for congestion notification.
- ECN enables end-to-end congestion notification between two endpoints on IP network
- ECN uses 2 LSB of Type of Service field in IP header



| ECN | ECN Behavior |
|-----|---------------------------|
| 00 | Non ECN Capable |
| 10 | ECN Capable Transport (0) |
| 01 | ECN Capable Transport (1) |
| 11 | Congestion Encountered |

Explicit Congestion Notification (ECN)

- IP Explicit Congestion Notification (ECN) is used for congestion notification.
- ECN enables end-to-end congestion notification between two endpoints on IP network
- ECN uses 2 LSB of Type of Service field in IP header
- In case of congestion, ECN gets transmitting device to reduce transmission rate using Congestion Notification Packet (CNP) without pausing traffic.

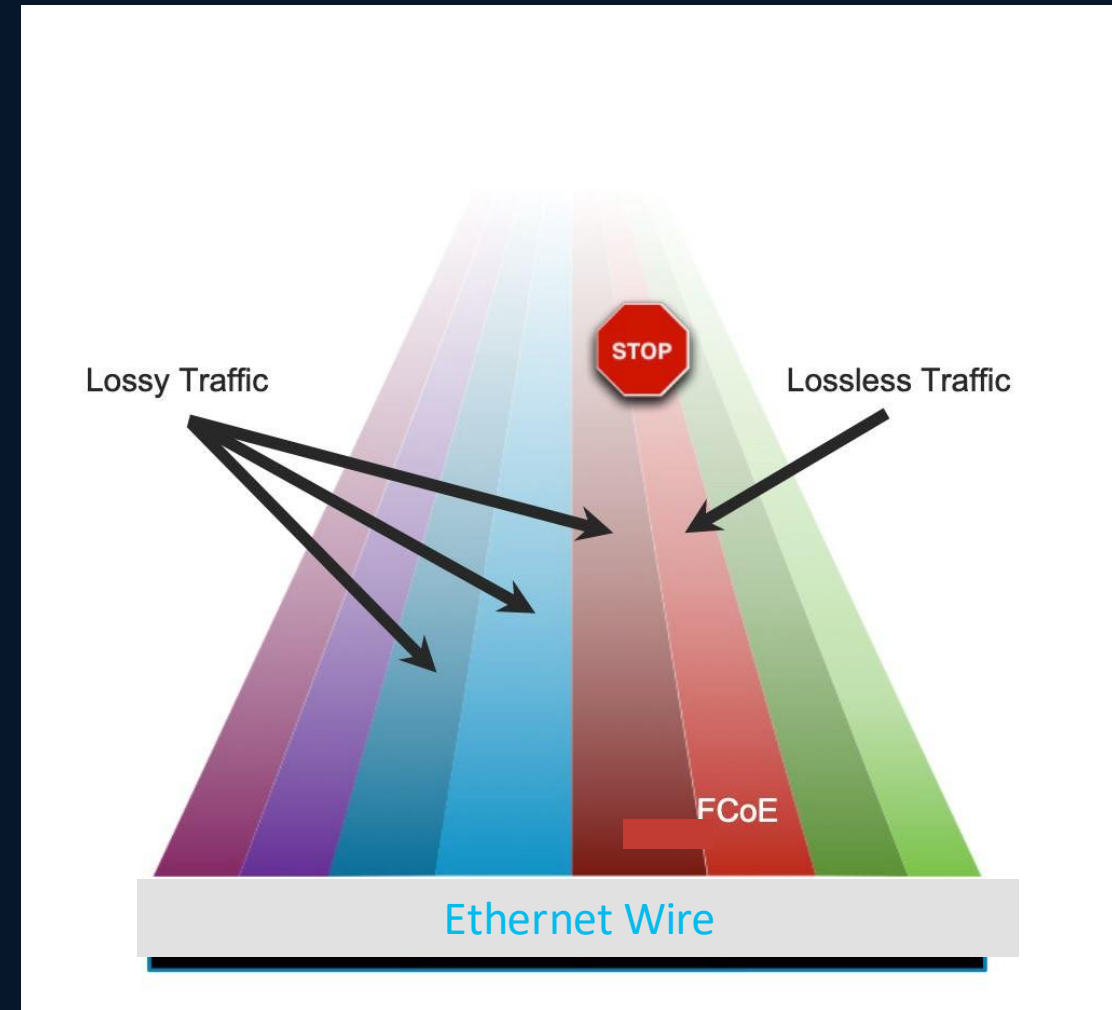


| ECN | ECN Behavior |
|-----|---------------------------|
| 00 | Non ECN Capable |
| 10 | ECN Capable Transport (0) |
| 01 | ECN Capable Transport (1) |
| 11 | Congestion Encountered |

Priority Flow Control


Flow Control Mechanism – 802.1Qbb

- A.k.a "Lossless Ethernet"
- PFC enables Flow Control on a Per-Priority basis
- PFC is also called Per-Priority-Pause
- Therefore, we have the ability to have lossless and lossy priorities at the same time on the same wire
- Allows traffic to operate over a lossless priority independent of other priorities
- Other traffic assigned to other priority will continue to transmit and rely on upper layer protocols for retransmission



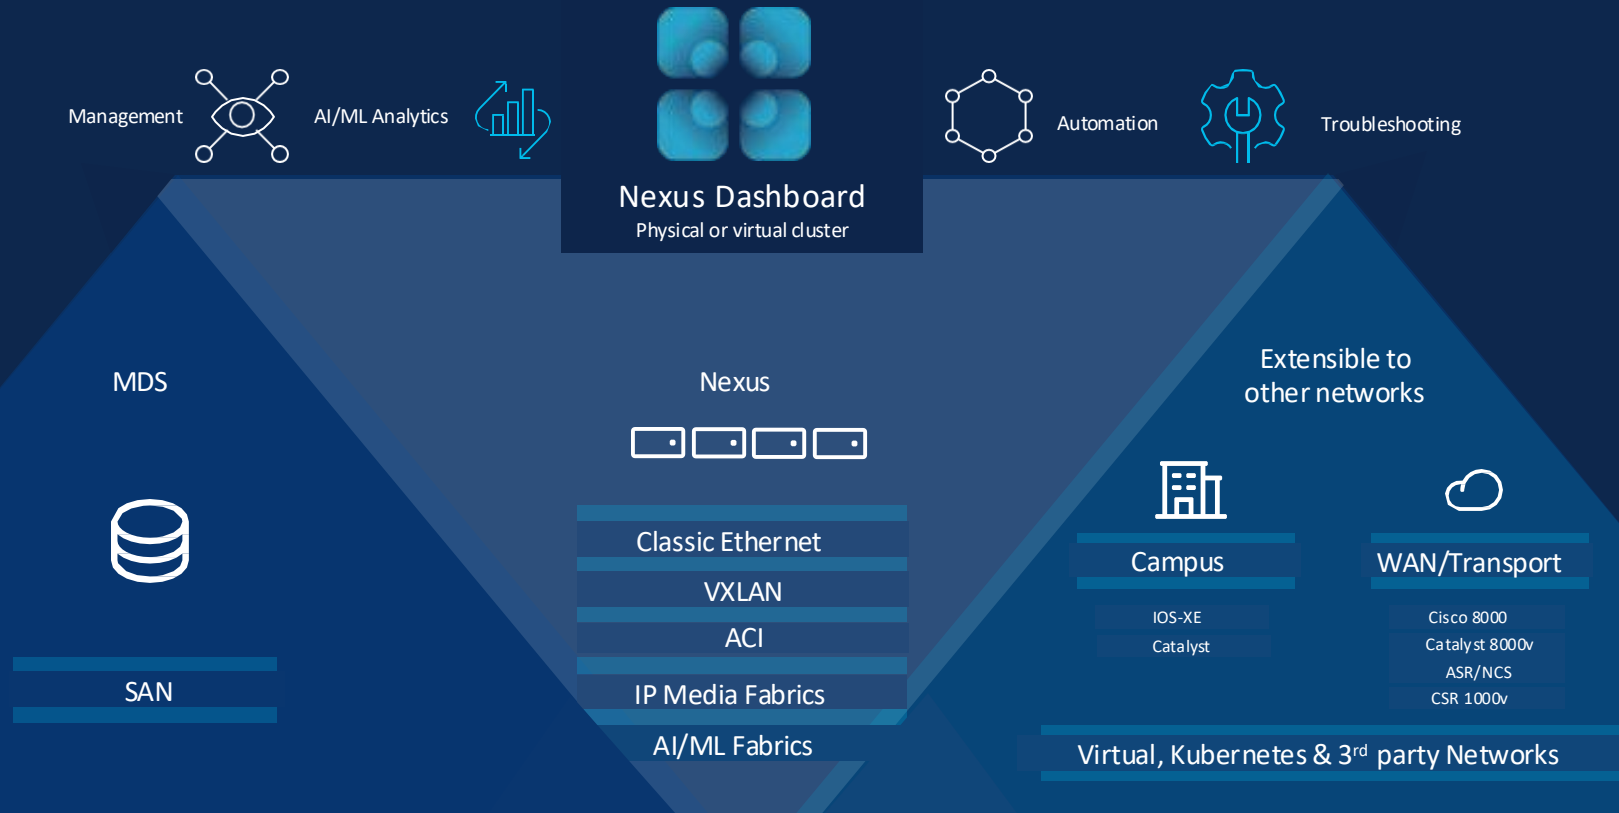
PFC – How pause frames are sent

- PFC sets thresholds in no-drop queue
- Headroom is present to accommodate “in flight” packets
- Under congestion, traffic is buffered in non-drop queue
- PFC frames are sent toward sender after queue utilization exceeds *xoff* threshold
- While draining the queue, and utilization is below *xon* threshold system will stop sending PFC frames



So how do we manage all of this?

Unify your operational experience



The **simplest** way to manage and operate data center connectivity

The power of one Nexus fabric experience



Common data plane - VXLAN

Common security policies – Security Groups

Common interconnect – Automated provisioning and Border Gateways (BGWs)



Nexus 9000



NX-OS and ACI



Nexus Dashboard delivers **ONE** fabric experience and **consistent operations** for NX-OS and ACI fabrics¹

Common Provisioning

AI/ML Data Center Blueprint for Networking

Addressing needs of the most demanding datacenter production networks



Best performing AI/ML networks



Intelligent buffering, low latency, telemetry/visibility, RoCEv2



Easily build lossless fabrics



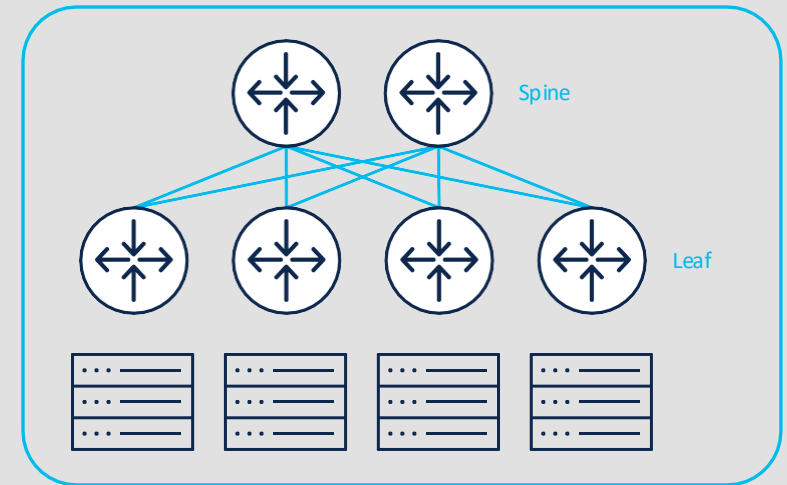
Efficient application performance



Lower TCO



Validated designs for network and ecosystem partners

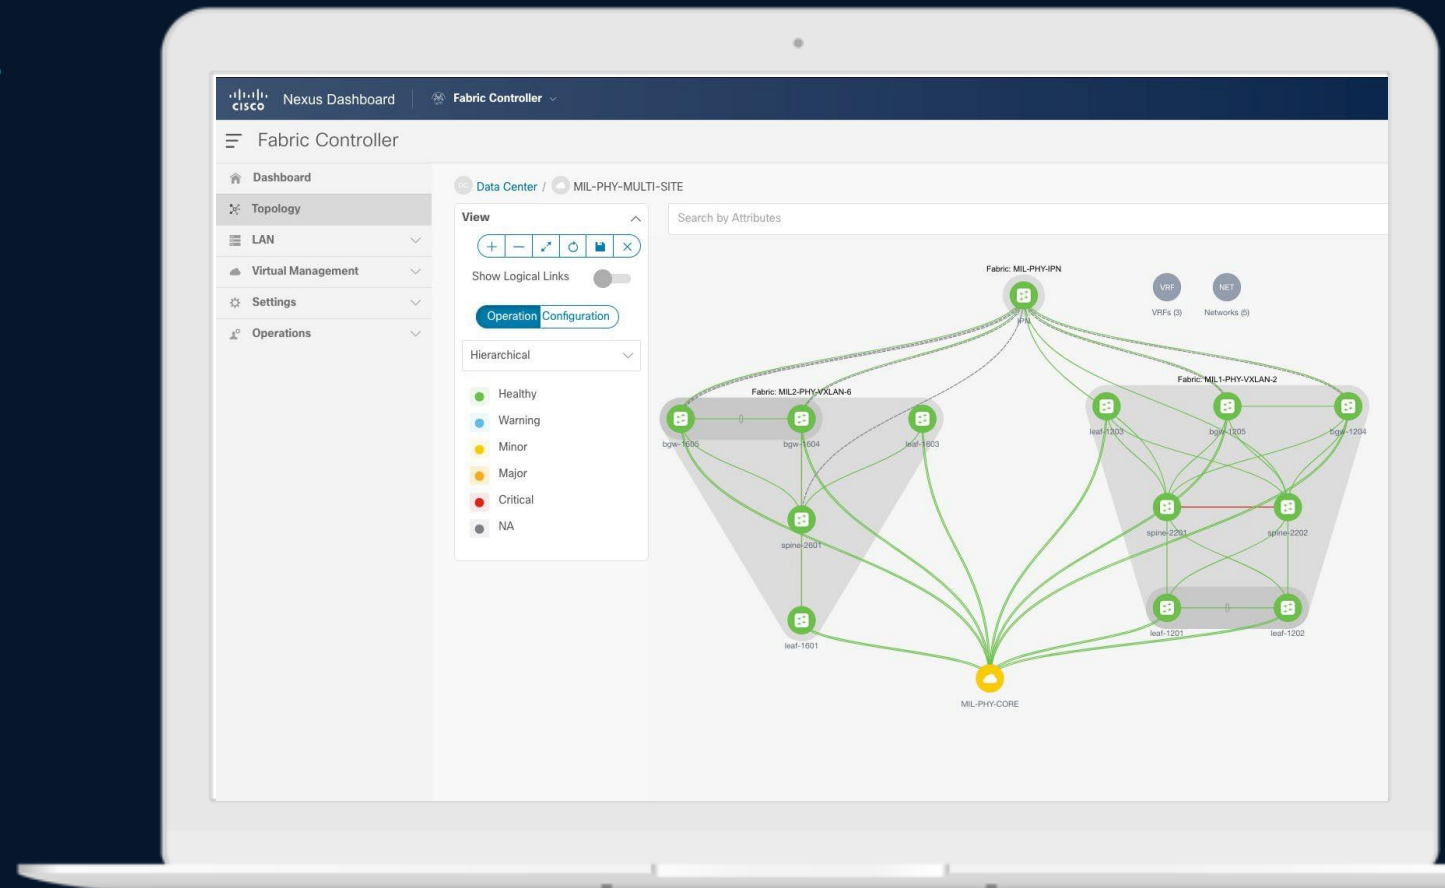


Nexus Dashboard Fabric Controller

Deploy using best practices with built-in AI/ML QoS policy templates

Key capabilities:

- Fabric provisioning and automation
- Underlay and Overlay definition
- Image management
- Change control
- Multi-site architecture
- Manages multiple fabrics, running multiple architectures, at the same time



Nexus Dashboard Fabric Controller – AI/ML Network

100Gb Leaf QoS Generated Config

```
class-map type qos match-any ROCEv2
  match dscp 26
class-map type qos match-any CNP
  match dscp 48
```

```
policy-map type qos QOS_CLASSIFICATION
  class ROCEv2
    set qos-group 3
  class CNP
    set qos-group 7
  class class-default
    set qos-group 0
```

```
policy-map type queuing QOS_EGRESS_PORT
  class type queuing c-out-8q-q6
    bandwidth remaining percent 0
  class type queuing c-out-8q-q5
    bandwidth remaining percent 0
  class type queuing c-out-8q-q4
```

100Gb Spine QoS Generated Config

```
class-map type qos match-any ROCEv2
  match dscp 26
class-map type qos match-any CNP
  match dscp 48
```

```
policy-map type qos QOS_CLASSIFICATION
  class ROCEv2
    set qos-group 3
  class CNP
    set qos-group 7
  class class-default
    set qos-group 0
```

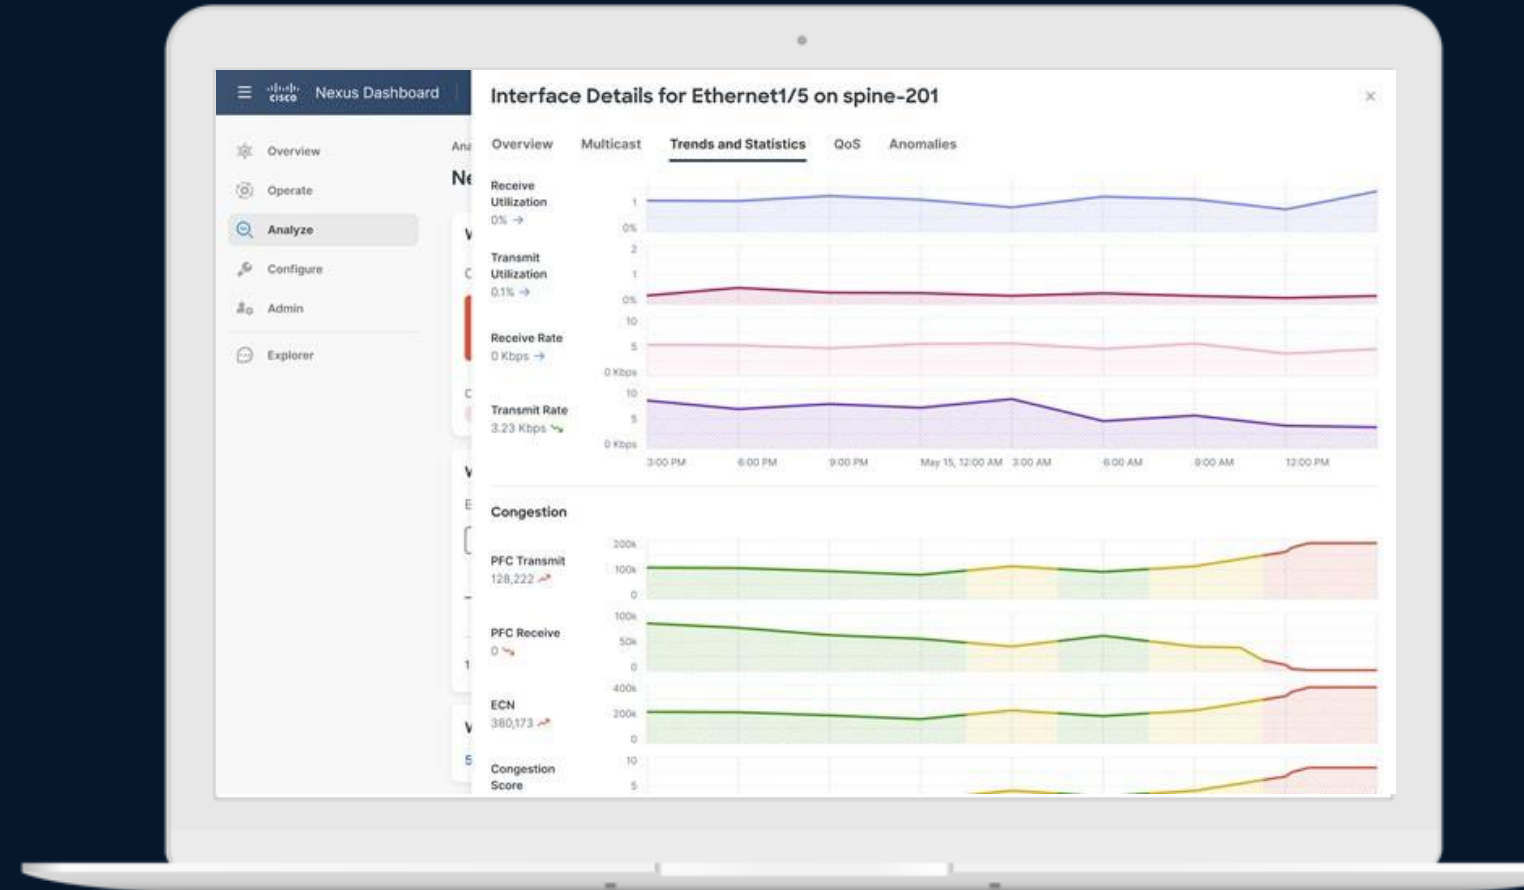
```
policy-map type queuing QOS_EGRESS_PORT
  class type queuing c-out-8q-q6
    bandwidth remaining percent 0
  class type queuing c-out-8q-q5
    bandwidth remaining percent 0
  class type queuing c-out-8q-q4
```

Nexus Dashboard Insights - AI/ML Networks






Nexus ASICs stream AI/ML DCQCN telemetry directly to Nexus Dashboard Insights



AI Network Visibility

- UI Dashboard
- Visibility – Lossless Ethernet
- Monitoring (ECN & PFC)
- Congestion Score
- Application to network performance correlation
- Telemetry and Netops



Nexus Dashboard Insights – Congestion Visibility

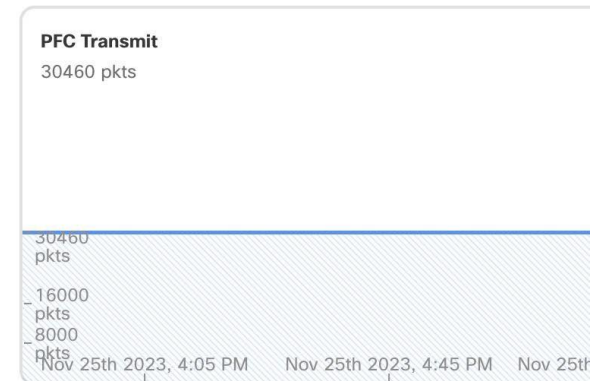
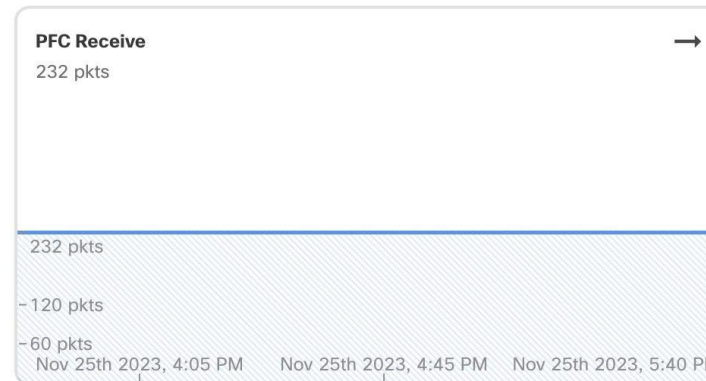

-  Overview
-  Operate
-  Analyze
-  Configure
-  Admin

-  Explore
-  Bookmarks

Interface Details for eth1/3 on RoCE-Spine-2

Overview Trends and Statistics Anomalies

Congestion



Microbursts

Microbursts by

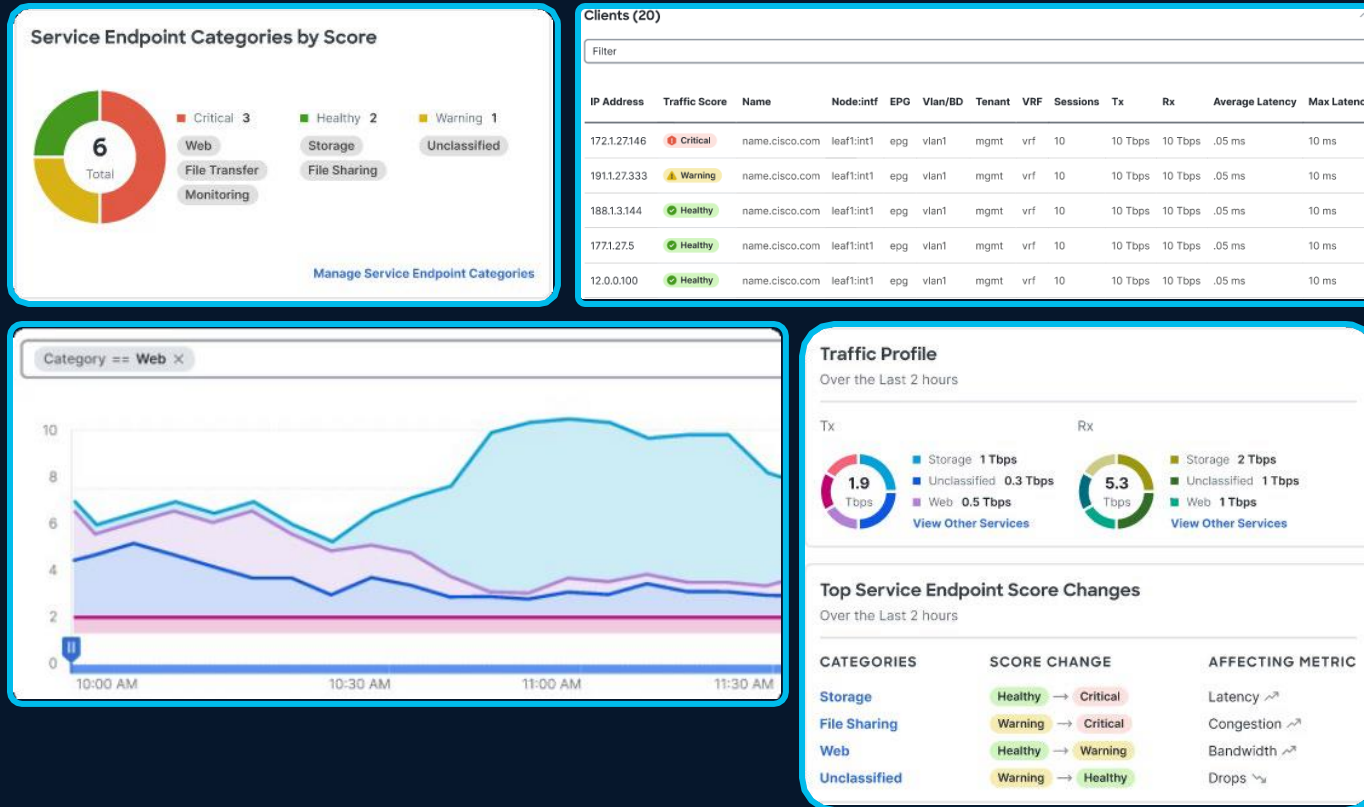
| Queue | Start Time | Number of Bursts | Max Duration (ns) | Avg Duration (ns) | Max Peak | Avg Peak |
|---------|-----------------------------------|------------------|-------------------|-------------------|-----------|-----------|
| queue-3 | Nov 25 2023 05:40:00.000000 PM | 60 | 466.61 ns | 185.26 ns | 2,231,424 | 1,424,765 |

10 Rows

Nexus Dashboard Insights

Now available

Traffic Analytics – ACI and NX-OS support



Congestion, Latency, Drops

Customize TCP categories

Detect potential issues, top talkers faster Auto-

discovered services (Web, FTP, etc) Pervasive -

Filters/Rules not needed

Requires: ACI – 6.1.1 and NX-OS 10.4(2F)

Benefits

One (physical), three, and six node cluster options. FX and newer platforms with one week of data retention

Takeaways

- Robust back-end network, and flexible front-end network
- Familiar data center fabric technologies, BGP or VXLAN
- Automate network for easier bring up and operation
- Visibility in network congestion, and bottleneck to troubleshoot and optimize

