The bridge to possible

Cisco Silicon One

Innovation at Speed and Scale

Andreas Bovopoulos
Director of Technical Marketing

BRKARC-1011

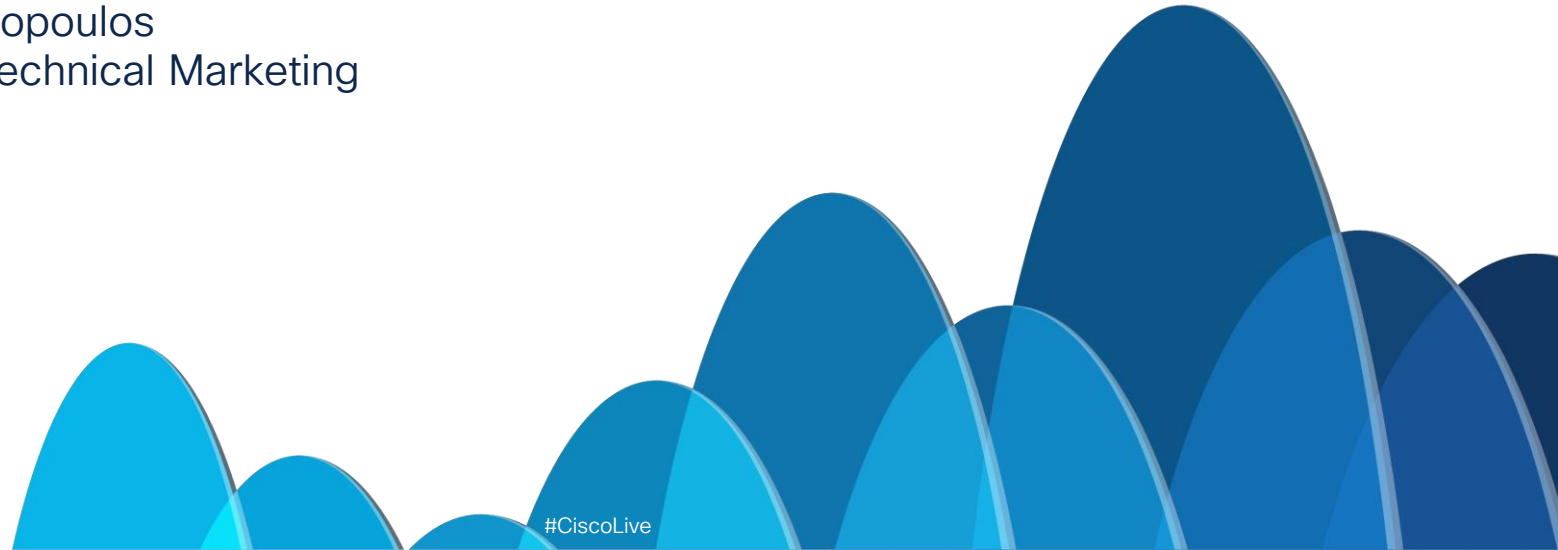

cisco Live!

Cisco Webex App

https://ciscolive.ciscoevents.com/ciscolivebot/BRKARC-1011

Questions?

Use Cisco Webex App to chat with the speaker after the session

How

1. Find this session in the Cisco Live Mobile App
2. Click “Join the Discussion”
3. Install the Webex App or go directly to the Webex space
4. Enter messages/questions in the Webex space

Webex spaces will be moderated by the speaker until June 7, 2024.

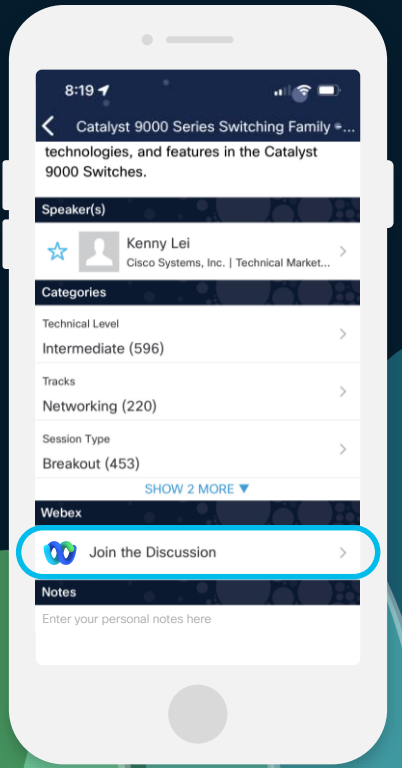

Agenda

- Silicon Industry and Silicon One
- Business Models
- Architecture Highlights
- Silicon One Product Line
- Cisco Silicon One & AI/ML
- Summary

Cisco Silicon One

Breaking the mold

If you approach a problem with the same organization and the same technology, you'll get the same outcome

The Fundamental Change

Then

Requirements drive unique architectures

Business models drive unique architectures

Merchant | Custom

Technical debt prohibits new architectures

=

20

Architectures

Now

CISCO Silicon One = One Architecture

Requirements drive unique Devices not architectures

Cisco enables multiple business models, Erase merchant vs. custom

- ✓ Silicon Only
- ✓ White Box
- ✓ Full System

Clean sheet, Investment, & Innovations enable One Architecture

NEW

Clean Sheet ~~Technical Debt~~

Invested over $1B

Hardware Innovations

Accelerated Innovation | Accelerated Deployment | Easier Maintenance

A better and easier to maintain Network

Cisco Silicon One: Accelerating Innovation

Massive Investment

Over 8 years of research and over a billion dollars

New Organization

Focused on building One Architecture **across the network** and **across business models**

New Technology

Invented **new** technology to allow convergence of routing and switching

Fundamental change in the industry

First truly scalable network silicon architecture

Agenda

- Silicon Industry and Silicon One
- ➢ Business Models
- Architecture Highlights
- Silicon One Product Line
- Cisco Silicon One & AI/ML
- Summary

Multiple Engagement Models

Cisco Silicon One – Multiple Business Models

3rd Party Provided

Cisco Provided

Full System | Hardware Only | Silicon Only

Agenda

- Silicon Industry and Silicon One
- Business Models
- ➢ Architecture Highlights
- Silicon One Product Line
- Cisco Silicon One & AI/ML
- Summary

Key Enabling Technology: Slice-Based

New Dataplane Architecture

Fully Scalable

New category of routing silicon

Common approach with switch silicon
Previously impossible for routing silicon

Routing at Switch Bandwidth

Key Enabling Technology: Network Processor

New Processing Architecture

Fully Scalable

Family Optimized P4 programmable NPU

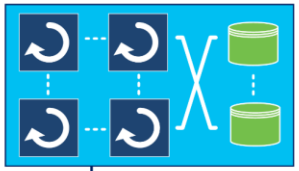

Hardware Accelerated Databases

Routing with Switching Efficiency

Switching with Router Programmability

Key Enabling Technology: Traffic Management and Buffering

New Memory Architecture

Fully Scalable

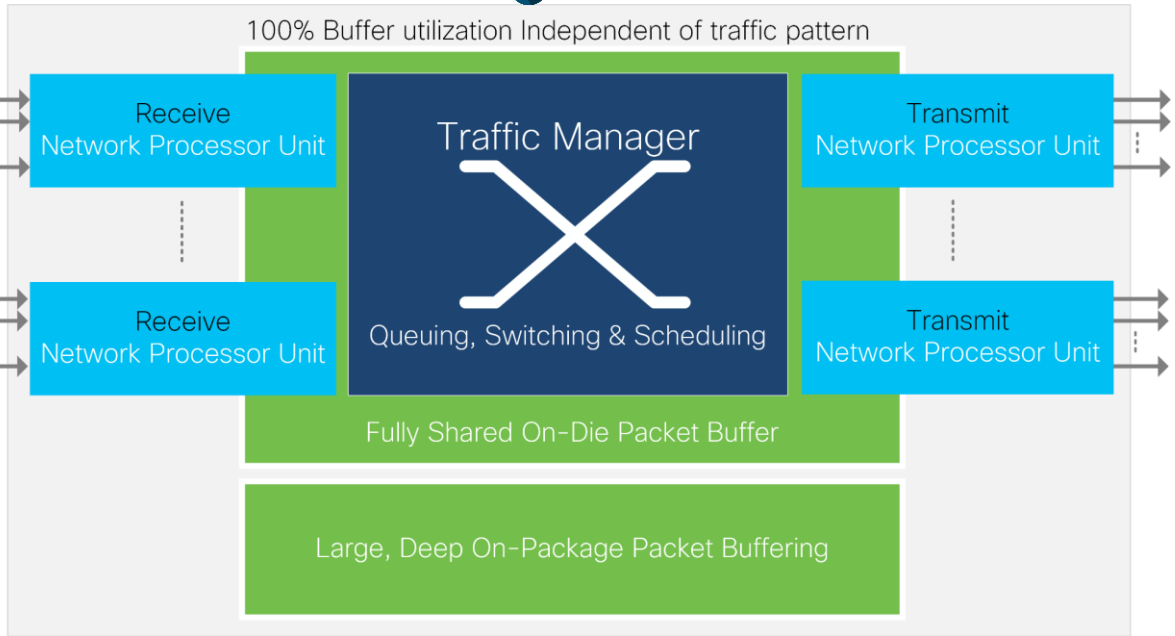

New levels of efficiency

No segmentation of packet buffer
Ultimate burst performance
Write once, read once, data doesn't move

Highly Efficient Router
Highly Efficient Switch

Cisco Silicon One: Endlessly Scalable Architecture

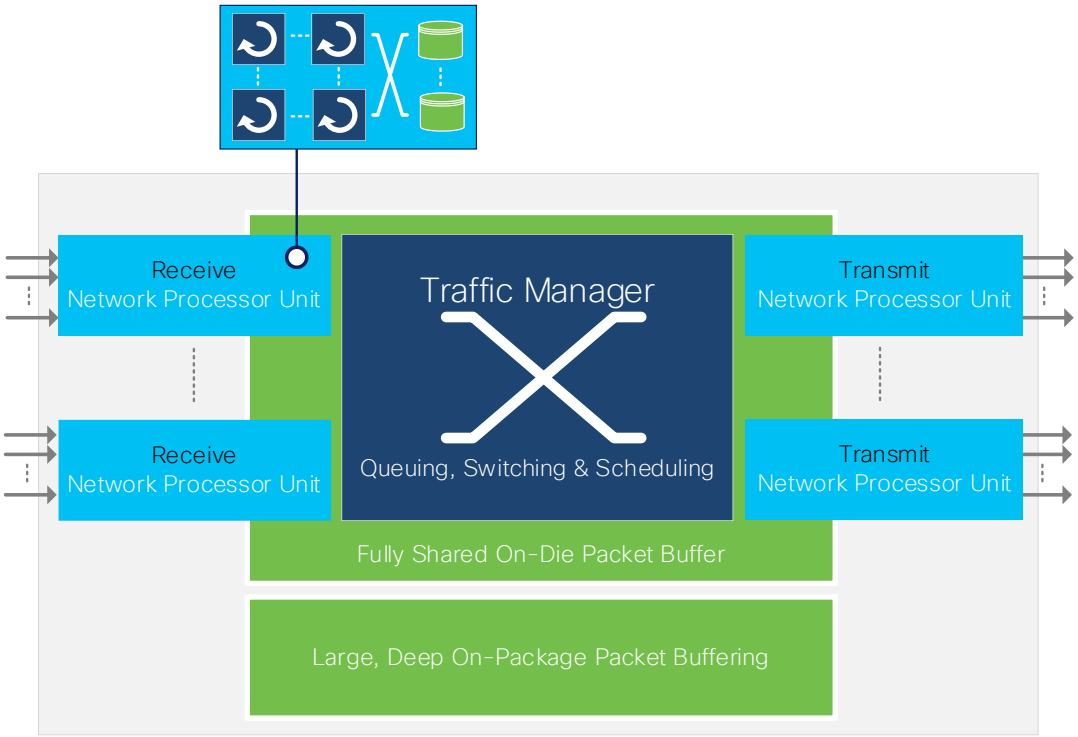

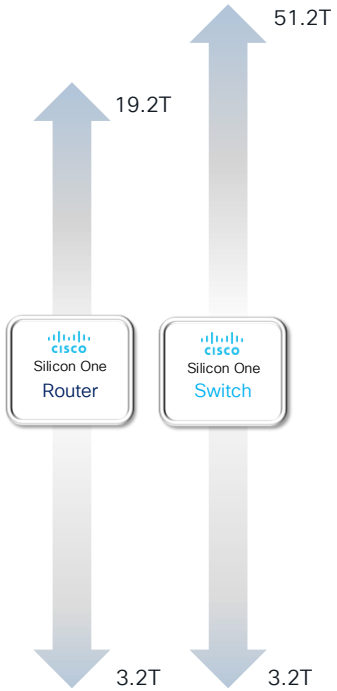

First truly scalable networking silicon

Agenda

- Silicon Industry and Silicon One
- Business Models
- Architecture Highlights
- ➢ Silicon One Product Line
- Cisco Silicon One & AI/ML
- Summary

Cisco Silicon One: One Architecture, Many Devices

Cisco Silicon One

Expanding over time

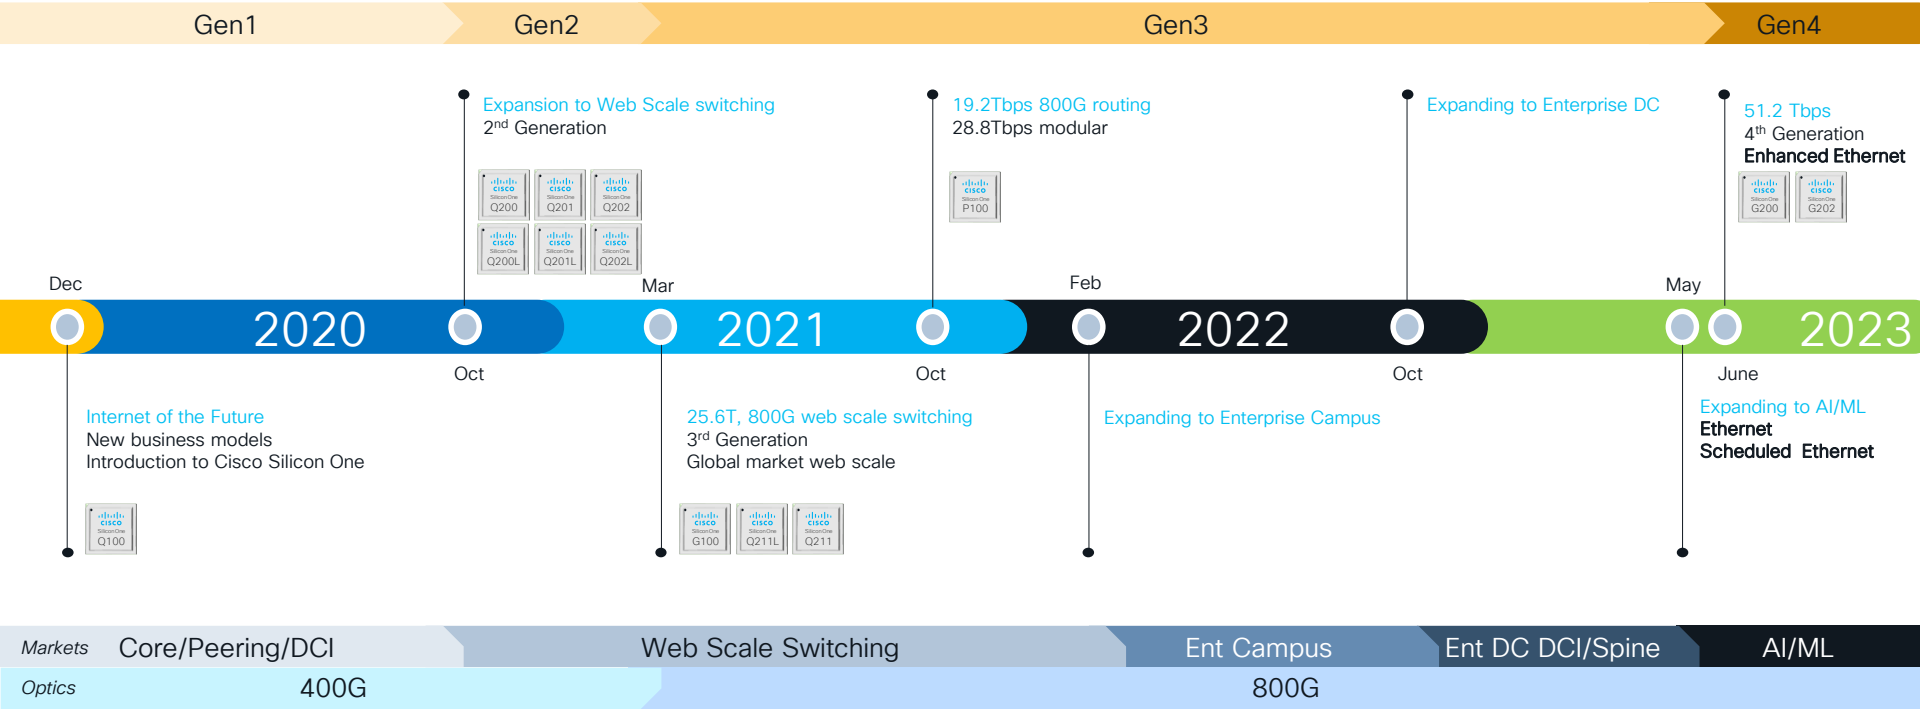

Cisco Silicon One Device Roadmap

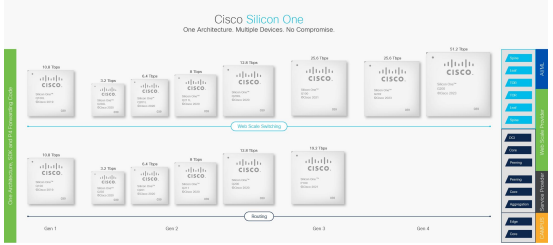

Gen1 | Gen2 | Gen3 | Gen4

16nm | 7nm | 5nm

Switching
Enhanced Ethernet AI

Routing
Scheduled Ethernet

Silicon One Q100/L

Silicon One Q2xx/L

Silicon One G100

Silicon One G200

Silicon One P100

56G SerDes
(400G optics)

112G SerDes
(800G optics)

Cisco Silicon One

Convergence without compromise

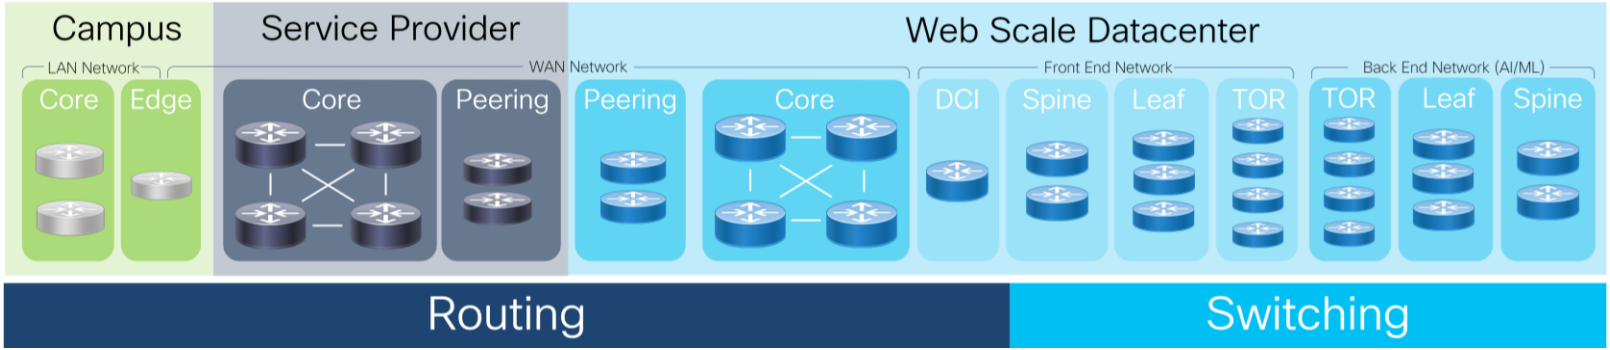

cisco
Silicon One

One architecture
One SDK
One form factor
One design

Unmatched programmability
Unmatched performance
Unmatched efficiency
Unmatched flexibility

Multiple devices
Multiple markets
Multiple business models

One network
One experience

One experience, across the network, without compromise

Enabled System Architectures

Fixed

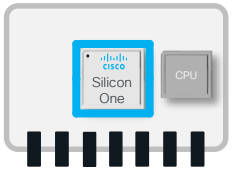

Standalone

Centralized

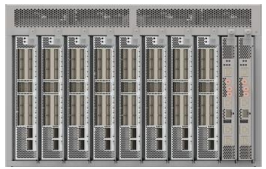

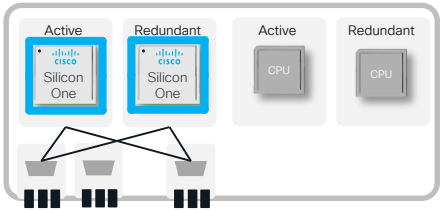

Standalone

Distributed

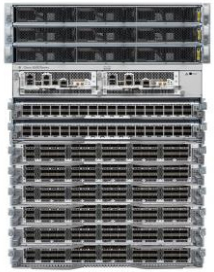

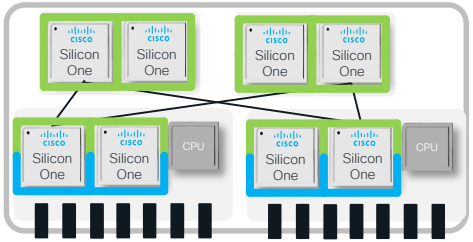

Fabric Element + Linecard

Disaggregated

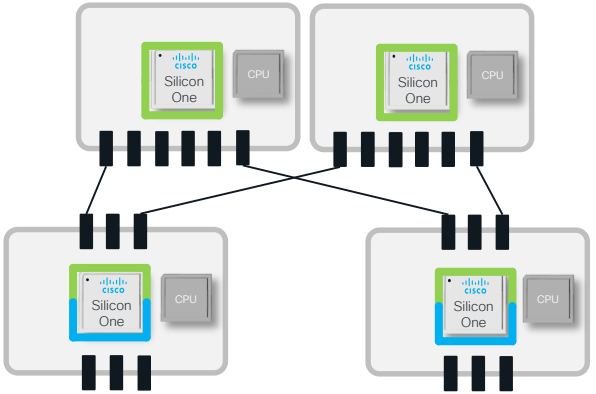

Fabric Element + Linecard

Cisco Silicon One – Modes of Operation

Standalone (SA)

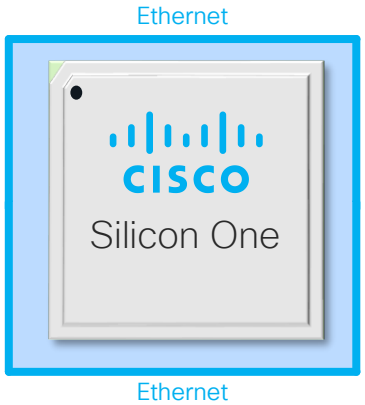

Linecard (LC)

Oversubscribed Linecard (LC)

Fabric Element (FE)

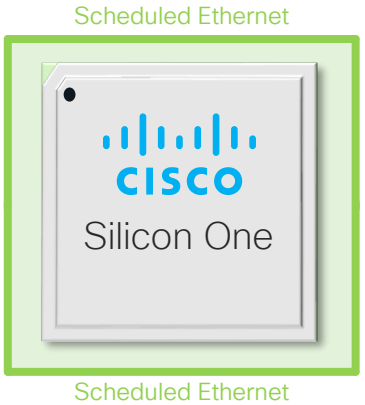

Device	Generation	Mode(s)
P100	3rd	LC, SA
Q200	2nd	LC, SA, FE
Q100	1st	LC, SA, FE
Q211	2nd	SA
Q201	2nd	SA
Q202	2nd	SA

Device	Generation	Mode(s)
G200	4th	SA
G202	4th	SA
G100	3rd	SA, FE
Q200L	2nd	LC, SA, FE
Q100L	1st	LC, SA, FE
Q211L	2nd	SA
Q201L	2nd	SA
Q202L	2nd	SA

Cisco Silicon One – Publicly Announced Systems

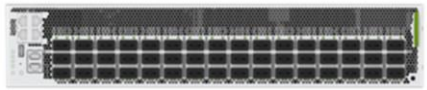

OCP Wedge400C by an ODM

Enterprise Datacenters (Nexus)

Cisco OS

Modular: 9800

SP and WebScale (8000 Series)

Cisco OS

SONiC

Modular: 8804/08/12/18
Fixed: 810x, 820x

Campus Core (Catalyst)

Cisco OS

C9500X

C9600X

Agenda

- Silicon Industry and Silicon One
- Business Models
- Architecture Highlights
- Silicon One Product Line
- ➢ Cisco Silicon One & AI/ML
- Summary

AI/ML Networks demand more from the network

AI/ML Workloads are unique

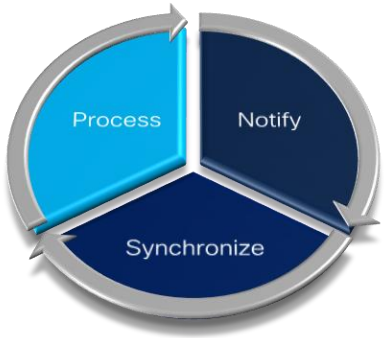

High Bandwidth flows
Long Lived flows
Synchronous
Massive Scale

Major Challenges to Solve

Load balance packets
How do you distribute packets across multiple paths

Congestion Reaction or Avoidance
How do you handle congestion caused by traffic pattern or load balancing decisions

Solving
Job Completion Time
(JCT)

Link Failure Avoidance
How do you recover from link failures?

Solving
Reliability at Scale

1 Ethernet

Standard but non-deterministic performance

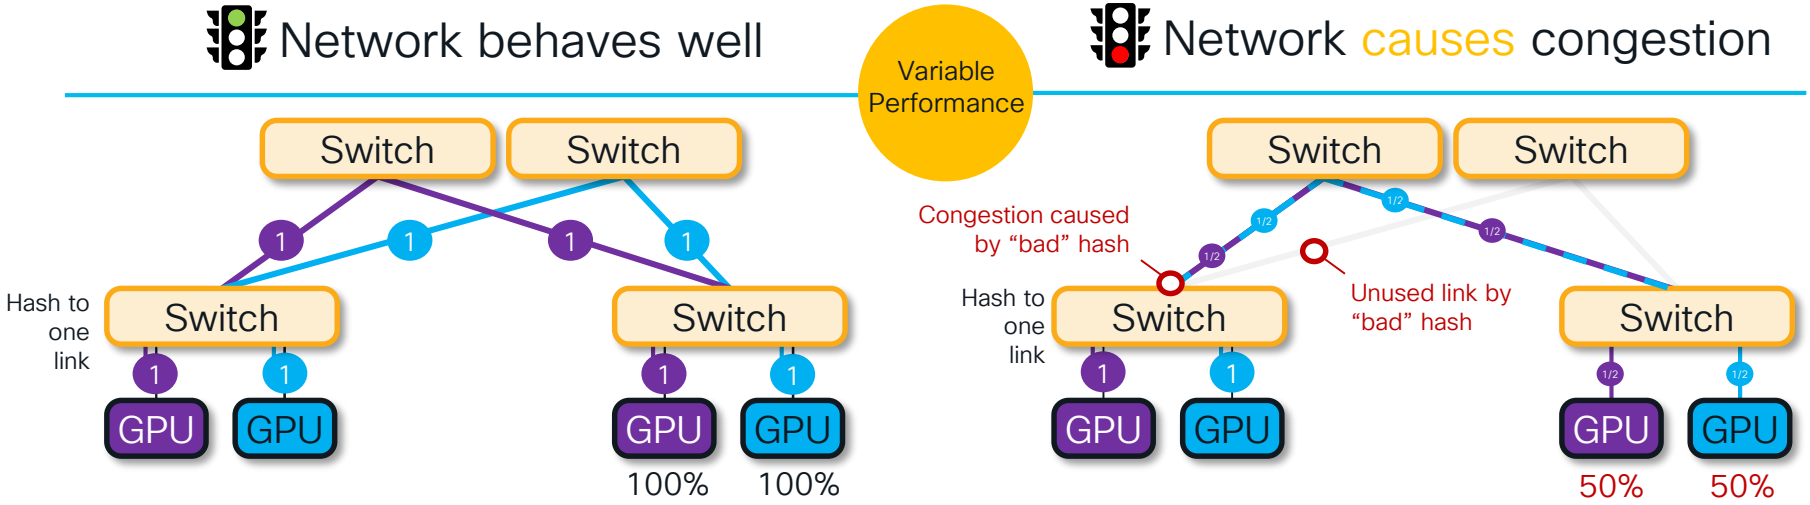

- Most traffic requires in-order delivery
- Ethernet ensures all packets of a flow go through one path
  - ECMP Hash packet header to identify flow → Map to outgoing link
  - Flow characteristics cause congestion even when unused BW is available
  - This problem gets worse when you have high bandwidth & long-lasting flows
- Training has high bandwidth & long-lasting flows, and GPUs stall for slowest path
  - One bad ECMP load balancing decision, stalls the entire cluster

Load Balancing causes congestion

AI/ML traffic characteristics magnify this situation

ECMP = Equal-Cost Multi-Path

3 Scheduled Ethernet / Distributed Scheduled Fabric (DSF)

Standard to the GPU but Proprietary between switches with Deterministic performance

Deterministic Performance

Network always behaves well

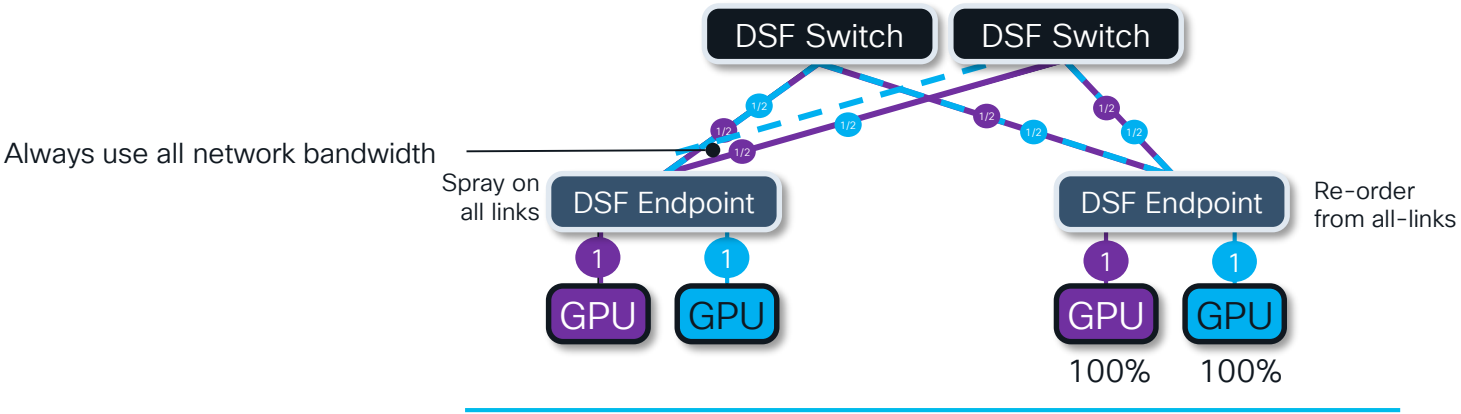

- Most traffic requires in-order delivery
- DSF **sprays packets** on all links to go through all paths
  - Ensures the network is not the bottleneck
  - DSF re-orders packets on the egress switch to deliver packets in-order to the GPUs
  - Enables wide radix deployments with smaller links to enable power optimized networks (Significant cost & power savings)
- DSF includes other important features
  - **Hardware based failure** detection and reachability
  - **Credit based scheduler** protects against incast events

Cisco Silicon One & AI/ML Infrastructure

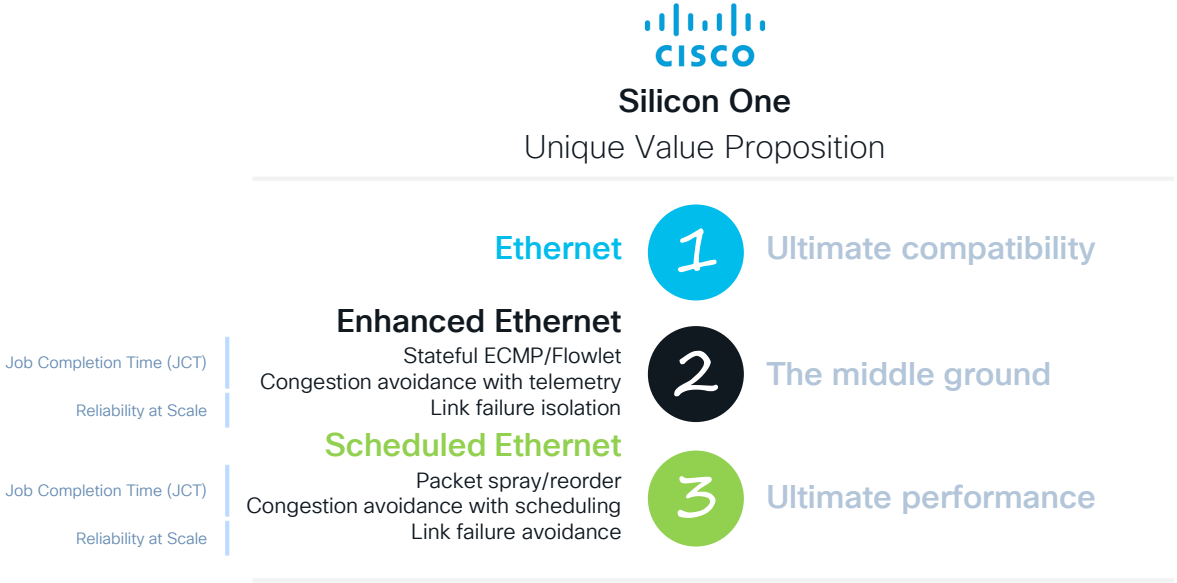

Cisco Silicon One convergence directly impacts AI/ML

One architecture, across 3 deployment models

Cisco Silicon One – AI/ML Fabric Options

	1 Ethernet	2 Enhanced Ethernet		3 Scheduled Ethernet
Load Balance	Stateless ECMP	Stateful Flow/Flowlet	Spray & re-order (optionally) in smartNIC	Spray & Re-order in leaf
Network Congestion Management	Congestion Reaction with ECN/PFC	Congestion Reaction with congestion score to adjust distribution	Congestion Reaction with congestion score to adjust distribution	Congestion Avoidance with fully scheduled Ethernet
Link Failure	Software	Hardware		
Job Completion Time	Good	Better	Best	
Tight coupling between NIC + Network	No	Yes	No	

Cisco Silicon One

Ethernet, Enhanced Ethernet, Scheduled Ethernet

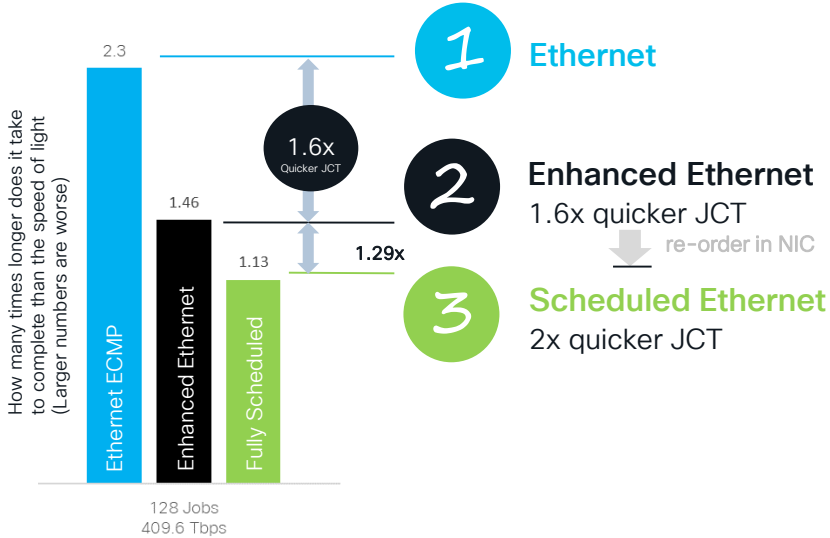

cisco
Silicon One

>2x
Compute for the same network

AI/ML Workload characteristics

High performance GPUs
Admissible traffic pattern
High bandwidth flows
Long lived flows
Job wide synchronization with barrier operation
Multiple jobs running concurrently on same infrastructure

Network performance is critical

One bad load balancing decision effects all GPUs in the job

Cisco AI Networking Solution

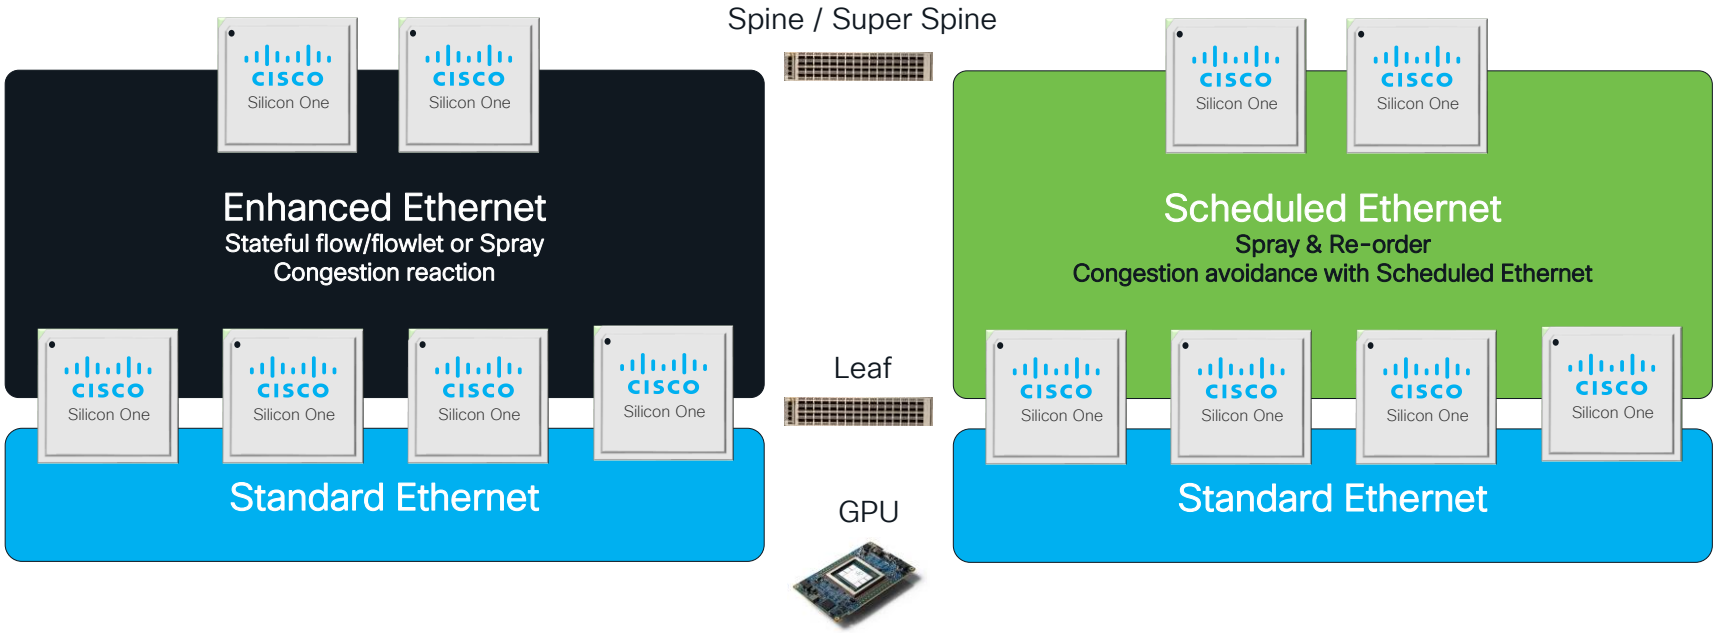

Agenda

- Silicon Industry and Silicon One
- Business Models
- Architecture Highlights
- Silicon One Product Line
- Cisco Silicon One & AI/ML
- ➢ Summary

Cisco Silicon One Value Proposition

Convergence without Compromise

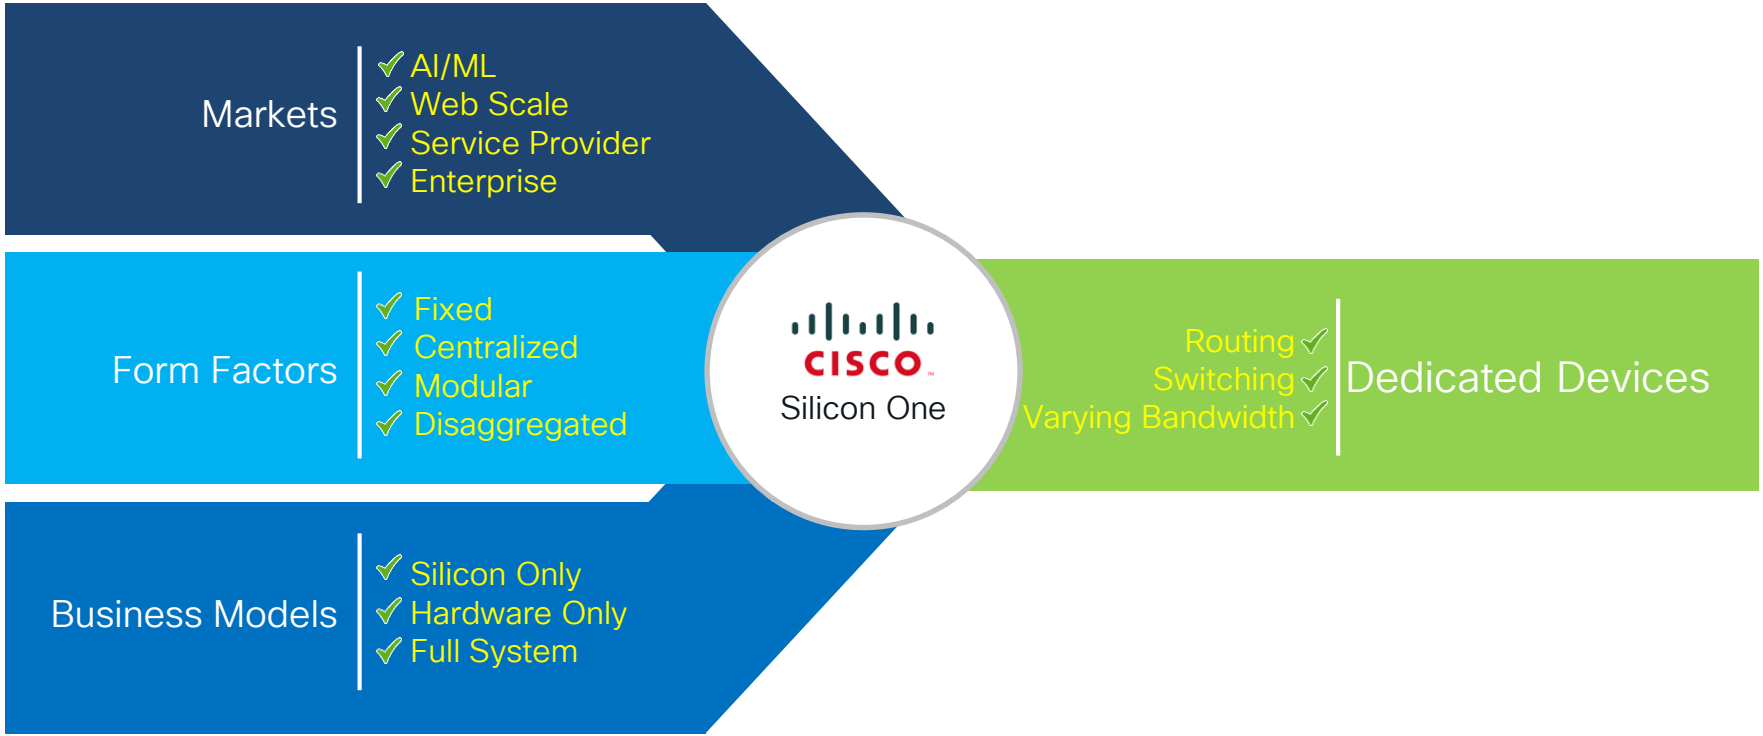

Why adopt Cisco Silicon One?

Minimize OPEX and CAPEX

- Convergence enables efficient qualification, deployment and operations
- Port the SDK/SAI Abstraction once, deploy everywhere
  - Across routing and switching
  - Across form-factor

✓ Deploy across whole network

✓ Simplify operations

Best Silicon Architecture

- Systems first design methodology
- Multiple, purpose built, devices
- Purpose built and Family optimized P4 Programmable NPU
- Rich Features and Telemetry
- Low Power
- Large and fully shared packet buffer with optional external packet buffer
- Scalable by design
- Only Cisco Silicon One extends Ethernet, to enhanced Ethernet and fully scheduled Ethernet
- Only Cisco Silicon One enables the same hardware systems and topology to run Ethernet, Enhanced Ethernet or Scheduled Ethernet

✓ Best Routing Silicon

✓ Best Switching Silicon

✓ Best for AI/ML

Massive Investment

- Largest investment in one networking silicon architecture
  - Fastest growing product line
  - Leadership position continuing to accelerate
- Huge investment in optics enabling future integration of optics

✓ Long term partner

Flexible Business Models

- Silicon Only, Hardware Only, Full System
- Commitment to your success

✓ Consumption on Your terms

Evolve Your Networks – Learn More

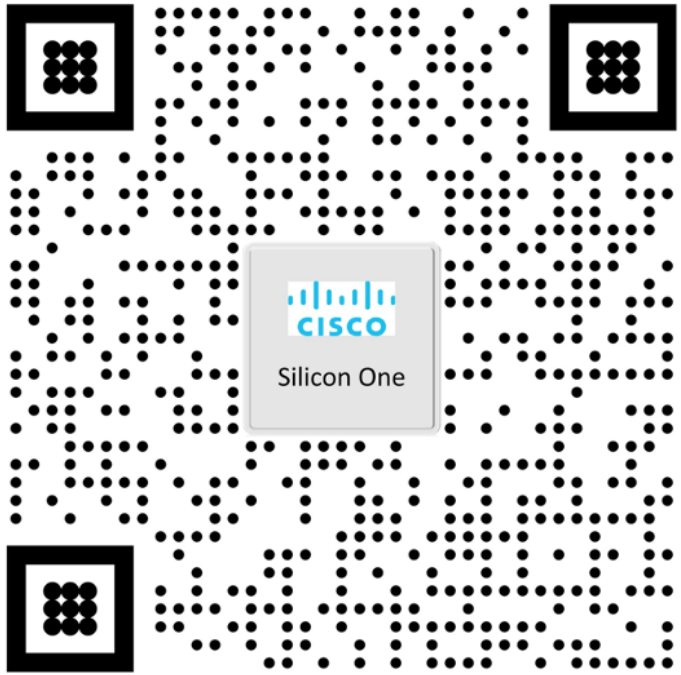

SiliconOne_inquiries@cisco.com

Continue your education

- Visit the Cisco Showcase for related demos
- Book your one-on-one Meet the Engineer meeting
- Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs
- Visit the On-Demand Library for more sessions at www.CiscoLive.com/on-demand

Contact me at: abov@cisco.com

Thank you