

Cisco Silicon One


Convergence without compromise

June 2023

Contents

Introduction	2
One architecture	3
One network	5
One form factor	6
Routing versus switching	7
Scheduled or unscheduled fabric	8
Machine learning and artificial intelligence	10
Conclusion	10
For more information	10

Cisco Silicon One



Introduction

Silicon is literally the foundation of every web-scale and service provider network. But up to now, customers have always had to choose a silicon architecture based on their specific requirements:

- Routing vs. web-scale switching
- Full featured vs. lean and mean
- Deep vs. shallow buffered
- Programmable vs. fixed function
- High vs. low scale
- Advanced vs. basic traffic management
- Fixed box vs. centralized box vs. modular line card vs. modular fabric card
- Scheduled vs. unscheduled fabric
- Customer type

Cisco Silicon One™ is a breakthrough technology that for the first time in history enables a single silicon architecture to erase these dividing lines and expand their reach to a massive portion of the networking market. Gone are the days where designers and network architects needed to invest in and learn multiple unique architectures in parallel. Support and operations teams

no longer need to train engineers on the behavior of all the unique silicon architectures and systems.

With our solution, network operators are no longer required to understand, qualify, deploy, and troubleshoot multiple disjointed architectures. Now you can learn and integrate a single architecture, design to a single Software Development Kit (SDK), and deploy it everywhere in the network faster and more simply. Support teams only need to understand one architecture, so they can troubleshoot issues more quickly. Network operations teams simplify facility designs and minimize electricity expenses with industry-leading power efficiency. This leads to significantly reduced Capital Expenditures (CapEx) and Operating Expenditures (OpEx) while cutting down time to market for new devices and services.

With Cisco Silicon One, our unique solution enables deployments all the way from the web-scale Top Of Rack (TOR) across the front-end and AI/ML back-end networks, through the service provider peering and core roles, through the enterprise campus network from core to the edge. No other architecture in the industry can span this space. Even more impressively, no other architecture can cover this space while still being best of breed in any one location in the network.

That's the power of Cisco Silicon One.

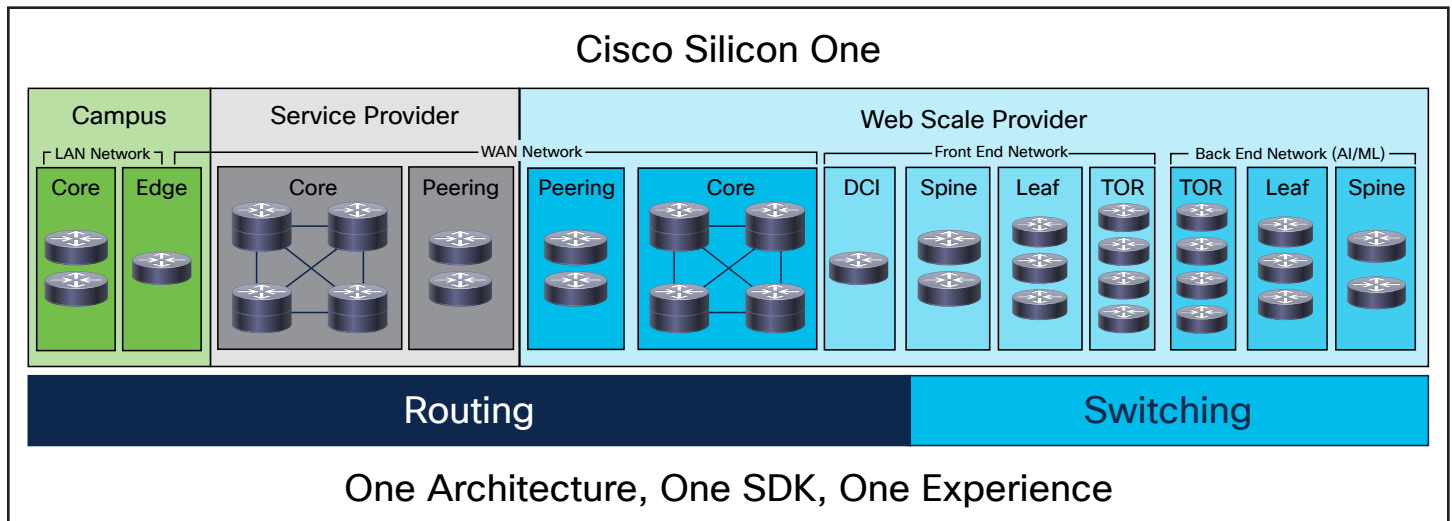


Figure 1. Network roles

One architecture

All Cisco Silicon One devices share a common set of blocks working together to create a common architecture that includes:

- One unified Silicon One SDK
- Large and fully shared on-die packet buffer
- High performance
- Low power
- Large scale
- Programmable forwarding engines
- Advanced features like tunnel termination and generation, ingress and egress Access Control Lists (ACLs), and Network Address Translation (NAT), all at line rate
- Advanced high-scale traffic management
- Advanced telemetry features

From our unified architecture, multiple devices are built to enable customers to trade off bandwidth, scale, cost, and power. This enables the same architecture to be deployed into both routing and web-scale switching roles, with software configuration change between modes required for different form factors of Line Card (LC), Standalone (SA), and Fabric Element (FE).

Table 1. Cisco Silicon One routing devices

Device	Generation	Ethernet Bandwidth	SerDes	Process	External Buffering	Mode(s)
P100	3 rd	19.2T	192x112G PAM4	7 nm	Yes	LC, SA
Q200	2 nd	12.8T	256x56G PAM4	7 nm	Yes	LC, SA, FE
Q100	1 st	10.8T	216x56G PAM4	16 nm	Yes	LC, SA, FE
Q211	2 nd	8T	160x56G PAM4	7 nm	Yes	SA
Q201	2 nd	6.4T	256x28G NRZ	7 nm	Yes	SA
Q202	2 nd	3.2T	128x28G NRZ	7 nm	Yes	SA

Table 2. Cisco Silicon One web scale switching devices

Device	Generation	Ethernet Bandwidth	SerDes	Process	External Buffering	Mode(s)
G200	4 th	51.2T	512x112G PAM4	5 nm	No	SA
G202	4 th	25.6T	512x56G PAM4	5 nm	No	SA
G100	3 rd	25.6T	256x112G PAM4	7 nm	No	SA, FE
Q200L	2 nd	12.8T	256x56G PAM4	7 nm	No	LC, SA, FE
Q100L	1 st	10.8T	216x56G PAM4	16 nm	No	LC, SA, FE
Q211L	2 nd	8T	160x56G PAM4	7 nm	No	SA
Q201L	2 nd	6.4T	256x28G NRZ	7 nm	No	SA
Q202L	2 nd	3.2T	128x28G NRZ	7 nm	No	SA

One network

Although any Cisco Silicon One device can be deployed anywhere in the network, traditional customer bandwidth, scale, cost, and power needs typically drive adoption of specific devices into specific roles. The P100, Q200, Q201, Q202, and Q100 devices are well suited for high-scale, deep-buffered routing deployments. Traditionally these are:

- Web-Scale and Data Center Interconnect (DCI)
- Web-Scale and Service Provider Core
- Web-Scale and Service Provider Peering
- Campus Edge
- Campus Core

The G200, G202, G100, Q200L, Q211L, Q201L, and Q202L are optimized for web-scale data center switching applications focused on highly efficient Ethernet switching. Traditionally these are:

- Web-Scale Top of Rack (TOR)
- Web-Scale Leaf
- Web-Scale and Enterprise Spine

With the explosive growth of Machine Learning (ML) and Artificial Intelligence (AI) the importance of the back-end network has expanded significantly. With Cisco Silicon One, web-scale switching devices can be used for a standard Ethernet-based deployment. If additional performance is needed, a fully scheduled fabric can be created by using the P100 or Q200/L devices as a TOR and the Q200L and G100 as the leaf and spine.

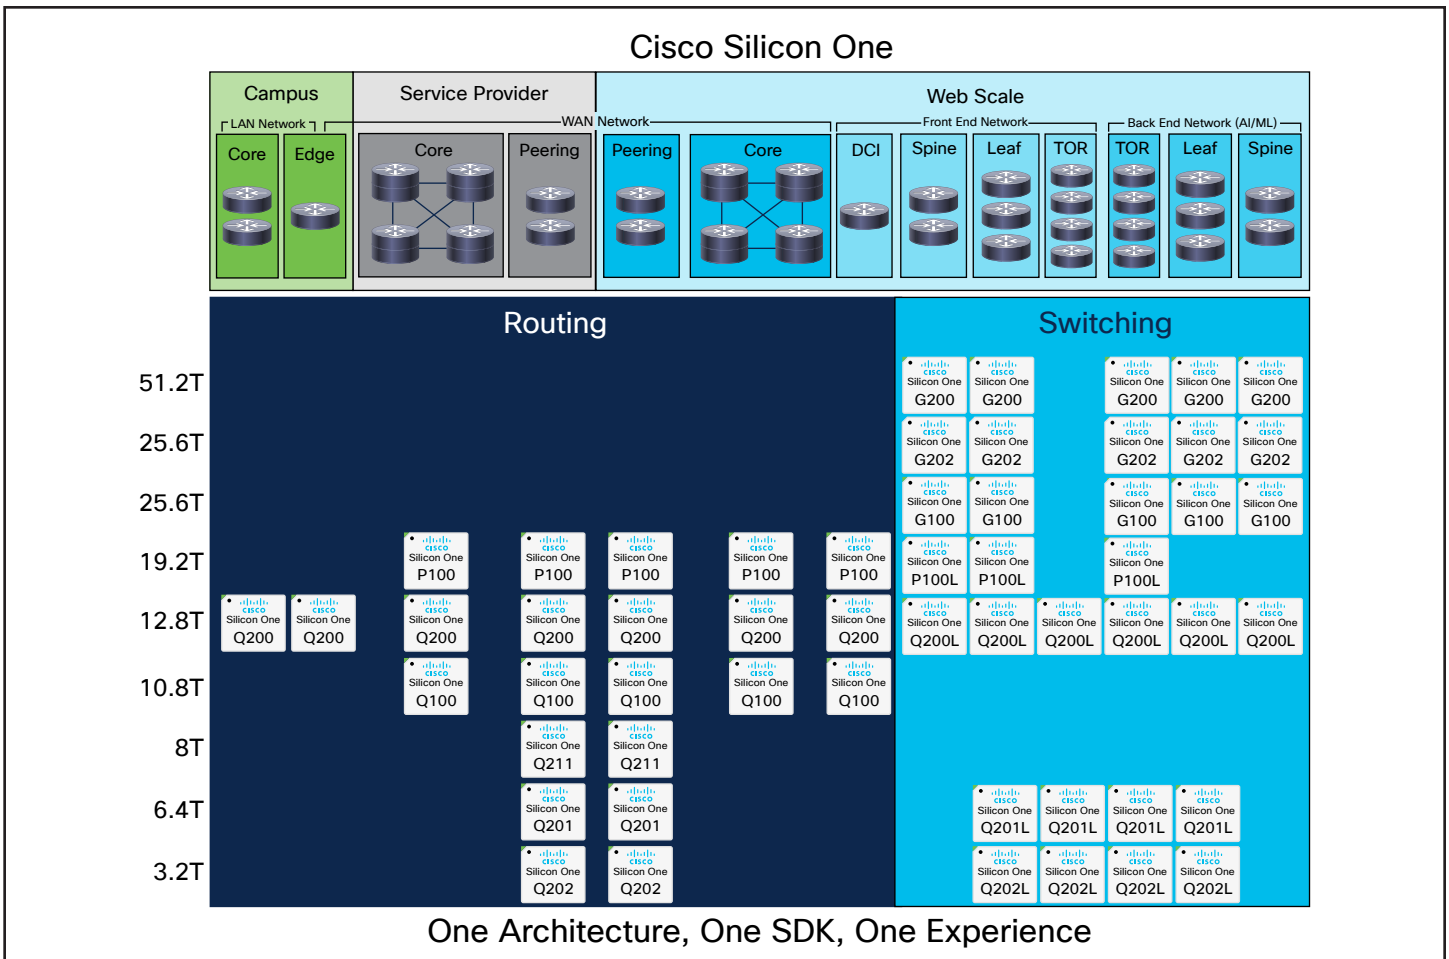


Figure 2. Cisco Silicon One across the network

One form factor

Not only can Cisco Silicon One devices be deployed anywhere in the network, but they can also be deployed in any form factor. The industry is accustomed to using different silicon architectures for standalone fixed boxes, standalone centralized boxes, modular line cards, modular fabric cards, disaggregated line cards

(leaf), and disaggregated fabric cards (spine), fracturing the development of features and behaviors based on the size of the system.

With our solution, a fully unified architecture can be deployed optimally across all these form factors.

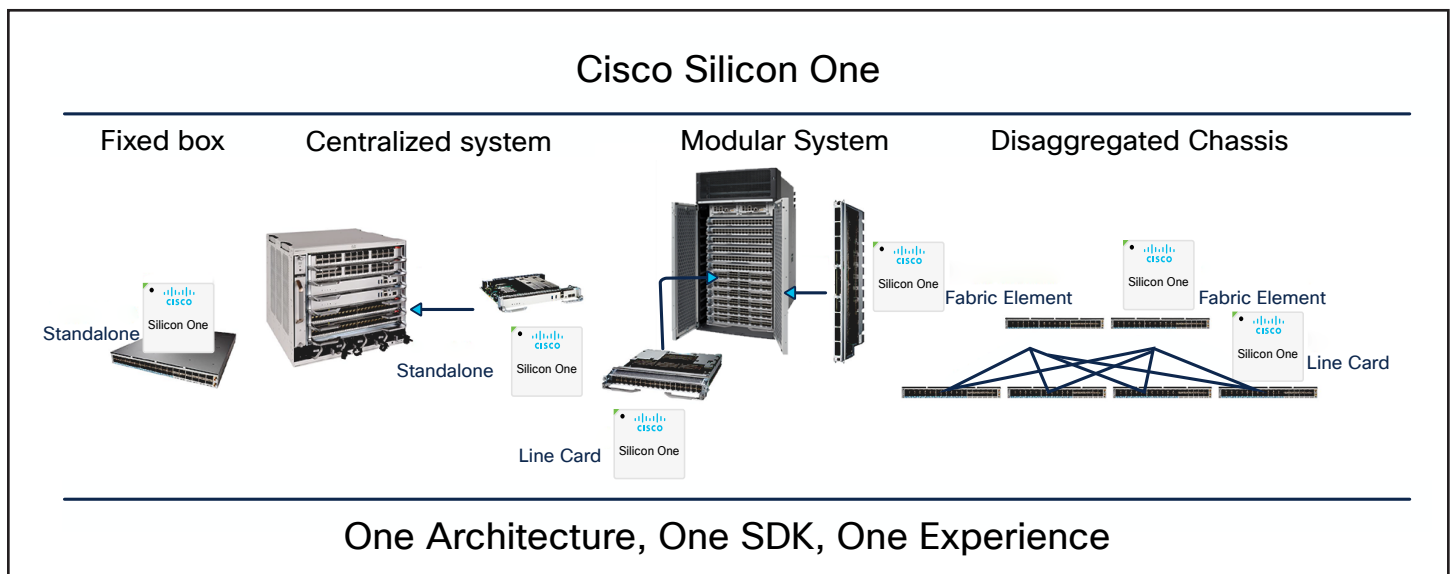


Figure 3. Cisco Silicon One across form factors

Cisco Silicon One offers a wide range of devices based on customer bandwidth, buffering, scale, and form factor needs.

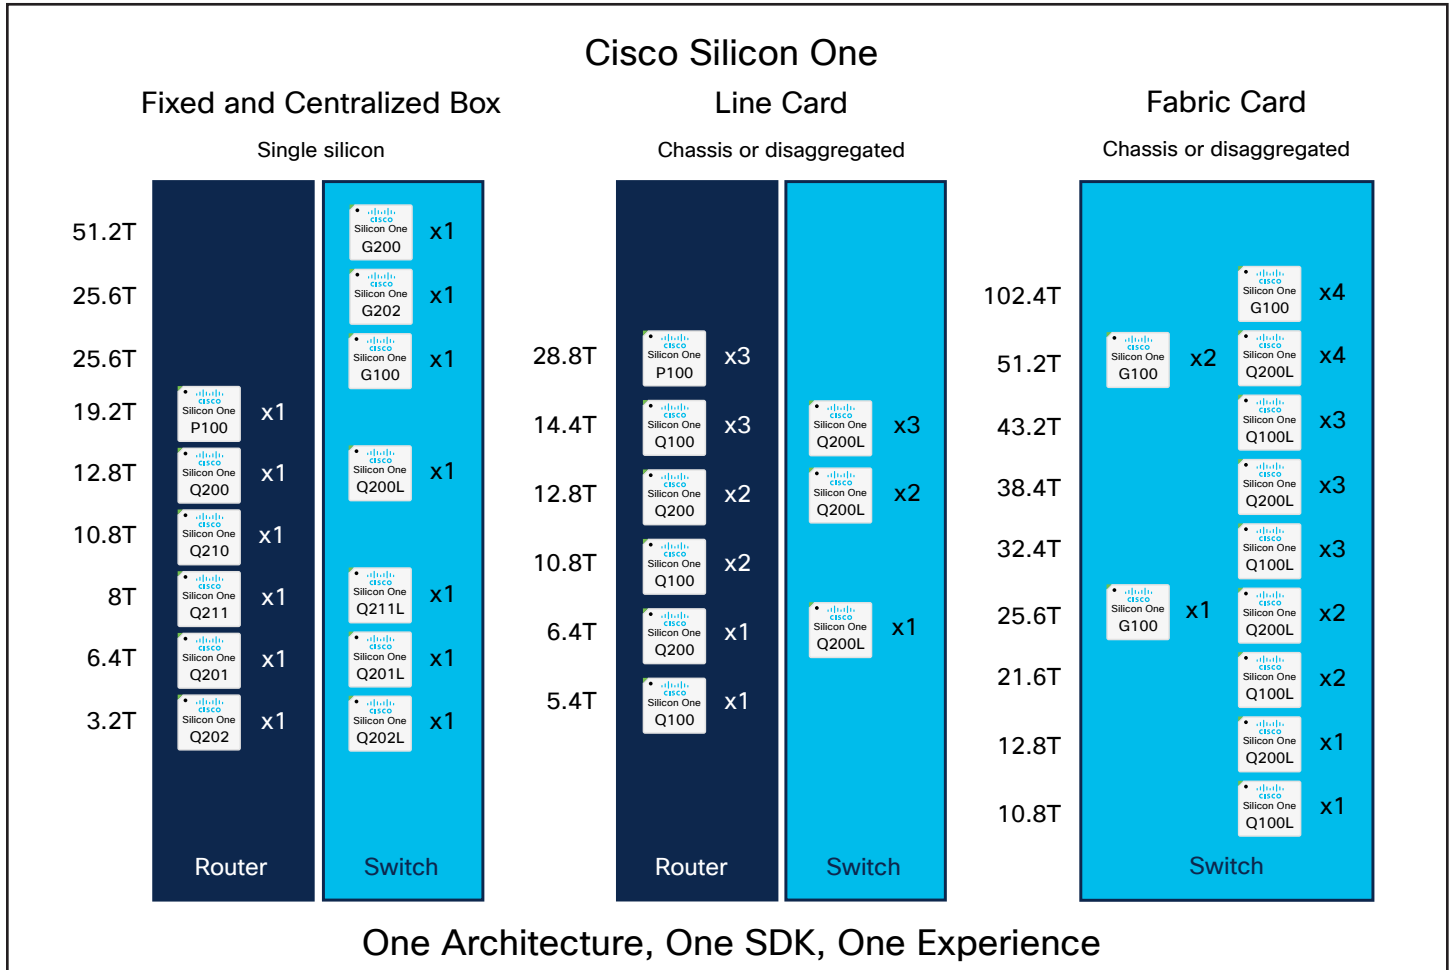


Figure 4. Cisco Silicon One devices across form factors

Routing versus switching

Cisco Silicon One allows equipment manufacturers to build a single piece of hardware that can accept pin-compatible Q200 routing silicon with deep buffers and Q200L switch silicon with a fully shared on-die buffer. This allows a single system design to become a class-leading 12.8-Tbps router or a 12.8-Tbps

switch. With footprint-compatible routing and switching silicon and a unified SDK, equipment manufacturers can accelerate time to market and network operators can decrease qualification time, enabling quicker deployment of the latest technologies.

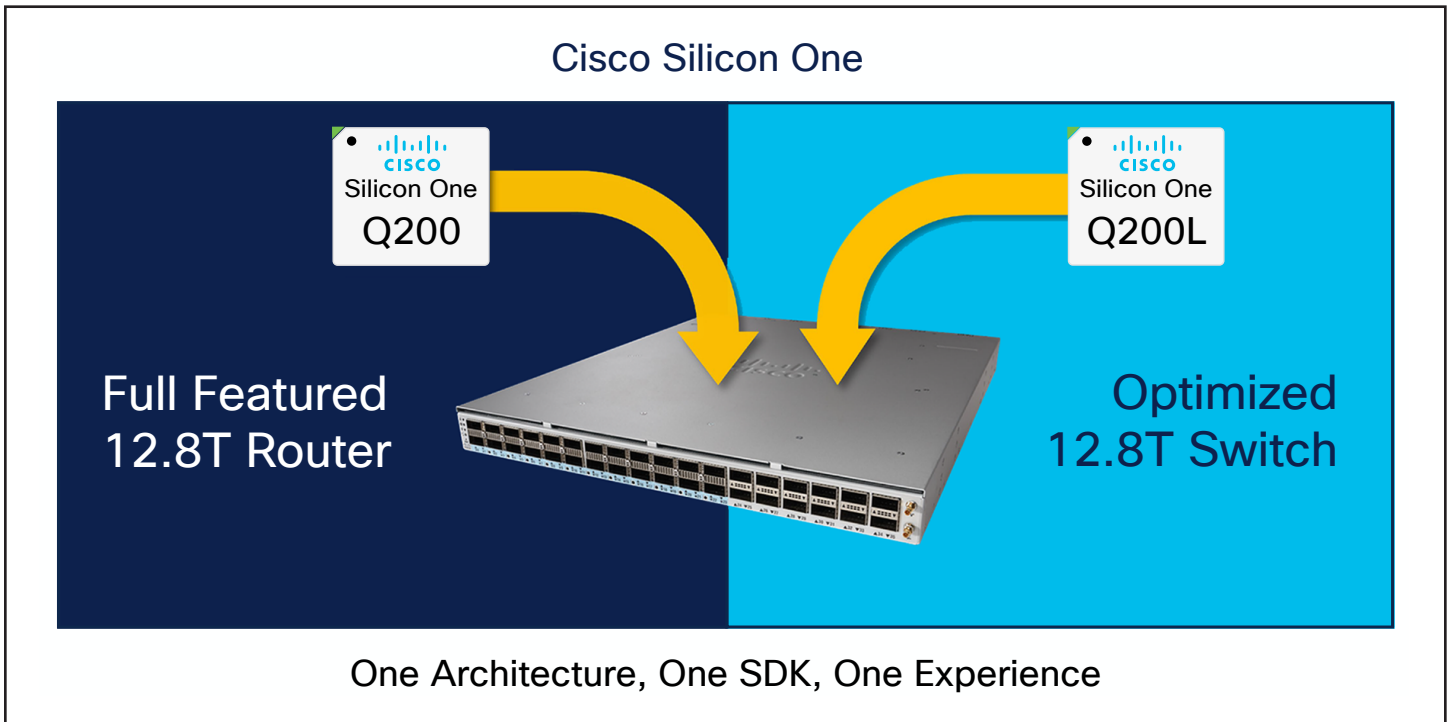


Figure 5. Cisco Silicon One universal hardware

Scheduled or unscheduled fabric

Our solution allows a common hardware platform to operate as individual routing and switching elements communicating over standard Ethernet with Equal-Cost Multi-Path (ECMP). Or, with simple software configuration

changes, it can operate as a fully scheduled fabric with ingress Virtual Output Queueing (VOQ) to create a distributed single routing or switching instance.

Table 3. Ethernet ECMP vs. scheduled fabric

Characteristic	Unscheduled Ethernet fabric	Fully scheduled fabric
Distribution method	ECMP hash	Spray and re-order
Link utilization	Low	High
Maximum flow limitations	Based on leaf and spine port BW	Based only on leaf port BW
Queueing	Queue per element	Ingress line-card Virtual Output Queue (VOQ)

Characteristic	Unscheduled Ethernet fabric	Fully scheduled fabric
Drop points	Ingress leaf, spine, egress leaf	Ingress leaf
Network view	Multiple unique routers and switches	One router or switch
Network OS complexity	Loose coupling	Tight coupling

This unique capability allows a modular chassis to take on multiple personalities depending on which operating systems are loaded. Similarly, a network operator can deploy a leaf-spine network of 12.8-Tbps fixed boxes with Q200 or Q200L, where each box works as a standard standalone device. But over time, they can

choose to convert these disjointed boxes into a fully scheduled fabric when their OS or network operations are ready. Similarly, the P100 and G100 devices can also be used to create even higher bandwidth for unscheduled or fully scheduled systems.

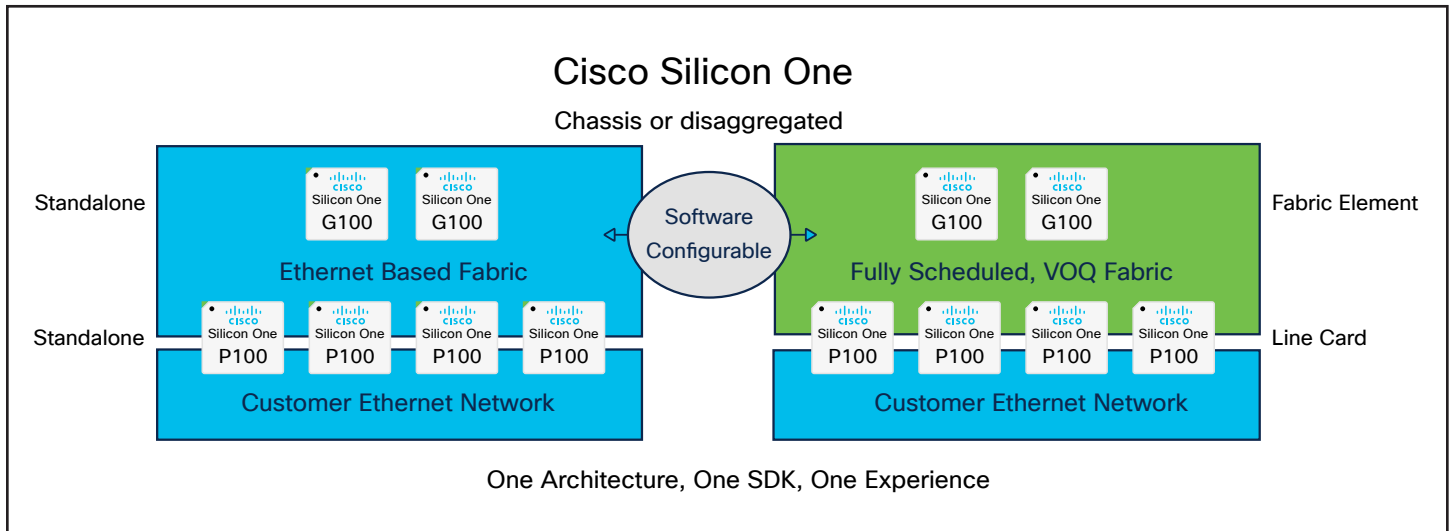


Figure 6. Cisco Silicon One scheduled or unscheduled fabric

Machine learning and artificial intelligence

In large-scale High-Performance Compute (HPC) environments, AI/ML network operators were forced to develop two incompatible, isolated networks. Many customers refer to the front-end network as the network that connects generic servers to one another, and to the outside world. This is what most people consider the traditional web-scale data center network. However, there is also a second network known as the back-end network. This network is designed to connect specialized compute and storage components to one another and has historically been built with proprietary interconnect technologies. The economics behind this type of network design can't keep up with the explosion of traffic that is occurring in AI/ML networks.

Network operators are searching for alternatives. The ideal economic solution is to use generic Ethernet as an interconnection topology. However, due to poor Equal-Cost Multiple Path (ECMP) load balancing decisions, the network can cause congestion even when the traffic pattern from the Graphical Processing Units (GPU) intentionally avoids congestion. This effect causes the network to slow down and increase the Job Completion Time (JCT).

The industry is investing in telemetry and intelligence in job placement to minimize these effects by detecting buildup early and rebalancing flows around the congestion points.

An alternate solution is to use a fully scheduled fabric to connect the GPUs. This approach sprays packets across all links and reorders the packets at the exit, so no network congestion builds up. Simply put, it allows for an ideal interconnect under all traffic conditions.

Unfortunately, most silicon architectures force network operators to choose a paradigm. One silicon architecture can use Ethernet; another one can use a proprietary interconnect; and one can use a proprietary fabric. This fork in the road forces customers to pick an interconnect technology today and hope that it works into the future with new ML/AI algorithms.

Advanced Cisco Silicon One technology has the unique ability to software configure the same network topology to use either standard Ethernet or fully scheduled fabric. Operators can deploy one network and evolve their choices over time, while maintaining maximum interoperability and performance.

Conclusion

Cisco Silicon One erases the hard dividing lines that have existed in the industry for decades, ushering in a new era of networking. Our unique solution is the only unified architecture that can span across routing and switching, from the web-scale data center TOR through the service provider and enterprise campus edge and core networks, and across all system design form factors. Customers can port the SDK once and deploy it everywhere.

For more information

Visit [Cisco Silicon One](#).