



Cisco Silicon One Product Family


Convergence without compromise

March 28, 2025 (Ver 6)

Contents

- Introduction – The Power of Cisco Silicon One 2
- One architecture – multiple devices 3
- One Network – many roles 5
- One architecture – any form factor 7
- One hardware design – multiple roles 8
- Scheduled or unscheduled fabric 9
- Machine Learning and Artificial Intelligence 10
- Conclusion 11
- For more information 11

Cisco Silicon One



Introduction – The Power of Cisco Silicon One

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

- Service provider vs. web-scale vs. enterprise deployments
- Routing vs. switching
- Full-featured vs. limited-featured
- Deep vs. shallow buffered
- Programmable vs. fixed function
- High vs. low scale
- Advanced vs. basic traffic management
- Fixed box vs. modular chassis vs disaggregated chassis configurations
- Scheduled vs. unscheduled fabric
- Customer type

Cisco Silicon One™ is a game changer for networking, eliminating long-standing divisions and unifying a significant portion of the networking market. Network designers, architects, operators and support teams no longer need to invest in, learn or deploy multiple unique architectures in parallel.

Adopting Cisco Silicon One-based systems, reduces Capital and Operating Expenditures (CapEx and OpEx) while accelerating new deployments and network or service upgrades. By removing the need to understand, qualify, deploy, and troubleshoot multiple distinct architectures, operators can focus on mastering a single architecture and a unified Software Development Kit (SDK), simplifying upgrades and streamlining operations. This approach also helps network operations teams to reduce facility design complexity and lower electricity costs, through industry-leading power efficiency. Additionally, support teams benefit from working with one architecture, enabling faster issue resolution.

Cisco Silicon One systems are now widely deployed, from web-scale Top of Rack (TOR) to front-end and AI/ML back-end networks, service provider peering, core, edge and access networks, as well as enterprise campus and Data centers. No other architecture in the industry matches this breadth of coverage. Even more impressively, Cisco Silicon One achieves this extensive reach while maintaining best-in-class performance at every network location, demonstrating its exceptional power and versatility.

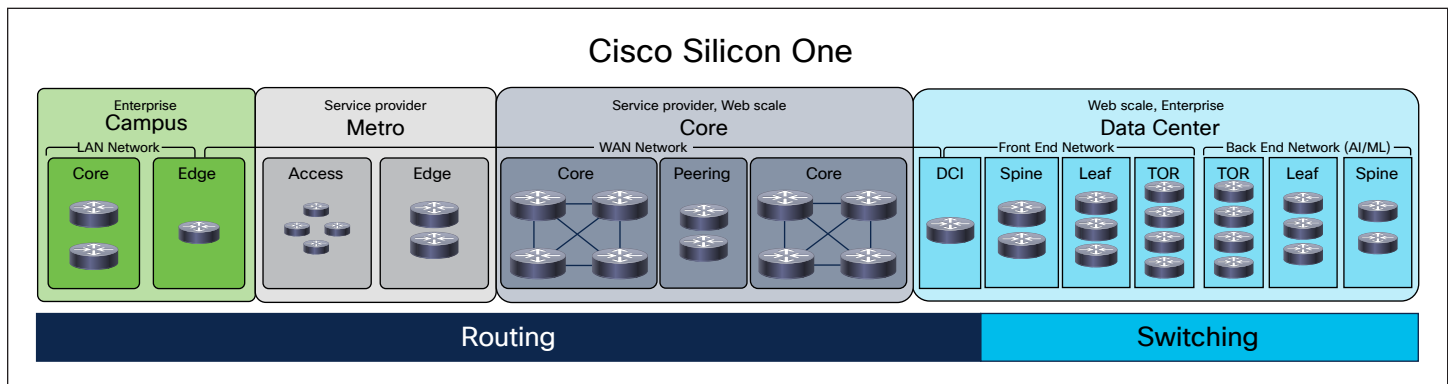


Figure 1. Network roles

One architecture – multiple devices

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

- A single, unified Silicon One SDK
- A large, fully shared on-die packet buffer
- High performance with lower power consumption
- Massive scalability
- 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

- High-scale, advanced traffic management
- Sophisticated telemetry features

The Cisco Silicon One unified architecture has already enabled the development of a broad portfolio of routing and switching devices, with many more in progress. Customers can select from a range of proven Silicon One solutions, each optimized for different bandwidth, scale, cost, and power requirements. Additionally, through software configuration, certain Silicon One devices, can operate in Line Card (LC), Standalone (SA) or Fabric Element (FE) mode.

Table 1. Cisco Silicon One Access Routing Devices

| Device | Generation | Ethernet bandwidth | SerDes | Process | External buffering | Mode(s) |
|--------|-----------------|--------------------|-------------|---------|--------------------|---------|
| A100 | 3 rd | 200G to 1.2T | 48x56G PAM4 | 16 nm | Yes | SA |

Table 2. Cisco Silicon One High-Features and Service Provider Edge Devices

| Device | Generation | Ethernet bandwidth | SerDes | Process | External buffering | Mode(s) |
|--------|-----------------|--------------------|--------------|---------|--------------------|---------|
| K100 | 3 rd | 6.4Tb | 192x56G PAM4 | 7 nm | Yes | SA, LC |

Table 3. Cisco Silicon One Enterprise Data center TOR and Aggregation Devices

| Device | Generation | Ethernet bandwidth | SerDes | Process | External buffering | Mode(s) |
|--------|-----------------|--------------------|--------------|---------|--------------------|---------|
| E100 | 3 rd | 6.4Tb | 192x56G PAM4 | 7 nm | No | SA |

Table 4. Cisco Silicon One High-Bandwidth 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 5. 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 – many roles

Cisco Silicon One devices are designed for versatility across the network, but specific customer requirements—such as bandwidth, scale, cost, and power—typically determine the best-fit solution for each role.

For high-scale, deep-buffered routing deployments, Cisco Silicon One currently offers the P100, Q200, Q201, Q202, and Q100 devices. These devices are typically deployed in:

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

For web-scale data center switching deployments, that prioritize highly efficient Ethernet switching, the G200, G202, G100, Q200L, Q211L, Q201L, and Q202L are preferred. These devices are typically deployed in:

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

The explosive growth of Machine Learning (ML) and Artificial Intelligence (AI) has significantly increased the importance of the back-end networks. With Cisco Silicon One, web-scale switching devices can seamlessly support standard Ethernet-based deployment. For higher performance, a fully scheduled fabric can be created by using P100 or Q200/L devices as a TOR and the Q200L and G100 as the leaf and spine layers.



For Service Provider Edge deployments, where features like high subscriber queues, flow analytics with telemetry, and large, flexible Layer 2 and Layer 3 tables are essential to support various tunneling protocols and programmability needs, the K100 is the ideal solution.

In enterprise TOR deployments, the E100 meets key requirements such as enterprise virtualization services, RDMA support, and seamless migration from

lower-speed networks (1Gbps, 10Gbps) to higher speeds—25Gbps, 50Gbps, and 100Gbps at the server level—and from 10Gbps and 40Gbps to 50Gbps, 100Gbps, and 400Gbps at the aggregation layer.

Designed to meet the demands of access networks, the A100 is ideally suited for Service Providers delivering both legacy and next-generation services to homes and businesses.

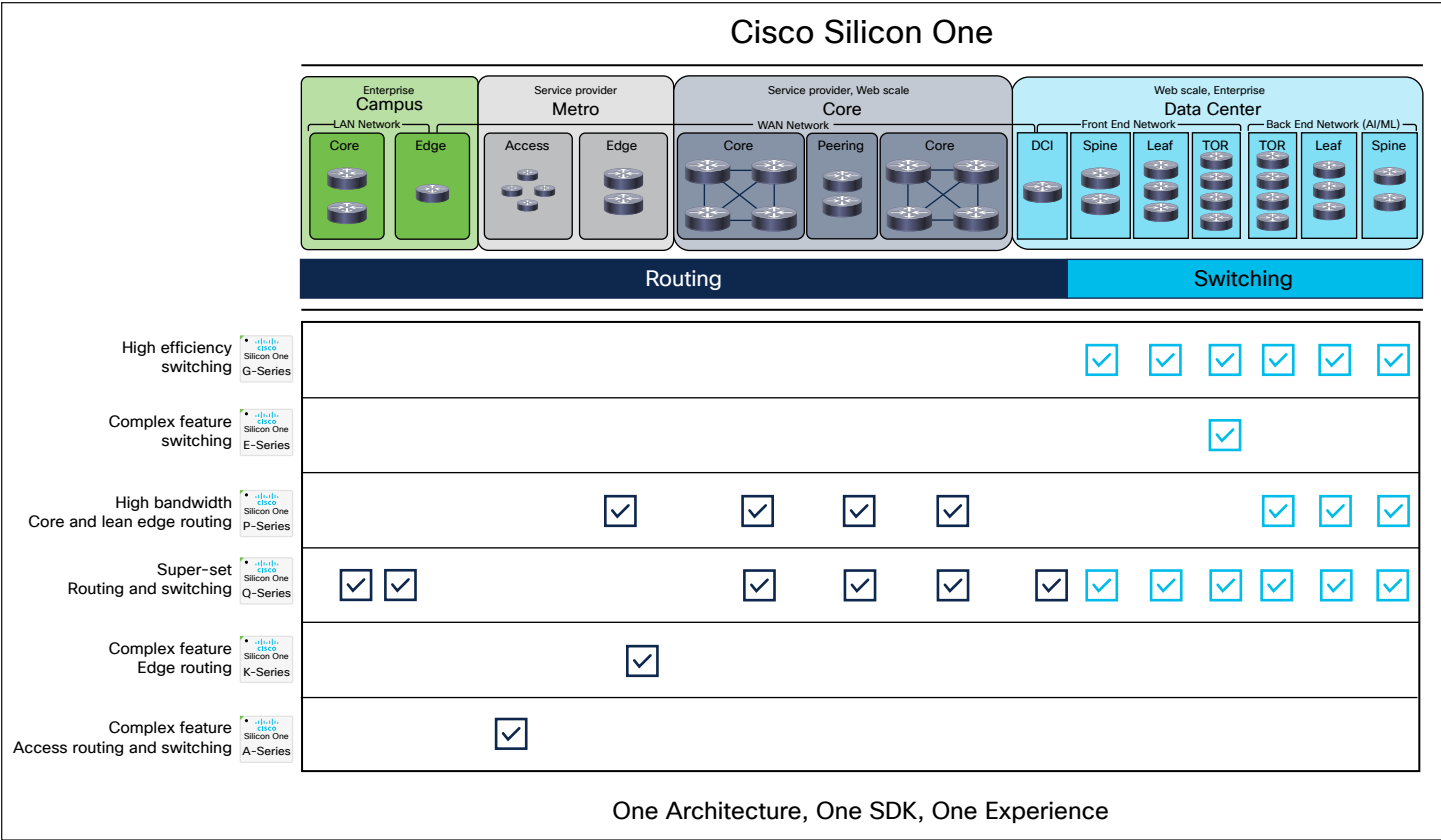


Figure 2. Cisco Silicon One Across the Network

One architecture – any form factor

Cisco Silicon One devices are designed for deployment anywhere in the network and in any form factor.

Traditionally, the industry has relied on different silicon architectures for standalone fixed boxes, standalone centralized systems, modular line cards, modular fabric cards, disaggregated line cards (leaf), and disaggregated fabric cards (spine). This fragmented approach has led to inconsistencies in feature development and system behavior across different form factors and system sizes.

In contrast, Silicon One's fully unified architecture eliminates these inconsistencies, enabling seamless deployment across the network in any form factor while optimizing both network performance and efficiency.

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

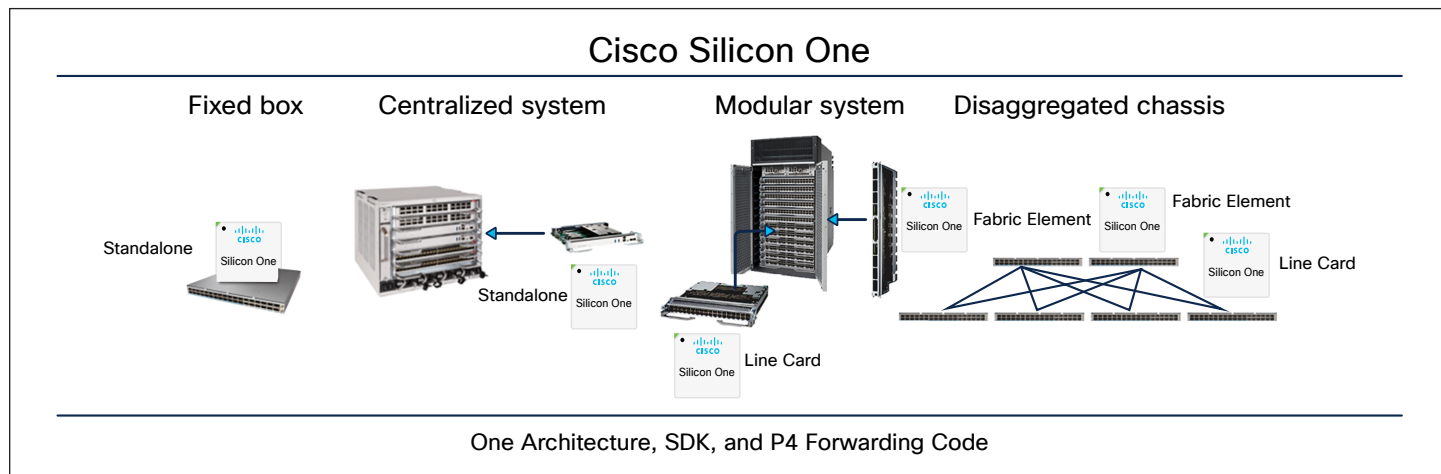


Figure 3. Cisco Silicon One: Deployed Across All Form Factors

One hardware design – multiple roles

In addition to the architectural commonality across devices, Cisco Silicon One maintains pin compatibility across devices serving different market segments, when possible. The pin compatibility allows equipment manufacturers to create a single hardware platform that can be populated with a variety of pin-compatible devices, allowing them to serve multiple roles within the network.

A prime example of this innovation is seen in the design of pizza-box systems, where the pin-compatible Q200 routing silicon with deep buffers and the Q200L switch silicon with a fully shared on-die buffer are interchangeable. The pin compatibility allows a single system design to serve either as 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, while network operators benefit from reduced qualification time, enabling faster deployment of the latest technologies.

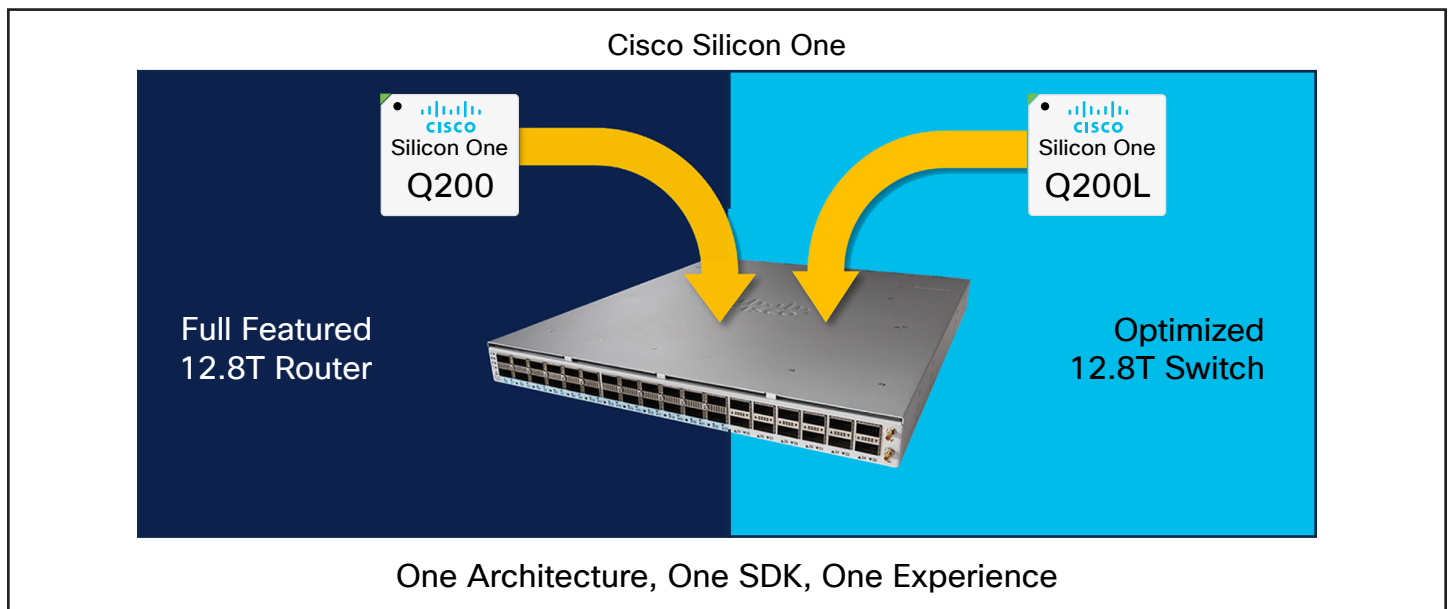


Figure 4. Cisco Silicon One: One Design, Universal Hardware

Scheduled or unscheduled fabric

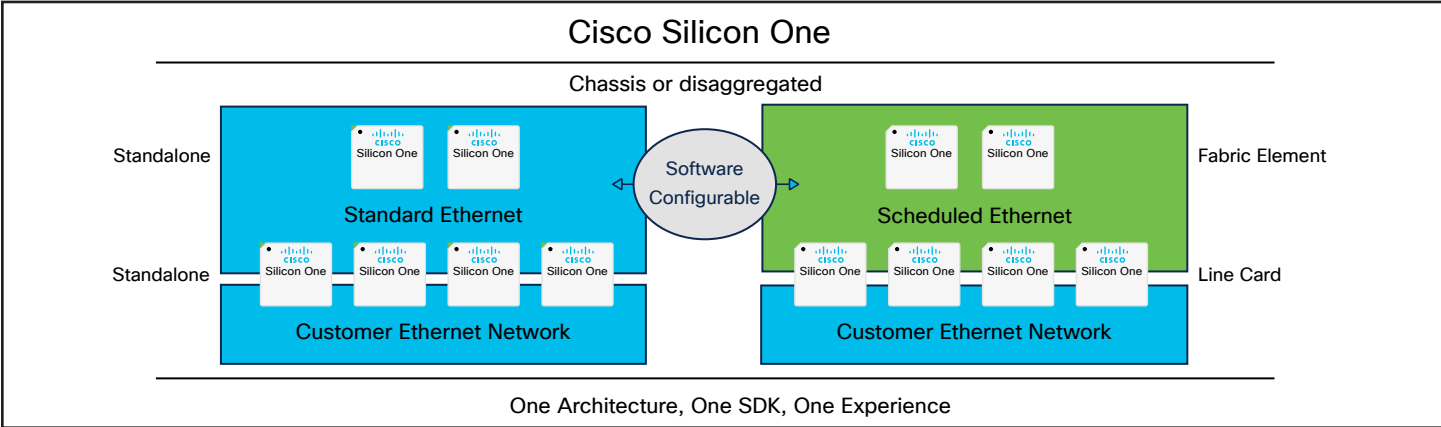


Figure 5. Cisco Silicon One: Scheduled or Unscheduled Fabric

Cisco Silicon One enables a fabric of interconnected devices to operate as individual routing and switching devices communicating over standard Ethernet with Equal-Cost Multi-Path (ECMP), or with a simple software configuration change, as a fully scheduled fabric with ingress Virtual Output Queueing (VOQ). Table 3 summarizes the key characteristics of these two deployment solutions.

Table 6. 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) |
| 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 |

The software reconfigurability of a deployed system is a unique advantage of Cisco Silicon One, allowing a modular chassis to take on multiple roles depending on the loaded operating modes. For example, a network operator can deploy a leaf-spine network based on the Q200 or Q200L using 12.8-Tbps fixed boxes, where each box operates as a standard standalone Ethernet device. Over time, the leaf-spine network can be reconfigured into a fully scheduled fabric. Similar in-network reconfigurability can also be achieved with P100 and G100 devices.

Machine Learning and Artificial Intelligence

In large-scale High-Performance Compute (HPC) environments, AI/ML network operators have traditionally been forced to develop two incompatible, isolated networks. The front-end network connects generic servers to one another and to the outside world—essentially a traditional web-scale data center network. The back-end network, on the other hand, connects specialized compute and storage components and has historically relied on proprietary interconnect technologies.

The economics behind a non-Ethernet based back-end network design can't keep up with the explosion of traffic that is occurring in AI/ML networks and thus, network operators are searching for alternatives. The ideal solution is to use generic Ethernet for interconnectivity. Yet, standard Ethernet faces challenges—poor Equal-Cost Multi-Path (ECMP) load balancing can create congestion, even when traffic patterns from GPUs are designed to

avoid it. This congestion slows down the network and increases Job Completion Time (JCT). To address these issues, the industry is investing in advanced telemetry and intelligent job placement, enabling real-time congestion detection and traffic rebalancing.

As discussed earlier, a proven alternative to standard Ethernet or proprietary interconnects is a fully scheduled fabric. This approach distributes packets evenly across all available links and reorders them at the exit, eliminating network congestion by design. Simply put, a fully scheduled fabric ensures an optimal interconnect under any traffic conditions. Unfortunately, all silicon architectures—except Cisco Silicon One—force network operators to choose between Ethernet and proprietary interconnects. This means customers must commit to an interconnect technology early on, hoping it remains viable as AI/ML workloads evolve.

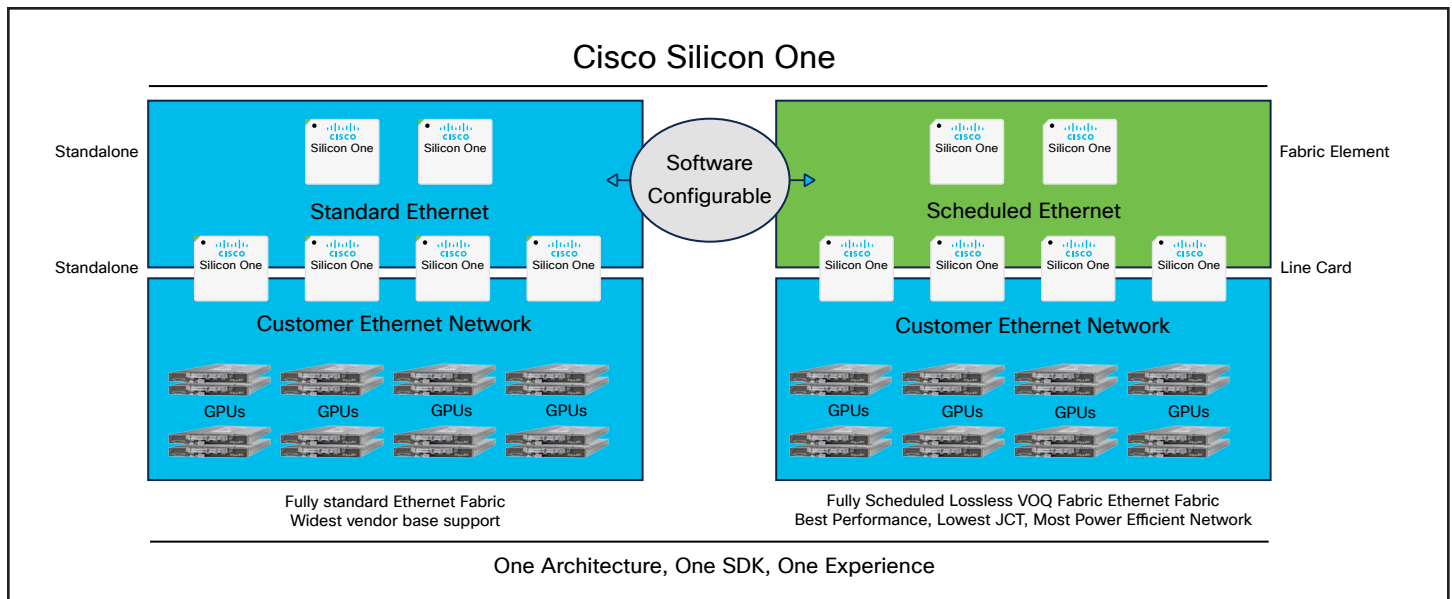


Figure 6. Cisco Silicon One Fabrics for GPU AI/ML Workloads

Cisco Silicon One's advanced technology uniquely allows operators to software-configure the same network topology to use either standard Ethernet or a fully scheduled fabric. This flexibility enables operators to deploy a single network while evolving their choices over time, ensuring maximum interoperability and performance.

Conclusion

Cisco Silicon One dismantles long-standing networking barriers, ushering in a new era of unified and flexible solutions. As the only architecture that seamlessly integrates both routing and switching, Cisco Silicon One supports a wide range of environments—from web-scale and enterprise data centers to service provider and enterprise campus networks, and all system design form factors. With a single SDK, customers can design once and deploy anywhere, driving greater efficiency and accelerating innovation across the network. This is the defining advantage of Cisco Silicon One.

For more information

Visit [Cisco Silicon One](#).