

Cisco Silicon One Product Family

Convergence without compromise

Contents

| | |
|---------------------------------------|---|
| Introduction | 2 |
| One architecture | 3 |
| One network | 4 |
| One form factor..... | 5 |
| Routing versus switching | 6 |
| Scheduled or unscheduled fabric | 6 |
| Conclusion | 7 |

Introduction

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

Cisco Silicon One™ is a break-through technology that for the first time in history enables a single silicon architecture to erase these dividing lines and span a massive portion of the networking market. Gone are the days where designers and network architects need to invest in and learn multiple unique architectures in parallel. Nor do support and operations teams 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 more rapidly and simply. Support teams only need to understand one architecture and 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.

Cisco Silicon One

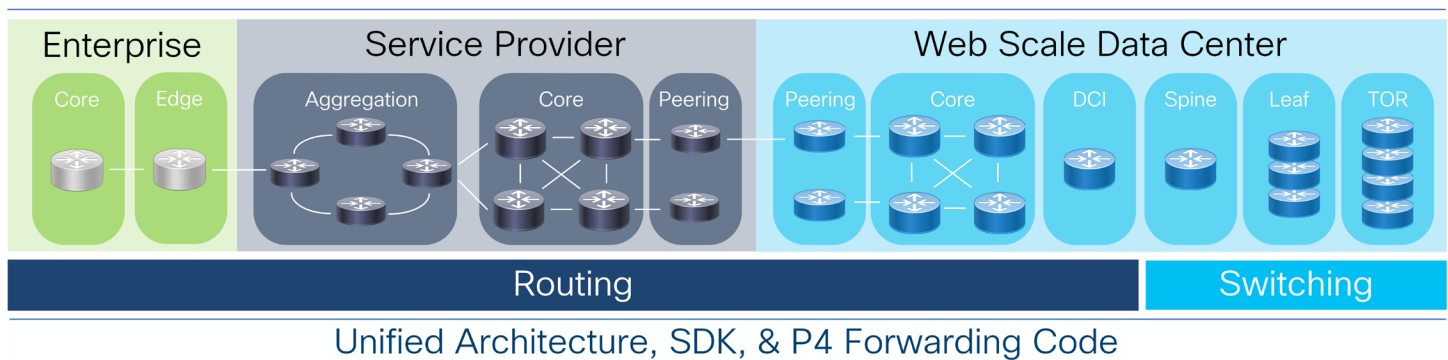


Figure 1. Network roles

With Cisco Silicon One, our unique solution enables deployments all the way from the web scale Top of Rack (TOR) switch, the service provider network to the aggregation sites, and the enterprise network from core to edge. No other architecture in the industry can span this space, but more impressively, no other architecture can cover this space while still being best of breed in any one location in the network.

One architecture

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

- One unified Silicon One SDK
- One unified P4 forwarding code
- Large and fully shared on-die packet buffer
- High performance
- Low power
- Large scale
- P4 run to completion programmable 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, enabling the same architecture to be deployed into both routing and web scale switching roles.

| | Generation | Ethernet Bandwidth | SerDes | Process | External Buffering |
|------|-----------------|--------------------|---------------|---------|--------------------|
| P100 | 3 rd | 19.2T | 192x112G PAM4 | 7nm | Yes |
| Q200 | 2 nd | 12.8T | 256x56G PAM4 | 7nm | Yes |
| Q100 | 1 st | 10.8T | 216x56G PAM4 | 16nm | Yes |
| Q211 | 2 nd | 8T | 160x56G PAM4 | 7nm | Yes |
| Q201 | 2 nd | 6.4T | 256x28G NRZ | 7nm | Yes |
| Q202 | 2 nd | 3.2T | 128x28G NRZ | 7nm | Yes |

Table 1. Cisco Silicon One routing devices

| | Generation | Ethernet Bandwidth | SerDes | Process | External Buffering |
|-------|-----------------|--------------------|---------------|---------|--------------------|
| G100 | 3 rd | 25.6T | 256x112G PAM4 | 7nm | No |
| Q200L | 2 nd | 12.8T | 256x56G PAM4 | 7nm | No |
| Q100L | 1 st | 10.8T | 216x56G PAM4 | 16nm | No |
| Q211L | 2 nd | 8T | 160x56G PAM4 | 7nm | No |
| Q201L | 2 nd | 6.4T | 256x28G NRZ | 7nm | No |
| Q202L | 2 nd | 3.2T | 128x28G NRZ | 7nm | No |

Table 2. Cisco Silicon One web scale switching devices

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 starting from the web scale Data Center Interconnect (DCI), through the peering roles into web scale, service provider and enterprise core networks, into service provider aggregation roles, and finally into the enterprise edge roles.

The G100, Q200L, Q201L, and Q202L are optimized for web scale data center switching applications focused on highly efficient ethernet switching from the TOR through the leaf and spine.

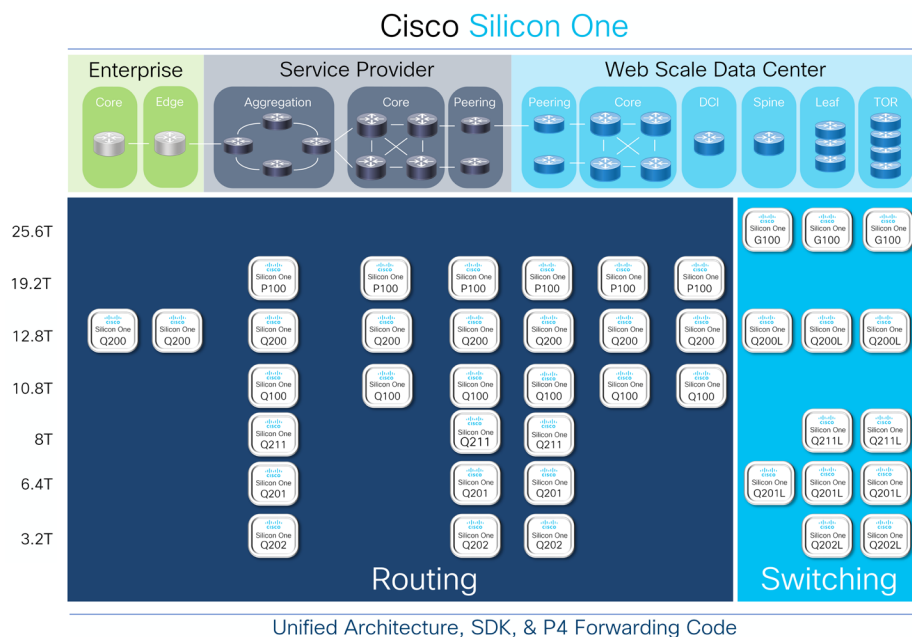


Figure 2. Cisco Silicon One across the network

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.

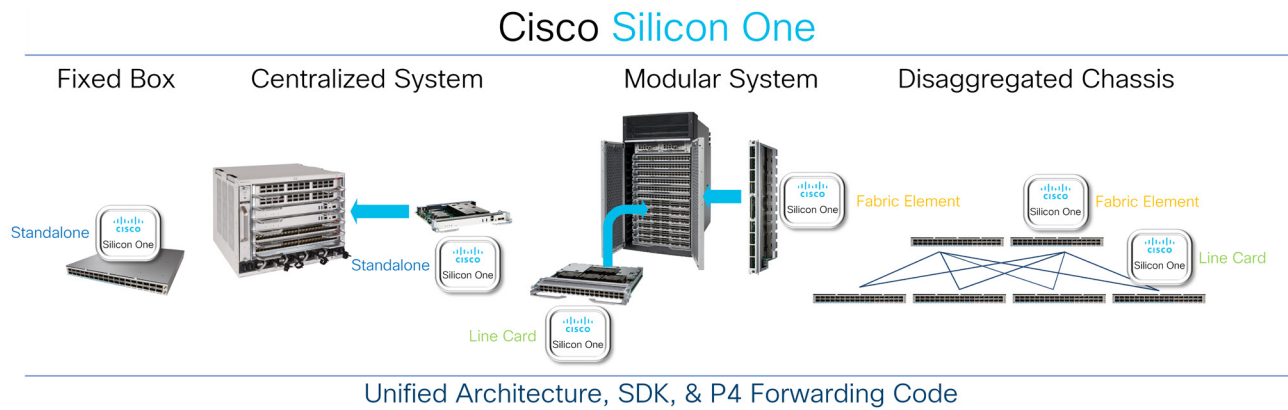


Figure 3. Cisco Silicon One across form factors

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

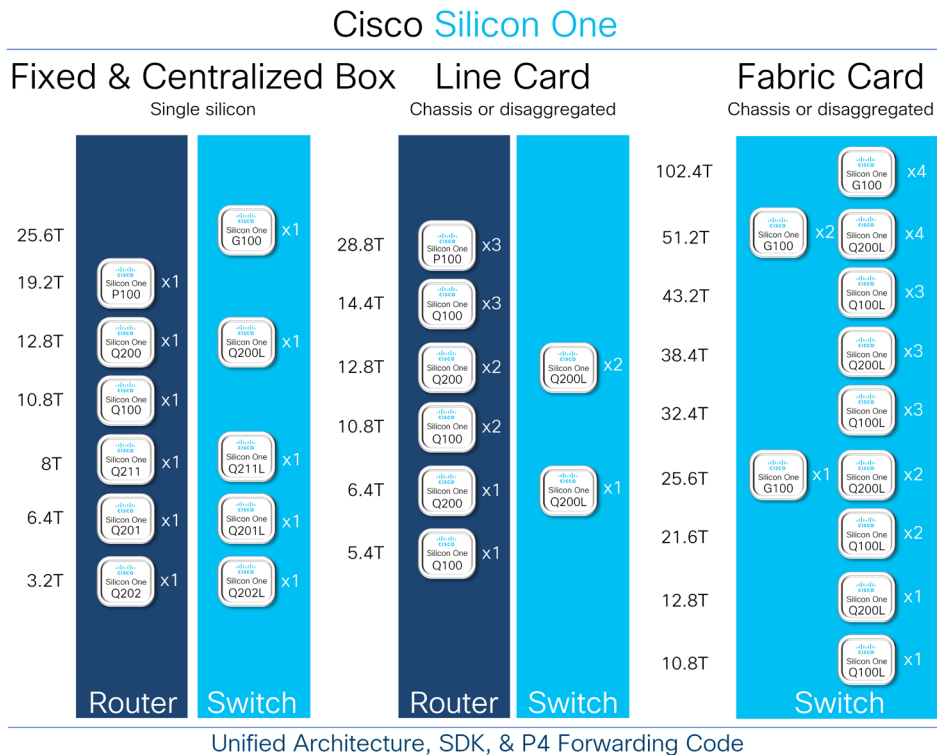


Figure 4. Cisco Silicon One devices across form factors

Routing versus switching

Cisco Silicon One allows equipment manufacturers to build a single piece of hardware which 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 12.8Tbps router or a 12.8Tbps 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.

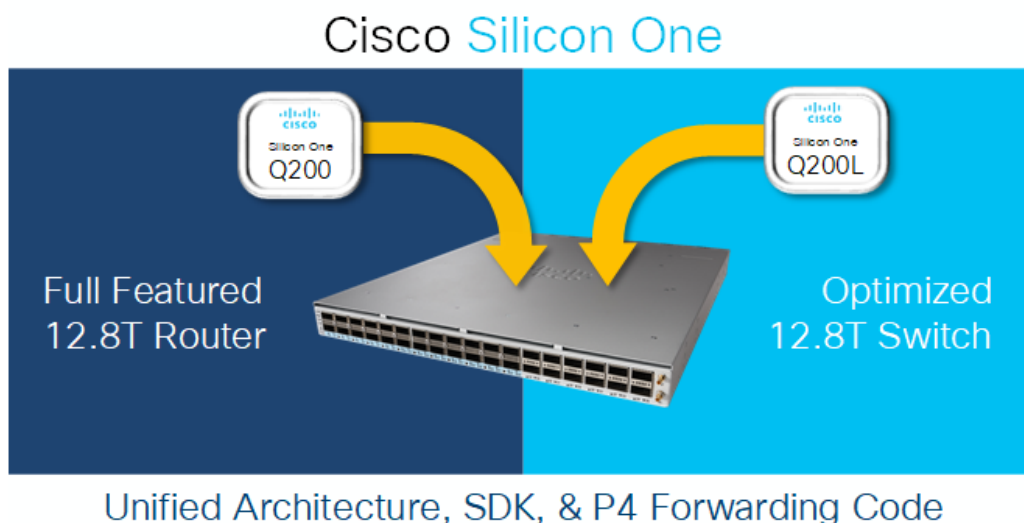


Figure 5. Cisco Silicon One universal hardware

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.

| | Unscheduled Ethernet Fabric | Fully Scheduled Fabric |
|--------------------------|--------------------------------------|--|
| Distribution Method | ECMP Hash | Spray & 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 |

Table 3. Ethernet ECMP vs. scheduled fabric

Learn more

For more information visit
[Cisco Silicon One](#).

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 12.8Tbps fixed boxes with Q200 or Q200L where each box works as a standard stand-alone device. But over time they can make the choice 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.

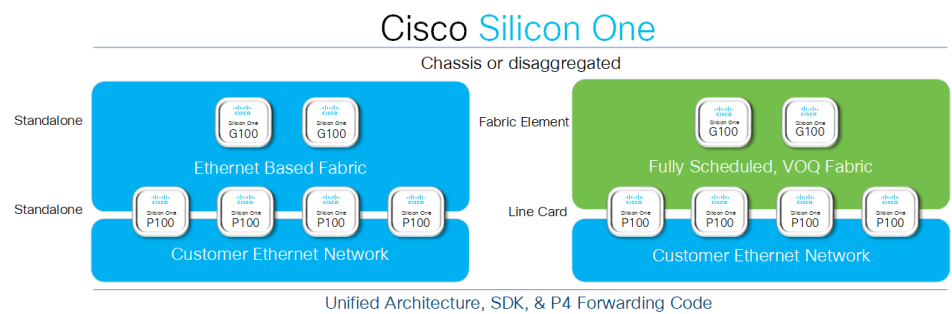


Figure 6. Cisco Silicon One scheduled or unscheduled fabric

Conclusion

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