

System Benefits of Cisco Silicon One P100

In a league all its own

Contents

Introduction	2
Building a fixed-box router	2
Fixed box latency savings	4
Fixed box power savings benefits	4
Building a modular router line card	6
Conclusion	8

Summary

In October of 2020, Cisco announced the innovative Q200 device. It was suddenly possible to build a 12.8 Tbps router with a single piece of silicon that would require between two and 10 pieces of silicon with other vendors. This efficiency gain enabled significant cost and power savings.

With the announcement of the Cisco Silicon One™ P100 device, we extend our leadership for the fixed box portfolio enabling a 19.2 Tbps router with a single piece of silicon, but we also enable the industry's first 28.8 Tbps, 36x800Gbps line-card for modular systems, providing a twofold increase in density over what is available from any other vendor.

Introduction

In this paper we'll study the impacts on both fixed and modular systems using the P100 and other silicon available on the market. We'll show that fixed systems can be built with significantly less silicon, resulting in large savings in power, space, latency, and cost. We'll also show that modular systems can be built with two times more bandwidth than other silicon, resulting in significant increases in bandwidth per system while simultaneously driving power and cost efficiency.

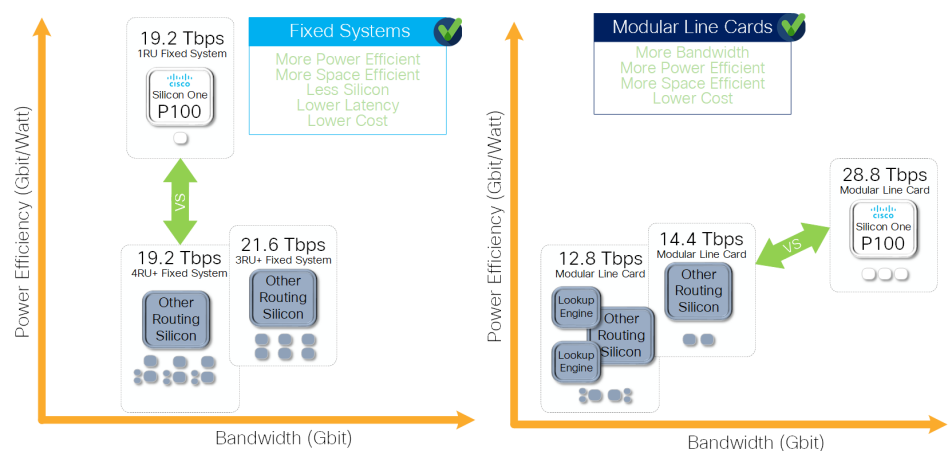


Figure 1. Cisco Silicon One P100 advantage

Building a fixed-box router

There are many pieces of routing silicon available on the market today from both third-party providers and systems companies. Currently the next highest performance routing silicon after Cisco Silicon One is a 7.2Tbps device. We believe that to get equivalent IPv4 Longest-Prefix-Match (LPM) scale, up to two external lookup engines may be needed.

With Cisco Silicon One P100 we can build a 19.2Tbps (48x400G or 24x800G) router with a single piece of routing silicon, while with other devices you can only build a lower scale 7.2Tbps (18x400G) system or a high scale 6.4Tbps (16x400G) system with external lookup engines.

Confusion around chip bandwidths arise when companies use misleading marketing tactics like using half-duplex math to advertise a 7.2 Tbps routing silicon as a 14.4 Tbps routing silicon. With Cisco Silicon One we don't need to play these games. When we advertise our devices as 19.2 Tbps it means we really can transmit 19.2 Tbps and receive 19.2 Tbps or put more simply, our device is 2.6 times higher bandwidth than any other routing silicon in the market at the time of the writing of this paper.

With the power efficiencies of Cisco Silicon One, it's possible to build a 24x800G using QSFP-DD800 optics in a 1RU form factor or a 48x400G using QSFP-DD56 optics in a 2RU form factor.

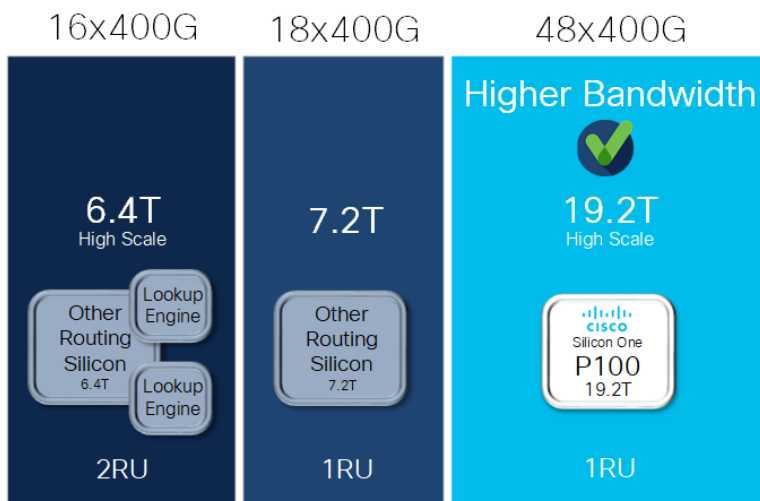


Figure 2. Fixed box with single device

To achieve higher bandwidth, other vendors need to use a minimum of six devices connected in a Clos configuration which may add up to three times the latency and rack units and five times the cost and power, magnifying the benefits of Cisco Silicon One and demonstrating how advanced our architecture is.

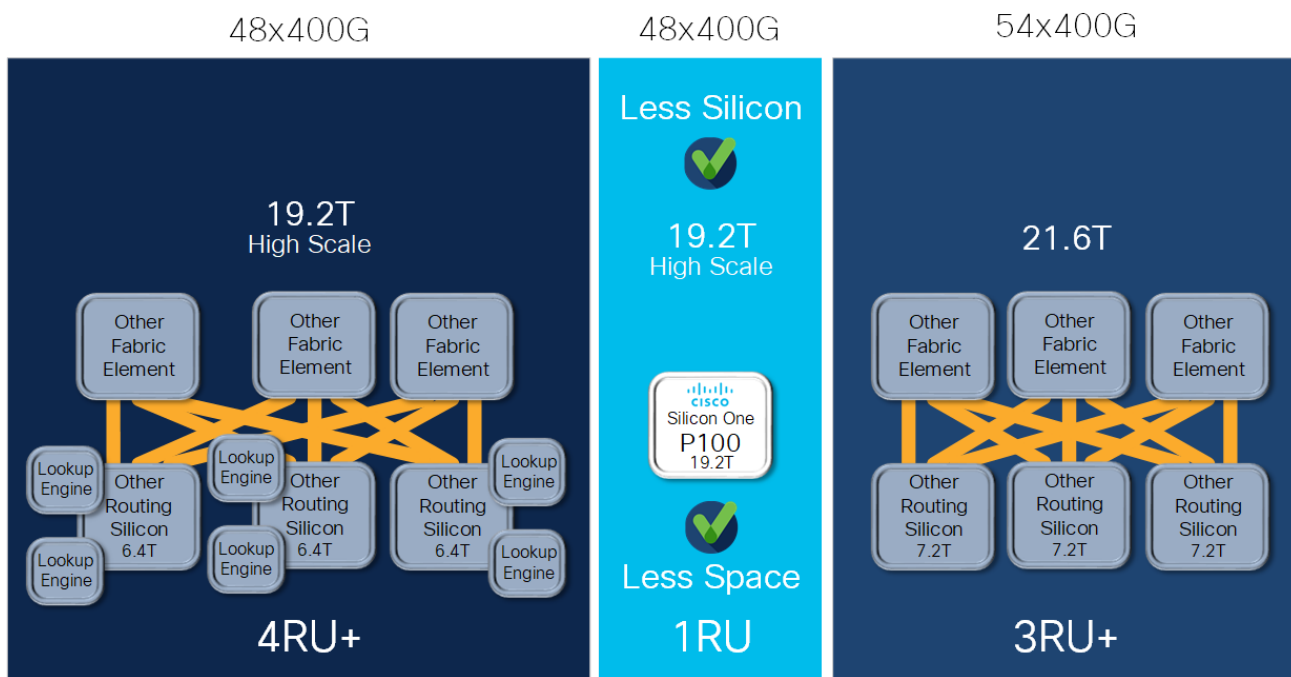


Figure 3. High bandwidth fixed box

Fixed box latency savings

A consequence of multiple devices is that packets need to travel over multiple hops as they traverse from an input port to an output port in the system. As packets arrive at the ingress device, they must be routed over Serializer/Deserializer (SerDes) to the fabric element and then again over another set of SerDes to the egress device, adding significant additional latency.

With Cisco Silicon One, the packets are switched locally within a device without ever going to another device. This means that Cisco Silicon One can forward packets with one third the number of hops, resulting in significant latency savings.

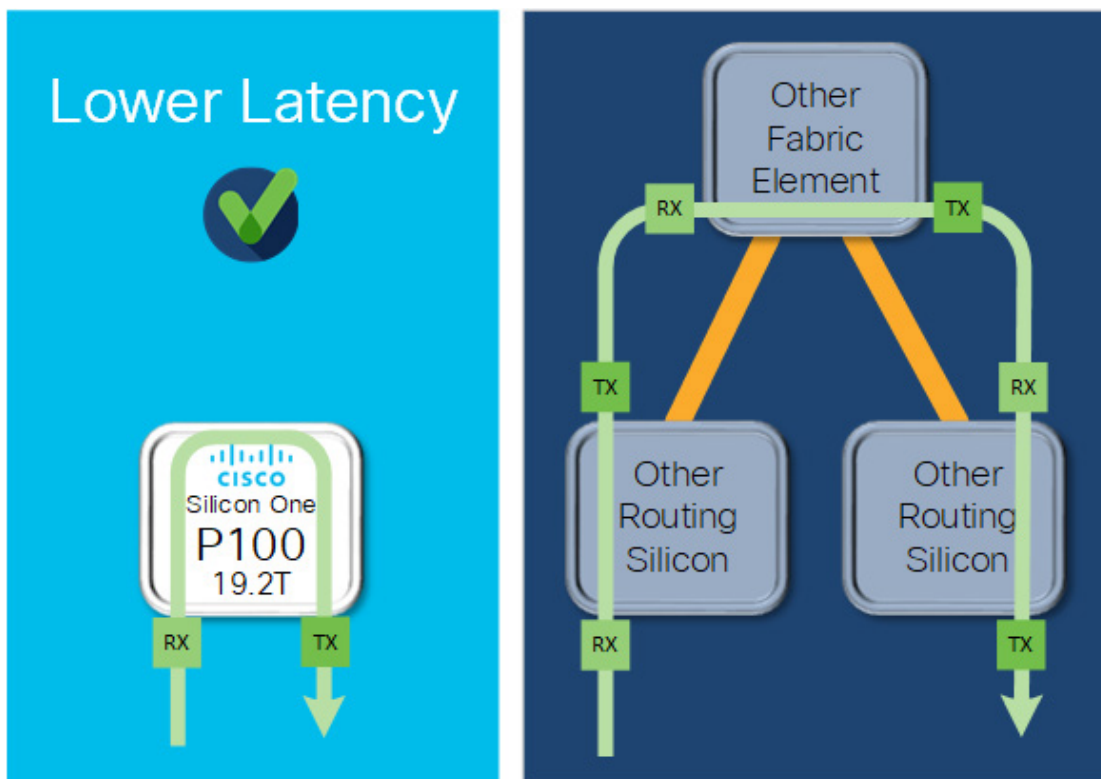


Figure 4. Fixed box latency savings

Fixed box power savings benefits

The power penalties of other architectures come from several places:

Impacts due to the silicon itself:

- Increase in silicon core die area (six to 12 devices versus one device)
- Additional SerDes to connect the chips together
- Reading the packet from buffers and moving the data multiple times

Secondary impacts:

- Increased loss in power planes delivering the higher current to the chips
- Increased loss in power supplies delivering the higher current to the system
- Increased fan speed to cool the higher power draw

These impacts translate to a very significant effect at the full system level. Systems built with the Cisco Silicon One P100 can provide 19.2 Tbps of routing performance in under 450W of power, while we estimate systems built with the next closest competitor consume more than 2,000W and with high-scale Longest Prefix Match (LPM) we estimate it requires more than 2,275W.

System	Pieces of Silicon	LPM Scale	Rack Units	Typical System Power	Power per 100G
19.2T P100 1RU System	1	High	1	< 450W	< 2.3W
21.6T Other Silicon 3RU System	6	Medium	3+	> 2,000W	> 9.2W
19.2T Other Silicon 4RU System	10	High	4+	> 2,275W	> 11.8W

Table 1. Fixed box system power

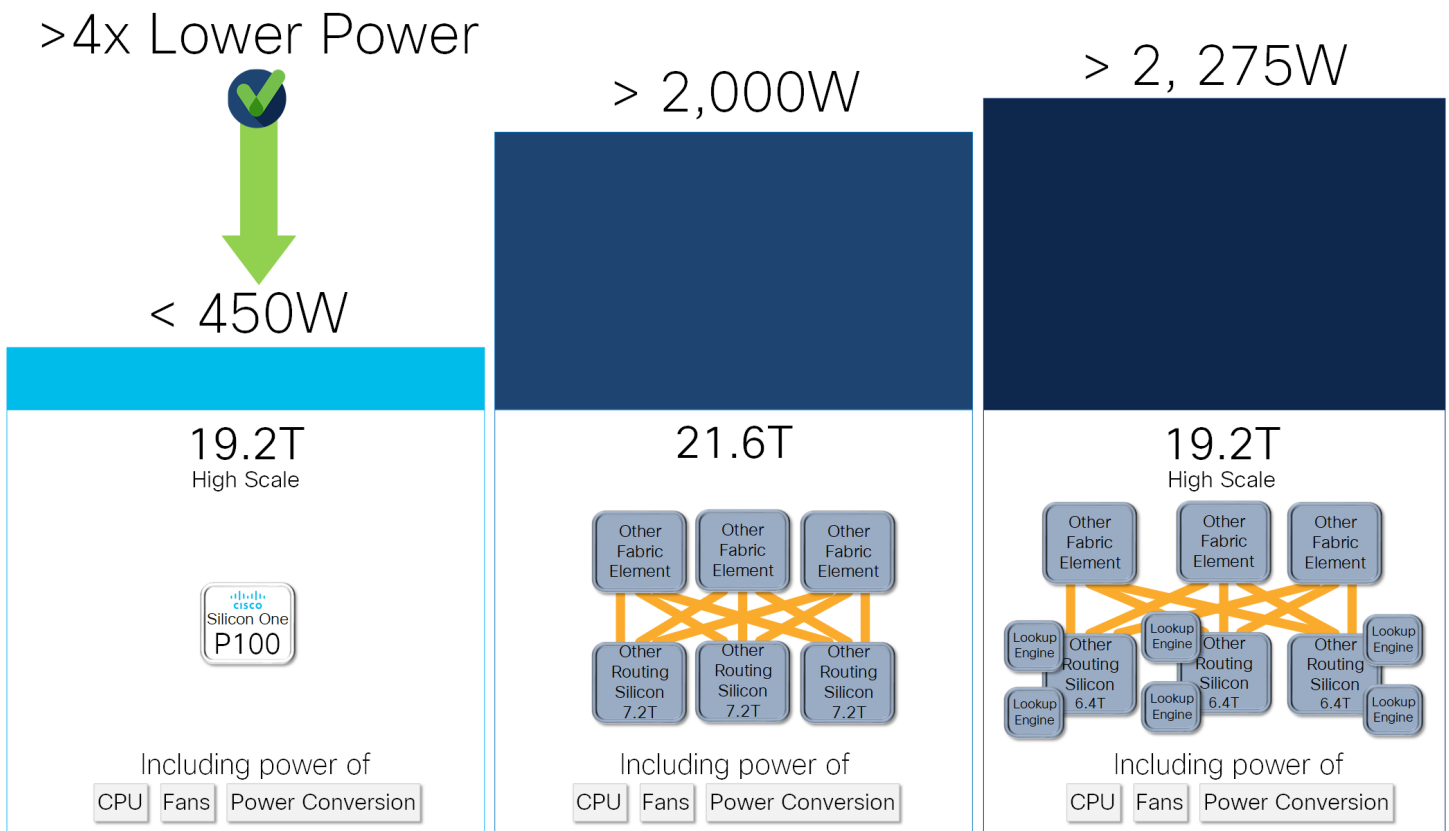


Figure 5. Fixed box system power (excluding optics)

Building a modular router line card

Modular systems are built with multiple pieces of silicon. Based on limitations on the faceplate, Printed Circuit Board (PCB), optics form factors, silicon, and orthogonal connector densities, most system vendors produce 36x400GE line cards with customer-facing 14.4 Tbps of bandwidth.

It's more complicated to understand how silicon bandwidth translates to customer facing bandwidth in modular systems than single silicon fixed boxed systems. To build modular systems, other routing silicon has a dedicated set of fabric interfaces on the back of the device, which are sometimes included in marketing material for chip bandwidth and sometimes they aren't.

Peeling away from the various marketing material available, we can see that with the next best routing silicon available on the market using two devices it's possible to build a 36x400, 14.4 Tbps line card. Like the high LPM scale fixed box architecture, we believe up to four external lookup engines are required on the line card, dropping the bandwidth for the line card slot to 12.8 Tbps while increasing the cost and power of the line card.

With Cisco Silicon One P100 we have flexible IO which can be used as fabric interfaces or generic Ethernet. It requires three P100s to build a 36x800, 28.8 Tbps line card which is two times higher bandwidth than any alternate solution, or two P100 devices to build a 25.6 Tbps modular line card.

Using our advanced architecture, coupled with 100G PAM4 SerDes technology, Cisco Silicon One enables line cards to push through previous limitations in front panel optics, PCB technologies, and orthogonal connectors, resulting in hardware that's two times higher bandwidth than other modular line cards.

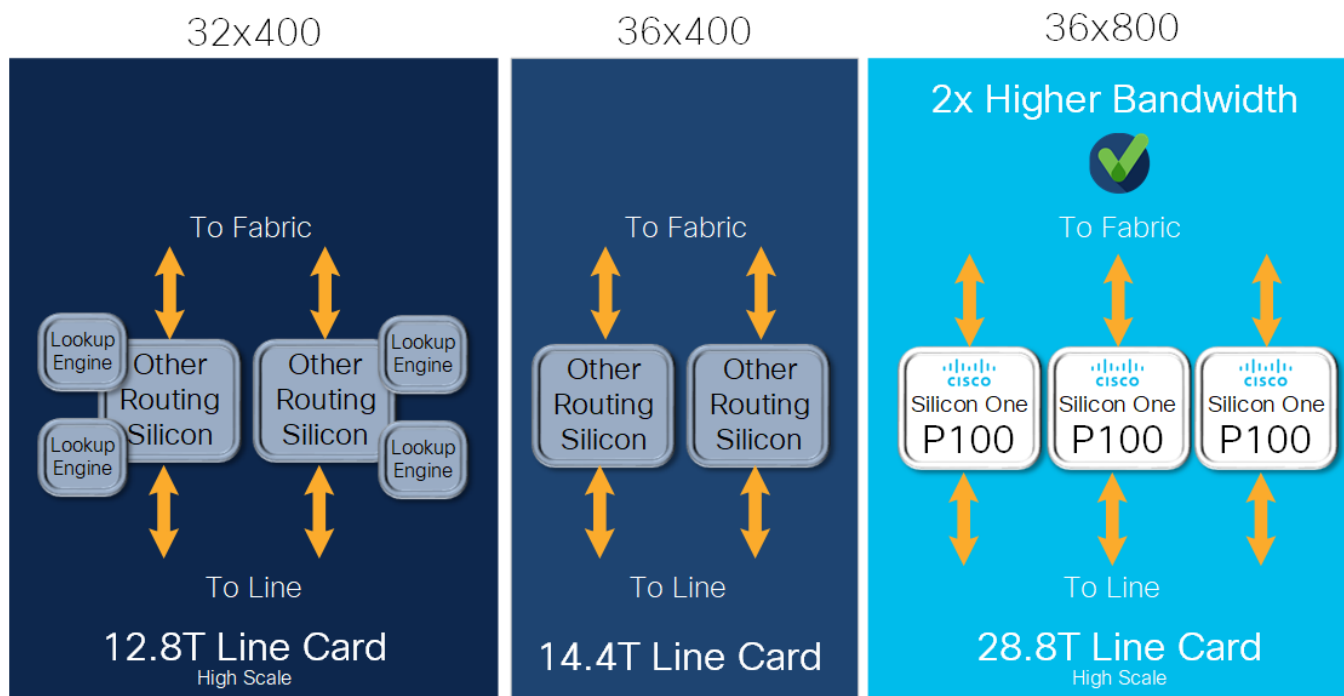


Figure 6. Modular line card

Utilizing 7nm, the same process node as other routing silicon on the market, the P100 benefits from the efficiency of the Cisco Silicon One architecture, resulting in massive power efficiency gains which we estimate to be 1.6 times more power efficient than others on the market.

System	Pieces of Silicon	LPM Scale	Power per 100G
28.8T P100 line card	1	High	< 4.2W
14.4T Other silicon line card	6	Medium	> 6.8W
12.8T Other silicon line card	10	High	> 8.8W

Table 2. Modular line card system power

>1.6-2x More Efficient

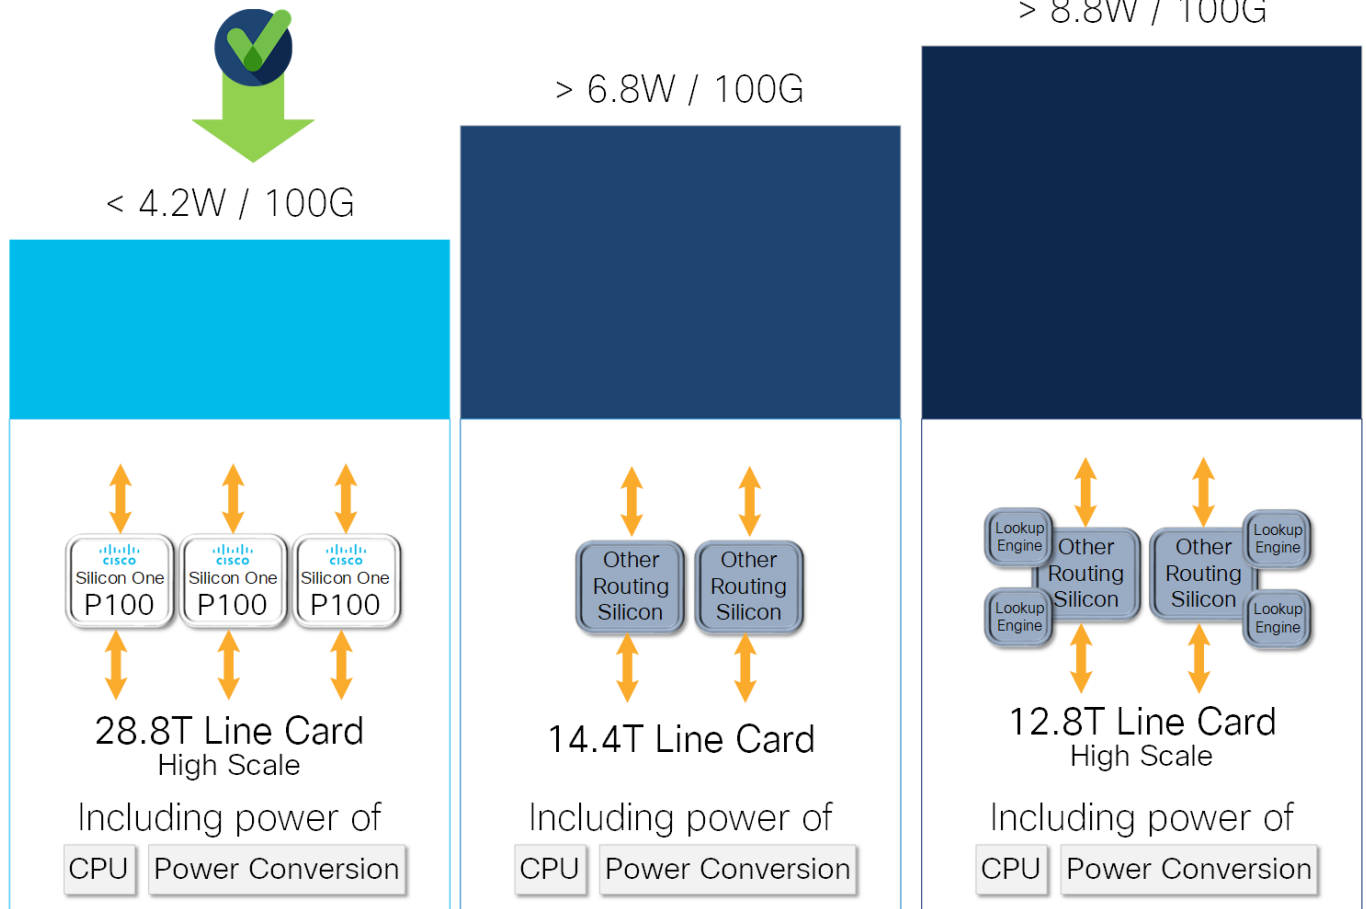


Figure 7. Modular line card system power

Learn more

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

Conclusion

Cisco Silicon One ushers in a new era of networking silicon, enabling densities and power efficiencies which were unheard of in the past. The radical reduction in silicon to build a 19.2 Tbps system impacts the cost, power, latency, and space for a router which can help provide significant cost savings to the network operators who deploy the equipment, while simultaneously enabling the industry's first 28.8 Tbps modular line card systems.