June 16, 2020 6 Comments

Data Center

400G Done Right: Cutting-Edge Nexus Innovation Continues

Thomas Scheibe

To help data centers scale with the increasing networking demands being placed on them, Cisco is making dramatic improvements in the capacity and manageability of our Nexus switching platforms.

Today we are announcing the next generation of Cisco's Cloud Scale ASIC, with bandwidth capacity of 25.6 terabits per second. We are also announcing availability of a new 400G line card for the Cisco Nexus 9500 modular chassis. Each new Nexus 9500 GX Line Card has 16 line-rate 400G ports.

We will also release new 32- and 64-port, single-ASIC Cisco Nexus 400G switches early next year, which take full advantage of the latest Cloud Scale ASIC.

It is critical that data centers scale up to support new capacities and densities. Cisco's Annual Internet Report estimates that there will be 29.3 billion networked devices and connections by 2023. The rapid growth in AI/ML workloads also leads to more distributed neural networks, which require higher peak bandwidth and greater densities in data centers. And, of course, the current shift to remote work has caused a surge in traffic between data centers and the access network.

Cisco-Designed Silicon

Our latest generation of the Cloud Scale ASIC has 4 times the capacity of current offerings. Its 7nm fabrication technology also means it's more power-efficient and more economical. This new ASIC supports key capabilities for the modern data center, including granular telemetry with real-time visibility into packets, flows, and performance. The chip has a larger L3 route table and more buffers than our previous ASICs, support for segment routing over the IPv6 dataplane (SRv6), and smart buffers to deliver optimal flow completion time in real-world networks.

The 25.6Tbps Cloud Scale ASIC was designed using a multi-die silicon technology, with one main die and multiple I/O chiplets. This decouples the switch logic from the SerDes and allows each to evolve independently. The newest Cloud Scale ASIC supports 512, 56 Gbps PAM4 SerDes. It is an ideal building block for 100G and 400G switch designs

The New 400G Products

The new Cisco Nexus 9500 GX line card, available now, provides 16 ports of line-rate 400G per slot in the existing Nexus 9500 Series chassis. For the first time in a Nexus line card, each port also supports 400G line-rate MACsec and Cloudsec. With this line card, customers can now deploy scalable mixed 100G/400G spine-leaf fabrics using a modular high-density 400G switch spine along with leaf switches with 100/400G uplinks.

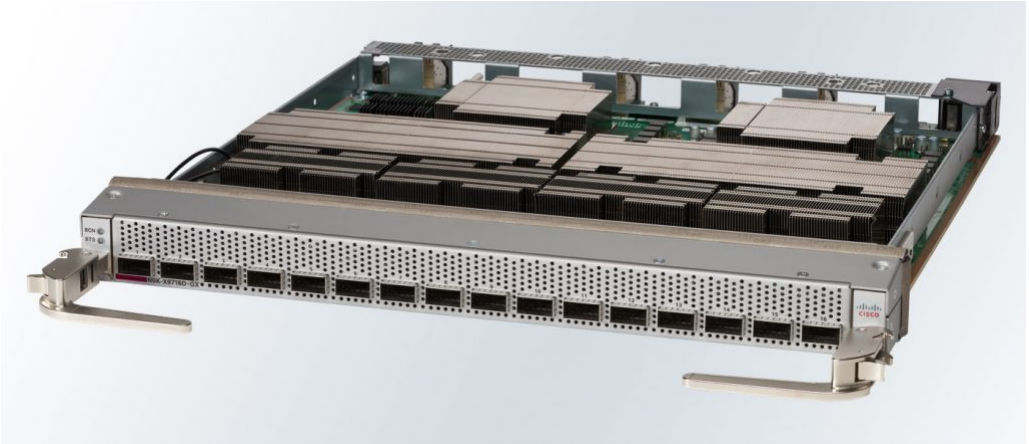

The new 400G Cisco line cards for the Nexus 9500 chassis support sixteen 400G ports.

The upcoming 64-port, 400G Nexus 9300-GX2 switch will deliver up to 25.6 Tbps of throughput, in a 2RU enclosure. A more compact 32-port model will deliver 12.8 Tbps of throughput in 1RU. These switches will empower customers to transition from 100G to 400G without changing their current fan-out designs.

The Cisco 400G Nexus portfolio provides the capabilities that early adopters need to pave the way for next-generation networking and enables large enterprises, service providers and 5G mobile operators to transform their cloud architectures to support increasing bandwidth demands, while reducing both CapEx and OpEx.

The upcoming Nexus 9300-GX2 switches will support up to 25.6 Tbps of throughput. 64-port model shown.

Both the new line card and the upcoming switches support flexible port configuration for lower speeds such as 10G, 25G, 50G, 100G and 200G, as well as backwards compatibility to QSFP28 or QSFP+ modules. The new line card provides further investment protection, as it is compatible with current generation 10G, 25G, 40G, and 100G Nexus 9500 Cloud Scale line cards.

Driving 400G Forward

Cisco's 25.6 Tbps ASICs, and our complete line of 400G line cards and switches, along with currently shipping 400G QSFP-DD optical modules, will support a transformation in how data center networks are architected and operated, with new efficiencies in cost, power, and fiber infrastructure.

The upcoming QSFP-DD-ZR/ZR+ Pluggable coherent modules will support 400G Ethernet interconnect payloads over Data Center Interconnect (DCI) links offering reaches up to 120km using higher order modulation.

We also continue to work with industry ecosystem and standards organizations to drive the standardization and interoperability of 400G technology building blocks. Beyond 400G, Cisco is also leading the industry with 800G Ethernet: Along with other industry leaders, we have formed the QSFP-DD800 MSA Group to increase the bandwidth of the QSFP-DD pluggable module form factor to 800Gbps. The new module standard includes backward compatibility with all other generations of QSFP modules, including QSFP+, QSFP28, QSFP56 and 400G QSFP-DD.

Our switches are built to last, as are the technologies, standards, and software that powers them. Cisco's next generation 400G Nexus switches feature strong price for performance, advanced flow data telemetry, and support for automation and Day-2 operations tools. They are designed for the future of cloud networking – standards-based, with investment protection for both backward and forward compatibility. This is 400G done right. I'm excited to see our customers deploy the Nexus products in their networks.

Learn more about High Capacity 400G Data Center Networking with Cisco Nexus Switches.

Share

Share:

6 Comments

- Yogesha MG, Technical Leader - CX says:
  June 16, 2020 at 7:04 pm

  Great to hear about this direction & Nexus has proven track record, this direction continues to keep us on leadership position.

Thanks
Yogesha MG
CX – Data Center

- Dave Hopkins says:
  June 16, 2020 at 9:40 pm

  This 400G bandwidth matches well with the bandwidth and low latency I/O characteristics of NVMe. The more these two technologies can be packaged together, the more likely they will be successful in delivering truly scaleable solutions.

  - Anonymous says:
    June 17, 2020 at 8:45 am

    Good job Mr Singh!

  - NoBS says:
    June 17, 2020 at 10:00 am

    I dont see the latency mentioned anywhere?

- Anonymous says:
  June 17, 2020 at 2:57 am

  Nice to hear! With the growing demand of ultra high speed Datacenter Networks, this definitely is a milestone for Cisco.

  Thanks
  Ashok
  CX DC

- Vern Stitt says:
  July 9, 2020 at 2:48 pm

  I'm very excited that NEXUS switches are getting new much faster line rates for backbone connectivity. All flash storage needs a backbone network that can support hundreds of servers at full speed.

Comments are closed.

CONNECT WITH CISCO