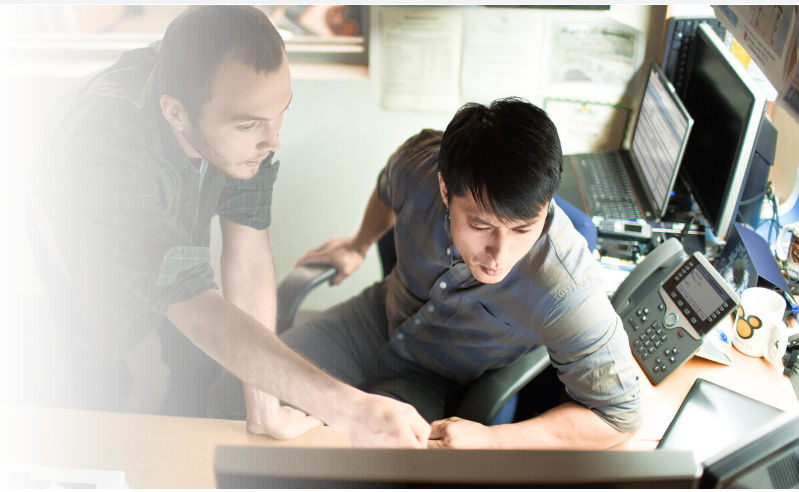

Cisco Nexus 9000 Series Switches with Cisco Cloud Scale Technology

Benefits

- **Smart buffers**: Large buffers coupled with advanced algorithms address real-world network congestion problems more efficiently, resulting in better application performance.
- **Deep telemetry and visibility**: Full packet and flow capture provide pervasive application visibility, real-time analysis, and troubleshooting.
- **Increased scale**: Increased router table sizes and end-host scale support scalable container-based deployments.
- **Price optimized**: Price optimized: Innovative application-specific integrated circuit (ASIC) design delivers feature-rich switches with 25-Gbps and 100-Gbps price points at the cost of 10- and 40-Gbps.

Overview

Data center networks have been rapidly evolving in recent years to address the changing nature of data center workloads. The modern data center workload is dominated by server-to-server traffic with intensive communication across hundreds and even thousands of servers. Additionally, shifts in application development and the associated growth in the use of Linux containers, microservices, and IP storage are also affecting many aspects of data center design.

With these shifts, network latency, throughput, scale, and visibility are extremely important to help ensure application performance and effective IT operations. A crucial design challenge in this networking environment is the selection and deployment of network switches that meet the performance requirements of applications and provide deep visibility into infrastructure and application behavior while reducing overall costs.

Cisco recognized that building the next-generation data center switch requires a next-generation application-specific circuit (ASIC), because the ASIC is the single largest factor influencing switch performance, capabilities, power profiles, and cost. To meet the new data center demands, Cisco developed the Cisco® Cloud Scale technology ASIC for next-generation switches.

By using a new generation of semiconductor device fabrication technology, Cisco Cloud Scale ASICs provide significantly greater density, more bandwidth, larger routing tables, larger buffers, more detailed visibility, and lower power consumption, all at a lower cost than competing Switches.

The following capabilities are among the main features in the new Cisco Nexus® 9000 Series Switches and modules with Cisco Cloud Scale ASICs:

- **Smart Buffering**: Cisco's new ASICs deliver larger internal buffers (40MB versus 16MB) plus several enhanced queuing and traffic management features not found in most merchant silicon switches. For more detail on these new features, see "Smart buffer vs Deep buffer" section at the end of this AAG.
- **Better Telemetry and Visibility**: The new Cisco switches with Cisco Cloud Scale ASICs support the Cisco Tetration Analytics™ platform for pervasive visibility into your applications and infrastructure. They are designed to capture every packet and every flow at line rate with no impact on the CPU. Switches with ASICs based on Cisco Cloud Scale technology also provide new sources of analytic information such as an enhanced flow table, buffer monitoring, and expanded counters to complement the diagnostic functions in the current generation of switches.

Figure 1. Cisco Tetration Analytics Platform

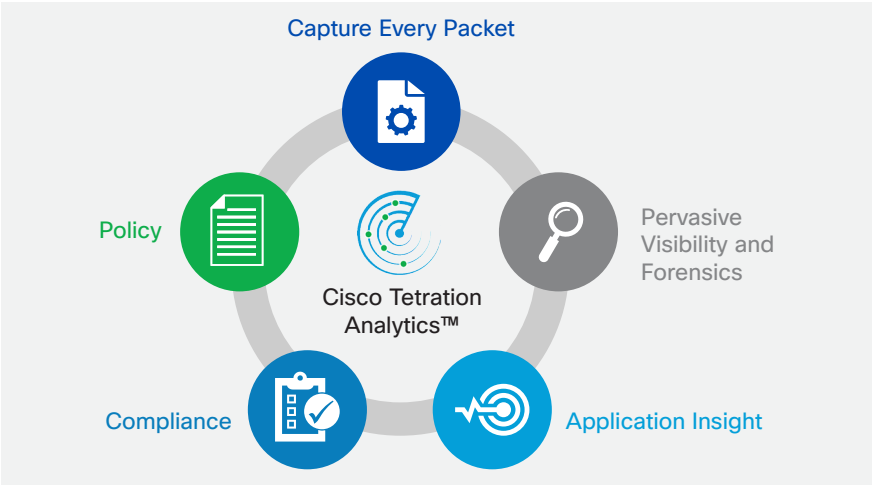

- **Greater Scalability:** The advances in application development and the accompanying use of containers and microservices all have repercussions for network scale. Cisco's switches with Cisco Cloud Scale technology offer increased route and end-host scale. The Cisco Nexus 9000 switches based on the new ASICs uses a new lookup-table architecture that can support up to 512,000 MAC address entries and up to 896,000 longest prefix match (LPM) entries: two to three times the number supported by popular Ethernet merchant silicon–based switches. The need for greater forwarding table capacity will continue to increase with the growth in IPv6 routing and containerization. Anticipating this demand, Cisco Nexus 9300 platform switches are designed to support 750,000 IPv6 routes compared with the 84,000 routes available with Ethernet merchant silicon.
- **Greater Bandwidth Capacity**: Demand for greater data center bandwidth will increase considerably in the next few years. Cisco's new switches offer more bandwidth per rack unit that is more cost effective than that provided by merchant silicon–based switches. The use of next-generation Cisco Cloud Scale fabrication technology has enabled Cisco to build a single switch-on-a-chip (SoC) ASIC that can support 3.6 terabits per second (Tbps) of line-rate routing capacity. It has also enabled Cisco to build the first 48 x 10- and 25-Gbps and 6 x 40- and 100-Gbps top-of-rack (ToR) switch in the industry with full Virtual Extensible LAN (VXLAN) support. By using the new ASIC, Cisco provides more capacity at lower cost.
- **Optimized Price Point**: By using a new generation of semiconductor device fabrication technology (16 nanometers [nm] versus 28 nm), Cisco new ASICs provide significantly greater transistor density and lower power consumption. With more transistors, switches support greater bandwidth, more ports, larger routing tables, more precise visibility, and larger buffers and enable you to implement advanced capabilities in the future. And with greater transistor density, you gain all these benefits at less cost.
- **Other Network Capabilities**: With the next-generation Cisco ASIC, capabilities such as segment routing, group-based security policy, network service headers, and full-featured VXLAN overlays are possible at much greater scale and at lower price points.

Although some of these features are available at much lower scale in merchant silicon–based switches, Cisco's new switches with Cisco Cloud Scale ASICs offer significantly greater scale and more features at an optimized price point.

Smart Buffer vs. Deep Buffer

Buffering is used to absorb traffic bursts, to avoid packet drops during network congestion. There are two architectural approaches to buffering:

Deep buffers (most merchant silicon) or Smart Buffers (Cisco Cloud Scale Technology). With a "deep buffer" approach, all data is placed in a common buffer on a first-come, first-served basis, and then moved out in the same order. It does not distinguish between flow sizes, both large flows and small flows are treated the same. Therefore, majority of the buffer gets consumed by large flows, starving the small flows. No matter how large the buffer is, the end result of deep buffers is added latency for both large and small flows, benefiting neither.

In contrast, Cisco Cloud Scale smart buffer approach offers two key innovations to manage buffers and queue scheduling intelligently by identifying and treating large and small flows differently.

Dynamic Packet Prioritization (DPP): DPP prioritizes small flows over large flows, helping ensure that small flows are not affected by larger flows due to excessive queuing.

Approximate Fair Drop (AFD): AFD introduces flow-size awareness and fairness to early-drop congestion-avoidance mechanism.

Both DPP and AFD ensure small flows are detected and prioritized and not dropped to avoid timeouts, while large flows are given early congestion notification through TCP to prevent over use of buffer space. As a result, smart buffers allow large and small flows to share the switch buffers in a much more fair and efficient manner. This provides buffer space for small flows to burst and large flows to fully utilize the link capacity with much lower latency times than a simple large buffer approach implemented in most merchant silicon switches.

Next Steps

To find out more about the new Cisco Cloud Scale switches, visit the Cisco Nexus 9000 Series website at www.cisco.com/go/nexus9000.

C45-737330-00 06/16