Cisco Silicon One G204

Contents

Value statement

Web-scale, enterprise, and service provider hardware is built around switching silicon, routing line card silicon, and routing fabric silicon. These three basic building blocks enable silicon and system vendors to create unique architectures tuned for individual markets and industries. However, forcing customers to consume and manage these disjointed, dissimilar products has also caused an explosion in complexity, CapEx, and OpEx.

The Cisco Silicon One™ architecture ushers in a new era of networking, enabling one silicon architecture to address a broad market space, while simultaneously providing best-of-breed devices. Cisco Silicon One doesn't mean one device across the network, but one architecture and many optimized devices across the network.

At 12.8 Tbps, the Cisco Silicon One G204 builds on the groundbreaking technology of the Cisco Silicon One G200. What's more, it fully optimizes the design for high bandwidth web-scale switching in front-end, storage networking and other applications, enabling a deterministic, low latency, and power efficient 16x800GE switch. The Cisco Silicon One G204 offers the same base level feature set as the G200 with key enhancements related to SerDes reach, power, latency, congestion management and telemetry at one-fourth the throughput.

Product overview

The Cisco Silicon One G204 processor is a 12.8-Tbps, full-duplex, standalone switching processor that can be used to build fixed form factor switches ideally targeted for web-scale data center top of rack and leaf switches serving front-end, storage networking and other applications.

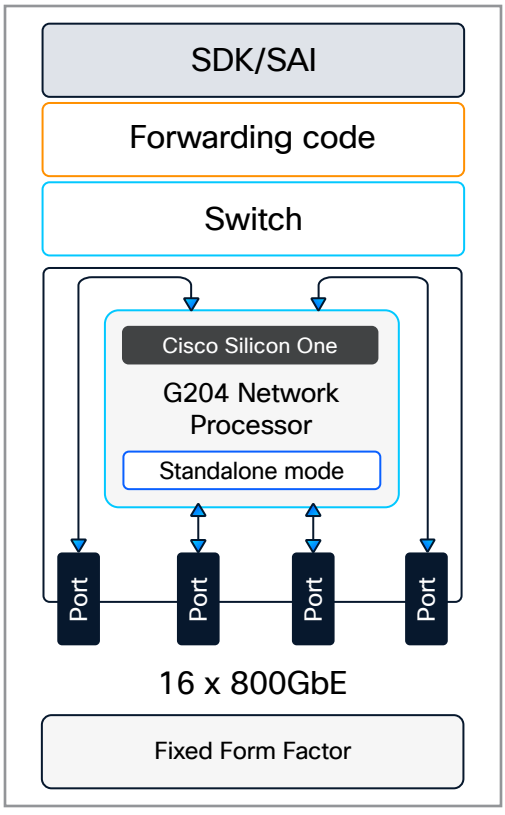

Features and benefits

Table 1. Architectural characteristics and benefits

Feature	Benefit
One architecture across multiple markets	One architecture greatly simplifies customer network infrastructure deployments, saving both OpEx and CapEx while simultaneously shortening qualification time.
One SDK across market segments and applications	One SDK provides a consistent point of integration for all applications across the entire network infrastructure, improving quality while reducing OpEx and CapEx for customers.
Latency optimized programmable network processor	Deterministic and low-latency programmable processor that offers additional run-to-completion flexibility for complex flows. This architecture uniquely addresses the requirements of web-scale providers' switching applications without sacrificing features and programmability.
Large and fully unified packet buffer	Fully shared on-die packet buffer allows any input or output port to consume the entire memory. This capability reduces packet loss and minimizes PFC events, thus maximizing network performance under varying traffic conditions and enabling low latency for RDMA and RoCEv2 protocols.
Unmatched telemetry and visibility	Support for standard and emerging web-scale, in-band telemetry protocols enables advanced congestion control and advanced flow tracking with temporal dynamics. Together with in-network trigger events, these capabilities enable post-event analysis in hardware time scales.
Advanced load balancing capabilities	Support for stateless and stateful congestion-aware load balancing techniques ensures optimal delivery of packets through the network. This helps to ensure optimal Flow Completion Time (FCT) for traditional web-scale networks.
Network resiliency assurance	Support for hardware-based link monitoring and rebalancing of traffic helps ensures optimized network utilization even under link failure conditions in large-scale networks.

Prominent feature

Flexibility and performance for next-generation web-scale front-end and storage networks.

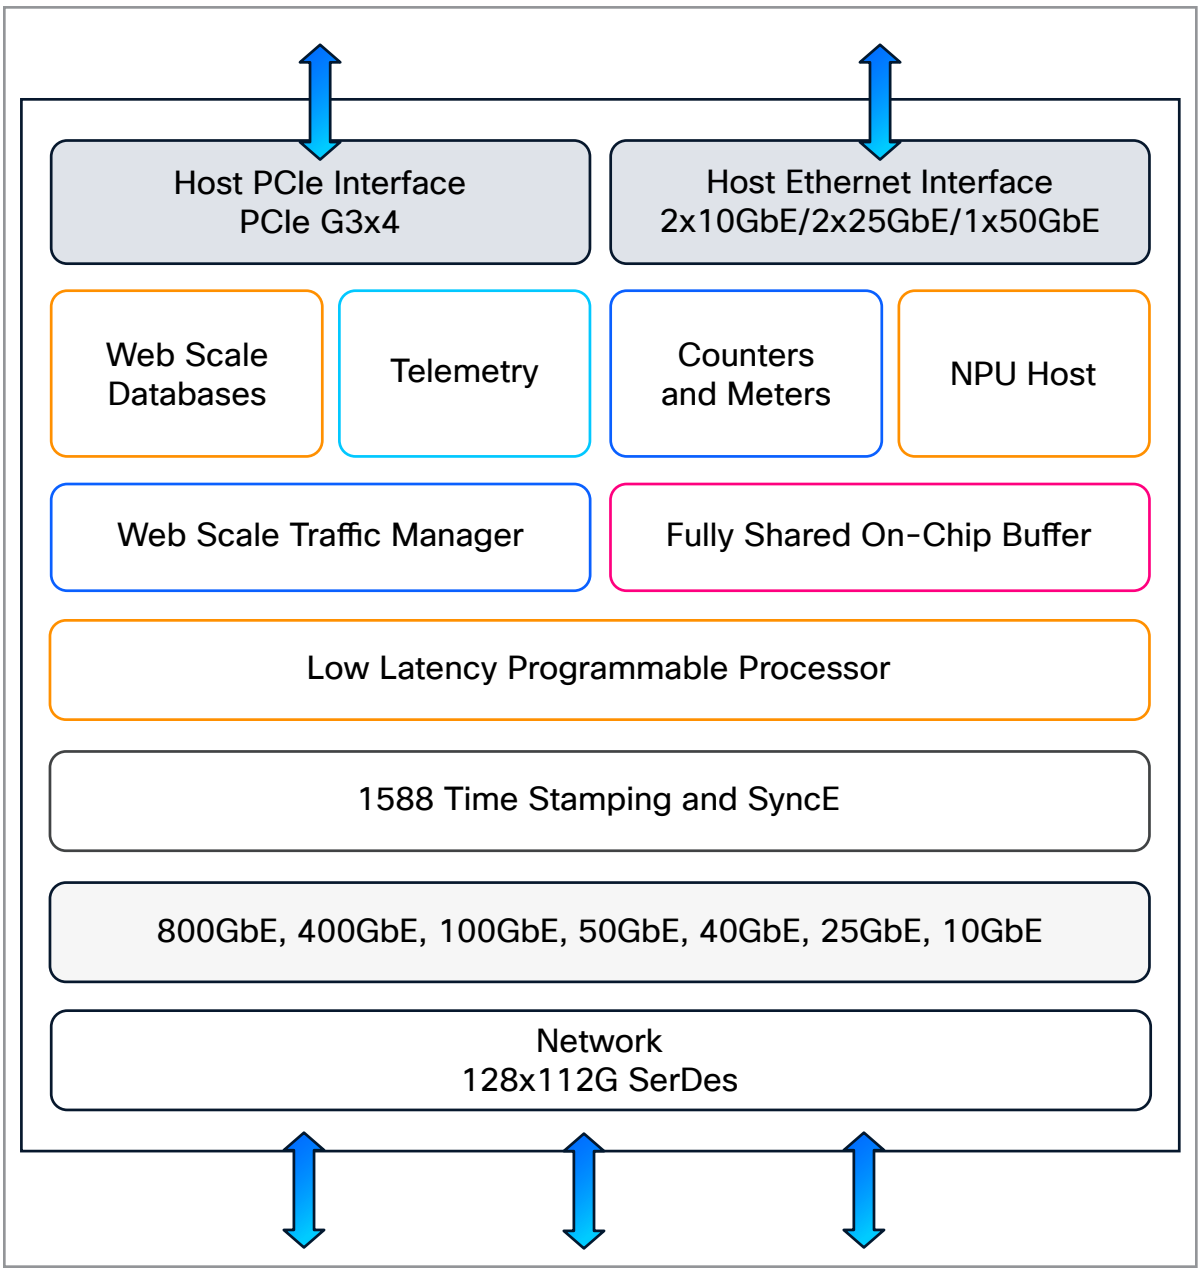

Features

- 128 112G long reach SerDes supporting NRZ and PAM4 modulation
- 128 Ethernet MACs enable maximum network scale for optimized network deployments
- Flexible port configuration supporting 10/25/40/50/100/200/400/800 Gbps
- Large, fully shared, on-die packet buffer
- 1588v2 and SyncE support with nanosecond-level accuracy
- On-chip, high-performance, programmable host Network Processing Unit (NPU) for high-bandwidth offline packet processing (for example, OAM processing, MAC learning)
- Multiple embedded processors for CPU offloading
- PCIe gen3 and two Ethernet interfaces to connect to the host CPU complex
- Advanced features for optimal load balancing, fault detection and recovery and telemetry

Traffic management

- Multiple output queues per output port support web-scale customers' future needs
- Support for ingress and egress traffic mirroring
- Support for link-level (IEEE802.3x) flow control
- Support for PFC priority-level (802.1Qbb) flow control
- Support for PFC watchdog
- Dynamic thresholds and policies help ensure optimal usage of the fully-shared packet buffer
- Support for probabilistic multi-color ECN marking
- Support for probabilistic multi-color WRED drop porfiles
- Switch CNP (congestion notification packet)
- Microburst detection

NPU

- Optimized, deterministic and low latency programmable processor with extensions to complete run-to-completion programmable processors for advanced features
- Web-scale optimized and fungible tables
- Achieves line-rate at small packet sizes and full web-scale feature sets running

Load balancing

- Flow load balancing using WECMP, ECMP or LAG with innovative non-correlated hashing functions to avoid polarization even across massive scale networks
- Support for WECMP without replicating entries
- Congestion-aware flowlet load balancing with ability to detect and handle elephant flows
- Congestion-aware packet spraying independent of flow characteristics
- Packet trimming with recycle

Instrumentation and telemetry

- Support for standard (P4-INT, IFA1.0, IFA2.0) and emerging web-scale in-band telemetry protocols
- Advanced flow scope with temporal dynamics and live network trigger for post-mortem analysis
- Programmable meters used for traffic policing and coloring
- Programmable counters used for flow statistics and OAM loss measurements
- Counters for port utilization, microburst detection, delay measurements, flow tracking, elephant flow detection, and congestion tracking
- Tail timestamping
- INT/CSIG (congestion signalling)
- Packet drop counters
- Traffic mirroring: (ER)SPAN on congestion and drop
- Support for sFlow and NetFlow

Front-end and storage networking

- Advanced load balancing techniques for optimal network performance
- Hardware based link failure isolation and re-routing to enable performance across large scale networks
- Advanced congestion control and telemetry benefiting both TCP and RoCEv2 RDMA traffic
- Deterministic low-latency performance

Software

- SDK APIs in C++, and in C
- Switch Abstraction Interface (SAI)
- Functional simulation environment
- SONiC refference on functional simulator and hardware platform
- Support for x86 and ARM host CPU complexes
- Distribution-independent Linux packaging
- Debug support: gRPC-based CLI and Python shell

Programmability

- Application development is handled by an IDE programming environment
- At compilation, the forwarding application generates low-level register/memory access APIs and higher-level SDK application APIs
- Provides application support for a wide range of data center, service provider, and enterprise protocols
- Ability to develop the SDK and applications running over the SDK over a simulated Cisco Silicon One device

Cisco application

Utilizing Silicon One's extensible programming toolkit, we are always adding features to address new markets and new customer requirements. A sample of features supported includes:

- IPv4/v6
- MPLS
- Ethernet Switching
  - 802.1d, 802.1p, 802.1q, 802.1ad
- IP Tunneling
  - IPinIP
  - GRE
  - VXLAN
- Segment Routing
  - SRv6 uSID
  - MPLS
- RDMA Support
  - Priority Flow Control (PFC) 802.1Qbb
  - Flow Control (802.3x)
  - Probabilistic multi-color ECN marking
  - Probabilistic multi-color WRED drop
- Integrated Routing and Bridging (IRB)
- HSRP/VRRP
- Policy-Based Routing
- Security and QoS ACLs
- ECMP and LAG (802.3ad)
- Multicast
  - IGMP
- Protection (Link/Node/Path and TI-LFA)
- QoS Classification and Marking
- Congestion Management
- Telemetry
  - NetFlow, sFlow
  - Inband Telemetry (P4-INT, IFA, and emerging protocols)
  - (ER)SPAN
  - Packet Mirroring with Appended Metadata
  - Lawful Intercept
- Warmboot
- DDoS Mitigation
  - Control-Plane Policing
  - BGP Flowspec
- Timing and Frequency Synchronization
  - SyncE
  - 1588

Product sustainability

Information about Cisco's Environmental, Social and Governance (ESG) initiatives and performance is provided in Cisco's CSR and sustainability reporting.

Table 2. Cisco Environmental Sustainability Information

Sustainability topic		Reference
General	Information on product-material-content laws and regulations	Materials
	Information on electronic waste laws and regulations, including our products, batteries and packaging	WEEE Compliance
	Information on product takeback and reuse program	Cisco Takeback and Reuse Program
	Sustainability Inquiries	Contact: csr_inquiries@cisco.com
Material	Product packaging weight and materials	Contact: environment@cisco.com

Cisco Capital

Flexible payment solutions to help you achieve your objectives

Cisco Capital makes it easier to get the right technology to achieve your objectives, enable business transformation and help you stay competitive. We can help you reduce the total cost of ownership, conserve capital, and accelerate growth. In more than 100 countries, our flexible payment solutions can help you acquire hardware, software, services and complementary third-party equipment in easy, predictable payments. Learn more.

For more information

Learn more about Cisco Silicon One.

C78-5456267-00 10/25