

Cisco Silicon One G300

Contents

Value statement2
Product overview2
Features and benefits.....	.3
Prominent features4
Cisco application.....	.8
Product sustainability9
For more information9

Value statement

Hardware for artificial intelligence (AI) networking and web-scale, enterprise, and service provider networks 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 ushered in a new era of networking, enabling one silicon architecture to address a broad market space while simultaneously providing best-in-class devices. Cisco Silicon One doesn't mean one device across the network, but one architecture and many optimized devices across the network.

At 102.4 Tbps, the Cisco Silicon One G300 processor builds on the groundbreaking technology of the Cisco Silicon One G200. What's more, it fully optimizes the design for high-bandwidth, AI networking and web-scale switching, enabling a deterministic, low-latency, and power-efficient 64x 1600 Gigabit Ethernet (GE) switch.

Product overview

The Cisco Silicon One G300 processor is a 102.4-Tbps, full-duplex, standalone switching processor that can be used to build fixed-form-factor switches ideally targeted for web-scale data center spine and leaf applications spanning front-end, back-end, scale-out, and scale-up networking roles serving AI applications.

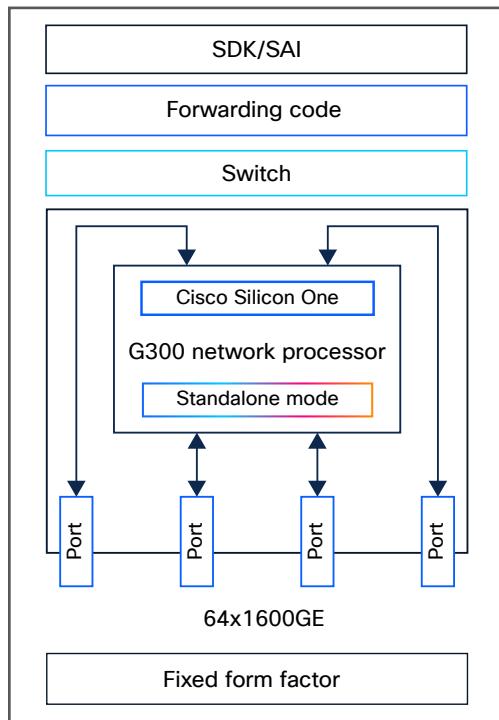


Figure 1. Form factor

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 software development kit (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 priority flow control (PFC) events, thus maximizing network performance under varying traffic conditions and enabling low latency for the Remote Direct Memory Access (RDMA) and RDMA over Converged Ethernet (RoCE) v2 protocols.
Unmatched telemetry and visibility	Support for standard and emerging web-scale, in-band telemetry protocols enable 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 helps ensure optimal delivery of packets through the network. This helps to ensure optimal flow completion time (FCT) for traditional web-scale networks and optimal job completion time (JCT) and tokens-per-second performance in massive-scale AI networks.
Network resiliency assurance	Support for hardware-based link monitoring and rebalancing of traffic helps ensure optimized network utilization even under link failure conditions in large-scale networks.

Prominent features

Flexibility and performance for next-generation, web-scale, front-end, and back-end AI networks

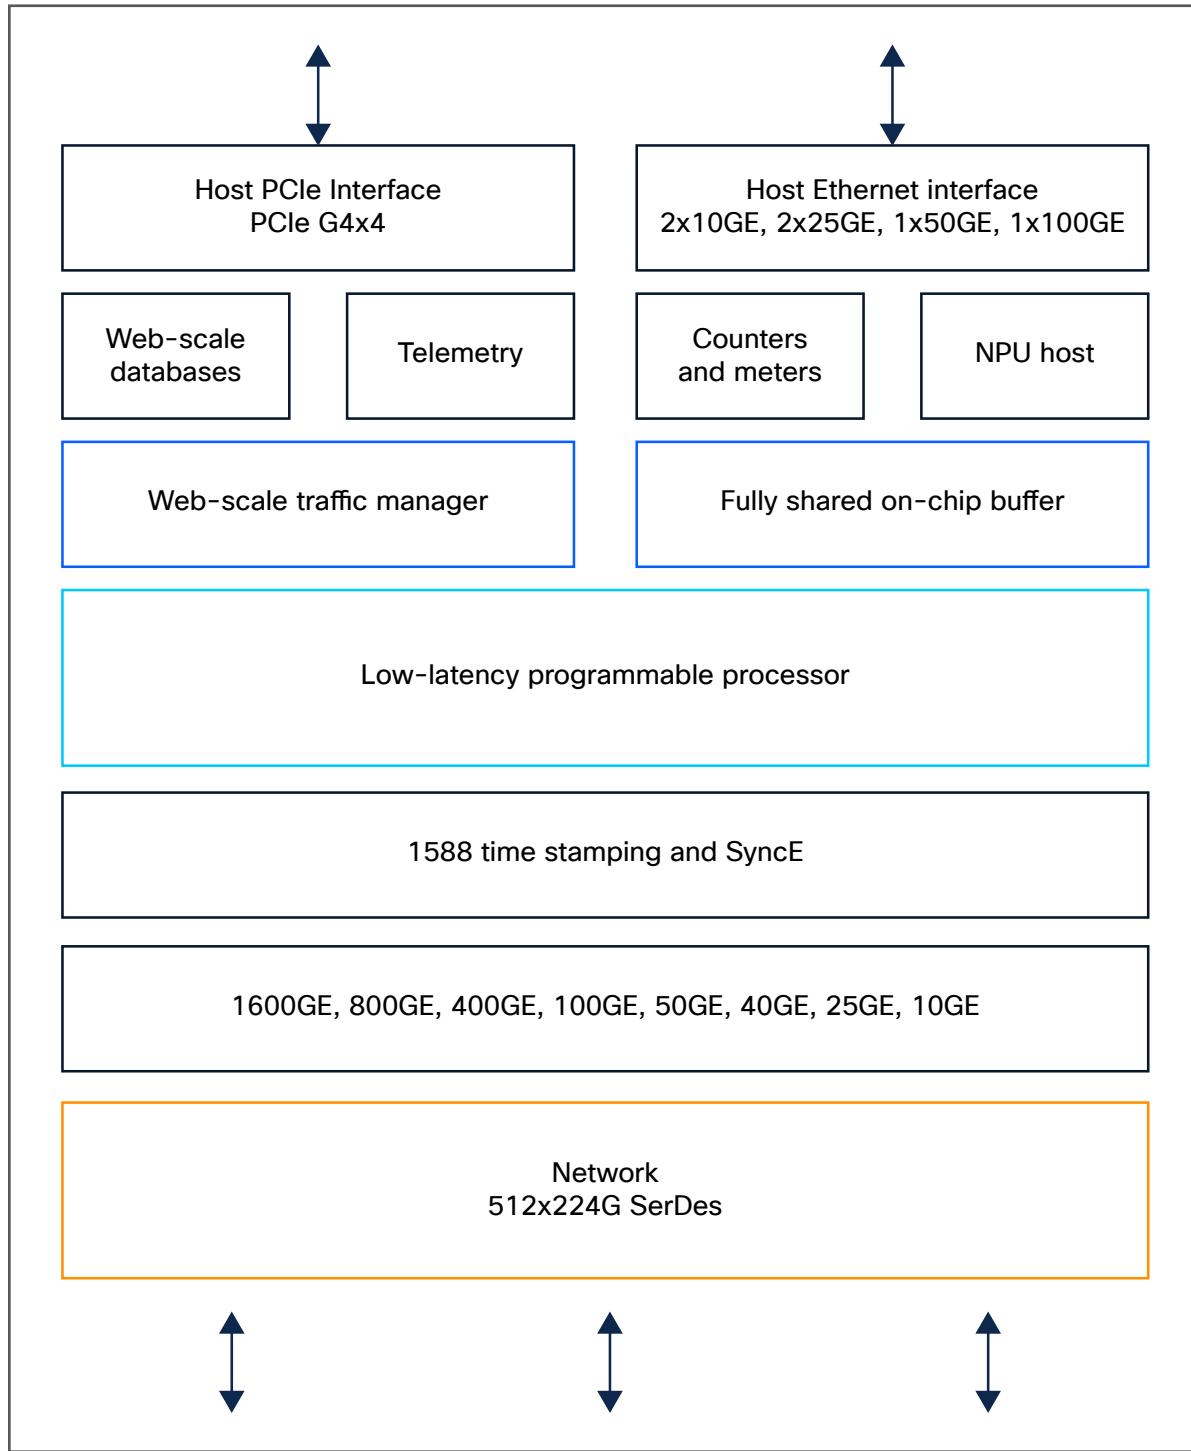


Figure 2. Block diagram

Features

- 512x 224-Gbps long-reach serializer/deserializer (SerDes) supporting non-return to zero (NRZ) and four-level pulse amplitude modulation (PAM4) to deliver up to 1.6 Tbps Ethernet
- 512x Ethernet MACs enable maximum network scale-out for low-latency-optimized network deployments
- Flexible port configuration supporting 10, 25, 40, 50, 100, 200, 400, 800, and 1600 Gbps
- Large, fully shared, on-die packet buffer provides reliable GPU data delivery in bursty traffic environments and over longer distances
- IEEE 1588v2 Precision Time Protocol (PTP) and Synchronous Ethernet (SyncE) support with nanosecond-level accuracy
- On-chip, high-performance, programmable host network processing unit (NPU) for high-bandwidth offline packet processing (for example, Ethernet Operations, Administration, and Maintenance [OAM] processing, MAC learning)
- Multiple embedded processors for CPU offloading
- PCIe Gen 4 and four Ethernet interfaces to connect to the host CPU complex
- Advanced features for AI deployments with optimal load balancing, fault detection, 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 (IEEE 802.3x) flow control
- Support for priority-level (802.1Qbb) flow control (PFC)
- Support for PFC watchdog
- Dynamic thresholds and policies help ensure optimal usage of the fully shared packet buffer
- Support for probabilistic multicolor explicit congestion notification (ECN) marking
- Support for probabilistic multicolor weighted random early detection (WRED) drop profiles

NPU

- Optimized, deterministic, and low-latency programmable processors enable co-innovation of new features with customers, accelerating time to revenue without requiring silicon re-spin
- AI networking and web-scale-optimized large and flexible tables scale, enabling high route and endpoint scale
- Achieves line rate at small packet sizes and with full AI networking and web-scale feature sets running

Load balancing

- Flow load balancing using weighted equal cost multi-path (WECMP), ECMP, or link aggregation group (LAG) with innovative noncorrelated hashing functions to avoid polarization, even across massive-scale networks
- Support for WECMP with dynamic load balancing (DLB)
- DLB: Congestion-aware flow and flowlet load balancing with ability to detect and handle elephant flows
- DLB: Congestion-aware packet spraying independent of flow characteristics
- Static pinning with fallback to ECMP
- Topology-aware load balancing using a hardware-assisted fabric routing protocol for extremely fast convergence
- Dynamically load-balanced flows using hardware-assisted fabric routing protocol information and sub-round trip time (RTT) signals
- Fast failover for ECMP flows
- Adaptive load balancing
- Native support for Segment Routing (SR) v6 micro segment identifier (uSID) for deterministic, scalable, and efficient traffic steering crucial for GPU-to-GPU communication

Instrumentation and telemetry

- Support for standard (P4-INT, inband flow analyzer [IFA] 1.0, IFA 2.0) and emerging web-scale, in-band telemetry protocols
- Switch congestion notification packet (CNP)
- Packet trimming (including back to sender)
- Congestion signaling (CSIG)
- Tail timestamp
- Unreachable destination notification packet (UDNP)
- Advanced flow scope with temporal dynamics and live network trigger for postmortem 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
- Traffic mirroring: (ER)SPAN (switched port analyzer, encapsulated remote SPAN) on congestion and drop
- Support for sFlow and NetFlow

AI networking

- Advanced load-balancing techniques for optimal network performance
- Hardware-based link failure isolation and rerouting to enable performance across large-scale networks
- Advanced congestion control and telemetry to optimize JCT and tokens-per-second performance
- Deterministic low-latency performance

Software

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

Programmability

- Application development is handled by an integrated development environment (IDE)
- 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 on a simulated Cisco Silicon One device



Cisco application

Using 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 those listed in the table below.

Table 2. Sample of supported features

<ul style="list-style-type: none"> ▪ IPv4/v6 ▪ MPLS ▪ Ethernet switching <ul style="list-style-type: none"> - 802.1d, 802.1p, 802.1q, 802.1ad ▪ IP tunneling <ul style="list-style-type: none"> - IP-in-IP - Generic Routing Encapsulation (GRE) - VXLAN ▪ Segment Routing <ul style="list-style-type: none"> - SRv6 uSID - Multiprotocol Label Switching (MPLS) ▪ RDMA support <ul style="list-style-type: none"> - Priority Flow Control (PFC) 802.1Qbb - Flow Control (802.3x) - Probabilistic multicolor ECN marking - Probabilistic multicolor WRED drop ▪ Integrated routing and bridging (IRB) ▪ Hot Standby Router Protocol (HSRP)/Virtual Router Redundancy Protocol (VRRP) ▪ Policy-Based Routing ▪ Security and quality-of-service (QoS) access control lists (ACLs) ▪ ECMP and LAG (802.3ad) ▪ Multicast <ul style="list-style-type: none"> - Internet Group Management Protocol (IGMP) ▪ Protection (Link/Node/Path and Topology-Independent Loop-Free Alternate [TI-LFA]) 	<ul style="list-style-type: none"> ▪ QoS classification and marking ▪ Congestion management <ul style="list-style-type: none"> - Priority-based flow control (PFC) - Enhanced congestion notification (ECN) ▪ Load balancing (LB) <ul style="list-style-type: none"> - Dynamic LB: {flow, flowlet, packet}-spray based - Weighted ECMP with DLB - Static pinning with fallback to ECMP - Topology-aware LB - Dynamically load-balanced flows - Global load balancing (GLB) - Fast failover for ECMP flows - Adaptive LB ▪ Telemetry <ul style="list-style-type: none"> - NetFlow, sFlow - In-band telemetry (P4-INT, IFA, and emerging protocols) - (ER)SPAN - Switch congestion notification packet (CNP) - Unreachable destination notification packet (UDNP) - Congestion signaling (CSIG) - Packet trimming (back to sender) - Packet mirroring with appended metadata - Lawful intercept ▪ Warmboot ▪ Distributed denial-of-service (DDoS) mitigation <ul style="list-style-type: none"> - Control-plane policing ▪ Border Gateway Protocol (BGP) Flow Specification ▪ Integrated tamper-resistant root of trust (RoT) mitigates vulnerability at the edge of the network ▪ Timing and frequency synchronization <ul style="list-style-type: none"> - SyncE - IEEE 1588 PTP
--	--

Product sustainability

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

Table 3. Cisco environmental sustainability information

Sustainability topic	Reference
General	Information on product-material-content laws and regulations
	Information on electronic waste laws and regulations, including our products, batteries, and packaging
	Information on product takeback and reuse program
	Sustainability inquiries
Material	Product packaging weight and materials

For more information

[Learn more](#) about Cisco Silicon One.