Cisco Silicon One K100 Processor

Version 3.0

October 2025

Contents

The silicon industry has always been plagued with the trichotomy of switching silicon, routing line card silicon, and routing fabric silicon. Using these three basic building blocks, silicon and system vendors created unique architectures tuned for individual markets and industries. Consequentially, forcing customers to consume and manage these disjointed and dissimilar products caused an explosion in complexity, CapEx, and OpEx for the industry.

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.

The Cisco Silicon One K100 is a 6.4Tbps, high-performance, flexible, power-efficient switching silicon for high-end access and edge deployments and for consolidating traffic and services into 100G/200G/400G aggregation networks. K100 and its carrier edge focus, provides a service provider the utmost flexibility in the deployment and evolution of edge services, enabling any service, anywhere-concept, the aggregation, and convergence of both wired-based and wireless access services such as residential IP broadband (OLT, IP, DSLAM), FTTx, BNG, cnBNG, PON, Carrier Ethernet, and 4G/5G xHaul into the access-agnostic IP/Ethernet Metro aggregation network. K100 addresses the service provider needs for increased capacity and scale, enabling deployment of edge SLA-compliant, QoS-enhanced, multi-tenant, multi-carrier edge services.

The Cisco Silicon One K100 builds on the groundbreaking technology of the Cisco Silicon One devices before it, while increasing the lead over all other routing silicon in the market with focus on high scale for service termination, port density, best-in-class timing, programmability, telemetry and analytics, and security provisions with its support for MACsec and IPsec.

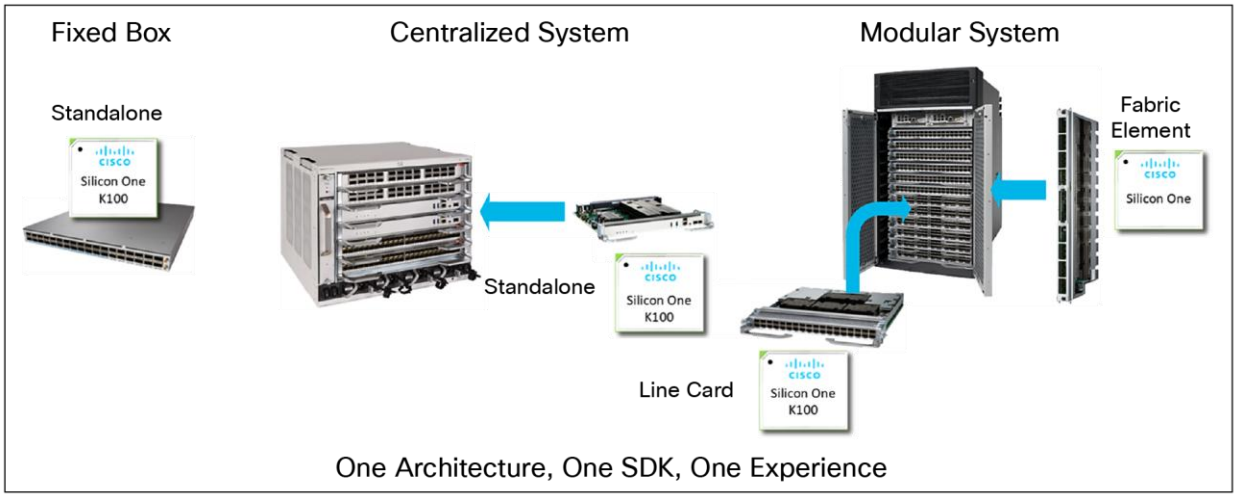

Figure 1. Cisco Silicon One K100 Configurations

As a standalone device, K100 finds its sweet spot in the service provider edge and typical bandwidth configurations in the range of 3.2 to 6.4 Tbps. Such systems offer the lowest Total Cost of Ownership (TCO) for cost and power and come in small form factors (1 to 3RU).

The maximum scalability and port fanout of a K100 system are achieved by configuring K100-based line cards within a Cisco Silicon One fabric-based chassis. These configurations offer several benefits, including scaling up the edge, providing, and if needed, N+1 redundancy, facilitating future capacity expansion and upgradability, and allowing the consolidation of edge and aggregation functionalities into a Silicon One fabric-based chassis.

Given the large variety of possible deployment scenarios, K100 supports per Serdes speed configuration, enabling the most flexibility in port speed configuration and deployment transition scenarios. For example, in service provider access and edge, a K100 enables seamless aggregation of a wide range of port speeds from 10 Mbps to 400 Gbps.

Beyond the flexibility of interface mix, K100 has large on-chip buffers expandable to even larger size via in-package HBM, which can be used by any and all of its large number of queues, supporting hierarchical QoS and ingress or egress traffic management. The efficient operation of the device in bursty traffic situations benefits from the ability of the device to identify mice and elephant flows and its support for flowlet load balancing.

K100 excels with its very large Layer 2 and Layer 3 tables, supporting highly flexible Layer 2 and Layer 3 configurations and tunneling protocols, and its support for large number of flows, flow analytics and telemetry (large ACL and NetFlow tables), and programmability that meets the performance needs of the diverse service provider edge environments.

With its IEEE 802.1ae MAC security support on all ports and all port speeds, it allows traffic encryption at the physical layer. In addition, it supports IPsec for P2P and P2MP IP tunnels, plus encryption across L2VPN and L3VXLAN tunnels.

The service provider edge market is a diverse market with hundreds of unique topologies and Layer 2, Layer 3, and tunneling service requirements. K100 facilitates deployment of such diverse networks both in protocol processing requirements and scale of services in the downstream and upstream directions. Thus, in the downstream direction, K100 addresses the needs of such scenarios with its large number of hierarchically scheduled and shaped queues that enable a large number of subscribers and services in the downstream direction. Similarly in the upstream direction, K100, with its large scale of meters and ACLs, can deal with the challenge of policing and aggregating a large amount of access traffic.

Product overview

K100 is a member of the “K” series of Cisco Silicon One processors and has a maximum I/O bandwidth of 9.6 Tbps. It can be configured as:

- A 9.6-Tbps (192x56G) I/O device with 6.4-Tbps full-duplex packet processing/forwarding routing/switching processing and large on-chip buffering, expandable via in-package HBM.
- A 6.4-Tbps (128x56G) network facing interface and a 3.-2Tbps (64x56G) fabric interface and large on -chip buffering. In this configuration, K100 processes up to 6.4-Tbps network facing full-duplex routing/switching traffic. The fabric interface enables connectivity to a Cisco Silicon One fabric for scalable fabric-based configurations.

Features and benefits

Table 1. Architectural Features and their Benefits

Features	Benefit
Unified architecture with the rest of the Cisco Silicon One devices	Greatly simplifies customer network infrastructure deployments
Unified SDK across market segments and applications	Provides consistent application programmability across the entire network infrastructure
High-bandwidth routing silicon	7-nm, 6.4-Tbps routing silicon with a wealth of features for service provider edge and aggregation applications
Power-efficient routing silicon	The power efficiency of 7 nm and the Cisco Silicon One architecture
Large and fully unified packet buffer	Fully shared on-die buffer and an external even larger in-package packet buffer
Routing efficiency with support of features at scale	Addresses the requirements of service provider edge and aggregation routing applications
Run-to-completion network programmable processor	Provides feature flexibility without compromising performance or power efficiency
Programmable	A programming processor to allow for rapid feature development
Encryption support	Line rate MACsec encryption across all device ports. IPsec, Cisco® ClearTag, and Cisco Cloudsec encryption support for IP network deployments

Technical details

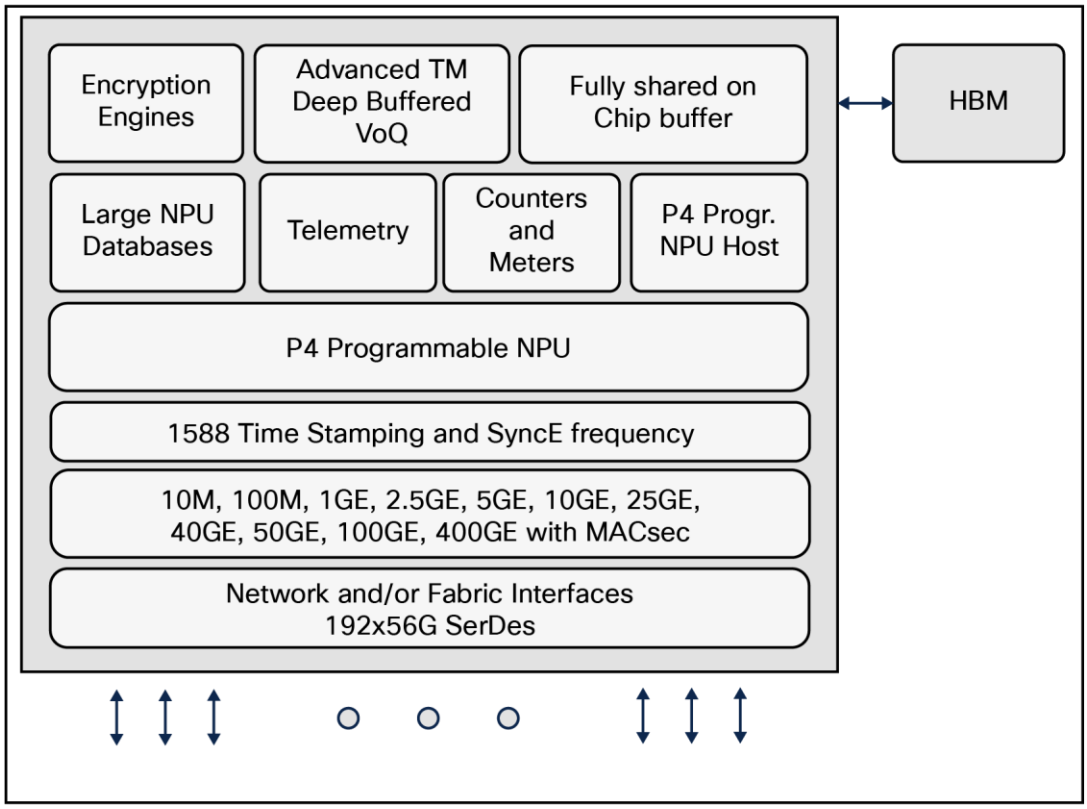

Figure 2. K100 Block Diagram

Features

- 192 56G SerDes; SerDes supporting NRZ and PAM4 modulation
- Flexible port and FEC configurations: supporting 1M/10M/100M/1G/2.5G/5G/10G/25G/50G/100G/200G/400G Ethernet interfaces
- Support for multi-MAC GSGMII and OSGMII MAC Interfaces and Cisco Multigigabit Technology (10G-mGig)
- Breakout support: 100G to 4x25G and 400G to 4x100G
- Large, fully shared, on-die packet buffer, expandable with in package HBM used for buffering and also for LMP table expansion
- Port level flow control: Link-level flow control, Priority Flow Control (PFC) with watchdog timer for deadlock avoidance
- TSN frame preemption for 10G and 25G ports
- 1-step and 2-step IEEE 1588v2 PTP support; Class-A, B timing for all speeds and Class-C timing on ports of 25G and above; SyncE support
- On-chip, high-performance, programmable host NPU for high-bandwidth offline packet processing (for example, OAM processing, MAC learning)
- Multiple embedded processors for CPU offloading

Traffic management

- Large pool of configurable queues, supporting DiffServ and 4-level hierarchical QoS
- Congestion Drop Control: per source port, per VoQ, delay based, multicast enqueue, Tx multicast, dynamic thresholding, WRED
- Support for link-level (IEEE802.3x), PFC priority-level (802.1Qbb) flow control, and ECN Marking (Delay based or VoQ-congestion based)
- Support for system-level, QoS, and scheduling for both unicast and multicast traffic
- 4-level hierarchical queuing and scheduling: supporting, services/customer, customers ports, customer/sub-interfaces, sub-interfaces per port
- Scheduling: 8 traffic classes, Weighted Fair Queuing (WFQ), strict priority, support for min bandwidth and shaping
- TM counters and meters
- Policers and meters: 1R2C and 2R3C policers, color-aware and color-blind policers, policer-based remarking
- Egress TM: via Virtual output Queues (VoQs) and egress policing
- Support for ingress and egress traffic mirroring

Network processor

- Run-to-completion programmable network processor
- Large and shared fungible tables
- Line rate performance even with complex packet processing
- Support for IPv4 fragmentation

Load balancing and link aggregation (LAG)

- 5-tuple based hashing
- Optional internal headers or flow label-based hashing
- Two-level ECMP
- Resilient hashing

Instrumentation and telemetry

- Programmable meters used for MEF-compliant traffic policing and coloring
- Programmable counters used for flow statistics and OAM loss measurements
- Programmable counters used for port utilization, microburst detection, delay measurements, flow tracking, and congestion tracking
- Comprehensive IP Measurements (IPM) to monitor, analyze, and optimize IP traffic, reliability and efficiency.
- Traffic mirroring: (ER)SPAN on drop
- Support for sFlow and NetFlow
- Micro-burst detection
- Elephant flow monitoring
- Export counters to remote collector

SDK

- Cisco Silicon One APIs provided in both C++ and Python
- Switch abstraction interface (APIs)
- Resource monitoring
- Express boot
- Distribution-independent Linux packaging
- Robust simulation environment enables rapid feature development
- CPU packet I/O through native Linux network interfaces
- Monitoring: Show configuration state, device health monitoring, device resource monitoring, counters

Serviceability

- Single CLI for state and serviceability
- C++/gRPC automation ready
- Packet trace: On-demand flow NPU processing analyzer
- Build in traffic generator
- Debug capabilities: Device level, port level, NPU behavior, TM, ASIC table reports

Programmability framework

- Forwarding application development as well as enhancements via an IDE programming environment
- At compilation, the forwarding application generates low-level register/memory access APIs and higher-level SDK Application APIs
- Device simulator that can be used for the development of the SDK and applications running over the SDK

K100 SDK

A sample of the features that are currently available and supported by the K100 SDK are listed below. Given that the device is programmable more features are added, as needed.

Table 2. Highlighted Key Features of the K100 SDK

Bridging • L2-Interfaces • Bridge domains • MAC table • MAC learning • Flooding • MC • Storm control • STP • Per protocol L2-interface counters • ECN support for bridged packets	**IP routing** • L3 interfaces • Switched virtual interfaces • IPV4/IPV6 LPM • Extend LPM to HBM • IPv4/IPv6 host routes • VRF • Next hop routers • Router MAC check • Per protocol L3-interface counters • Strict and loose unicast RPF • RX/TX SVI
VLAN services • Per port default VLAN selection • VLAN anywhere (dense mode) • 1 / 2 VLAN tags-based service identification • Programmable Ether-types • Support for QinQ / 802.1ad • Encapsulate 1 / 2 VLAN tag per L2/L3 interface • Private VLAN	**IP multicast** • IPv4/IPv6 multicast • PIM-SM, PIM-SSM, PIM-ASM • IGMP snooping • Multicast RPF • Directly connected sub net • Efficient LAG, vPC pruning
MPLS • LDP, LDP over TE • MPLS-segment routing • BGP LU (over LDP and SR) • TE, RSVP-TE, SR-TE • RSVP TE • MPLS label encapsulation • QoS map/EXP-based ECMP • L2-VPN/L3-VPN service • Point-to-point PWE service • Pseudowire headend (PWHE)	**SRv6** • Base format • uSID (F3216 and F4816) formats • UPD, USD, USP, PSP disposition • H.Encap.red up to 2 SIDs (11 uSIDs) • WLIB support for uSID • SR-TE • L2-VPN services, including EVPN • L3-VPN services

IP tunnels and VxLAN

- IP-in-IP, GRE, GUE, VxLAN tunnel encapsulations
- IPV4/IPv6 underlay support
- Point to point tunnels
- P2MP tunnels
- EVPN control plane support
- Active-active multi-homing
- Head-end replication and multicast underlay
- L2 and L3 VNI mapping
- Extended VxLAN tunnels with recycle
- Tunnel split horizon
- Tenant routed multicast
- Support for ECMP of tunnels
- ECN propagation for IP tunnels

Mirroring

- Rx and Tx mirror sessions with multiple simultaneous mirror copies per packet
- ERSPAN v2 encapsulation
- sFlow
- Generic UDP header encapsulation
- Mirroring based of L2 / L3 logical ports
- Statistical sampling of mirrored copies
- ACL-based mirroring

Crypto

- 802.1AE MACsec
- IPsec

Security and QoS ACLs

- CTS/SGT
- NAC
- FHS
- MAC-based ACLs/Time-based, Object-group

Visibility and Telemetry

- IPM (Internet protocol measurement)
- Streaming telemetry
- Microburst detection
- Copy on drop/delay
- NetFlow
- Packet Trace
- Built-in packet generator
- CFM, BFD
- Histograms

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	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

For more information

Learn more about the Cisco Silicon One.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **https://www.cisco.com/go/offices**.

Printed in USA C78-4379514-01 10/25