

Cisco Silicon One E104

Contents

Product overview	3
Features and benefits.....	3
Technical details.....	4
Product sustainability	8
For more information	8

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 E104 is a 3.2 Tbps, high-performance, flexible, power-efficient routing/switching silicon for enterprise data centers enabling smooth speed and services upgrades to meet the changing needs of virtualized data centers and automated cloud environments.

The Cisco Silicon One E104 builds on the groundbreaking technology of the Cisco Silicon One devices before it, while increasing the lead over all other switching silicon in the market with focus on high capacity, port density, unmatched scale of services, best-in-class timing, programmability, telemetry and analytics, and security provisions with its support for MACsec and IPsec.

Enterprise data centers are unique and are distinguished by the variety of configuration and evolution paths. E104 has been conceived to enable this non-homogeneity and offers flexibility in the ways it can be used and deployed.

In an enterprise data center, systems are predominantly deployed in fixed form factor single-chip systems, and E104 serves such deployments. E104 finds its sweet in data center spine and leaf architectures and typical bandwidth configurations up to 3.2 Tbps. Such systems offer the lowest Total Cost of Ownership (TCO) for cost and power.


The E104 offers versatile deployment options with its large number and individually configured SerDes, allowing seamless transition for data centers from 1 and 10 Gbps speeds to higher speeds like 25 Gbps, 50 Gbps, and 100 Gbps at the server level, and from 10 Gbps and 40 Gbps to 50 Gbps, 100 Gbps, and 400 Gbps at the aggregation layer. Additionally, it boasts

large on-chip buffering shared among numerous per port queues. Its efficient operation in fluctuating traffic conditions is enhanced by its ability to differentiate between small and large flows and its support for flowlet load balancing.

E104 stands out with its large Layer 2 and Layer 3 tables supporting highly flexible Layer 2 and Layer 3 configurations and tunneling protocols, its support for large number of flows and flow analytics and telemetry (large ACL and NetFlow tables), and its programmability that meets the performance needs of the ever-evolving virtualized data centers 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.

E104 offers protocol support for lossless transport of RDMA over Converged Ethernet with its support of DCB protocols: (1) Priority-based Flow Control – (PFC) to prevent drops in the network and pause frame propagation per priority class; (2) Enhanced Transmission Selection – (ETS) to reserve bandwidth per priority class in network contention situation, and (3) Explicit Congestion Notification – (ECN) that provides end-to-end notification per IP flow by marking packets that experienced congestion, without dropping traffic.



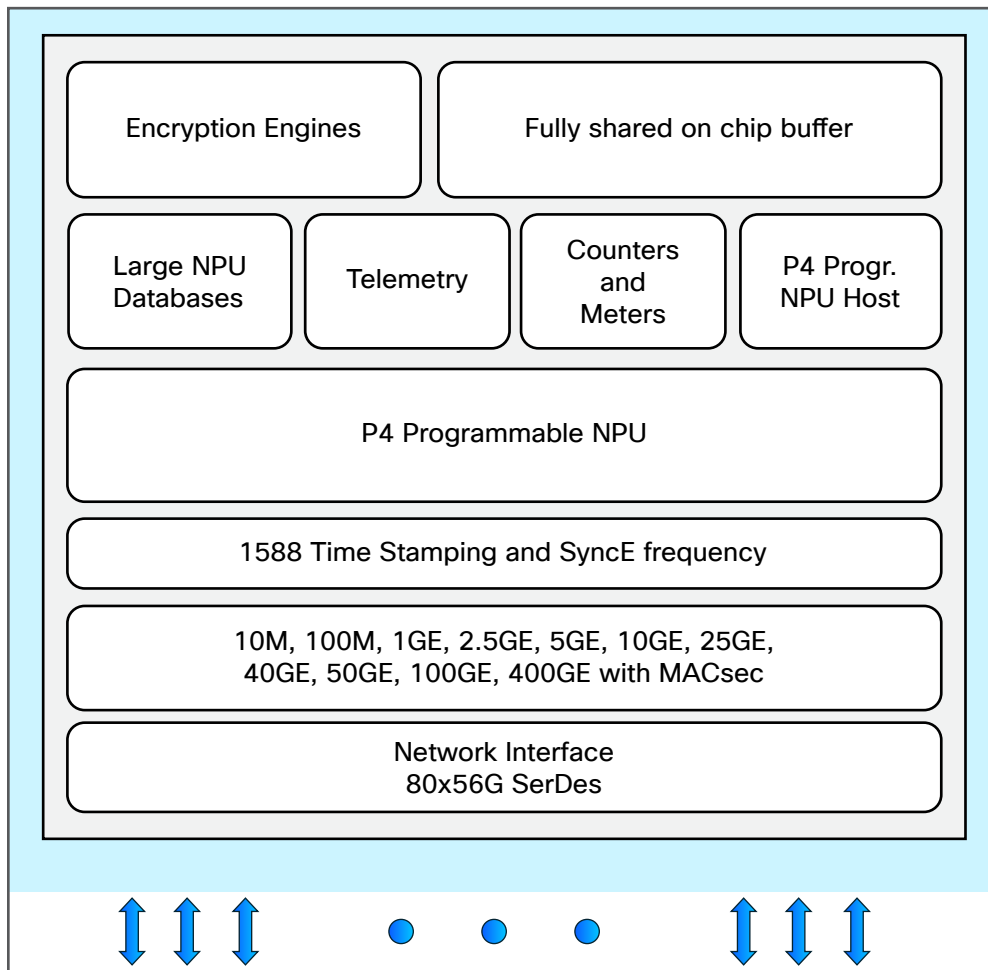
Product overview

E104 is a member of the “E” series of Cisco Silicon One processors with 80x56G I/O SerDes and 3.2 Tbps full-duplex packet processing/forwarding routing/switching processing and large on-chip buffering.

Features and 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 switching/routing silicon	7-nm, 3.2 Tbps switching silicon with a wealth of features for enterprise data center applications
Power-efficient routing/switching silicon	The power efficiency of 7 nm and the Cisco Silicon One architecture
Large and fully unified packet buffer	Fully shared on-die buffer shares by any and all device queues
Switching and routing efficiency with support of features at scale	Addresses the requirements of enterprise data center leaf (TOR) 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 encryption support for IP network deployments
RoCEv2 support	E104 addresses the data center need for lossless transport for RDMA over Converged Ethernet

Technical details



Features

- 80x56G 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
- 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
- Support for lossless transport for RDMA over Converged Ethernet
- 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 and RoCE RDMA support over Converged Ethernet

- Support for ingress and egress traffic mirroring
- Support for link-level (IEEE802.3x), PFC priority-level (802.1Qbb) flow control, and ECN marking
- Support for enhanced transmission selection (ETS)
- TM counters and meters
- Policers and meters: 1R2C and 2R3C policers, color-aware and color-blind policers, policer-based remarking

Network processor

- Run-to-completion programmable network processor
- Large and shared fungible tables
- Line rate 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 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
- 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

E104 SDK

A sample of the features that are currently available and supported by the E104 SDK are listed below. Given that the device is programmable more features are added, as needed.

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

NAT

- Static NAT
- Dynamic NAT
- PAT
- Inside/outside NAT
- TCP FIN/RST detection
- NAT application-level gateway
- Hardware flow aging
- NAT domains

Product sustainability

Information about Cisco's environmental, social, and governance (ESG) initiatives and performance is provided in Cisco's CSR and sustainability [reporting](#).

Table 1. 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