

Asterfusion CX864E-N

Open Ultra Ethernet Switch



Highlights

- 51.2T switch with 64x800G OSFP or 128x400G/512x100G in 2RU
- World's fastest switch with port-to-port latency under 560ns for 800GE ports
- Max TDP of 2200W with 64x800G SR8 ports under full traffic load
- Large on chip buffer of 200+MB for better RoCE (RDMA over Converged Ethernet) performance
- 10ns PTP and SyncE performance supports strictly synchronized AI parallel computing
- Advanced INT (in-band network telemetry) for packet delays, drops, and path traversal, enabling more advanced congestion control algorithms
- Open AsterNOS based on SONiC with best SAI support, more robust and reliable.
- Compatible with heterogeneous GPUs and SmartNICs from all leading vendors
- Line-rate programmability to support evolving UEC (Ultra Ethernet Consortium) standards

Overview

Engineered to meet the demands of AI training/inference, high-performance computing (HPC) and cloud computing/storage, the Asterfusion CX864E-N Open Ultra Ethernet Switch is a cutting-edge data center solution. Boasting a massive switching capacity of 51.2T and port density of 64x800G, it supports ultra-scale data centers and delivers end-to-end performance on par with InfiniBand networks.

Compliant with UEC standards, it features comprehensive open APIs, facilitating seamless integration for data centers and HPC clusters. As a vendor-neutral network device, it is compatible with heterogeneous GPUs and network cards from various vendors. With industry-leading low latency and high reliability, the Asterfusion CX864E-N Open Ultra Ethernet Switch stands as the premier choice for data centers in the AI era.

Product Applications

- AI Training Fabric
- AI Inference Fabric
- Data Center Fabric
- Ethernet Storage Fabric
- HPC (High Performance Computing)
- Supercomputer

Hardware

- Industry-leading MARVELL TERALYNX 10 (TL10) ultra-low latency ASIC chips with 51.2Tbps switching capacity
- Up to 28.7Bpps of L2/L3 full wire-speed forwarding
- 512 x 112G Long Reach (LR) best-in-class SerDes, supporting 64x800G, 128x400G, or 512x100G interfaces via optical modules/splitters
- Load-aware per flowlet/packet balancing to prevent congestion and ensure efficient utilization of available bandwidth
- Intel Xeon 4/8-core CPU for line-rate programmability, enabling future networking protocols through software upgrades

- 200+ MB of on-chip buffer, dynamically shared by every port, ensures superior network quality with fewer packet drops
- Active queue management avoid traffic congestion by ECN
- Support power and fan redundancy to ensure continuous business operation
- Air flows in from the front panel, and exhausts from the rear panel

Networking

- 8 x CX864E-N support 512 GPUs interconnection with 400G per port speed
- 192 X CX864E-N support 8192 GPUs interconnection with 400G per port speed
- 192 X CX864E-N support 128K ML/AI nodes interconnection with 100G per port speed
- Support for all leading collective communication libraries, including NCCL, RCCL, oneCCL, BCL, MPI, and Gloo
- Optimize collective operations (Allreduce, Allgather, All-to-all, Broadcast) with built-in multicast
- Support various network topologies, including spine-leaf, fat-tree, dragonfly, butterfly, 3D mesh/torus, hyper-x and hypercube

Operating System

- Pre-installed AsterNOS is an enterprise version of SONiC with SAI as the kernel
- Integrated rich L2/L3 network features, complete support for network virtualization, QoS policies and other services
- Advanced functional containerized & event-driven system architecture to accelerate network service development/customization
- Provides an open REST API interface that can be centrally managed and invoked by third-party applications such as cloud management platform
- Provides KLISH command line for traditional network engineers

High Availability

- 1+1 hot-swappable power supplies and 3+1 hot-swappable fans

- BFD for BGP and OSPF in 50ms
- BGP multi-homing for multiple connections to a server with automatic load balancing and failover
- Multi-chassis link aggregation group (MC-LAG) for active/active L2 multipathing
- Up to 8192-way equal-cost multi-path (ECMP) routing for load balancing and redundancy

Automated O&M

- Integrated with Python and Ansible to support automated operation and maintenance
- Supports ZTP which enables devices to automatically obtain and load deployment files

Specifications

Product Model		CX864E-N
Network interface	800GE	64
	400GE (via splitter)	128
	200GE (via splitter)	256
	100GE (via splitter)	512
Switching chip	Chip model	MARVELL TERALYNX 10
	Switching capacity	51.2Tbps
Forwarding performance	Packet forwarding rate	28.7Bpps of L2/L3 forwarding
	Cut-through latency	560ns
	Packet buffer	200+MB
Control and manage CPU (COMe)	CPU	Intel Xeon 4/8-core
	RAM	32GB SODIMM
	SSD	256GB M.2 SATA/NVME
Management interface	USB	1 x USB2.0
	Console	1 x Console RJ45
	MGMT	1 x MGMT GE RJ45

Electrical characteristics	Fan redundancy (Hot-pluggable)	3+1
	Power redundancy (Hot-pluggable)	1+1
	Input voltage	200-240V AC HVDC 200V~320V
	Maximum power consumption	2200W (with full ports of 800G-SR8)
Dimensions	Height	2U
	Dimensions (WxHxD, mm)	440x87x660
Operating conditions	Operating temperature	0 - 40°C
	Relative humidity	10% - 90% (non-condensing)

Software Features


Layer-2 features	4k VLANs
	VLAN-Based MAC learning
	LAG/LACP
	MAC dynamic learning
	MAC priority
	MAC migration
	LLDP neighbor discovery and aging
	LLDP custom TLV
	STP
	MSTP
	IGMP
	IGMP over VXLAN
	IGMP Snooping
	9K Jumbo frame

Layer-3 features	IPv6 NDP
	IPv6 ND proxy
	Static ARP
	Dynamic ARP
	Gratuitous ARP
	ARP proxy
	Routing protocols: OSPF, BGP, MP-BGP
	Static routes
	8192-path Equal Cost Multipath Routing (ECMP)
	ECMP elastic hash
	DHCPv4/v6
	VRF
Virtual network	Uplink and downlink association (monitor-link)
	VXLAN encapsulate and decapsulate
	VXLAN layer-2 forwarding
	VXLAN layer-3 gateway
	Distributed gateway
	ARP suppression
	VXLAN L2/L3 tunnel automatically establish
	Dynamic virtual network routing
	Symmetric IRB
High reliability	VM migration awareness
	Containerized system
	MC-LAG
	BFD: BGP association, OSPF association
	SLA


	Track
	ICMP-echo
	BGP multi-homing
Security & Control	ACL for IPv4/v6
	DSCP mapping, marking, classify
	IACL, EACL
	IPv4/v6 CoPP
	Policy speed limited
	TACACS+
QoS	Multi-Queue scheduling
	Flow classification
	DSCP mapping queue
	Speed limit policy
	Bandwidth guarantee
	Tail drop
	WRED
	Data center network: PFC, ECN, ETS, DCBX
	Congestion control: QCN, DCQCN, DCTCP
	RoCEv2
Visibility	In-band network telemetry (INT)
	SPAN
	ERSPAN
	Detailed telemetry on packet delays (HDC), packet drops (BDC) and path traversed by a packet (IPT)
Time Synchronization	PTP (IEEE 1588v2)
	NTP
System management	CLI, WEBUI, controller

	DevOps: Ansible, Python
	Out-of-Band management interface
	System information monitoring
	ONIE install
	ZTP
	SONiC upgrade
	Warm reboot
	Fast reboot
	SNMP
	Syslog
	REST API
	gRPC
	CRM

Hardware Panel



System Architecture



Order Information

Part Number	Description
CX864E-N	Ultra Ethernet Switch, 64x800G OSFP, hot-pluggable 1+1 PSU, hot-pluggable 3+1 FANS, front-to-power airflow, 2RU height
Optional components and spares	
OSFP-800G-2*SR4-MM	800G-SR8/2*SR4, OSFP, Dual MPO12, MMF 850nm, 100m/OM4
OSFP-800G-2*DR4-SM	800G-DR8/2*DR4, OSFP, Dual MPO12, SMF 1310nm, 500m/OS2
OSFP-800G-2*FR4-SM	800G-FR8/2*FR4, OSFP, Dual duplex LC, SMF 1310nm, 2km
OSFP-800G-2*LR4-SM	800G-LR8/2*LR4, OSFP, Dual duplex LC, SMF 1310nm, 10km
PWR-CRPS-AC-3000-AF	3000W AC Slim Power
FAN-AF-III	FAN
SVC-Basic-1Y-CX-E	Basic H/W Service and Warranty
SW-AsterNOS-Upgrade-CX-E	AsterNOS upgrade subscription

Warranty and Service Support

Asterfusion CX864E-N switches come with 2-year Basic H/W service and warranty, preloaded perpetual licensed AsterNOS and 1-year AsterNOS upgrade subscription.

To acquire more info about company, products, and solutions: www.cloudswit.ch
Sales: bd@cloudswit.ch
Business Hotline: +86-0512-62982976