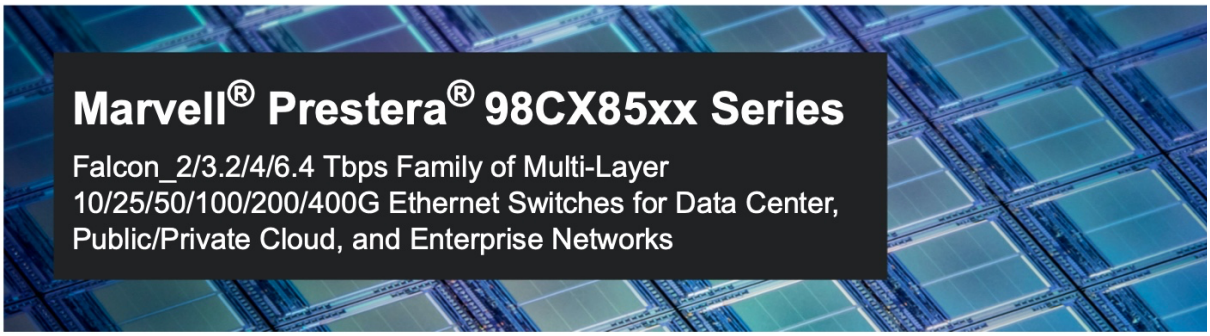

Prestera® 98CX85xx Series

Falcon_2/3.2/4/6.4 Tbps Family of Multi-Layer 10/25/50/100/200/400G Ethernet Switches for Data Center, Public/Private Cloud, and Enterprise Networks

PRODUCT OVERVIEW

The Prestera® CX Falcon 2T/3.2T/4T/6.4T Family is a new generation of highly-integrated packet processors enabling single and multi-chip solutions for public and private data center networks. The Falcon 2T/3.2T/4T/6.4T Family is ideally suited for data center transition from 10/25/100GbE towards 50/100/200/400GbE port speeds.

The Falcon 2T/3.2T/4T/6.4T Family supports advanced data center features, such as network analytics including in-band network telemetry, virtual overlay networking with programmable tunnel header encapsulation, NFV service function chaining, low latency cut-through switching, dynamic load balancing, and advanced congestion mechanisms.

The Falcon 2T/3.2T/4T/6.4T Family provides an ideal platform for SDN applications, utilizing the multiple TCAM lookups for generic match/action rules, large L2/L3 forwarding tables, and programmable elements embedded in the data path. The Falcon 2T/3.2T/4T/6.4T Family also supports MPLS and IPV6 based segment routing.

Falcon 2T/3.2T/4T/6.4T supports secure Infrastructure as a Service (IaaS), using multiple tunneling capabilities—a mandatory feature in multi-tenant virtualized cloud data centers to ensure privacy among tenants.

Falcon 2T/3.2T/4T/6.4T has sophisticated QoS features to optimize QoS for data center cloud applications, such as metering, counting, scheduling, and shaping. It also supports advanced traffic monitoring features, such as sFlow, IPFIX, and remote port analysis. The hardware OAM engines auto-generate and receive OAM traffic to monitor cloud end-to-end service connectivity and provide high-accuracy delay and packet loss measurement. The OAM engine can be used to enable higher accurate network analytics.

To facilitate efficient network analytics and troubleshooting, the Falcon 2T/3.2T/4T/6.4T Family has extensive traffic counters across all the processing and forwarding/filtering engines.

Falcon 2T/3.2T/4T/6.4T supports the Marvell® extended-bridging (eBridge) architecture, a unified architecture implementing a hardware-based virtualization of interfaces and switching domains. This architecture enables the key data center technologies, such as the standard virtual overlay encapsulations (for example, VXLAN, NVGRE).

The programmable header editor allows for supporting new emerging encapsulations, such as VXLAN-GPE, Geneve, and NSH.

Falcon 2T/3.2T/4T/6.4T includes a large shared L2 /L3-host Forwarding Database (FDB) and Longest Prefix Match (LPM) tables for high scaling of multi-tenancy data center L2/L3 services.

Data center networks support huge traffic and potential multiple uBurst. Falcon 2T/3.2T/4T/6.4T supports large shared packet buffers with dynamic thresholds, allowing efficient use of packet buffer resources.

Modern data centers must support multiple virtual network domains, where each domain may have a different data plane encapsulation. A common use case is VXLAN routing, where a packet arrives from one VXLAN tenant system and must be routed and forwarded to a different VXLAN tenant system.

For data center interconnect (DCI), an MPLS-based network (for example, VPLS, EVPN) may be used to interconnect remote data centers. In this use case, Falcon 2T/3.2T/4T/6.4T serves as the gateway between the VXLAN domain and the MPLS domain, processing packets in a single-pass.

The Falcon 2T/3.2T/4T/6.4T Family supports two-step and one-step Precision Time Protocol (PTP) per the IEEE 1588v1/v2 standard. This provides high-accuracy clock synchronization required by a data center time center application, such as financial trading and time-based updates in SDN.

The Falcon 2T/3.2T/4T/6.4T Family supports advanced flow hash calculation and dynamic load-balancing mechanisms. The hash function utilizes a programmable key and provides balanced distribution over any number of members in the group. The hash scheme supports symmetric hashing allowing bi-directional traffic in a flow to travel over the same path. Polarization is avoided by allowing each load balancing mechanism (for example, L2 ECMP, L3 ECMP, LAG) to utilize a different hash value.

The Falcon 2T/3.2T/4T/6.4T Family is feature and software compatible with previous generations of the Prestera® CX family of devices including support for open APIs via Switch Abstraction Interface (SAI) and Telemetry and Monitoring (TAM) APIs.

The Falcon 2T/3.2T/4T/6.4T Family includes devices with up to 132 high-speed, low-power, multi-rate SERDES:

- Up to 128x50G PAM4 SERDES supporting:128 ports of 1/10/25/50 GbE support FC-FEC (2112) and (2080) and RS-FEC (528,514) and (544,514)
- 64 ports of 100 GbE over 2x50G SERDES supporting RS-FEC (544,514)
- 32 ports of 100 GbE over 4x25G SERDES with RS FEC (528,514)
- 32 ports of 200 GbE over 4x50G SERDES supporting RS-FEC (544,514)
- 16 ports of 400 GbE over 8x50G SERDES supporting RS-FEC (544,514)
- 164x10G SERDES supporting 4 ports of 1/10 GbE

Falcon 2T/3.2T/4T/6.4T Device Characteristics

Device	Falcon 2T 98CX8512	3.2T 98CX8520	** 3.2T 98CX8522	4T 98CX8530	** 6.4T 98CX8540
SERDES Lanes	80x25G + 2x10G	64x50G + 2x10G	128x25G + 4x10G	80x50G + 2x10G	128x50G + 4x10G
1/10G Ports	2	2	4	2	4
1/10/25G Ports	80	64	128	80	128
50G Ports	40	64	64	80	128
100G Ports	20	32	32	40	64
200G Ports	N/A	16	N/A	20	32
400G Ports	N/A	8	N/A	10	16
PCIe Interface	x2 Gen3	x2 Gen3	x2 Gen3	x2 Gen3	x2 Gen3
PCIe Serial DMA Channels	32	32	64	32	64
Maximum BW	2 Tbps	3.2 Tbps	3.2 Tbps	4 Tbps	6.4 Tbps
Forwarding Rate	2.8 Bpps	2.8 Bpps	2.8 Bpps	2.8 Bpps	2.8 Bpps
Package	60 x 60 mm HSBSFC BGA, 1 mm pitch	60x60 mm HSBSFC BGA, 1 mm pitch	60x60 mm HSBSFC BGA, 1 mm pitch	60x60 mm HSBSFC BGA, 1 mm pitch	60x60 mm HSBSFC BGA, 1 mm pitch

Falcon 2T/3.2T/4T/6.4T Family Block Diagram

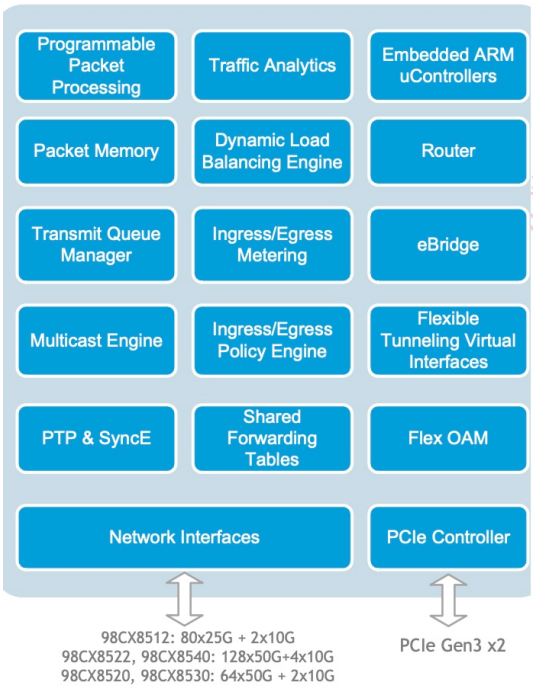

Features and Benefits

- High Integration
  - Large shared packet buffer memory
  - Large shared forwarding tables
  - Large queuing pools
  - Up to 1K remote physical ports
  - High-speed SERDES
- Multi-Rate FlexLink Ports
  - 1/10/25/50G Ethernet ports with a single SERDES
  - 50/100G Ethernet ports with dual SERDE
  - 40/100/200G Ethernet ports with Quad SERDES
  - 200/400G Ethernet ports with Octal SERDES
- SERDES Capabilities
  - Rx and Tx/Rx Training capability per SERDES lane for
  - optimizing the SERDES performance and error rate
  - Built-in test capabilities (Loopback and PRBS) per SERDES lane
  - Eye Opening Monitor per SERDES lane for tuning and diagnostics
- Layer-2 Wire-Speed Switching Engine
  - 802.1Q-compliant bridging
  - Large Forwarding Database (FDB) providing exact-match lookup for MAC entries, IGMPv3/MLDv2 IP Multicast, FCoE entries, and Router Host entries
  - Learning and forwarding based virtual ports (ePorts) and virtual bridge domains (eVLANs) for data center L2 overlay services
  - L2ECMPandLinkAggregationGroups(LAGs)
- Layer-3 Wire-Speed Routing Engine
  - Longest Prefix Match (LPM) for IPv4/6 and IP Multicas
  - Policy-based routingVRF, VRF-Lite, BGP/MPLS IP VPNs
  - Multicast routing supporting PIM-SM/DM and PIM-bidirectional routing protocols
  - ECMP routing for load-balancing traffic
  - NetworkAddressTranslation(NAT44,66)
  - IPV6 Segment routing
- Data Center
  - Low-latency Cut-through switching
  - Virtual Overlay Networking – NVGRE, VXLAN-GPE, GENEVE, SPB, TRILL, GRE
  - Server Virtualization – IEEE 802.1Qbg EVB, 802.1BR Port Extender
    NFVServiceFunctionChaining– NetworkServiceHeader (NSH)
    Dynamic load balancing for optimizing flow distribution
  - FCoE forwarding utilizing the L2 FDB
  - Data Center Bridging (DCB) standards – Priority-based Flow Control (PFC), Congestion Notification (CN), and Enhanced Transmission Selection (ETS)
  - Explicit Congestion Notification (ECN) marking, including support for DC-TCP and Phantom Queues
- Network Analytics
  - Mirror packets on Drop or Congestion
  - Max Queue fill level
  - uBurst Duration
  - INT and in-situ OAM
  - Latency measurements and statistics
  - Approximate Fair Dropping (AFD) and Dynamic Packet Prioritization (DPP)

- Data Center Interconnect (DCI)
  - Universal gateway one-pass bridging/routing to support interworking between virtual network domains
  - MPLS-based services, such as VPWS, VPLS, EVPN
  - EVPN/VPLS/MPLS-over-GRE services
  - Ethernet-over-GRE services
- TCAM and Exact Match Classifications
  - 12 ingress TCAM lookups and 4 Egress TCAM lookups
  - Configurable TCAM keys size ranging from 10B to 70B, and fully configurable key content based on packet metadata or first 160B of the packet header
  - Lookup based on exact match Hash tables
  - SAFE,StorageAwareFlowEngineforNVMe-oFabricflow
  - Management
- Quality of Service (QoS)
  - 1KQoSProfiletoflexiblyassignapackettrafficclass,drop precedence, and packet CoS marking
  - 3levelsofschedulingsupportforStrictPrioritydShaped Deficit Weighted Round-Robin (SDWRR) per level
  - Ingress/Egress Policers with Single/Two Rate 3-Color Marking and bandwidth sharing
  - FASTER, Enabling high port count systems with End-to-End QoS
  - Virtual Output Queues (VoQ) for head of line blocking prevention in a chassis fabric
  - Centralized traffic management with Remote Physical Ports, where the device performs queuing and shaping on behalf of fan-out line card devices
- Scalability
  - Cascade up to 1K Prestera®-CX/DX/EX devices
  - eDSA/DSA-compatible with Prestera®-CX/DX/EX devices
  - Large number of ECMP/LAG groups and group members
- Network Management and Statistics
  - Ingress and Egress port MIB counters
  - sFlow (RFC 3176)
  - IPFIX (RFC 3917)
  - Large counter pool with flexible assignment
  - uBurst and Elephant flows detection
  - Dynamic Load Balancing support for optimizing link utilization and goodput
- Synchronization and Precision Time Protocol (PTP)
  - High accuracy one-step and two-step PTP compliant with IEEE 1588v1/v2
  - SyncE compliant
- Hardware-based Operations, Administration, and Maintenance (OAM)
  - Support for IEEE 802.1ag, ITU-T Y.1731 continuity check, delay measurement, and loss measurement
  - MPLS OAM, including MPLS Bidirectional Forwarding Detection (BFD) and ITU-T G.8113.1
- Miscellaneous
  - Adaptive Voltage Scaling (AVS) support
  - I/O voltage levels – 1.8V or 3.3V
  - Power-save/down/scaling capabilities, along with the Marvell Alaska® low-power PHY devices
  - LED interface for port activity status
- Development Kit
  - Prestera Software Suite (CPSS) driver support
  - CompatiblewithPrestera®-CX/DX/EXAPIsandconcepts