Emulex Engine (XE)100 Series (Skyhawk) of Ethernet Controllers

High Performance Networking for Enterprise Virtualization and the Cloud

XE100 series of Ethernet Controllers

Overview

The Emulex Engine™ (XE)100 series (Skyhawk™) of Ethernet Controllers provides a single chip, which can support single-port 40Gb Ethernet (40GbE) or up to four ports of 10GbE/1GbE capability delivering multiple benefits for the enterprise cloud, including:

- Increasing data center IT agility and scalability through deployment of a secure multi-tenant cloud
- Driving scalability and flexibility in converged infrastructures
- Optimizing server hardware utilization by scaling high density virtualization

The XE100 series of 10GbE and 40GbE controllers are designed for the high bandwidth and flexibility demands of tier 1 enterprise applications with storage protocol (Fibre Channel over Ethernet (FCoE) and iSCSI) offloads, more scalable virtualization, support for RDMA over Converged Ethernet (RoCE) fabric, enhanced Single Root I/O Virtualization (SR-IOV), Network Interface Card (NIC) port partitioning, and cloud optimization using overlay network technology. The controllers also features innovative, purpose-built hardware architecture to carry network and storage traffic over a single full-duplex 40GbE port or four 10GbE/1GbE ports.

Extensible protocol offload architecture

The Emulex XE100 series of Ethernet Controllers utilizes an innovative combination of high performance hardware-based data path elements for protocol data movement and industry-standard embedded processors for protocol control path. Hardware-based isolation between offloaded protocols ensures precise control over bandwidth allocation and sharing across network and storage traffic. Protocol offload features, bandwidth allocation and management features are exposed through industry standard interfaces.

Emulex Virtual Network Exceleration™ (VNeX) overlay network offloads for multi-tenant cloud networking

Scaling existing technologies for private or public multi-tenant infrastructures requires networking solutions that can enable virtual machine (VM)-to-VM communication and virtual workload migration across Layer 2 and Layer 3 boundaries without impacting connectivity or performance.

At the same time, these solutions need to ensure isolation and security for thousands or millions of tenant networks. However, with existing technology, the available 4094 virtual Local Area Network (VLAN) IDs are insufficient to isolate and secure each tenant in a private or hybrid cloud data center.

Target applications

- Single chip solution for LAN on Motherboard (LOM)
- Storage devices support iSCSI and FCoE initiator, target modes, NIC and RoCE in embedded environments

Key benefits

- Maximizes server hardware return on investment (ROI) with high VM density
- Simplifies deployment of secure, scalable multi-tenant cloud infrastructures
- Minimizes TCO through deployment of heterogeneous workloads on Converged Infrastructure
- Accelerates applications and storage performance
- Provides the bandwidth needed for slot constrained server platforms
- Reduces complexity through the deployment of a common network platform
- Reduces management, infrastructure and energy costs

Key features

- Integrated 40GBASE-KR4, 10GBASE-KR and 1000BASE-KX PHY interfaces for LOM and backplane applications
- Superior network convergence—storage and network traffic over a common 10GbE or 40GbE infrastructure
- SR-IOV
- Data acceleration with RoCE support
- Powerful hardware offloads for:
  - Overlay networks (NVGRE & VXLAN)
  - Storage protocols: iSCSI and FCoE
  - Stateless TCP
- Greater bandwidth with PCI Express (PCIe) 3.0
- VMware vSphere NetQueue support
- Microsoft Windows Server VMQ & Dynamic VMQ support

Emulex Engine (XE)100 Series (Skyhawk) of Ethernet Controllers

Virtual Extensible Local Area Network (VXLAN), supported by VMware, and Network Virtualization using Generic Routing Encapsulation (NVGRE), supported by Microsoft, are next generation overlay networking solutions that address these requirements. These solutions are a frame-in-frame data packet encapsulation scheme enabling the creation of virtualized Layer 2 subnets that can span physical L3 IP networks. Traffic from each VM is tunneled to a specific virtual network; the packets are then routed transparently over the existing physical infrastructure.

Emulex VNeX offload technology powered by a multi-core ASIC engine accelerates the performance of network virtualization by preserving legacy stateless TCP offloads and scaling methods on encapsulated packets, providing full native network performance in a virtual network environment.

RDMA support

The XE100 Controllers leverage RoCE enabling server to server data movement directly between application memory without any CPU involvement providing high throughput and data acceleration on a standard Ethernet fabric without the need for any specialized infrastructure or management.

Flexible workload storage connectivity with FCoE and iSCSI offloads

XE100 Controllers support FCoE hardware-based offload using the same enterprise-class Emulex drivers that work with Emulex LightPulse® Fibre Channel Host Bus Adapters (HBAs). XE100 controllers also support iSCSI hardware-based offload, providing performance that is superior to iSCSI solutions based on software initiators and basic NICs. Finally, XE100 series of Ethernet Controllers also have the ability to support simultaneous iSCSI and FCoE hardware offloads concurrently with NIC traffic on the same port (i.e. concurrent storage).

Optimized host virtualization density with SR-IOV support

SR-IOV optimizes I/O for VMs, enabling higher host server virtualization ratios to deliver maximum server ROI. SR-IOV provides a more cost-effective solution than multiple, physical ports.

SR-IOV enables multiple VMs to directly access the XE100 series of Ethernet Controllers I/O resources, thus allowing VM networking I/O to bypass the host and take a path directly between the VM and the controller, eliminating redundant I/O processing in the hypervisor. This, in turn, allows higher I/O performance, lower CPU utilization and reduced latency as compared to the alternative of software-emulated NIC devices that are implemented in the hypervisor.

Optimized bandwidth allocation with Emulex Universal Multi-Channel™ port partitioning

Emulex Universal Multi-Channel (UMC) is ideal for virtualized server environments because bandwidth allocation can be optimized to support VM migration, management and I/O intensive applications. UMC allows multiple PCIe physical functions (PFs) to be created on each controller port. Each XE100 controller can be configured with up to sixteen functions.

The key benefits of deploying Emulex UMC technology include:

Lower total cost of ownership (TCO)

- Consolidates multiple 1GbE adapters, associated cables and switch ports
- Higher VM workload bandwidth allocation to drive higher VM density on host servers
- Lower per-Gb bandwidth cost compared to deploying multiple 1GbE adapters

Optimized I/O utilization

- Granular bandwidth provisioning minimizes wasted idle bandwidth and waste of dedicated 1GbE adapters
- Enables Service Level Agreement (SLA)-based provisioning and deployment

Simplified deployment

- UMC is not dependent on specialized operating system (OS) support
- Works with any 10GbE or 40GbE switch

UMC is ideal for virtualized server environments because bandwidth allocation can be optimized to support I/O intensive applications, VM migration, and management functions.

Simplified management OneCommand® Manager application

The OneCommand Manager application provides centralized management of products based on the XE100 series of Ethernet Controllers throughout the data center from a single management console. The OneCommand Manager application provides a graphical user interface and a scriptable command line user interface. OneCommand Manager for VMware is fully integrated with VMware vCenter to simplify management for virtual server deployments.

Fourth generation platform delivers enterprise-class reliability and performance

Leveraging generations of advanced, field-proven controller and adapter technology, XE100 Controllers meet the robust interoperability and reliability requirements of enterprise and scale-out data centers.

Emulex Engine (XE)100 Series (Skyhawk) of Ethernet Controllers

Ethernet standards

- Single IEEE 802.3ba 40GBASE Ethernet port (40GBASE-KR4/40GBASE-SR4/40GBASE-CR4)
- Up to quad IEEE 802.3-2008 10GBASE Ethernet ports (10GBASE-SR/10GBASE-LR/10GBASE-CR)
- Up to quad 1GBaseX/SGMII auto negotiation
- IEEE 802.3an-2006 (10GBASE-T) with external PHY
- IEEE 802.1Q VLAN
- IEEE 802.3x Flow control with pause frames
- IEEE 802.1Qbg Edge Virtual Bridging
- IEEE 802.1Qaz Enhanced Transmission Selection (ETS); Data Center Bridging Capability Exchange (DCBX)
- IEEE 802.1Qbb Priority Flow Control (PFC)
- IEEE 802.3ad Link Aggregation/LACP
- IEEE 802.1AB Link Layer Discovery Protocol (LLDP)

Ethernet Network Interface (Layer 2 NIC) and TCP/IP

- NDIS 6.0, 6.2, 6.3-compliant Ethernet functionality
- IPv4/IPv6 TCP, UDP checksum offload
- IPv4/IPv6 Receive Side Scaling (RSS)
- IPv4/IPv6 Large Receive Offload (LRO)
- IPv4/IPv6 Large Send Offload (LSO)
- Dynamic VMQ (Windows Server 2012 Hyper-V) and NetQueue (VMware vSphere)
- Programmable MAC and VLAN addresses
- 128 MAC/VLAN addresses per port
- Support for hash-based multicast MAC address filters
- Support for hash-based Broadcast frame filters per port
- VLAN offloads (insertion and extraction)
- Jumbo frame support up to 9000 Bytes

I/O virtualization

- Stateless L2, L3 and L4 offloads for frame-in-frame encapsulation (VXLAN, NVGRE)
- PCI-SIG Address Translation Service (ATS) v1.0
- Support for up to 512 hardware queues
- Virtual Switch Port Mirroring for diagnostic purposes
- Virtual Ethernet Bridging (VEB)
- Virtual Ethernet Port Aggregator (VEPA)
- OneConnect Universal Multi-Channel, support for up to 16 PCIe PFs per ASIC
- NIC SR-IOV
  - up to 63 Virtual Functions (VFs) per port
- QoS for controlling and monitoring bandwidth assigned to and used by virtual entities
- Configurable control of network bandwidth by physical port, queue, or protocol
- Traffic shaping and QoS across each VF and PF

Fibre Channel over Ethernet (FCoE) offload

- Hardware offload for FCoE protocol
- ANSI T11 FC-BB-5 compliant
- Programmable World Wide Name (WWN)
- Support for FIP and FCoE Ether Types
- Concurrent Logins (RPI): up to 8K per controller (FCoE adapter-only mode)
- Open Exchanges (XRI): up to 4K per controller (FCoE adapter-only mode)
- FCoE initiator and target modes
- Support up to 255 N_Port ID Virtualization (NPIV) interfaces per port
- Concurrent FCoE and iSCSI support on each port

Internet Small Computer System Interface (iSCSI) offload

- Hardware offload for iSCSI protocol
- Header and data digest support
- Up to 4K outstanding commands (iSCSI adapter-only mode)
- Up to 512 offloaded iSCSI connections (iSCSI adapter-only mode)
- FCoE initiator and target modes
- Up to 3072 outstanding commands (target models, iSCSI adapter-only mode)
- Up to 1792 offloaded iSCSI connections (target models, iSCSI adapter-only mode)
- iSCSI initiator and concurrent initiator/target modes
- Support for multipath I/O
- Operating system-agnostic INT13-based iSCSI boot and iSCSI crash dump support
- RFC 4171 Internet Storage Name Service (iSNS)
- Support for both IPv4 and IPv6 connections
- MTU packet size support up to 8342 bytes
- Concurrent iSCSI and FCoE support on each port

Converged Enhanced Ethernet (CEE) and Data Center Bridging (DCB)

- IEEE 802.1Qbb PFC
- IEEE 802.1Qaz Enhanced Transmission Selection (ETS)
- IEEE 802.1Qaz DCBX
- Absolute per-priority rate control option/ configuration

Remote Direct Memory Access (RDMA)

- Direct data placement in application buffers without CPU intervention
- Supports IBTA RoCE specifications
- Linux Open Fabrics Enterprise Distribution (OFED) support
- Low latency queues for small packet sends and receives
- Windows Server SMB Direct (SMB over RDMA)

PCI Express (PCIe) interface

- PCIe 3.0 x8 (8, 5.0, and 2.5 GT/s per lane) compliant interface:
  - Up to 64 Gbps full duplex bandwidth
  - Configurable width and speed to optimize power versus bandwidth
- Support for up to 16 PCIe PFs
- Support for x1, x2, x4, and x8 links widths
- NIC SR-IOV
  - up to 63 VFs per port
- Message Signal Interrupts Extended (MSI-X)
- Advanced Error Reporting (AER)
- Completion Timeout (CTO)
- Function Level Reset (FLR)
- Alternative Routing ID Interpretation (ARI)

Emulex OS driver support

- Windows
- Red Hat Enterprise Linux
- SUSE® Linux Enterprise Server
- Oracle Linux
- VMware vSphere
- CentOS
- Debian (NIC mode only, storage not supported)
- Ubuntu (NIC mode only, storage not supported)
- FreeBSD (NIC mode only, storage not supported)

Emulex Engine (XE)100 Series (Skyhawk) of Ethernet Controllers

System level design

- vCenter management plugin support
- Role-based management, integrated with Active Directory and LDAP
- Flexible personality definition for networking and storage protocols
- Multi-channel configuration and bandwidth control
- UEFI and x86 remote boot support including PXE v2.1, UEFI 2.3.1, iSCSI and FCoE
- MAC statistics gathering (SNMP, Ethernet MIB, MIB2, RMON, RMON2)
- Offline and online firmware updates
- IEEE 1149-1 compliant JTAG support
- Network Controller Sideband Interface (NC-SI) support
- ACPI v2.0 compliant power management
- IPMI pass-through allows baseboard management controller (BMC) access to the network
- Integrated Thermal Sensor works with management utilities

Package type

- 563-ball Flip Chip Ball Grid Array (FCBGA)
- 27 mm x 27 mm package size
- RoHS (lead free) compliant, including China RoHS

Electrical specifications

- Power supply 1.8V, 1.5V, 1.2V, 0.9V

Ordering information (see port configuration table below)

XE104-P1
- XE104 40Gb CNA Controller (NIC, standard RoCE, hardware offloaded iSCSI and FCoE)

XE104-N-P1
- XE104 40Gb Ethernet Controller (NIC, standard RoCE protocols only)

XE102-P1
- XE102 20Gb CNA Controller (NIC, standard RoCE, hardware offloaded iSCSI and FCoE)

XE102-N-P1
- XE102 20Gb Ethernet Controller (NIC, standard RoCE protocols only)

Port configurations

Port configuration	XE102	XE104
4x 1Gbps	Yes	Yes
1x 10Gbps	Yes	Yes
2x 10Gbps	Yes	Yes
2x 10Gbps + 2x 1Gbps	Yes	Yes
4x 10Gbps	No	Yes
1x 40Gbps	No	Yes

This data sheet represents the underlying functionality of the Emulex I/O Controller. It does not guarantee feature enablement through firmware and drivers. Please contact your Emulex sales representative for details on your particular feature requirements.

EMULEX®

World Headquarters 3333 Susan Street, Costa Mesa, CA 92626 +1 714 662 5600
Bangalore, India +91 80 40156789 | **Beijing, China** +86 10 84400221
Dublin, Ireland +35 3 (0) 1 652 1700 | **Munich, Germany** +49 (0) 89 97007 177
Paris, France +33 (0) 158 580 022 | **Tokyo, Japan** +81 3 5325 3261 | **Singapore** +65 6866 3768
Wokingham, United Kingdom +44 (0) 118 977 2929 | **Brazil** +55 11 3443 7735

www.emulex.com