

Voltaire® 10 Gigabit Ethernet Switching for Scale-Out Data Centers



Scaling-Out Data Centers Using 10 Gigabit Ethernet: A New Paradigm

Data center architectures are constantly evolving, with major changes now emerging in several key areas:

- **Data center consolidation:** Building larger shared (private or public) data centers instead of many smaller ones
- **Focus on application and business services:** Moving away from manual IT processes
- **Virtualization anywhere:** Servers, I/O, storage, networks, and applications are virtualized and decoupled from physical hardware
- **Fabric convergence:** Networking, storage, and inter-process communication (IPC) from multiple applications traveling over the same physical wire

These trends have significant impact on the fabric architecture of the data center. Fabrics must now support larger-scale Layer 2 (L2) networks since server virtualization and mobility, new storage protocols, and low latency messaging must reside on the same L2 domain. Today's data center infrastructure also has more complex management requirements and must address virtualization from the ground up.

As data centers become denser and more virtualized, the loads over the underlying interconnect fabric increases significantly. Factors that add to the increased IO load include:

- Extensive use of shared external file (NAS) or block (SAN) storage
- Multi-socket, high density server platforms on which virtualization is applied
- More communication between different physical segments/racks due to server and application mobility

Limitations of Traditional Networks

There are many dimensions to consider when discussing data center interconnect scalability: from interconnect bandwidth and latency to price and power consumption.

Legacy data center networks were built to connect application silos, with each silo containing a set of servers or a rack that ran a specific application. The application had little communication with the external world since most of its intensive transactional, messaging, and data/storage traffic ran within the rack and only a fraction of that traffic was delivered to consumers outside of each rack. For this reason, there was a very small amount of uplink bandwidth from the access (rack) layer to the aggregation and core layers. Also, the aggregation and core layer switches often developed L3/4 routing functionality based on a deep buffered store-and-forward chipset architecture as they connected separate application subnets. This resulted in messaging latencies as high as 50 microseconds for each switch, which adds up to 300 microseconds of application latency over the entire fabric. This level of performance simply does not meet the bandwidth and end-to-end fabric latency requirements of the next generation virtualized data center.

Voltaire 10GbE Unique Value Proposition

Scale-out switching

- High capacity, low latency
- Linear scale-out topology
- Enabling application mobility

Data center optimized

- Central fabric management
- FCoE, CEE support

Application IO acceleration

- Multicast messaging offload
- MPI messaging acceleration
- RDMA Ethernet support




Figure 1: Traditional Data Center Topology




Figure 2: Scale-out Data Center Fabric Topology

Voltaire® 10 Gigabit Ethernet Switching for Scale-Out Data Centers

As traditional three-tier architectures of switches scale in port count, the price tag grows substantially. For every port added in the access layer, cost, power and space requirements of additional uplink aggregation and core ports increase in order to maintain the topology. These ports, which are engineered as router and service ports, have substantial price and power overhead and are redundant for the desired functionality, as shown in the following table.

Platform	Power/Port	Price/Port	Latency	Max wire speed ports in 1U
Access	7 - 10W	\$400-900	> 5 us	48 in 1U
Aggregation/Core	35 - 100W	\$2,600-\$5,000	> 40 us	6 in 1U
Difference	~ 5-10X	~ 5-10X	~ 5-10X	> 8x

Voltaire Scale-out Data Center Fabric Architecture

Voltaire’s new scale-out fabric architecture based on Converged Enhanced Ethernet (CEE), delivers a scalable and efficient fabric for servers and storage in the data center.

Efficiency: Cost, Power, Latency & Density

Voltaire’s 10GbE fabric solution comprises very high-density switches that can switch many 10GbE ports at wire speeds and without oversubscription. In addition, the switch architecture provides several unique traffic management capabilities that reduce the latency to less than one microsecond, while at the same time guaranteeing application performance and lossless or lossy behavior for the various traffic classes.

The scale-out design of the switch couples less complex and less expensive hardware elements with a powerful scale-out software stack, enabling lower costs, lower power consumption, and higher switching density than other aggregation switches available today. Furthermore, multiple switches can be meshed together to form enormous topologies without losing these advantages.

Linear Scalability

Voltaire’s scale-out solutions are not just wired in a mesh topology—they also allow the traffic to flow over more than a single path between every two switches. Traditional aggregation networks use a Spanning Tree protocol, which deactivates most of the links and does not allow such a high degree of fabric path utilization. The immediate benefit of the Voltaire fabric is that there is more available bandwidth within the fabric to service the servers, eliminating the need to add additional core switches when the traffic load increases.

In addition, Voltaire’s fabric management software—Voltaire Unified Fabric Manager™ (UFM)—allows multiple switches to appear as one very large Ethernet switch to any external device. This transparent architecture allows simple integration in heterogeneous switching environments while delivering on the core benefits of scale-out architectures.

Compared to traditional Ethernet aggregation switches arranged in a hierarchical design, the Voltaire-based solution combined with any off-the-shelf top of rack (TOR) switch demonstrates a clear advantage in both cost and efficiency in larger scale configurations. The Voltaire solution allows users to build larger consolidated server and storage farms at much lower costs, while maintaining the highest performance levels.





Figure 3: Voltaire’s scalability (green) vs. traditional networking solutions (red)

For example, when comparing the price and power consumption of a 1000 node fabric solution using the Voltaire Vantage 8500 and the Cisco 7000/5000, the Voltaire solution costs half the price, consumes one-third of the power and has ten times better latency with a similar server interconnect bandwidth.

Switch	Cost	Power Consumption	Latency
Cisco Nexus 7000 core	\$2,375,000	36,900 Watts	16 μs
Voltaire Vantage 8500 core	\$1,100,000	10,000 Watts	1.5 μs

1000 server cluster cost and power comparison (with 1:3 oversubscription)

Voltaire® 10 Gigabit Ethernet Switching for Scale-Out Data Centers

Introducing the Voltaire Vantage™ 8500

The Voltaire Vantage 8500 switch is a completely modular, high-performance, high density Layer-2 core switch, optimized for enterprise data center and cloud computing environments. It features the industry's most power-efficient and lowest latency capabilities on Ethernet, enabling new levels of efficiency, scalability and real-time application performance, while at the same time consolidating multiple/redundant network tiers and significantly reducing infrastructure expenses.

The Voltaire Vantage 8500 is the industry's largest non-blocking 10 Gigabit Ethernet switch. With up to 288 ports of 10GbE connectivity, the Voltaire Vantage 8500 switch features an impressive 11.52 Tb/s non-blocking backplane, allowing applications to perform at maximum bandwidth and efficiency. As a focused, L2 core switch, it seamlessly integrates with existing switch infrastructures, and enables application performance at the lowest bandwidth-to-power ratio— with only 600-1200 nanoseconds of port-to-port latency and power consumption as low as 10 watts per port.

In addition, the Voltaire Vantage 8500 implements the IEEE data center bridging lossless mode of traffic optimized for Fibre Channel over Ethernet (FCoE) and RDMA traffic types, which require a reliable and lossless transport layer.

Additional CEE features are implemented by the switch across the fabric. These include Ethernet multi-pathing, I/O virtualization, and fabric-wide congestion management and QoS. The switch is interoperable with standards-based switches, 1GbE and 10GbE network adapter cards, and third-party management and virtualization solutions enabling a best-of-breed, cost effective data center network.

Using Voltaire's Vantage technology, up to 12 Voltaire Vantage 8500 switches can be clustered together to form a 3400 port non-blocking 10GBE switching fabric—an important capability for virtualized data centers.

Virtualized from the Ground Up

In response to the proliferation of server virtualization in data centers, the Vantage 8500 was designed from the ground up to support virtual machine switching environments. It supports virtual ports and applies policies in the granularity of a virtual port in service of each virtual server machine. The Vantage 8500 also allows users to seamlessly manage roaming virtual machine ports, while maintaining their unique policies as they relocate across the fabric and monitoring them as if they were physical ports. The Vantage 8500's powerful operating system, VT-OS, also allows multiple simultaneous sessions and granular access control of virtual or physical switch elements.

VT-OS comprises a full Ethernet Layer-2 protocol stack and advanced chassis management capabilities to automate regular administrative tasks while offering a user-friendly CLI, GUI, and SNMP interface.

Complementing VT-OS is Voltaire's Unified Fabric Manager™ (UFM™) software management platform, which enables fabric-wide, application-centric cluster management of the Voltaire Vantage 8500 cluster. UFM software also provides management for third-party switches in conjunction with other data center management applications.



Voltaire Vantage™ 8500

Highest Capacity Chassis

- Voltaire 11.5Tb/s chassis
- Highest slot capacity of 960 Gb/s
- Total 10GbE-port options: 288 (full rate)
- Signal quality driving 12m DAC (Twinax)
- Flat Layer 2 Scalability to 3,400 nodes (full rate)

Key Features

- Full L2 stack, L3 ready, L2-4 ACLs
- Enhanced Ethernet (CEE) feature set
- Ethernet L2 Multi-pathing
- Virtual I/O port configuration per server VM

Voltaire® 10 Gigabit Ethernet Switching for Scale-Out Data Centers

TECHNICAL SPECIFICATIONS

Voltaire Vantage 8500

- 12-slot, 19" rack mountable chassis, 15U high
- Redundant power supplies (sPSUs) and fan units (sFUs)
- Bisectonal bandwidth: 11.52 Tb/s
- Port-to-port latency: 600-1200 nanoseconds
- Up to 12 10GbE line boards supported per chassis, with a total of 288 ports
- Each slot may accommodate one 24-port 10GbE Line Board
- Line board indicators: Link up/down and Traffic activity LEDs per port, power and info LEDs
- Up to 4 fabric boards supported per chassis
- Fabric board indicators: Link up/down and traffic activity LEDs per line board link port, power, info and hot-swappable LEDs

Management

- Supported management:
 - Web based management
 - CLI
 - SNMP v2/v3
 - Telnet
 - SSH v2
 - SSL
 - RADIUS
 - AAA
 - Syslog
 - RMON, XRMON
 - File download via HTTP, HTTPS, SCP, FTP, SFTP
 - NTP
 - Ethernet like MIB
 - SNMP MIB II
 - RBAC
- Indicators: Power, chassis manager active, info and hot-swap LEDs
- Additional rear control panel
- Connectors: EIA/TIA-232 Console DB-9, 10/100 Ethernet RJ-45
- Indicators: chassis management activity (2 LEDs), Fabric Boards (4 LEDs) and temperature LED
- Management/device reset button

Layer 2

- 4K VLANs
- Jumbo frames (9K)
- IGMP snooping
- Port mirroring
- QoS (e.g. shaping, policing, and marking)
- Up to 16K Unicast MAC addresses

Policy and Layer 3

- L2/3/4 ACLs
- Actions: forward/drop, QoS assignment, counting/metering/policing and mirroring

IEEE Compliance

- 802.1
- 802.1Q VLAN
- 802.1P CoS
- 802.1d STP
- 802.1w Rapid STP
- 802.1s MSTP
- 802.1ab LLDP
- 802.3
- 802.3ae
- 802.3ad Link aggregation/LACP (16 ports/channel)
- 802.3x Flow Control

DCB (Datacenter Bridging) Compliance

- 802.1Qbb PFC
- 802.1Qau QCN
- 802.1Qaz ETS
- VEB (Virtual Ethernet Bridge)

Scale-out Layer 2 Advanced Features

- Virtual I/O Port Object
- Layer 2 Multi-pathing

Cooling

- Airflow: front-to-back
- Heat dissipation: 14336 BTU/hour max. for full configuration

Physical/ Electrical

- 288 10GbE ports SFP+ cages:
 - 10GBASE-SR (30m to 300m)
 - 10GBASE-LR (10Km)
 - 10GBASE-LRM (220m with FDDI grade)
- SFP+ Twinax cable (copper)
- Power supply redundancy
- Max. power consumption: 4.3 KW

- Max. power consumption measured: 3.3 KW
- Typical power consumption: 10Watt/Port

Physical Characteristics

- 19-inch rack-mountable
- Dimensions (H x W x D): 25.6 in. (660 mm) x 17.5 in. (444 mm) x 22.75 in. (578 mm)
- Optional front or rear rack mounting
- Weight: 128 to 163 lb (58 to 74 Kg), depending on configuration and packaging

Environmental

- Operating
 - Ambient temperature: 32° to 113° F (0° to 45° C)
 - Humidity: 15 to 80%, non-condensing
 - Altitude: 0 to 9843 ft (3000m)
- Storage
 - Temperature: -13° to 158° F (-25° to 70° C)
 - Humidity: 5 to 90 non-condensing
 - Altitude: 0 to 15,000 ft (4570m)

Certifications

- Safety (Voltaire Typical)
 - UL60950
 - CB IEC60950
 - CSA-C22.2 No.60950-00
- EMC (Voltaire Typical)
 - 47CFR FCC part 15
 - EN55022:98/EN55024:98/EN61000-3-2:00/EN61000-3-3:95
 - VCCI

Ordering Information

Part Number	Product Description
Vantage 8500 Components	
VLT-30023	8500 SWITCH BASIC BUNDLE (CHASSIS + FANS + X2 PS + X1 SMB)
VLT-30022	8500 SMB MANAGEMENT MODULE
VLT-30012	8500 SFB FABRIC MODULE
VLT-30024	8500 SLB 24 PORTS LINE MODULE
SFP+ copper - DAC	
CBL-00239	ETH CABLE SFP+ (10GE) 1M 30AWG
CBL-00240	ETH CABLE SFP+ (10GE) 3M 28AWG
CBL-00242	ETH CABLE SFP+ (10GE) 5M 26AWG
CBL-00244	ETH CABLE SFP+ (10GE) 7M 24AWG
CBL-00245	ETH CABLE SFP+ (10GE) 10M 24AWG

Part Number	Product Description
SFP+ Optic Transceivers	
OPT-00006	ETH SFP+ SR Transceiver (up to 300m)
OPT-00005	ETH SFP+ LR Transceiver (up to 10km)
OPT-00010	ETH SFP+ LRM Transceiver (up to 220m)



Contact Voltaire to Learn More

1.800.865.8247
info@voltaire.com
www.voltaire.com

©2009 Voltaire Inc. All rights reserved. Voltaire and the Voltaire logo are registered trademarks of Voltaire Inc. Grid Director is a trademark of Voltaire Inc. Other company, product, or service names are the property of their respective owners. Information in this document is subject to change without notice. Voltaire assumes no responsibility for any errors that appear in this document. All statements regarding Voltaire's future direction and intent are subject to change or withdrawal without notice.