

Solution Brief: Addressing Cloud Computing and Virtualized Data Center Challenges with Voltaire



Key Challenges in Virtualized and Cloud Environments

I/O is the new bottleneck

- CPU capacity is growing exponentially, memory performance is now quadrupled with Intel Xeon 5500, and dozens of virtual machines (VMs) are squeezed into the same server while the I/O capacity remains the same.
- Software-based I/O virtualization and switching adds significant overhead, and degrades application performance.
- Networks are heavily used for storage, application messaging, and VM migration, so network oversubscription is no longer desired.

Fabrics are Expensive, Inefficient, and Don't Scale

- Each server is connected to multiple networks and single purpose fabrics, requiring multiple adapters and switches.
- 10GbE networks are extremely expensive and power consuming when scaling to thousands or even just hundreds of nodes.

No Fabric Isolation or SLA

- Network SLA cannot be monitored or guaranteed per VM or traffic class, since everyone is sharing the same NIC and physical wires.
- Network policies are hard to enforce in dynamic server environments.
- There is a lack of security and isolation at L2 due to operational challenges.

Lack of Visibility and Automation

- Monitoring and troubleshooting tools are designed for physical elements and fixed relations between applications and physical infrastructure.
- Network provisioning consists of manual and expert configuration, and is not integrated as part of the overall service/cloud provisioning.

Most data centers are starting to realize the benefits of virtualization, cloud computing, and automation. However there are still some key challenges creating significant inefficiencies and inhibiting wider adoption. Insufficient I/O capacity, network architectures that are inherently expensive and inefficient, the lack of fabric isolation and SLA, and inadequate visibility and automation are all issues plaguing today's virtualized and cloud computing environments.

New data center and cloud optimized connectivity solutions are required to address these challenges. Voltaire's family of scale-out 10GbE and InfiniBand switches and software enable cloud computing solutions that deliver:

- Faster fabrics with 10 or 40 Gb/s connectivity to address performance challenges
- The ability to offload I/O virtualization tasks to the adapter and switch hardware, thus reducing I/O overhead caused by software-based I/O virtualization and switching
- Converged fabrics that can run multiple isolated virtual networks over the same wire with guaranteed security and service levels, replacing multiple physical networks
- Flatter fabrics using a scale-out architecture that eliminates expensive core network and storage switches – removing potential network bottlenecks and allowing server mobility while maintaining application SLA
- Central fabric resource management that can enforce fabric policy in real-time across virtual I/O, virtual switches, edge or blade switches, and core switches
- Dynamic installation of policies to meet desired application SLA, guaranteed application or traffic class isolation, and virtual server migration
- Fabric monitoring for real-time, application-centric information on traffic, I/O and bottlenecks

Voltaire Fabric Solutions

Voltaire provides an integrated approach to cloud networking that delivers the best performance, efficiency, and isolation over a shared fabric, as well as the simplest and most automated fabric management—all with the lowest solution costs and power consumption.

The building blocks of Voltaire's Fabric Solutions include:

1. **Voltaire® Vantage™ and Grid Director™ Switches** – High-density, lossless, data center optimized L2 core 10GbE and 40Gb/s InfiniBand switches
2. **Voltaire Unified Fabric Manager™ (UFM™) Software** – Fabric management software that orchestrates Voltaire switches and third party physical or virtual switches to deliver aggregated monitoring, dynamic fabric partitioning, and fabric policy automation
3. **Voltaire I/O and Application Acceleration Software** – A set of software solutions to accelerate application performance and latency
4. **Voltaire Storage Accelerator (VSA) and Gateway** – Tools to accelerate and simplify storage access through fast caching and SAN I/O Virtualization
5. **Solution Integration** – Delivery of an integrated cloud solution that can be managed as a unified computing system working with key OEM, reseller, and ISV partners

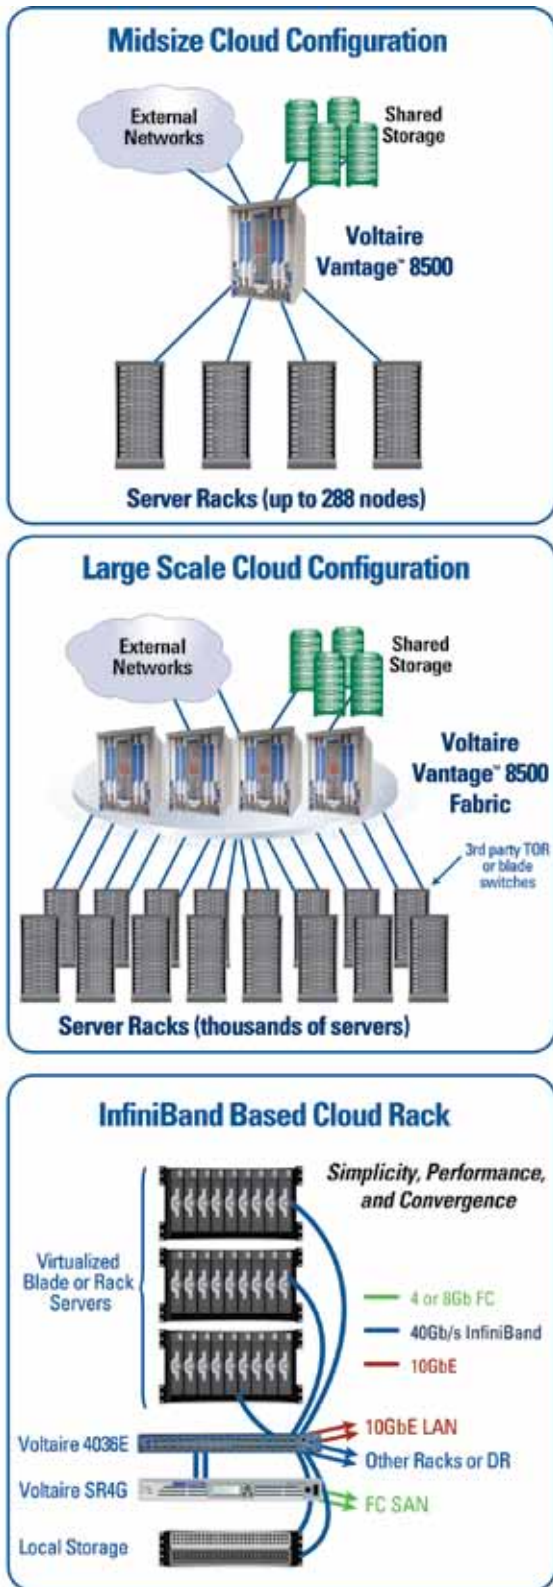


Figure 1: Midsize and Large Scale Cloud Configurations

Data Center Topology Using Voltaire Vantage 8500

At the core of the Voltaire-based data center solution is the Vantage 8500 switch. One or more switches are connected to servers, external networks, and shared storage (Figure 1). In large scale configurations, the first switching tier may consist of any 1GbE or 10 GbE top-of rack (TOR) or blade switch connected to one or more Voltaire Vantage 8500 switches that serve as the fabric aggregation layer.

With Voltaire's Vantage 8500, configurations require no more than two switching tiers. Even for configurations consisting of thousands of nodes, the price and power per node remain constant, and latency is significantly reduced.

Alternatively, data center architects can use Voltaire's Grid Director™ 40Gb/s InfiniBand switches— specifically the Grid Director™ 4036E, which offers a combination of 40Gb/s InfiniBand server fabric ports and 10GbE uplink/LAN ports in a dense 1U form factor—as a faster and more economical alternative to Converged Ethernet solutions.

Accelerating and Isolating the I/O in Virtual Environments

Virtualization technologies today implement I/O virtualization and switching in software, software bridges and NIC emulators. When these technologies are used, they add significant overhead and latency.

Technologies such as para-virtualization help reduce the I/O emulation overhead by using an efficient API between guest virtual machines and the VM hypervisor. Para-virtualization can improve the bandwidth of I/O significantly, but it still has significant overhead. Many enterprise and real-time applications are not running efficiently over VMs due to software-based I/O virtualization overhead.

New technologies have been introduced that offload the I/O switching and partitioning to hardware, relieving the I/O challenges. In this scenario, a NIC must be able to generate multiple Virtual I/O instances (e.g. using PCI SR-IOV), and a switch must be able to monitor or enforce policy for each Virtual I/O as an independent entity, and potentially even switch packets arriving from a source port back to its origin (i.e., when two VMs from the same device communicate with each other).




Figure 2: I/O Virtualization Modes

When deploying many VMs on the same server, storage I/O becomes a key bottleneck. Furthermore, since many VMs produce very random storage access patterns, traditional storage systems may not be able to sustain the load. As a result, faster storage access and storage caching solutions are required. Voltaire Storage Accelerator (VSA) software combined with solid state storage delivers the speed required, and can also be used as a fast I/O virtualization gateway to FC, eliminating the need for expensive FC HBAs and FC switches.

In cloud environments, multiple applications and tenants are deployed over the same physical server and share the same network wire, creating additional challenges. Furthermore, each VM or application may have multiple connections to different fabrics (such as LAN, storage, IPC, management). These connections may have different service attributes such as lossless vs. lossy, low-latency, etc. Therefore, it is essential to partition the I/O and fabric resources to multiple isolated virtual end-points (ports) and virtual fabrics with unique resource assignment and service attributes.

The **Voltaire Vantage 8500** switch supports hardware-based I/O virtualization and allows users to create, monitor, and provision Virtual I/O Port Objects (VIPO). Users can assign specific policies to each Virtual I/O. These policies will be maintained even when the VM migrates to a different physical port, and each Virtual I/O port is connected to a specific virtual fabric that may have unique policy and behavior (e.g. low-latency, lossless, etc.). Users can manage virtual or physical network and I/O objects using **Voltaire Unified Fabric Manager (UFM)** software.

UFM provides mechanisms to define application and service I/O policies that are automatically pushed to the various switches (physical or virtual). Traffic is constantly monitored, and statistics and alarms on both the physical and virtual connectivity are provided. UFM is integrated with virtual machine hypervisors (such as VMware and Linux KVM) to automatically provision virtual switches and to synchronize the policy and state of the fabric in the case of any change or migration event.

Voltaire also delivers I/O and storage acceleration solutions that significantly boost application performance, as well as I/O gateways that enable fabric convergence and isolation.

Managing the Fabric as a Service with UFM

Integration and automation are key challenges in cloud computing environments. Specifically, network resources are ignored or handled in a very basic manner, leading to inefficient use of resources, lack of isolation, and inadequate service levels. Voltaire's breakthrough fabric optimization and management technology allows managing fabrics as an application service with easy integration into cloud management solutions and much simpler fabric provisioning.

UFM maintains a powerful model that accepts high-level application or service requirements, and then translates the requirements into specific policies that are dynamically enforced on the I/O and switching elements in order to maintain the application SLA. In addition, UFM collects information from a variety of sources and remaps or aggregates them to the application model and to the virtual or physical infrastructure elements.

The application or service requirements can be inserted manually, or policies or monitoring can be automated through a rich and open web-services based API.

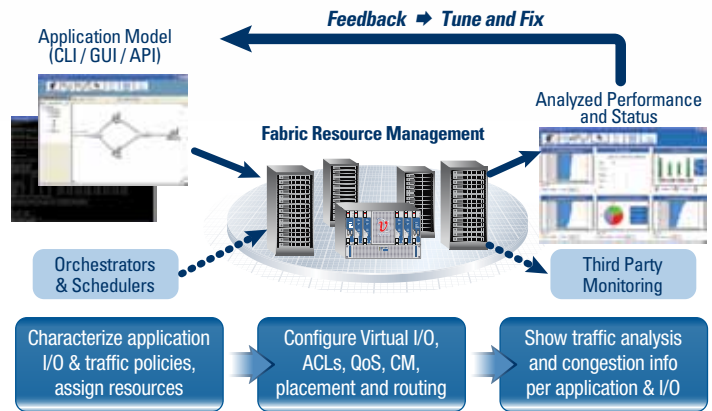


Fig 4: UFM Screenshots



Monitoring and Troubleshooting	<ul style="list-style-type: none">■ Enables automated fabric and resource discovery■ Collects and analyzes statistics from virtual or physical switches, optional host agents, and pluggable user or application-specific collectors■ Quickly identifies health problems, root cause, and application bottlenecks■ Identifies and visualizes fabric congestion and spreading■ Maps alarms, triggers, and logs to virtual or physical infrastructure
Dynamic Fabric Provisioning to Maintain Fabric Isolation and Application SLA	<ul style="list-style-type: none">■ Creates and manages multiple isolated virtual data centers on top of a shared fabric■ Automates VLAN assignment and I/O partitioning (virtualization)■ Applies QoS (class, priority, limits, guarantee) per Virtual I/O or traffic flows■ Guarantees HA (multi-rail, multi-path configuration, and policy synchronization)■ Suggests optimal VM placement based on fabric allocation or load■ Provides congestion isolation, control/throttling, and monitoring
Simplified and Unified Administration	<ul style="list-style-type: none">■ Manages multiple switches through a unified management console■ Groups operations, policies, and software upgrades■ Tracks relations between physical, virtual, and application elements in real-time

Summary

Voltaire's cloud computing solution, consisting of Voltaire's Vantage and Grid Director switches, Unified Fabric Manager (UFM) software, host software, and partner products, addresses the key challenges facing today's cloud and virtual data center environments.

Voltaire solutions deliver:

The Best Performance

- High performance of 10-40Gb/s per server with minimal oversubscription
- Hardware-assisted I/O virtualization for best application performance
- Storage access acceleration
- Automated traffic and congestion management

The Most Cost Effective and Efficient at Scale

- Affordable and power-efficient non-blocking switches
- Linear scalability with fixed price/port and performance
- Fabric convergence (LAN, SAN, and IPC) to a single wire

Optimal Isolation and Fabric SLA

- Multiple independent traffic classes delivered over the same fabric
- Partitioning of the fabric to multiple isolated virtual data centers
- Enforcement of SLA and QoS to applications and virtual fabric elements

Simplified Management and Integration

- Application and VM driven policies mapped dynamically to actual resources
- Extensive monitoring and analysis at different layers
- Central fabric maintenance and troubleshooting
- Open and extensible architecture for third party and user integration



Contact Voltaire to Learn More

1.800.865.8247
info@voltaire.com
www.voltaire.com

©2010 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.