

Voltaire® Storage Accelerator (VSA)

Delivering fastest storage access using commodity hardware



Overview

Recent exponential growth in CPU capacities and core counts means that many applications are now inhibited by I/O speed. Furthermore, increased data center consolidation and virtualization creates random storage access patterns, resulting in poor storage performance and slow response times. With the emergence of faster and denser solid state storage solutions, the hard disk and its electro-mechanical limitations is no longer the bottleneck.

New SAS (Serial SCSI) technologies allow accessing a single disk at 6 Gb/s, and some new SSD disks provide more than 10 Gb/s bandwidth per drive. With just a couple of disks in a box, any array or Fibre Channel (FC) link capacity can easily become saturated. In addition, access times to SSD disks are 100-1000X faster than traditional hard drives, and so the number of I/Os per second can reach into the millions per array vs. thousands with older technologies. With this enormous increase in speed, new tiered solutions utilizing faster transports and a scale-out fabric architecture are needed to enable the full potential of emerging storage technologies, and eliminate the application bottleneck resulting from slow storage I/O.

Introducing Voltaire Storage Accelerator (VSA) Software

To address these I/O performance challenges, Voltaire provides a highly scalable, high performance, and low-latency software solution for tier one storage and gateways. **Voltaire Storage Accelerator (VSA) software** can be deployed across industry standard storage servers or appliances to provide ultra-fast remote block storage access over a low-latency 40 Gb/s InfiniBand or 10GbE fabric. VSA leverages iSCSI or iSER (iSCSI RDMA Extension) to maximize throughput. VSA is designed to process many transactions in parallel and with very little I/O overhead, leading to unprecedented speeds. Multiple VSA-enabled appliances can also be managed as one fault-tolerant, scalable unified system. A VSA storage cluster can process many millions of transactions per second, deliver hundreds of GB/s of data, and speed access time by a factor of 100X over traditional storage, leading to far greater application performance at lower costs and power.

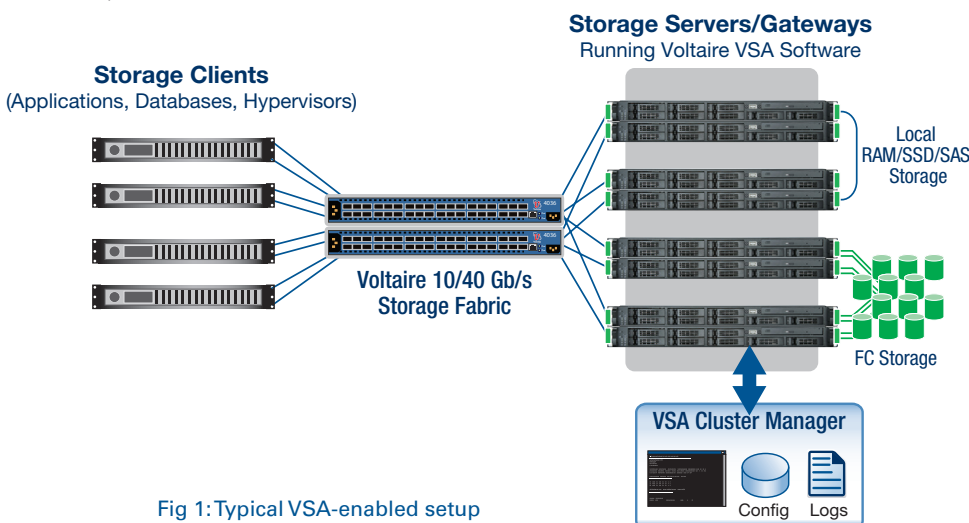


Fig 1: Typical VSA-enabled setup

Features

- Highest speed block storage target & gateway software
- Up to a million I/Os per second and 5,000 MB/s per storage target
- Only 0.07 millisecond access time
- Read and write storage cache
- Raid 0/1, Multipath, and Replication
- High-availability and scale-out design with integrated cluster management

Benefits

- Significantly improves application performance and response time
- Much greater performance and scalability with lower costs and power consumption
- Enables fabric convergence and lowers infrastructure costs
- Easy to manage and monitor
- Preserves existing storage investments

Applications

- High-performance storage clusters
- Databases
- Server virtualization and cloud computing
- Data caching and logging
- Billing systems
- InfiniBand to FC, or 10GbE to FC storage gateways

Voltaire® Storage Accelerator (VSA)

VSA reduces data center costs and complexity by enabling fabric convergence. 10GbE and/or 40Gb/s InfiniBand-based clients can use VSA as a fast FC gateway and storage cache, and to access remote FC storage without FC HBAs or FC access switches. Hundreds of virtual FC HBAs/ports can be provisioned with a single VSA-based gateway appliance (allocated to virtual or physical 10GbE or InfiniBand clients), while also boosting storage access performance.

VSA also discovers storage resources in storage server platforms, as well as any FC storage attached to it. It can configure user/client access to storage resources, and monitor the system or storage traffic (i.e. bandwidth, latency, input/output operations per second or IOPS, etc.).

VSA is installed on a standard Linux storage server and can front end DAS (SAS/SATA), SAN (FC/iSCSI), Flash, or RAM based storage. VSA ensures that the storage devices can be accessed in parallel over the fabric at speeds that are a few magnitudes faster than any SAN technology, while providing iSCSI management semantics for ease of use. In addition, it can be deployed as a tier one storage cache in front of slower DAS or FC SAN storage systems.

VSA is designed to maintain the highest transport reliability and availability, and to support link bonding, multi-pathing, mirroring, and fail-over. It also allows users to maximize performance and availability while using commodity, low-cost hardware elements.

VSA relies on third party or OS-based solutions to provide file system, volume management, backup and snapshot functionality, including traditional or cluster file systems such as GPFS, Ixrix, Oracle ASM, or VMware VMFS.

Voltaire Storage Accelerator Architecture

Unlike existing systems that can support a few 100K I/Os per second at the most, VSA-enabled systems can deliver millions of I/Os at low latency with gigabytes of throughput. VSA's zero-copy architecture leverages RDMA, allowing it to saturate the server I/O buses and deliver more than 5GB/s of bandwidth per storage server/gateway. The software incorporates multiple ultra-fast datamover engine processes. These communicate directly with the NIC/HCA hardware and leverage hardware offload capabilities to process the storage transactions. This enables unprecedented IOPs performance and a very low access time of only a few dozen microseconds end-to-end. VSA datamover engines communicate with a services layer that delivers caching, RAID, multipath, and I/O virtualization. Users can deploy a cluster of VSA servers with a combination of DAS, SAN, and Flash storage, and achieve the aggregate performance and capacity.

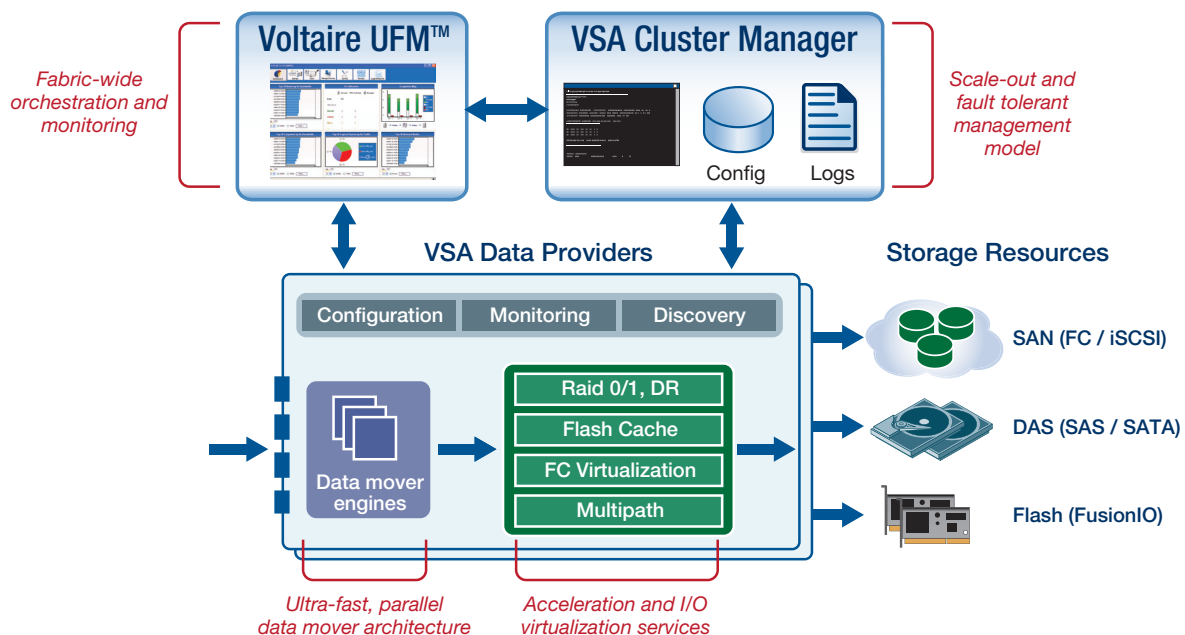


Figure 2: VSA Software Architecture

Manageability and Usability

Multiple VSA appliances are managed as a cluster with unified CLI and management. The VSA cluster manager provides central resources and hierarchy discovery, automated configuration for network and storage elements, central monitoring and logging, and secure role-based remote access.

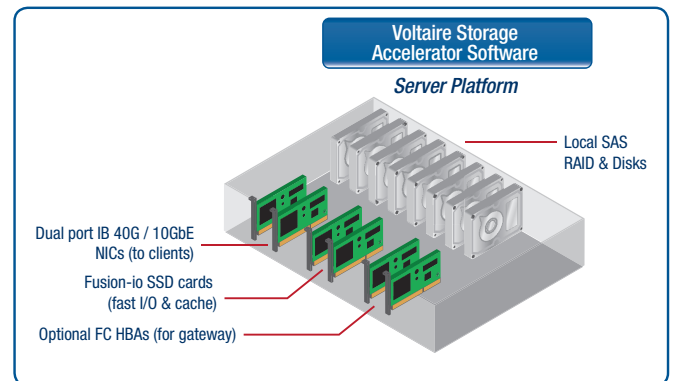
VSA is integrated with Voltaire's Unified Fabric Manager™ (UFM™) software allowing central remote discovery, monitoring, and configuration. With UFM, users benefit from end-to-end visibility into fabric resources and can control fabric policy through an application and service-oriented view.

VSA is provided with automated installation tools and detailed user documentation.

Configuration Examples

A VSA server or gateway is based on a standard x86 server running the Linux operating system and equipped with one or more of the following storage options:

- RAM storage (based on system memory)
- Fusion-io Flash Disks (PCIe)
- SAS/SATA RAID controllers and disks
- Qlogic 4/8Gb FC Adapters (with NPIV support) for gateways



Application Example 1: Scale-out Tier One Storage and FC Gateway

Multiple storage servers can be used in a cluster to deliver massive application capacity, performance of millions of transactions per second, and hundreds of Gigabytes per second at far lower cost, space, and power than alternative solutions.

VSA can also be placed as a Flash cache in front of traditional FC storage systems.

Application Example 2: Virtual Machine Booster (Rack level cache)

When deploying a rack of servers, each hosting a couple of dozen virtual machines, the result is 1000 VMs generating independent I/O transactions to storage. Because of the high density and random I/O patterns involved, most traditional storage systems cannot cope with the load and produce extremely slow performance and response time, leading to underutilized CPU resources and degraded application performance.

VSA software installed on a standard server:

- Uses internal fast SSD or Flash for boot, temporary data, or storage cache
- Significantly increases VM performance and the VMs per server ratio
- Delivers easy-to-use FC I/O port virtualization, enabling fabric convergence

Performance

The performance results described below were conducted using the following storage server hardware:

- Software: Voltaire VSA installed over CentOS 5.5
- Storage Server: HP DL 380 G6, with 2 x Intel 5520 CPUs (Quad core), 12GB Memory
- Fabric: Voltaire 40Gb/s InfiniBand
- Storage: 2 x FusionIO ioDrive Duo PCIe cards

Note that faster performance can be achieved when using more storage resources and/or with faster CPUs.

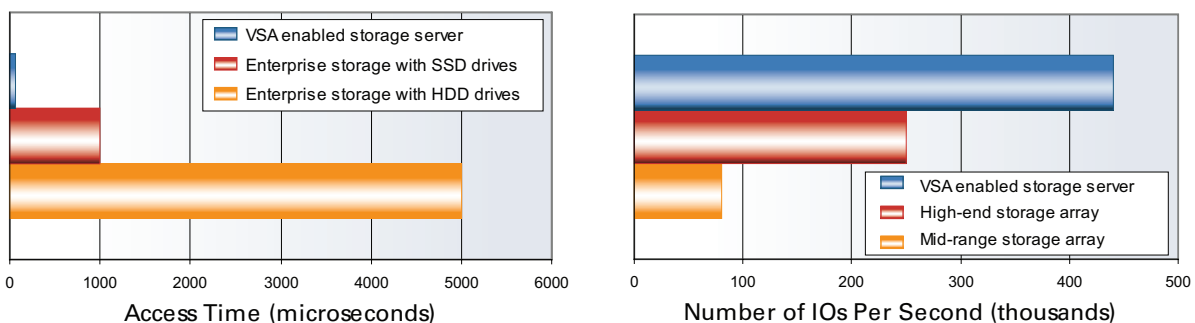


Fig 3: VSA IOPs and access time (latency) compared to traditional storage systems

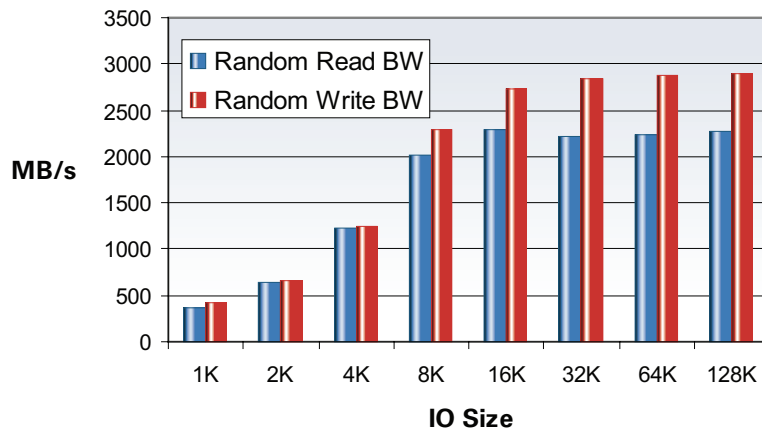


Fig 4: VSA performance with various block IO sizes for random read/write tests

TECHNICAL SPECIFICATIONS

VSA Software Pre-requisites

- VSA Server or Gateway configuration
 - ▶ x86_64, using dual socket each with 4 cores is recommended for high performance
 - ▶ 6GB RAM Minimum, 12GB Recommended
 - ▶ 20GB Available Disk Space for OS
 - ▶ HCA: Mellanox ConnectX DDR/QDR
 - ▶ 10GbE NIC: Mellanox, Intel
 - ▶ OS: RedHat 5.x or CentOS 5.x or OEL 5.x or RedHat 6.0
- Storage Devices and HBAs
 - ▶ FC HBA: Qlogic 4/8Gb FC
 - ▶ SAS/SATA HBA: LSI or HP
 - ▶ Flash: Fusion-io
 - ▶ iSCSI devices via 1/10GbE or InfiniBand
- Storage protocols (to client)
 - ▶ iSCSI over 1/10GbE or InfiniBand (iPoIB)
 - ▶ iSER (iSCSI RDMA), supported with Linux clients

Ordering Information

- VSA is licensed per storage server or gateway and their configuration.
- VSA is offered in various packages. Packages differ in functionality and price.
- For more details please visit www.voltaire.com or contact info@voltaire.com.

Features and Functionality

- Storage and Access Services
 - ▶ LUN mapping and masking
 - ▶ RAID0, RAID1, Multipath
 - ▶ Flash cache (read and write)
 - ▶ FC I/O Virtualization (using NPIV)
 - ▶ Multipath
- Configuration and Management
 - ▶ Auto create targets & Luns, ACLs
 - ▶ Auto configure networking
 - ▶ Auto create Virtual FC Ports for clients
 - ▶ Multipath, Global Balancing
- Discovery and monitoring
 - ▶ Auto discover disks, virtual volumes, and raid devices
 - ▶ Auto discover multiple paths
 - ▶ Discover FC fabric elements
 - ▶ Monitor system, network, storage, and session performance
- Maintenance, Security, and High-Availability
 - ▶ Installer, Persistent configuration, watch-dogs
 - ▶ Local/remote logging
 - ▶ Multiple Users, Secure Access, Audit Logs