



Product Highlights

High Performance

- Parallel storage system providing throughput aggregation
- Tunable parameters to meet application needs / profile

Scalability

- Easily increase compute and storage capacity
- Up to 128 nodes and 32 Petabytes

Cost control

- Commoditized architecture independent from hardware vendors
- Drastically cut storage costs by pooling and reusing internal disks

Simplicity

- Shared virtual storage pools
- Natural compliance with standards
- Non intrusive

Reliability

- Data protection and continuity of service ensured through RAIN X
- Fastest rebuild time after server or disk failure (up to 1 TB in 45 minutes)

Dynamicity

- Ability to add more disk in the cluster without rebuilding the storage

Meeting the Challenge of Commoditized Server Farm Storage

High-performance compute clusters and commoditized server farms place extreme demands on storage systems. Performance requirements are typically in the multi gigabyte-per-second range and it is critical that the storage environment can be easily and effectively managed. All this would not be so complex if costs for high end storage systems were in par with those of computing servers.

Seanodes[®] Shared Internal Storage and its flagship product Exanodes[®] 5.0 for Linux answer those questions with a smart, world shattering, and affordable solution.

The First Shared Internal Storage Solution

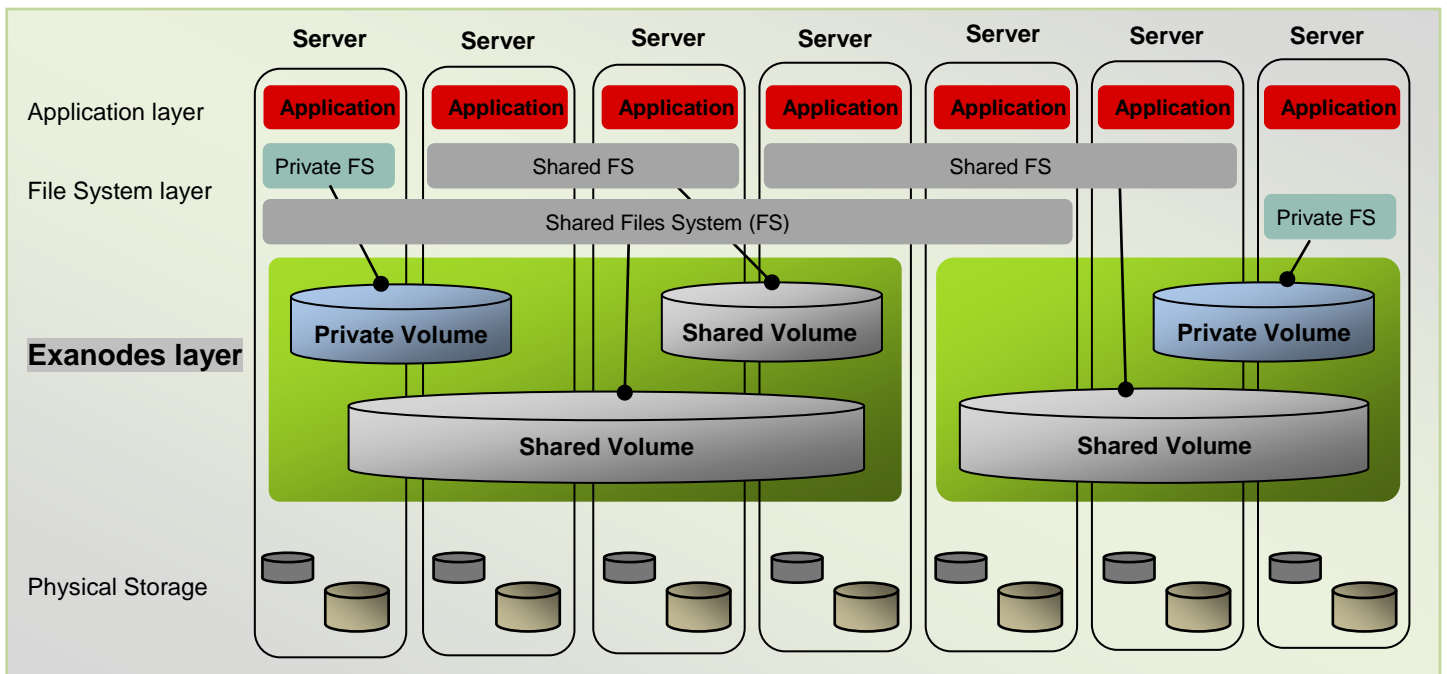
Exanodes is a software only solution which creates a high-performance and affordable shared storage system using application servers' own hard disks, which are generally under-used.

The internal storage is clustered, making it possible to share data among the servers (or nodes) of a cluster in a reliable, high-performance way. It is perfectly suited to the needs of high-performance cluster environments as it is based on their architecture and philosophy (see figure below).

Because of its versatility and high performance, Exanodes is the ideal storage solution for commoditized server farms dedicated to High Performance Computing, Hosted Services or other areas with large numbers of Linux boxes.

Exanodes is non-intrusive and is easily implemented alongside other storage solutions. Its integration is therefore quick and easy and in no way affects previous solutions.

Exanodes is a pure software solution that doesn't require any additional hardware to enable storage virtualization (beside disc drives) and high availability on the cluster.



Product Highlights

High-performance storage

Aggregation of bandwidth

With Exanodes, any computing server is also a storage server. The processing of inputs / outputs is done simultaneously on a large number of servers. The application therefore benefits from cluster parallelism for computing as well as storage throughput.

Natural scalability

Exanodes allows the storage architecture to scale naturally with the computing architecture : adding a server to the cluster increases computing and storage performance at the same time.

Cost Control

Considerable reduction of external storage needs (SAN or NAS)

Exanodes drastically reduces acquisition costs of a storage infrastructure. Indeed, external storage needs (SAN or NAS) are greatly reduced or even eliminated.

External storage is confined to less demanding uses: archiving of computing results, data distribution for workflows involving several clusters, management of backups, etc.

Reduction of operating costs

The simplicity of the Exanodes solution (integration and day-to-day management) makes it possible to increase the amount of storage and computing nodes that a person can administer.

In addition to acquisition costs, substantial savings can thus be made throughout the use of the cluster.

Reliable storage

Redundant Array of Inexpensive Nodes (RAIN)

To increase data availability and data protection, Exanodes provides a RAIN storage system. All data is replicated on disks belonging to different nodes.

When a node fails, Exanodes' RAIN system is able to carry on with the processing of the requests thanks to data replication.

Applications thus have uninterrupted access to their data and no data is lost when a node or a disk fails.

RAIN of RAID

Exanodes also virtualizes any internal RAIDs on the cluster nodes. By combining the protection of internal RAIDs with that of the RAIN, Exanodes provides highly reliable storage which can withstand numerous disk failures.

RAIN X

RAIN X is a data layout scheme that allows Exanodes to support cumulated disk and server failures in the same system. It provides distributed spare space and has the ability to rebuild data with an unmatched efficiency.

Simplicity

Shared storage system

Exanodes provides all the advantages of a storage system shared by the nodes of a cluster. Data management is thereby simplified.

Within the same cluster, applications become truly mobile as it is no longer necessary to load the data to the server where they are used. Vice versa, computing results are accessible from any node, removing the need for staging out data after the computing phases.

Natural compliance with standards

Exanodes has an extensive compatibility matrix. It supports all types of block devices as well as high-performance networks, Linux private file systems and clustered file systems.

Users thus retain the freedom to choose the software and hardware technologies best suited to the requirements of their applications in order to optimize each layer of their infrastructure.

No change to existing system

Exanodes is non-intrusive. Applications access Exanodes' high-performance storage area, without any need for reprogramming.

Moreover, Exanodes is installed without any need to modify the Linux kernel and requires no additional hardware, no need for external SAN storage or fabrics and no specific storage competencies.

Exanodes 5.0 for Linux Specifications	Exanodes 5.0 for Linux Supported Systems	
	Hardware Platform	Software Platform
<p>Exanodes cluster :</p> <ul style="list-style-type: none"> From 2 to 128 servers From 2 to 512 block devices (full disks, disk partitions or RAID) Up to 64 block devices per node Up to 32 Petabytes <p>Exanodes disk group limits :</p> <ul style="list-style-type: none"> Up to 512 block devices (full disks, disk partitions or RAID) Up to 2 Petabytes <p>Exanodes reliability :</p> <ul style="list-style-type: none"> Up to 16 spares 	<p>Supported CPU architecture :</p> <ul style="list-style-type: none"> x86_64 1Ghz min 512MB min <p>Supported disk technology :</p> <ul style="list-style-type: none"> Any disk drive/partition technology supported by Linux and seen as a block device (SATA, SAS, SCSI , SSD, RAID) <p>Supported High Performance Interconnect :</p> <ul style="list-style-type: none"> Gigabit Ethernet 10Gigabit <p>Other :</p> <ul style="list-style-type: none"> Contact us 	<p>Linux Kernel Version :</p> <ul style="list-style-type: none"> 2.6 <p>Linux Distribution :</p> <ul style="list-style-type: none"> RHEL 5 and derivatives CentOS 5 <p>Shared File systems :</p> <ul style="list-style-type: none"> Seanodes FS <p>Local File systems :</p> <ul style="list-style-type: none"> Ext3 XFS <p>Other :</p> <ul style="list-style-type: none"> Contact us

Seanodes Prizes & Awards

