

Executive summary

Compuverde has recognized the necessity to see storage as a strategic dimension in order to reach full effect and potential now and onwards. Storage is the last part of many total system setups to be transformed from fixed to flexible, from rigid to pragmatic. Compuverde is now, as a result of nearly 20 years experience, changing the storage industry fundamentally. Delivering a hyper-converged software solution with flexibility and scalability, storage is with Compuverde a proper foundation to build connectivity and distributed services on – a strategic asset for forward-thinking businesses big and small. Developed from the ground and to meet the future needs of storage platforms, Compuverde has created an architecture that is completely symmetrical.

The storage system is massively scalable, software defined and completely hardware agnostic. Compuverde provides complete capabilities for virtualized environments and vNAS. When handling scale-out NAS, Compuverde delivers a single file system spanning over all servers. The solution supports all major protocols, multiple operating systems, cache, flash and ensures redundancy with support of replicas, RAID 5 and Erasure-coding. Virtualization with Compuverde provides a hyper-converged storage where a hypervisor of choice and VMs run on the same server as the storage layer delivering enterprise-class performance, scalability and availability.



Telecom graded solution



Hardware agnostic



Allows for cloud storage



Extreme performance



Extreme scalability



Global patent portfolio



Support for virtual machines



Fully flash compliant

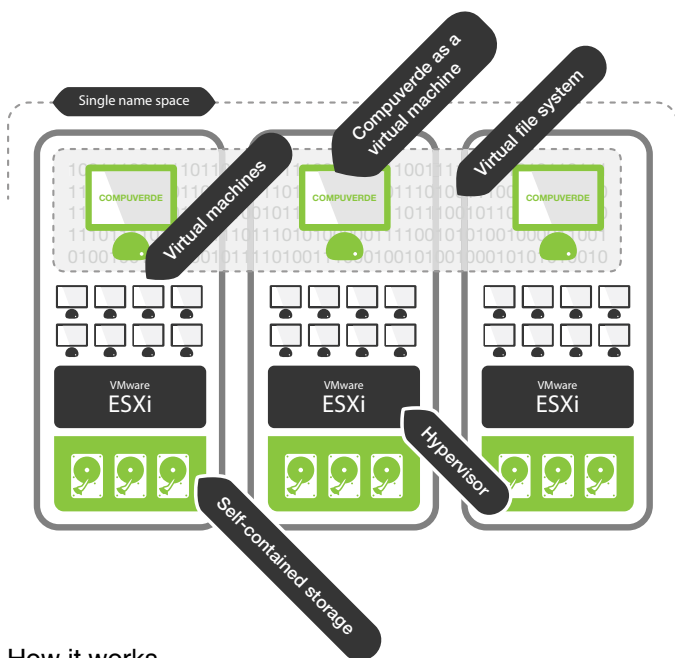
Products

Being hardware agnostic and completely software defined, Compuverde offers a scale-out NAS and scale-out VM setup that you won't have to estimate the size of beforehand nor invest in any specialized and expensive external systems. Compuverde is a solution that manages to combine high levels of flexibility, redundancy and performance in one package.

Virtualization

Compuverde supports virtualization and offers full integration with VMware ESXi, Windows Hyper-V, KVM and XEN. Compuverde also has full support for cache synchronization with other VM hosts.

- Run hypervisor and VM's in same server as storage
- Use of cache for performance
- Synchronized cache for fast live migration
- Synchronized to storage for redundancy and availability
- Scalable up to hundreds of nodes



How it works

Compuverde is run as a virtual machine within the host computer and is given ownership of the physical discs, taking control of the setup. The Compuverde servers are then forming a virtual file system that spans over all nodes in the cluster. The guest virtual machines are using Compuverde as storage for their own VM-images. While the machines are booting up, each machine is in contact with the virtual storage layer in Compuverde through the implemented protocol.


With Compuverde, you are also able to work with files and documents from any virtual machine as well as share files and documents between them. When a file is created by one virtual machine, it is at once also available from any other virtual machine in the cluster using the shared file system that Compuverde is providing.

With this, Compuverde is allowing for a fully VDI-compatible environment – meaning that you with Compuverde's virtual NAS-solution, have the possibility to share and work any way you want.

Scale-out NAS / vNAS

Compuverde has full support for scale-out NAS through both the CIFS and NFS protocol. The tight integration between these protocols and the storage core, ensures full scale-out as well as high availability features through NFS4.1, pNFS and SMB 3.

- Single file system spanning over all servers
- Native support of protocols
- Flash support for high performance
- Flexible scale-out by adding nodes



How it works

The scale-out NAS solution from Compuverde contains a layer with direct-attached storage. And on top of this storage layer sits the gateway layer. This is where Compuverde's software is forming a virtual file system that spans over all nodes in the cluster. Each server sees the same synchronized view of the filesystem and can access the same files.

When a user reads or writes data, the data is first stored in the local cache. Each cache is synchronized horizontally with the others. For full protection, the data is then replicated across multiple nodes before de-staging data from the cache. The cache is also at all times synchronized vertically with the persistent storage layer.

To ensure redundancy, Compuverde supports copies of data, RAID 5 for the smallest setups and Erasure coding for setups from six up to hundreds of nodes. With Erasure coding you will also get the benefit of creating a significantly smaller footprint within your data center.


Even if you should lose a node in the cluster, the data is available from any other node. This means that you are able to write data to one node in your cluster, while at the same time reading the same data from another.

Key features & benefits

Extreme scalability

By using a decentralized and symmetric architecture all storage nodes share the same functionality and responsibilities. No special purpose nodes or central nodes are needed which greatly simplify deployment and maintenance and remove any potential bottlenecks or single points of failure.

Because the storage cluster is built on a symmetric architecture, linear scaling up to hundreds of petabytes and billions of files is achieved simply by adding more storage nodes to the cluster. Adding storage nodes and increasing capacity can be carried out during runtime and does not interrupt any ongoing operations in the cluster.



Hardware agnostic

The Compuverde software runs on standard commodity hardware and servers without lock-in to a specific vendor and/or technology. Customers can add additional hardware of choice as needed to scale performance and capacity over time.

Supports integrated cache

Compuverde supports integrated cache consisting of NVRAM and SSD. The cache is at all times synchronized with the storage nodes and also synchronized horizontally with the other gateways' cache, giving extreme performance wins.

High availability

Compuverde software ensures high availability at all times by synchronizing the cache horizontally between gateways and vertically down to the object store. Therefore whichever gateway you access, the data is available. By supporting NFS 4.1, pNFS and SMB 3, Compuverde even more ensures availability, scalability and redundancy.

Redundancy

Compuverde support copies of files, RAID 5 and Erasure-coding. Compuverde also has full support for policies on file- and object level. The redundancy parameters can all be overridden by using policies or use of an API. The placement of replicas or pieces of data can be completely controlled by said policies.

Telecom grade

Compuverde secures five nines. In addition, fault recovery should be extremely fast and allow for almost complete redundancy. Compuverde software both meets, and for some areas even exceeds these definitions.

Greener data

The software developed by Compuverde will help businesses lower their energy costs with up to 50% - by making it possible for organizations to choose environmental friendly green hardware.

Supports all major protocols

The solution supports scale-out file systems with support of all major protocols, e.g. NFS, CIFS, CDMI, OpenStack, Amazon S3. Compuverde also supports NFS 4.1, pNFS and SMB 3 to ensure high availability.

Hypervisor support

Compuverde makes it easy to support your existing VM environment by fitting seamlessly with the industry's leading hypervisors, including VMware ESXi, Windows Hyper-V, KVM and Xen.

Multi O/S support

Compuverde Gateway has a very small O/S footprint and runs on both Linux and Windows. The small footprint makes the product more easily portable between O/S.

Handling of metadata

Metadata is stored within the storage cluster rather than in the gateways. By doing so, the system is made much more scalable. Metadata is also very sensitive information and needs to be protected from data loss. Compuverde store and handle metadata using the same algorithms as normal files. This creates a more robust and less complex architecture, for both software and hardware.

Fully flash compliant

Intelligent use of cache through flash storage compatibility is one of Compuverde's great advantages in the solution.

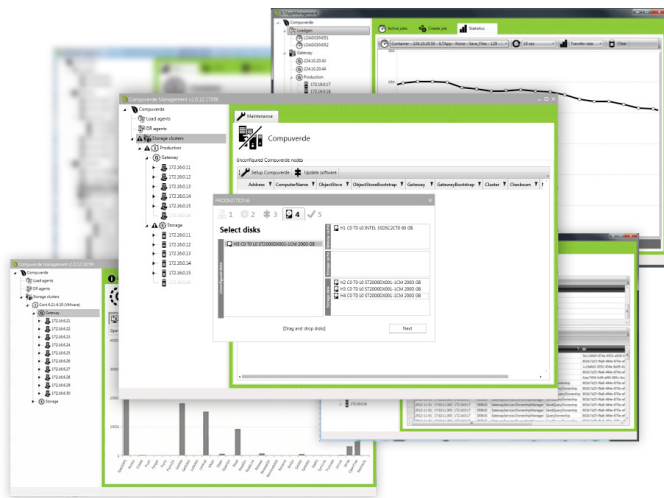

Cloud enabling

By using Compuverde, moving your data storage, processing and distribution to a cloud environment is no longer a far-fetched vision but a viable and effective alternative. Through virtualization you are no longer bound by physical servers, data halls and more, but can instead focus on delivering on your business promises.

Set-up and management

Installation process

Compuverde has an easy-to-use graphical interface - user friendliness being an overall priority for Compuverde, we put a lot of effort in to making both installation and management as straight forward as possible. Our installation process is easy to follow and is done in 30 minutes or less.



Management tool

Manage, monitor and control the entire cluster from any node through a simple graphical user interface. Through this, tasks like reading statistics, changing configurations, monitoring logs and rolling out new firmware updates, is made easy and possible during runtime without any effects on availability or performance of the storage cluster.

The Compuverde management and monitoring console is running in an Windows environment and makes it easy regardless if the storage cluster is of a size at 100TB, 1PB or 10+PB.

It is also possible to monitor and manage the storage clusters with other commercial network management systems using the Compuverde management API.

Performance

Overview

Scaling up by adding more storage nodes provides a linear increase in average performance in the cluster. The lack of single entry points enables a parallel load towards the cluster which in turn ensures an extremely high aggregated performance in the cluster.

The selfhealing process also benefits from this since all the nodes in the cluster share the work of creating lost data, the time for recreating data from lost nodes is dramatically shortened. The more storage nodes in the cluster, the shorter time it takes to complete the task.

Test information

The performance tests between Compuverde and two well-known competitors was performed on identical clusters with low to midrange hardware specifications using SPECsfs2008.

SPECsfs

SPECsfs2008 is the latest version of the Standard Performance Evaluation Corporation benchmark suite measuring file server throughput and response time, providing a standardized method for comparing performance across different vendor platforms. SPECsfs2008 results summarize the server's capabilities with respect to the number of operations that can be handled per second, as well as the overall latency of the operations. The suite is a follow-on to the SFS97_R1 benchmark, adding a CIFS workload, an updated NFSv3 workload, support for additional client platforms, and a new test harness and reporting/submission framework.

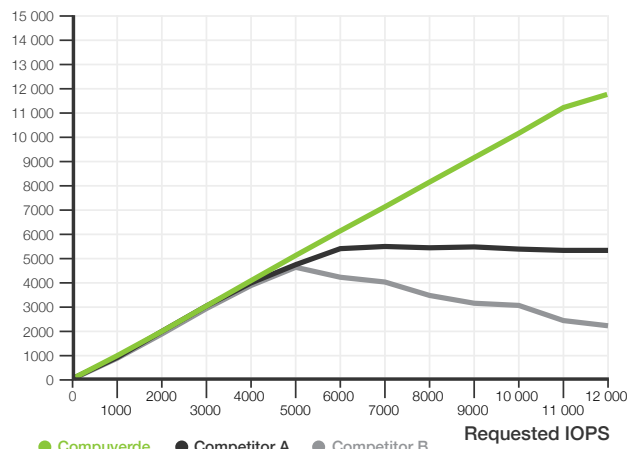
SPECsfs CIFS test setup

6x

CPU: Intel Xeon E5-2620 (2.0 GHz)
RAM: 32 Gb
Network: 2 Gbit
Cache: 1 x Intel 330 (SSD/60 GB)
Storage: 8 x Seagate SV35 (SATA/3 TB) JBOD

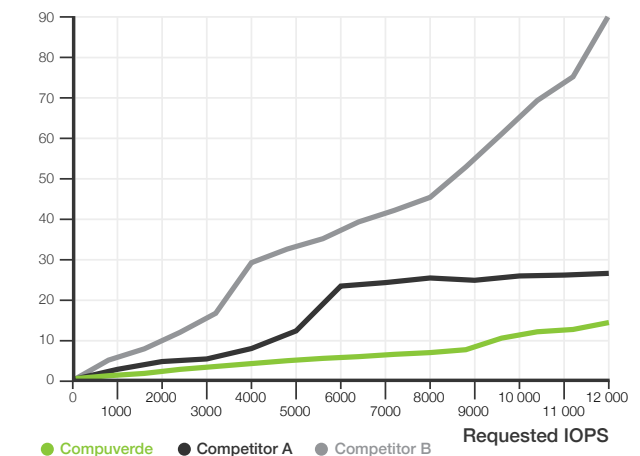
IOPS

Actual IOPS
(More is better)




Response time

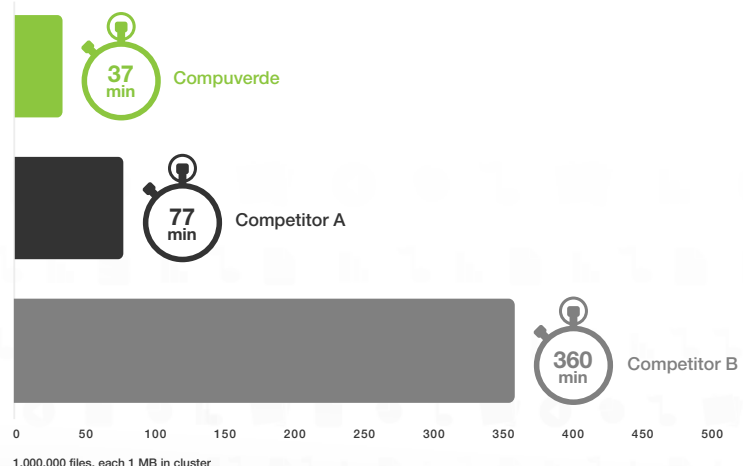
Milliseconds
(Less is better)



Replication - Time to recreate missing data



Time to recreate missing data (minutes)



Specifications - Structured data

System features	
Operating system support	Linux / Windows
Hardware independent	Yes
System architecture	Active / Active N-node Cluster
Data addressing	Structured
Scalability	Linear / Multiple Object Store clusters each 100+ billion files = several exabytes
Metadata storage	Within Object Store cluster (no local metadata)
High availability	> 99,999%
Elasticity	Runtime changes of cluster size
Policy & redundancy granularity	Per directory

Cluster features	
Protocol support	
NFS	v3 / v4.0 / v4.1 / pNFS
CIFS	Yes (SMB1 / SMB2 / SMB3)
CDMI	Yes (Domain / Container / Object)
OpenStack	Yes (Swift)
Amazon S3	Yes
Cache support	NVRAM / SSD
Cache synchronization	Yes
Cache replication	Yes
Configuration distribution	Automatic
Multi-tenant	Yes
Snapshot support	Yes (1024 snapshots per directory)
Off-site replication	Yes
Quota	Yes*

File features	
Policy	
Redundancy level	Yes
Retention	Yes
Immutability	Yes
Tier	Yes
Large file support	Yes
Partial update of file	Yes
Directory nestling	Deep

Management features	
Alarm	SNMP / E-mail
Statistics	SNMP / SMI-S*
Logging	System / Traffic
Management tool	Yes
Syslog server	Yes
Management API	Yes

Specifications - Unstructured data

System features	
Operating system support	Linux/Windows
Hardware independent	Yes
System architecture	Symmetric
Data addressing	Unstructured 128-bit
Scalability	Linear / 100+ PB / 100+ billion objects
High availability	> 99,999%
Elasticity	Runtime changes of cluster size
Policy & redundancy granularity	Per object

Cluster features	
Unstructured protocol	API (flat 128-bit) C++ / Java / .NET / PHP / Python / Ruby
Geographical disparity	Latitude, Longitude, Altitude / Rack designation
Self-healing cluster support	Yes
Automatic detection	Node failure / Disk failure / Data inconsistency
Healing mode	Background job
High performance	Yes, many to many replication pattern
Cache support	None / SSD
Configuration distribution	Automatic
Snapshot support	Yes (1024 snapshots)
Write patterns	Commit all copies / Write and replicate

Object features	
Versioning support	1-15 versions
Data redundancy	Object copies / Erasure coding
Object copies	2-15 copies
Erasure coding	Yes
RAID 5	2+1, 4+1, 6+1, 8+1, 10+1, 12+1
Cauchy-Reed Solomon	4/6, 6/8, 9/12, 10/16
Retention	R/W and WORM data
Tier support	Yes
Policy	Redundancy level / Versioning / Retention / Tier
Partial update of object	Yes

Management features	
Firmware update	Runtime upgrade of firmware / Rolling upgrade
Alarm	SNMP / E-Mail
Statistics	SNMP / SMI-S*
Logging	System/Traffic
Management tool	Yes
Syslog server	Yes
Management API	Yes