



Enterprise Virtualization OS (eEVOS)

Building a highly available total solution for server virtualization

euroNAS Technical White Paper Rev. 16-11-30

Copyright © 2015 euroNAS GmbH. All Rights Reserved.

euroNAS believes that the information in this publication is accurate of its publication date. The information in this publication is provided “as it is”. euroNAS GmbH makes no warranties of any kind with respect to the information in this publication. euroNAS GmbH specifically disclaims implied warranties or fitness for a particular purpose.

Index

Introduction	3
Why you should use high-availability environment?.....	3
Advantage of euroNAS not being locked to specific hardware	3
Supported Guest Operating Systems	4
Architectural Overview.....	5
eEVOS features and benefits.....	6
High-Availability cluster components and features.....	9
Hardware recommendations	11
Performance recommendations.....	12
Limitations.....	12

Introduction

Virtualization has become extremely important in many IT environments. It allows multiple operating systems to share the same hardware.

This does not only bring enormous cost savings and consolidation but also an increased need for proper disaster recovery scenarios.

In the classic IT environment, you would have dedicated server for each operating system and application. Virtualization enables you to put several operating systems and run multiple different applications on the same server.

This document describes how to use and build a highly-available environment for your VMs that ensures business continuity even in the event of a total server failure.

euroNAS Enterprise Virtualization OS (eEVOS) provides you a powerful, scalable high-availability solution that is both affordable and easy to deploy. It brings to small and midsize companies the enterprise functionality, usually affordable only to large companies.

eEVOS does this by combining the best of both worlds. It provides storage proven in thousands of installations and easy to manage virtualization based on high performance KVM. Integrated VM-Backup software makes sure that your VMs are safe even in case of some catastrophic event.

Why you should use high-availability environment?

In the non-virtualized environment, if one server would crash – only one application would go down (for example mail server, website, database etc.). Today, if the server breaks, the danger of total disaster where all applications will die at once is much bigger.

It is more than ever important to make sure that your virtual machines continue to run even in the case of total disaster. On the market there are many virtualization solutions providing high speeds and redundancy. They can protect you with redundant power supply, UPS or RAID.

eEVOS goes a step further by protecting you even from a total failure of one of the servers and also enabling you to simultaneously keep the same data on 2 different places.

High-Availability will protect your data in case of some catastrophic events like fire, earthquake or theft and provide business continuity at the same time.

It will minimize downtime of the virtual machines by monitoring the nodes. Should one node fail, virtual machines will automatically restart on working node.

Advantage of not being locked to specific hardware

eEVOS runs on standard x64 servers. This enables your total freedom in choosing the hardware that fits to your performance, scalability and budget.

Legacy storage manufacturers often use proprietary hardware and an operating system which will only run on the hardware in question. This makes it difficult to expand in the future and forces you to buy new units. euroNAS will grow with the needs of your company. It is not locked to proprietary hardware – this makes it a highly flexible and scalable solution.

By selecting euroNAS storage software you are not only getting highest possible redundancy in case of server failure but also hardware tolerant operating system that runs on any hardware.

Supported Guest Operating Systems

Currently (December 2015) following Guest Operating Systems are supported

Guest Operating System	64-Bit Version	Hardware Acceleration (Virtio)	Guest Driver
Microsoft Windows Server 2012 R2	Yes	Yes	*
Microsoft Windows Server 2012	Yes	Yes	*
Microsoft Windows Server 2008 R2	Yes	Yes	*
Microsoft Windows Server 2008	Yes	Yes	*
Microsoft Windows 10	Yes	Yes	*
Microsoft Windows 8.1	Yes	Yes	*
Microsoft Windows 8.	Yes	Yes	*
Microsoft Windows7	Yes	Yes	*
Red Hat Enterprise Linux 7.x	Yes	Yes	Built-in
Red Hat Enterprise Linux 6.x	Yes	Yes	Built-in
CentOS 7.x	Yes	Yes	Built-in
CentOS 6.x	Yes	Yes	Built-in
Fedora Server 23	Yes	Yes	Built-in
Fedora Workstation 23	Yes	Yes	Built-in
Debian Jessie	Yes	Yes	Built-in
Debian Wheezy	Yes	Yes	Built-in
Ubuntu Server 14.04 LTS	Yes	Yes	Built-in
Ubuntu Desktop 14.04 LTS	Yes	Yes	Built-in
Ubuntu Desktop 15.10	Yes	Yes	Built-in

* Must be installed during the install

Architectural Overview

Enterprise Virtualization OS (eEVOS) is based on following components

euroNAS Storage OS

Award winning euroNAS Storage OS represents one of the most performant and most reliable storage solutions on the market. It provides perfect environment for storing the VMs.

Built-in Virtualization

Virtualization is embedded natively in euroNAS Storage OS. It is based on industry standard KVM Hypervisor. Intuitive web-based management helps you centrally manage your VMs.

Simple web-based management interface

As all euroNAS products, eEVOS is perfect for IT Professionals. You are not only gaining a high available storage solution with continuous replication but also an efficient and simple to use management interface. Whole configuration can be done within minutes.

Monitoring and Statistics

eEVOS will show you extensive information about the server usage and performance.

eEVOS features and benefits

euroNAS Virtualization is a complete storage and virtualization solution with great performance, high-availability, and manageability features.

High-availability

When using standard Storage your network applications depend on this single point of failure.

Whereas many appliances will provide you with redundant controllers and components – most of them will not protect you from a total unit failure.


Even with proper backup policy it takes a while to recover and manual intervention to get server and VMs back online.

When using euroNAS eEVOS your VMs will quickly recover, should a hardware or software fault occur, it is detected through intelligent features and the software automatically starts the VMs on the other server.

Virtual Machines can be shared by mirroring them between the servers or setting both virtualization servers to share the same storage (shared disk cluster).

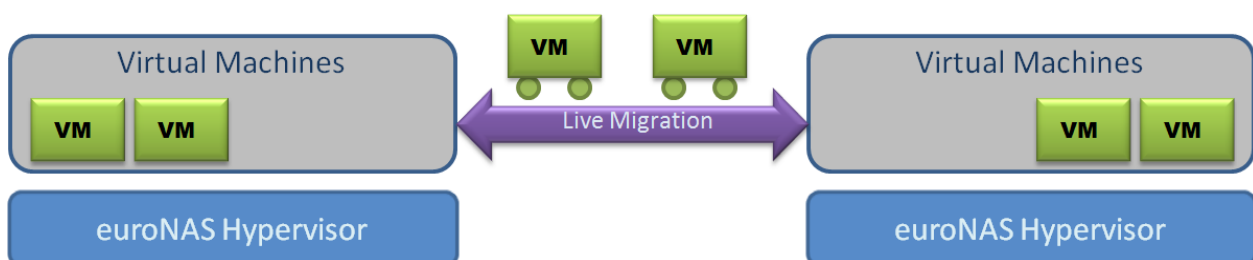
Server Mirror

euroNAS advanced technology creates a cluster of two mirrored servers with real time, continuous data replication and synchronization. Both servers contain identical data securing redundancy in case of server failure.




Move VMs on the fly within seconds to another node

euroNAS Hypervisor helps you to migrate running virtual machines from one cluster node to another one without downtime. VM will keep its connections; RAM and all settings will be moved. Whole migration process takes just several seconds. It is very useful functionality if you wish to shut down or reboot one of the nodes for maintenance or upgrade.



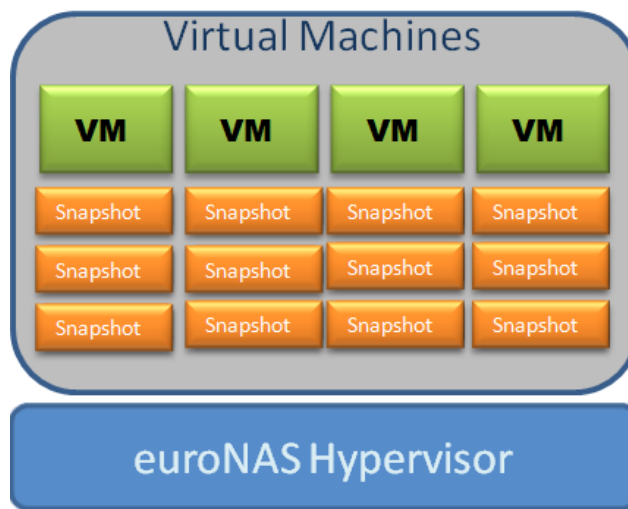
VM Failover in case of node failure

Should a hardware or software fault occur, it is detected through intelligent features. Functioning node will automatically take over and auto-restart the VMs.



Cluster Aware VM Snapshots

VM Snapshot work on any node even if the file is used and opened on another node. This allows you to create snapshots of your VMs and run these snapshot VMs on a different node than the original VM.



Intelligent VM Cloning

Create fully independent clones of your original virtual machine at its current state.

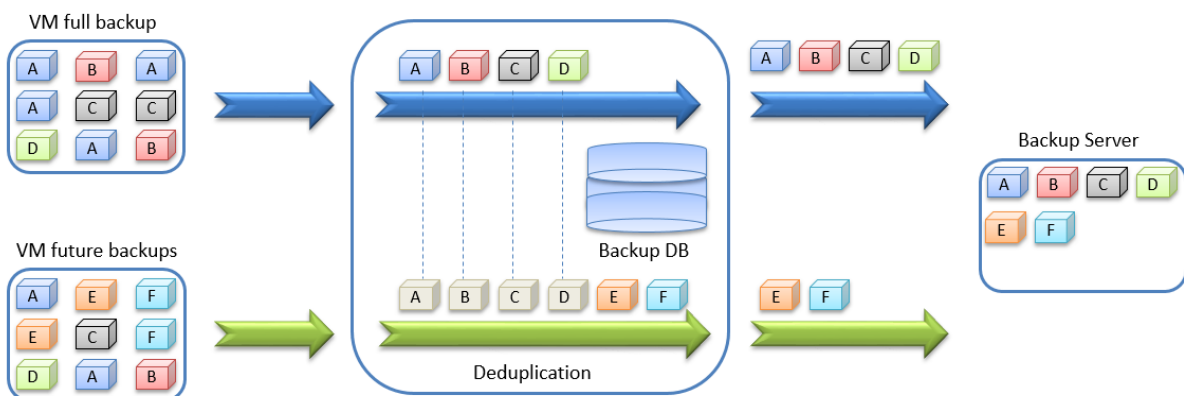


Thanks to the “copy-on-write” file system clones will use original VM as a reference and will require much less disk space to store.

Deduplicated VM Backups

Backup your VMs reliable and scheduled without downtime. Support for Volume Shadow Copy Service (VSS) cares for the best possible data consistency of your Windows VMs. With this you can also copy files that would be usually blocked during write I/O (for example Registry or system data bases).


As target you can use a standard NAS server. VMs are backed up using deduplication. The greatest advantage is that double data blocks will only be copied only once. This ensures the most effective use of your backup server. Checksum check makes sure that backup is consistent. Backup is not tied to the existing installation – you can restore the backup on any eEVOS installation.



euroNAS “Hybrid Storage Cache” helps you to speed up your VMs

euroNAS “Hybrid Storage Cache” improves the performance of your storage server by dynamically moving frequently used data to a faster, smaller device such as a SSD.

This way you can build a hybrid virtualization server that will provide you with higher I/O performance and larger capacity at the same time.



euroNAS eEVOS components and features

euroNAS consists out of following parts

Cluster node

Cluster node represents individual server that is member of the cluster. All nodes are equal – from each node you can monitor and configure cluster services.

Cluster mirror

Cluster mirror represents a mirror of 2 drives on each individual server. This 2 drives can be a single disk, software or raid array. They are mirrored in real-time. When creating cluster drive you have the option to create a mirror.

Shared Storage

Shared Storage represents an external storage that is being accessed by all nodes simultaneously. It can be connected via iSCSI, FC or SAS. It provides higher I/O performance than cluster mirror but data is not being mirrored (lower redundancy).

Cluster drive

On this drives virtual machines (VMs) are installed.

This drives don't necessary need to use full size of the available space. It is possible to create multiple cluster drives on the same disk or RAID array.

Network Test IP

The purpose of this IP(s) is to prevent storage desynchronization, also known as “split brain” situations. Such situation may occur when both nodes are running, but one or both are disconnected from the network.

HA Cluster nodes will ping this address constantly and check if still is able to reach the network. It is recommended to define more than just one IP address.

We recommend any IP address of some device that is constantly on. Most customer use switch, router, gateway, printer, mail server etc.

If one of the nodes is unable to ping this IP it will automatically know that it cannot reach the network and push all virtual machines to other node.

If a node cannot reach all defined test IP addresses it will shut down services to make sure that data remains consistent.

Virtual machines (VMs)

Virtual machines (VMs) are installed a cluster drive and are active on both systems. A full list of supported guest operating systems can be found on www.euronas.com

Replication network

Internal network used for heartbeat and data replication between the nodes.

Network for guest OS (VMs)

VMs can be connected to the networking using following methods

Bridge

Bridge can be configured in network configuration. VMs are directly passed through the NICs defined in the bridge. It is possible to create more than one bridge. This way it is possible to create dedicated networks for individual VMs.

NAT

In NAT the VMs are connected through their own private network. They will not be accessible from the outside world. They are however able to reach outside world (internet).

Hardware recommendations

CPU

For best performance multicore CPU or several CPUs should be used. For the best performance we recommend Intel XEON or AMD Opteron CPUs. CPUs must support hardware-assisted virtualization (VT-x and AMD-V). All nodes should have identical CPUs in order to support the move of the VMs from one node to other.

Installation Disk

For best performance the Enterprise Virtualization OS should be installed on an SSD with a minimum size of 20 GB. Installation disk is kept separate from data – it is not possible to install VMs on it.

Data Disk

Since installation disk is exclusively used for the OS, for data at least one additional disk is needed. Data disk can be simple disk or disks managed by the hardware RAID controller.

System RAM

The server should have at least 16 GB of RAM

Networking (cluster high-availability)

There must be one or more reliable IP Addresses on the Corporate Network that both servers can reach for proper operation of the High-Availability/Cluster features

An example would be the router which is on the Corporate Network or connects the servers to the Corporate Network

For best performance and high availability the server should have 4 individual Network Interface Cards with a minimum bandwidth of 1 Gigabit (10 Gigabit recommended). They should be configured as two separate bonded networks. One of the bonds should be connected to the Corporate Network. The other bond should be connected directly to the other server on a private Replication Network

If the servers are in the same location they should be cabled directly to each other without any switches between them if possible.

For top performance the Replication Network should be 10 Gigabit

If dual-port cards are used each port on the card should be assigned to a separate bond so that the failure of one card does not break the entire bond

DHCP should not be used unless static leases are provided

Using DHCP without static leases will cause the cluster to fail on IP address change

Performance recommendations

For the most optimal performance we recommend use of server CPUs such as Intel Xeon or AMD Opteron.

First generation of such CPUs has only supported CPU virtualization, second has additionally support for memory management unit (MMU) virtualization. For the best performance we recommend using second-generation or higher. The bottleneck of euroNAS Cluster is the disk I/O and the speed of the Ethernet controller. We recommend the use of SAS or SSD disks and hardware RAID controller.

We strongly recommend using identical hardware (especially CPU). Due to server replication in real time, the performance will depend on the weakest link. Using slower disks on one server will slow the server with faster disks as well.

Limitations

eEVOS detects automatically many possible fail scenarios and moves resources automatically. There are some scenarios where cluster service is unable to assure business continuity.

CPUs must be identical on all nodes

If the CPUs are not identical, switchover of the VMs to another node without downtime can lead to the crash of the VM

Replication network failure – corporate network available from both nodes

In this case server can reach network test IP from both nodes so both nodes think that they are online and available. On the other hand, replication network is unavailable. In order to prevent possible data loss server replication and resources are halted until replication network becomes available again. In order to prevent this we recommend network card teaming (port failover)

Powering down active node during sync process

During the sync process the node with more current data must not be powered off. In this case server replication service will stop the drives in order to prevent data corruption.

Only server with less current data can be powered off. This does apply if both nodes are in sync. In this case either of the servers can be powered off or rebooted.

When one node immediately fails, VMs will be automatically restarted on working node

eEVOS is capable of moving VMs on a running system without downtime. However, in case of total node failure there is no time left to move the RAM memory to other node. This means that VM will be freshly started on the running node. Content of RAM Memory will be lost.