



Virtuozzo Hybrid Server 7

Installation Guide

January 22, 2025

Virtuozzo International GmbH

Vordergasse 59

8200 Schaffhausen

Switzerland

Tel: + 41 52 632 0411

Fax: + 41 52 672 2010

<https://virtuozzo.com>

Copyright ©2016-2025 Virtuozzo International GmbH. All rights reserved.

This product is protected by United States and international copyright laws. The product's underlying technology, patents, and trademarks are listed at <https://www.virtuozzo.com/legal.html>.

Microsoft, Windows, Windows Server, Windows NT, Windows Vista, and MS-DOS are registered trademarks of Microsoft Corporation.

Apple, Mac, the Mac logo, Mac OS, iPad, iPhone, iPod touch, FaceTime HD camera and iSight are trademarks of Apple Inc., registered in the US and other countries.

Linux is a registered trademark of Linus Torvalds. All other marks and names mentioned herein may be trademarks of their respective owners.

Contents

1. Preparing for Installation	1
1.1 Requirements for Standalone Installations	1
1.1.1 System Limits	2
1.2 Planning Infrastructure for Virtuozzo Storage with GUI Management	2
1.2.1 Understanding Virtuozzo Storage Architecture	2
1.2.1.1 Storage Role	3
1.2.1.2 Metadata Role	3
1.2.1.3 Network Roles (Storage Access Points)	4
1.2.1.4 Supplementary Roles	4
1.2.2 Planning Node Hardware Configurations	5
1.2.2.1 Hardware Requirements	5
1.2.2.2 Hardware Recommendations	6
1.2.2.3 Hardware and Software Limitations	9
1.2.2.4 Minimum Configuration	10
1.2.2.5 Recommended Configuration	11
1.2.2.6 Raw Disk Space Considerations	15
1.2.3 Planning Network	16
1.2.3.1 General Network Requirements	16
1.2.3.2 Network Limitations	17
1.2.3.3 Per-Node Network Requirements	17
1.2.3.4 Network Interface Roles	21
1.2.3.5 Network Recommendations for Clients	22
1.2.3.6 Sample Network Configuration	23
1.2.4 Understanding Data Redundancy	24
1.2.4.1 Redundancy by Replication	26
1.2.4.2 Redundancy by Erasure Coding	28

1.2.4.3	No Redundancy	29
1.2.5	Understanding Failure Domains	29
1.2.6	Understanding Storage Tiers	29
1.3	Planning Infrastructure for Virtuozzo Storage with CLI Management	30
1.3.1	Understanding Virtuozzo Storage Architecture	30
1.3.2	Virtuozzo Storage Configurations	32
1.3.2.1	Minimum Configuration	33
1.3.2.2	Recommended Configuration	34
1.3.3	Hardware Requirements	35
1.3.4	Network Requirements	37
1.4	Obtaining the Distribution Image	38
1.5	Preparing for Installation from USB Storage Drives	38
2.	Installing Virtuozzo Hybrid Server	41
2.1	Starting Installation	41
2.1.1	Choosing Installation Type	42
2.1.2	Enabling Forced Detection of SSDs	42
2.1.3	Changing the Swap Size Calculation Method	43
2.1.4	Installing on USB Drive	43
2.2	Installing with GUI Management	44
2.2.1	Accepting EULA	45
2.2.2	Setting Date and Time	46
2.2.3	Selecting the Keyboard Layout	47
2.2.4	Configuring Network	48
2.2.4.1	Creating Bonded and Teamed Connections	49
2.2.5	Installing Virtuozzo Automator Components	51
2.2.5.1	Installing Management Panel and Compute on the First Server	52
2.2.5.2	Installing Compute on the Second and Other Servers	53
2.2.6	Installing Virtuozzo Storage Components	54
2.2.6.1	Installing Management Panel and Storage on the First Server	55
2.2.6.2	Installing Storage on the Second and Other Servers	56
2.2.7	Choosing the Installation Destination	57
2.2.8	Finishing the Installation	58
2.2.8.1	Activating a License	58
2.2.8.2	Configuring ReadyKernel	59
2.2.8.3	Completing Installation	60

2.2.8.4	Building the Infrastructure	60
2.3	Installing with CLI Management	61
2.3.1	Accepting EULA	61
2.3.2	Setting Date and Time	62
2.3.3	Selecting the Keyboard Layout	63
2.3.4	Configuring Network	64
2.3.4.1	Creating Bonded and Teamed Connections	65
2.3.4.2	Creating VLAN Adapters	69
2.3.5	Choosing the Storage Type	71
2.3.6	Setting Virtuozzo Storage Installation Options	71
2.3.6.1	Creating a New Virtuozzo Storage Cluster	72
2.3.6.2	Joining an Existing Virtuozzo Storage Cluster	74
2.3.6.3	Virtuozzo Storage Server Roles	75
2.3.7	Partitioning the Hard Drives	78
2.3.8	Finishing the Installation	82
2.3.8.1	Activating a License	83
2.3.8.2	Configuring ReadyKernel	84
2.3.8.3	Completing Installation	84
2.4	Installing in the Basic Graphics Mode	85
2.5	Installing via VNC	85
2.6	Installing in the Text Mode	85
2.7	Configuring Server Ports	86
2.7.1	Opened Ports on Standalone Servers	86
2.7.2	Opened Ports on Servers in Virtuozzo Storage Clusters	87
3.	Exploring Additional Installation Options	89
3.1	Booting into the Rescue Mode	89
3.2	Enabling the Expert Mode	90
3.3	Running Virtuozzo Hybrid Server in Virtual Machines	90
3.3.1	Restrictions and Peculiarities	91

CHAPTER 1

Preparing for Installation

This chapter describes how to plan the deployment of Virtuozzo Storage with GUI and CLI management. It also explains how to obtain the Virtuozzo Hybrid Server distribution image and prepare for installing Virtuozzo Hybrid Server from a USB flash drive.

1.1 Requirements for Standalone Installations

The recommended hardware requirements for running Virtuozzo Hybrid Server 7 as a standalone installation are as follows:

- x86-64 platform with hardware virtualization support: Intel VT-x (with “unrestricted guest”) or AMD Virtualization (tested only on AMD EPYC).

Note: To check if the Intel processor supports the “unrestricted guest” feature: 1) Download `vmxcap.py` from <https://github.com/qemu/qemu/blob/master/scripts/kvm/vmxcap>, 2) Run `python vmxcap.py | grep -i unrest`. The result must be `yes`.

- CPUs: at least 4 cores, a 64-bit processor is required for running 64-bit guest operating systems.
- RAM: 4 GB or more.
- HDD: the minimal required disk size is the sum of the minimal `/vz` and root partitions, as well as the swap size:

```
min_hdd_size = min_vz_size + min_root_size + swap_size
```

Where `min_root_size` is 12 GiB, `min_vz_size` is 30 GiB, and `swap_size` depends on the RAM size. See

Partitioning the Hard Drives on page 78 for more details.

- SSD (optional): at least 30 GiB (at least 32 GiB with /boot).
- Network: an Ethernet network adapter and a valid IP address.

The actual number of virtual machines and containers you can run on a physical server and their performance depend on resources they will consume.

1.1.1 System Limits

The table below lists the current hardware limits for Virtuozzo Hybrid Server 7 servers:

Hardware	Theoretical	Certified
RAM	64 TB	1 TB
HDD	1 EB	50 TB

1.2 Planning Infrastructure for Virtuozzo Storage with GUI Management

To plan your infrastructure for Virtuozzo Storage managed via the web-based management panel, you will need to decide on the hardware configuration of each storage node, plan the storage networks, decide on the redundancy method (and mode) to use, and decide which data will be kept on which storage tier.

Information in this section is meant to help you complete all of these tasks.

1.2.1 Understanding Virtuozzo Storage Architecture

The fundamental component of Virtuozzo Storage is a cluster: a group of physical servers interconnected by network. Each server in a cluster is assigned one or more roles and typically runs services that correspond to these roles:

- Storage role: chunk service or CS
- Metadata role: metadata service or MDS
- Network roles:

- iSCSI access point service (iSCSI)
- S3 gateway (access point) service (GW)
- S3 name service (NS)
- S3 object service (OS)
- Web CP
- SSH
- Supplementary roles:
 - management
 - SSD cache
 - system

Any server in the cluster can be assigned a combination of storage, metadata, and network roles. For example, a single server can be an S3 access point, an iSCSI access point, and a storage node at once.

Each cluster also requires that a web-based management panel be installed on one (and only one) of the nodes. The panel enables administrators to manage the cluster.

1.2.1.1 Storage Role

Storage nodes run chunk services, store all the data in the form of fixed-size chunks, and provide access to these chunks. All data chunks are replicated and the replicas are kept on different storage nodes to achieve high availability of data. If one of the storage nodes fails, remaining healthy storage nodes continue providing the data chunks that were stored on the failed node.

Only a server with disks of certain capacity can be assigned the storage role (see *Hardware Requirements* on page 5).

1.2.1.2 Metadata Role

Metadata nodes run metadata services, store cluster metadata, and control how user files are split into chunks and where these chunks are located. Metadata nodes also ensure that chunks have the required amount of replicas and log all important events that happen in the cluster.

To ensure high availability of metadata, at least five metadata services must be running per cluster. In this

case, if up to two metadata service fail, the remaining metadata services will still be controlling the cluster.

1.2.1.3 Network Roles (Storage Access Points)

Storage access points enable you to access data stored in Virtuozzo Storage clusters via the standard iSCSI and S3 protocols.

To benefit from high availability, access points should be set up on multiple node.

The following access points are currently supported:

- **iSCSI**, allows you to use Virtuozzo Storage as a highly available block storage for virtualization, databases, office applications, and other needs.
- **S3**, a combination of scalable and highly available services (collectively named Virtuozzo Object Storage) that allows you to use Virtuozzo Storage as a modern backend for solutions like OpenXchange AppSuite, and Dovecot. In addition, to developers of custom applications Virtuozzo Object Storage offers an Amazon S3-compatible API and compatibility with the S3 libraries for various programming languages, S3 browsers, and web browsers.

NFS, SMB, and other access point types are planned in the future releases of Virtuozzo Storage.

The following remote management roles are supported:

- **Web CP**, allows you to access the web-based user interface from an external network.
- **SSH**, allows you to connect to Virtuozzo Storage nodes via SSH.

1.2.1.4 Supplementary Roles

- **Management**, provides a web-based management panel that enables administrators to configure, manage, and monitor Virtuozzo Storage clusters. Only one management panel is needed to create and manage multiple clusters (and only one is allowed per cluster).
- **SSD cache**, boosts chunk write performance by creating write caches on selected solid-state drives (SSDs). It is recommended to also use such SSDs for metadata, see [Metadata Role](#) on page 3. The use of write cache may speed up write operations in the cluster by two and more times.
- **System**, one disk per node that is reserved for the operating system and unavailable for data storage.

1.2.2 Planning Node Hardware Configurations

Virtuozzo Storage works on top of commodity hardware, so you can create a cluster from regular servers, disks, and network cards. Still, to achieve the optimal performance, a number of requirements must be met and a number of recommendations should be followed.

1.2.2.1 Hardware Requirements

The following table lists the minimal and recommended hardware for a single node in the cluster:

Type	Minimal	Recommended
CPU	Dual-core CPU	Intel Xeon E5-2620V2 or faster; at least one CPU core per 8 HDDs
RAM	4GB	16GB ECC or more, plus 0.5GB ECC per each HDD
Storage	System: Approximately 100GB SATA HDD (see Partitioning the Hard Drives on page 78 for more details on the value that depends on the RAM size) Metadata: 100GB SATA HDD (on the first five nodes in the cluster) Storage: 100GB SATA HDD	System: Approximately 250GB SATA HDD (see Partitioning the Hard Drives on page 78 for more details on the value that depends on the RAM size) Metadata+Cache: One or more recommended enterprise-grade SSDs with power loss protection; 100GB or more capacity; and 75 MB/s sequential write performance per serviced HDD. For example, a node with 10 HDDs will need an SSD with at least 750 MB/s sequential write speed (on the first five nodes in the cluster) Storage: Four or more HDDs or SSDs; 1 DWPD endurance minimum, 10 DWPD recommended
Disk controller	None	HBA or RAID

Continued on next page

Table 1.2.2.1.1 -- continued from previous page

Type	Minimal	Recommended
Network	1Gbps or faster network interface	Two 10Gbps network interfaces; dedicated links for internal and public networks
Sample configuration		Intel Xeon E5-2620V2, 32GB, 2xST1000NM0033, 32xST6000NM0024, 2xMegaRAID SAS 9271/9201, Intel X540-T2, Intel P3700 800GB

1.2.2.2 Hardware Recommendations

The following recommendations explain the benefits added by specific hardware in the hardware requirements table and are meant to help you configure the cluster hardware in an optimal way:

General hardware recommendations:

- At least five nodes are required for a production environment. This is to ensure that the cluster can survive failure of two nodes without data loss.
- One of the strongest features of Virtuozzo Storage is scalability. The bigger the cluster, the better Virtuozzo Storage performs. It is recommended to create production clusters from at least ten nodes for improved resiliency, performance, and fault tolerance in production scenarios.
- Even though a cluster can be created on top of varied hardware, using nodes with similar hardware in each node will yield better cluster performance, capacity, and overall balance.
- Any cluster infrastructure must be tested extensively before it is deployed to production. Such common points of failure as SSD drives and network adapter bonds must always be thoroughly verified.
- It is not recommend for production to run Virtuozzo Storage on top of SAN/NAS hardware that has its own redundancy mechanisms. Doing so may negatively affect performance and data availability.
- At least 20% of cluster capacity should be free to avoid possible data fragmentation and performance degradation.
- During disaster recovery, Virtuozzo Storage may need additional disk space for replication. Make sure to reserve at least as much space as available on a single storage node.

Storage hardware recommendations:

- Using the recommended SSD models may help you avoid loss of data. Not all SSD drives can withstand enterprise workloads and may break down in the first months of operation, resulting in TCO spikes.
 - SSD memory cells can withstand a limited number of rewrites. An SSD drive should be viewed as a consumable that you will need to replace after a certain time. Consumer-grade SSD drives can withstand a very low number of rewrites (so low, in fact, that these numbers are not shown in their technical specifications). SSD drives intended for Virtuozzo Storage clusters must offer at least 1 DWPD endurance (10 DWPD is recommended). The higher the endurance, the less often SSDs will need to be replaced, improving TCO.
 - Many consumer-grade SSD drives can ignore disk flushes and falsely report to operating systems that data was written while it in fact was not. Examples of such drives include OCZ Vertex 3, Intel 520, Intel X25-E, and Intel X-25-M G2. These drives are known to be unsafe in terms of data commits, they should not be used with databases, and they may easily corrupt the file system in case of a power failure. For these reasons, use to enterprise-grade SSD drives that obey the flush rules (for more information, see <http://www.postgresql.org/docs/current/static/wal-reliability.html>). Enterprise-grade SSD drives that operate correctly usually have the power loss protection property in their technical specification. Some of the market names for this technology are Enhanced Power Loss Data Protection (Intel), Cache Power Protection (Samsung), Power-Failure Support (Kingston), Complete Power Fail Protection (OCZ).
 - Consumer-grade SSD drives usually have unstable performance and are not suited to withstand sustainable enterprise workloads. For this reason, pay attention to sustainable load tests when choosing SSDs.
- The use of SSDs for write caching improves random I/O performance and is highly recommended for all workloads with heavy random access (e.g., iSCSI volumes).
- Running metadata services on SSDs improves cluster performance. To also minimize CAPEX, the same SSDs can be used for write caching.
- If capacity is the main goal and you need to store non-frequently accessed data, choose SATA disks over SAS ones. If performance is the main goal, choose SAS disks over SATA ones.
- The more disks per node the lower the CAPEX. As an example, a cluster created from ten nodes with two disks in each will be less expensive than a cluster created from twenty nodes with one disk in each.
- Using SATA HDDs with one SSD for caching is more cost effective than using only SAS HDDs without such an SSD.

- Use HBA controllers as they are less expensive and easier to manage than RAID controllers.
- Disable all RAID controller caches for SSD drives. Modern SSDs have good performance that can be reduced by a RAID controller's write and read cache. It is recommend to disable caching for SSD drives and leave it enabled only for HDD drives.
- If you use RAID controllers, do not create RAID volumes from HDDs intended for storage (you can still do so for system disks). Each storage HDD needs to be recognized by Virtuozzo Storage as a separate device.
- If you use RAID controllers with caching, equip them with backup battery units (BBUs) to protect against cache loss during power outages.
- Make sure that the sum of all SSD and HDD speeds does not exceed the speed of the PCI bus to which the HBA/RAID controller is attached.
- If one of the HDDs is slower than the rest, it will hold back the entire storage cluster. The cluster will work as fast as its slowest drive.

Network hardware recommendations:

- Use separate networks (and, ideally albeit optionally, separate network adapters) for internal and public traffic. Doing so will prevent public traffic from affecting cluster I/O performance and also prevent possible denial-of-service attacks from the outside.
- Network latency dramatically reduces cluster performance. Use quality network equipment with low latency links. Do not use consumer-grade network switches.
- Do not use desktop network adapters like Intel EXPI9301CTBLK or Realtek 8129 as they are not designed for heavy load and may not support full-duplex links. Also use non-blocking Ethernet switches.
- To avoid intrusions, Virtuozzo Storage should be on a dedicated internal network inaccessible from outside.
- Use one 1 Gbit/s link per each two HDDs on the node (rounded up). For one or two HDDs on a node, two bonded network interfaces are still recommended for high network availability. The reason for this recommendation is that 1 Gbit/s Ethernet networks can deliver 110-120 MB/s of throughput, which is close to sequential I/O performance of a single disk. Since several disks on a server can deliver higher throughput than a single 1 Gbit/s Ethernet link, networking may become a bottleneck.
- For maximum sequential I/O performance, use one 1Gbit/s link per each hard drive, or one 10Gbit/s link per node. Even though I/O operations are most often random in real-life scenarios, sequential I/O is important in backup scenarios.

- For maximum overall performance, use one 10 Gbit/s link per node (or two bonded for high network availability).
- It is not recommended to configure 1 Gbit/s network adapters to use non-default MTUs (e.g., 9000-byte jumbo frames). Such settings require additional configuration of switches and often lead to human error. 10 Gbit/s network adapters, on the other hand, need to be configured to use jumbo frames to achieve full performance.

1.2.2.3 Hardware and Software Limitations

Hardware limitations:

- Each physical server must have at least two disks with the assigned three roles: System, Metadata, and Storage. The System role can be combined with the Metadata or Storage role, if the system disk capacity is greater than 100GB.

Note:

1. It is recommended to assign the System+Metadata role to an SSD. Assigning both these roles to an HDD will result in mediocre performance suitable only for cold data (e.g., archiving).
 2. The System role cannot be combined with the Cache and Metadata+Cache roles. The reason is that is I/O generated by the operating system and applications would contend with I/O generated by journaling, negating its performance benefits.
-

- Five servers are required to test all the features of the product.
- The system disk must have at least 100 GBs of space.
- The maximum supported physical partition size is 254 TiB.

Software limitations:

- The default storage mount point must not be changed if the GUI is used.
- The maintenance mode is not supported. Use SSH to shut down or reboot a node.
- One node can be a part of only one cluster.
- Only one S3 cluster can be created on top of a storage cluster.
- Only predefined redundancy modes are available in the management panel.

- Thin provisioning is always enabled for all data and cannot be configured otherwise.

Note: For network limitations, see *Network Limitations* on page 17.

1.2.2.4 Minimum Configuration

The minimum configuration described in the table will let you evaluate Virtuozzo Storage features:

Node #	1st disk role	2nd disk role	3rd and other disk roles	Access points
1	System	Metadata	Storage	iSCSI, S3 (private and public)
2	System	Metadata	Storage	iSCSI, S3 (private and public)
3	System	Metadata	Storage	iSCSI, S3 (private and public)
4	System	Metadata	Storage	iSCSI, S3 (private)
5	System	Metadata	Storage	iSCSI, S3 (private)
5 nodes in total		5 MDSs in total	5 or more CSs in total	Access point services run on five nodes in total

Note: SSD disks can be assigned metadata and cache roles at the same time, freeing up one more disk for the storage role.

Even though five nodes are recommended even for the minimal configuration, you can start evaluating Virtuozzo Storage with just one node and add more nodes later. At the very least, a Virtuozzo Storage cluster must have one metadata service and one chunk service running. However, such a configuration will have two key limitations:

1. Just one MDS will be a single point of failure. If it fails, the entire cluster will stop working.
2. Just one CS will be able to store just one chunk replica. If it fails, the data will be lost.

1.2.2.5 Recommended Configuration

The recommended configuration will help you create clusters for production environments:

Node #	1st disk role	2nd disk role	3rd and other disk roles	Access points
Nodes 1 to 5	System	SSD; metadata, cache	Storage	iSCSI, S3 (private and public)
Nodes 6+	System	SSD; cache	Storage	iSCSI, S3 (private)
5 or more nodes in total		5 MDSs in total	5 or more CSs	All nodes run required access points

Even though a production-ready cluster can be created from just five nodes with recommended hardware, it is still recommended to enter production with at least ten nodes if you are aiming to achieve significant performance advantages over direct-attached storage (DAS) or improved recovery times.

Important: To ensure high availability of metadata, at least five metadata services must be running per cluster in any production environment. In this case, if up to two metadata service fail, the remaining metadata services will still be controlling the cluster.

Following are a number of more specific configuration examples that can be used in production. Each configuration can be extended by adding chunk servers and nodes.

HDD Only

This basic configuration requires a dedicated disk for each metadata server.

Nodes 1-5 (base):

Disk No.	Disk Type	Disk Role(s)
1	HDD	System
2	HDD	MDS
3	HDD	CS
...		
N	HDD	CS

Nodes 6+ (extension):

Disk No.	Disk Type	Disk Role(s)
1	HDD	System
2	HDD	CS
3	HDD	CS
...		
N	HDD	CS

HDD + System SSD (No Cache)

This configuration is good for creating capacity-oriented clusters.

Nodes 1-5 (base):

Disk No.	Disk Type	Disk Role(s)
1	SSD	System, MDS
2	HDD	CS
3	HDD	CS
...		
N	HDD	CS

Nodes 6+ (extension):

Disk No.	Disk Type	Disk Role(s)
1	SSD	System
2	HDD	CS
3	HDD	CS
...		
N	HDD	CS

HDD + SSD

This configuration is good for creating performance-oriented clusters.

Nodes 1-5 (base):

Disk No.	Disk Type	Disk Role(s)
1	HDD	System
2	SSD	MDS, cache
3	HDD	CS
...		
N	HDD	CS

Nodes 6+ (extension):

Disk No.	Disk Type	Disk Role(s)
1	HDD	System
2	SSD	Cache
3	HDD	CS
...		
N	HDD	CS

SSD Only

This configuration does not require SSDs for cache.

Nodes 1-5 (base):

Disk No.	Disk Type	Disk Role(s)
1	SSD	System, MDS
2	SSD	CS
3	SSD	CS
...		
N	SSD	CS

Nodes 6+ (extension):

Disk No.	Disk Type	Disk Role(s)
1	SSD	System
2	SSD	CS
3	SSD	CS
...		

Continued on next page

Table 1.2.2.5.9 -- continued from previous page

Disk No.	Disk Type	Disk Role(s)
N	SSD	CS

HDD + SSD (No Cache), 2 Tiers

In this configuration example, tier 1 is for HDDs without cache and tier 2 is for SSDs. Tier 1 can store cold data (e.g., backups), tier 2 can store hot data (e.g., high-performance virtual machines).

Nodes 1-5 (base):

Disk No.	Disk Type	Disk Role(s)	Tier
1	SSD	System, MDS	
2	HDD	CS	1
3	SSD	CS	2
...			
N	HDD/SSD	CS	1/2

Nodes 6+ (extension):

Disk No.	Disk Type	Disk Role(s)	Tier
1	SSD	System	
2	HDD	CS	1
3	SSD	CS	2
...			
N	HDD/SSD	CS	1/2

HDD + SSD, 3 Tiers

In this configuration example, tier 1 is for HDDs without cache, tier 2 is for HDDs with cache, and tier 3 is for SSDs. Tier 1 can store cold data (e.g., backups), tier 2 can store regular virtual machines, and tier 3 can store high-performance virtual machines.

Nodes 1-5 (base):

Disk No.	Disk Type	Disk Role(s)	Tier
1	HDD/SSD	System	

Continued on next page

Table 1.2.2.5.12 -- continued from previous page

Disk No.	Disk Type	Disk Role(s)	Tier
2	SSD	MDS, T2 cache	
3	HDD	CS	1
4	HDD	CS	2
5	SSD	CS	3
...			
N	HDD/SSD	CS	1/2/3

Nodes 6+ (extension):

Disk No.	Disk Type	Disk Role(s)	Tier
1	HDD/SSD	System	
2	SSD	T2 cache	
3	HDD	CS	1
4	HDD	CS	2
5	SSD	CS	3
...			
N	HDD/SSD	CS	1/2/3

1.2.2.6 Raw Disk Space Considerations

When planning the Virtuozzo Storage infrastructure, keep in mind the following to avoid confusion:


- The capacity of HDD and SSD is measured and specified with decimal, not binary prefixes, so “TB” in disk specifications usually means “terabyte”. The operating system, however, displays drive capacity using binary prefixes meaning that “TB” is “tebibyte” which is a noticeably larger number. As a result, disks may show capacity smaller than the one marketed by the vendor. For example, a disk with 6TB in specifications may be shown to have 5.45 TB of actual disk space in Virtuozzo Storage.
- Virtuozzo Storage reserves 5% of disk space for emergency needs.

Therefore, if you add a 6TB disk to a cluster, the available physical space should increase by about 5.2 TB.

1.2.3 Planning Network

Virtuozzo Storage uses two networks (e.g., Ethernet): a) an internal network that interconnects nodes and combines them into a cluster, and b) a public network for exporting stored data to users.

The figure below shows a top-level overview of the internal and public networks of Virtuozzo Storage. One network interface on each node is also used for management: through it, administrators can access the node from the management panel and via SSH.



1.2.3.1 General Network Requirements

- Make sure that time is synchronized on all nodes in the cluster via NTP. Doing so will make it easier for the support department to understand cluster logs.

Note: In Virtuozzo Hybrid Server 7, time synchronization via NTP is enabled by default using the

chronyd service. If you want to use ntpdate or ntpd, stop and disable chronyd first.


1.2.3.2 Network Limitations

- Nodes are added to clusters by their IP addresses, not FQDNs. Changing the IP address of a node in the cluster will remove that node from the cluster. If you plan to use DHCP in a cluster, make sure that IP addresses are bound to the MAC addresses of nodes' network interfaces.
- Fibre channel and InfiniBand networks are not supported.
- Each node must have Internet access so updates can be installed.
- MTU is set to 1500 by default.
- Network time synchronization (NTP) is required for correct statistics.
- The management role is assigned automatically during installation and cannot be changed in the management panel later.
- Even though the management node can be accessed from a web browser by the hostname, you still need to specify its IP address, not the hostname, during installation.

1.2.3.3 Per-Node Network Requirements

Network requirements for each cluster node depend on roles assigned to the node. If the node with multiple network interfaces has multiple roles assigned to it, different interfaces can be assigned to different roles to create dedicated networks for each role.

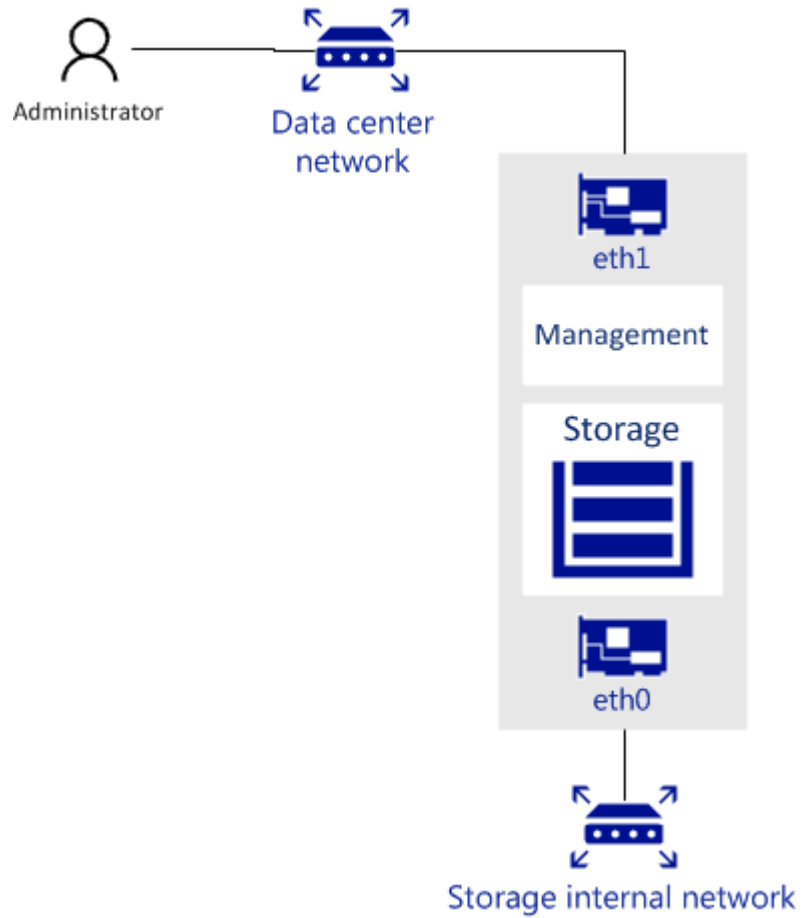
- Each node in the cluster must have access to the internal network and have the port 8889 open to listen for incoming connections from the internal network.
- Each storage and metadata node must have at least one network interface for the internal network traffic. The IP addresses assigned to this interface must be either static or, if DHCP is used, mapped to the adapter's MAC address. The figure below shows a sample network configuration for a storage and metadata node.



- The management node must have a network interface for internal network traffic and a network interface for the public network traffic (e.g., to the datacenter or a public network) so the management panel can be accessed via a web browser.


The following ports need to be open on a management node by default: 8888 for management panel access from the public network and 8889 for cluster node access from the internal network.

The figure below shows a sample network configuration for a storage and management node.

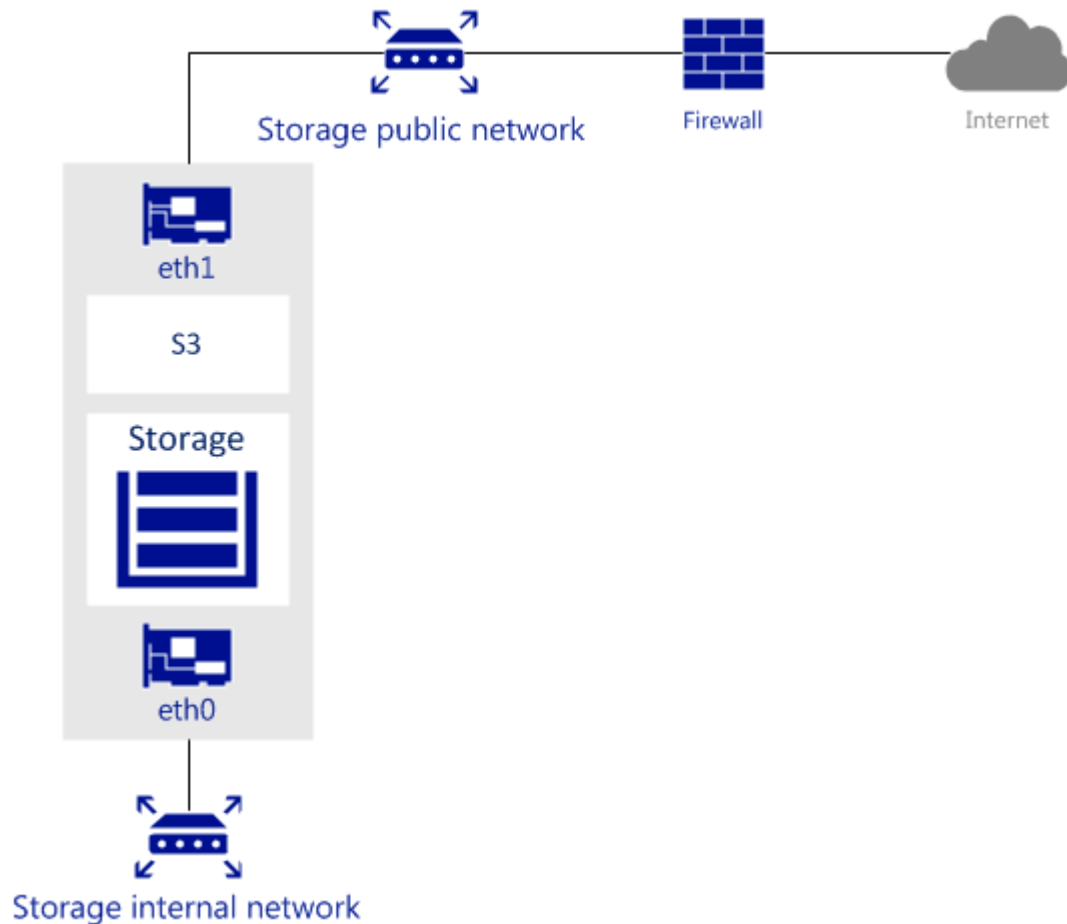


- A node that runs one or more storage access point services must have a network interface for the internal network traffic and a network interface for the public network traffic.

The figure below shows a sample network configuration for a node with an iSCSI access point. iSCSI access points use the TCP port 3260 for incoming connections from the public network.



The next figure shows a sample network configuration for a node with an S3 storage access point. S3 access points use ports 443 (HTTPS) and 80 (HTTP) to listen for incoming connections from the public network.



Note: In the scenario pictured above, the internal network is used for both the storage and S3 cluster traffic.

1.2.3.4 Network Interface Roles

For a Virtuozzo Storage cluster to function, network interfaces of cluster nodes must be assigned one or more roles described below. Assigning roles automatically configures the necessary firewall rules.

- **Internal.** If one or more internal roles are assigned to a network interface, traffic on all ports is allowed to and from said interface.
 - **Management.** The network interface will be used for communication between the nodes and the management panel. To perform this role, the network interface must be connected to the internal network. This role must be assigned to at least one network interface in the cluster.

- **Storage.** The network interface will be used for transferring data chunks between storage nodes. To perform this role, the network interface must be connected to the internal network. This role must be assigned to one network interface on each storage node.
- **Object Storage.** The network interface will be used by the S3 storage access point. To perform this role, the network interface must be connected to the internal network. This role must be assigned to one network interface on each node running the S3 storage access point service.
- **Public.** If one or more public roles (and no internal roles) are assigned to a network interface, only traffic on ports required by the public role(s) is allowed to and from said interface.
 - **iSCSI.** The network interface will be used by the iSCSI storage access point to provide access to user data. To perform this role, the network interface must be connected to the public network accessible by iSCSI clients.
 - **S3 public.** The network interface will be used by the S3 storage access point to provide access to user data. To perform this role, the network interface must be connected to the public network accessible by S3 clients.
 - **Web CP.** The network interface will be used to transfer web-based user interface data. To perform this role, the network interface must be connected to the public network.
 - **SSH.** The network interface will be used to manage the node via SSH. To perform this role, the network interface must be connected to the public network.
- **Custom.** These roles allow you to open specific ports on public network interfaces.

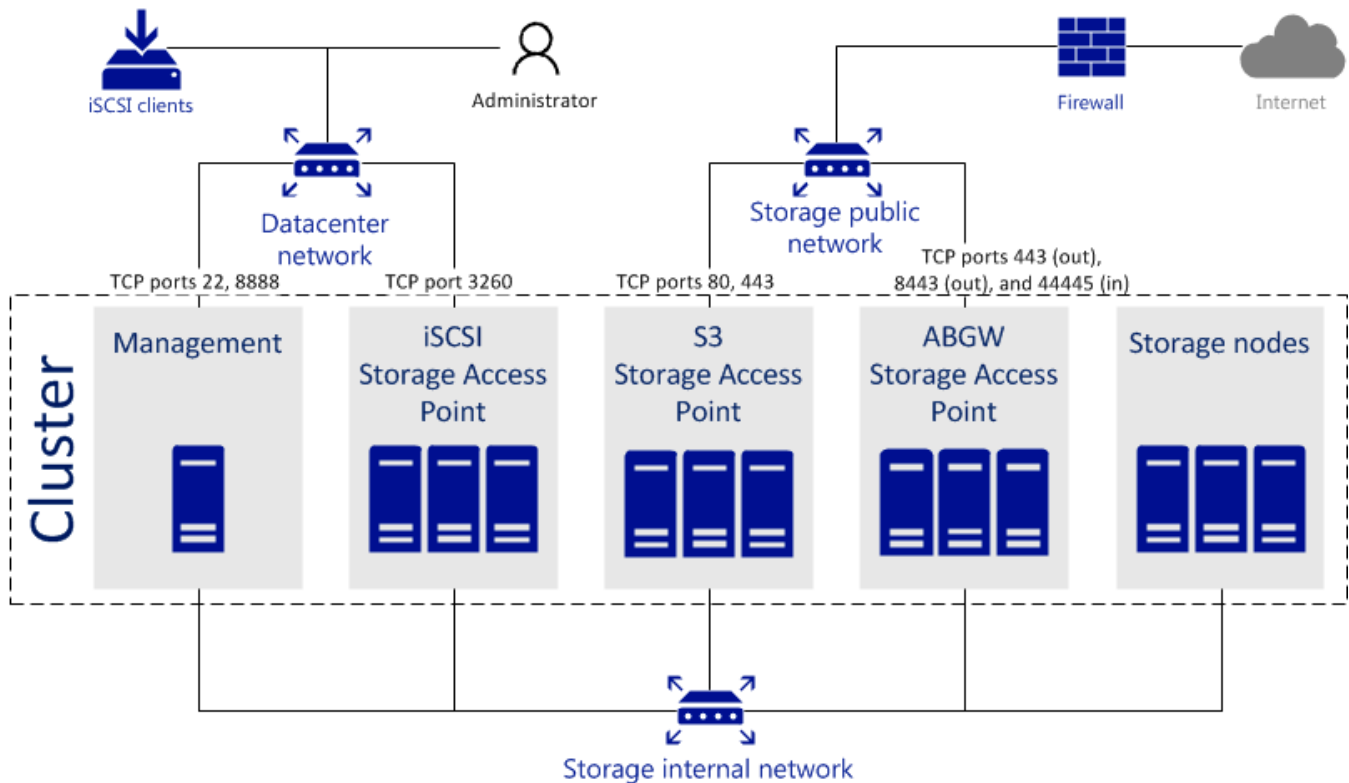
1.2.3.5 Network Recommendations for Clients

The following table lists the maximum network performance a Virtuozzo Storage client can get with the specified network interface. The recommendation for clients is to use 10Gbps network hardware between any two cluster nodes and minimize network latencies, especially if SSD disks are used.

Storage network interface	1Gbps	2 x 1Gbps	3 x 1Gbps	10Gbps	2 x 10Gbps
Entire node maximum I/O throughput	100MB/s	~175MB/s	~250MB/s	1GB/s	1.75GB/s
Single VM maximum I/O throughput (replication)	100MB/s	100MB/s	100MB/s	1GB/s	1GB/s
Single VM maximum I/O throughput (erasure coding)	70MB/s	~130MB/s	~180MB/s	700MB/s	1.3GB/s

1.2.3.6 Sample Network Configuration

The figure below shows an overview of a sample Virtuozzo Storage network.



In this network configuration:

- The Virtuozzo Storage internal network is a network that interconnects all servers in the cluster. It can be used for the management, storage (internal), and S3 (private) roles. Each of these roles can be moved to a separate dedicated internal network to ensure high performance under heavy workloads.

This network cannot be accessed from the public network. All servers in the cluster are connected to this network.

Important: Virtuozzo Storage does not offer protection from traffic sniffing. Anyone with access to the internal network can capture and analyze the data being transmitted.

- The Virtuozzo Storage public network is a network over which the storage space is exported. Depending on where the storage space is exported to, it can be an internal datacenter network or an external public network:

- An internal datacenter network can be used to manage Virtuozzo Storage and export the storage space over iSCSI to other servers in the datacenter, that is, for the management and iSCSI (public) roles.
- An external public network can be used to export the storage space to the outside services through S3 storage access points, that is, for the S3 public role.

1.2.4 Understanding Data Redundancy

Virtuozzo Storage protects every piece of data by making it redundant. It means that copies of each piece of data are stored across different storage nodes to ensure that the data is available even if some of the storage nodes are inaccessible.

Virtuozzo Storage automatically maintains the required number of copies within the cluster and ensures that all the copies are up-to-date. If a storage node becomes inaccessible, the copies from it are replaced by new ones that are distributed among healthy storage nodes. If a storage node becomes accessible again after downtime, the copies on it which are out-of-date are updated.

The redundancy is achieved by one of two methods: replication or erasure coding (explained in more detail in the next section). The chosen method affects the size of one piece of data and the number of its copies that will be maintained in the cluster.

The general rule is as follows:

- Choose replication for highly loaded VMs, Windows VMs, and other workloads that generate lots of IOPS.
- Choose erasure coding for lightly loaded Linux VMs and backups.

Virtuozzo Storage supports a number of modes for each redundancy method. The following table illustrates data overhead of various redundancy modes.

Note: The numbers of nodes listed in the table indicate only the requirements of each redundancy method but not the number of nodes needed for the Virtuozzo Storage cluster. The minimum and recommended cluster configurations are described in *Minimum Configuration* on page 10 and *Recommended Configuration* on page 11, respectively.

Redundancy mode	Nodes required to store data copies	Nodes that can fail without data loss	Storage overhead, %	Raw space required to store 100GB of data
1 replica (no redundancy)	1	0	0	100GB
2 replicas	2	1	100	200GB
3 replicas	3	2	200	300GB
Encoding 1+0 (no redundancy)	1	0	0	100GB
Encoding 1+2	3	1	200	300GB
Encoding 3+2	5	2	67	167GB
Encoding 5+2	7	2	40	140GB
Encoding 7+2	9	2	29	129GB
Encoding 17+3	20	3	18	118GB

The 1+0 and 1+2 encoding modes are meant for small clusters that have insufficient nodes for other erasure coding modes but will grow in the future. As redundancy type cannot be changed once chosen (from replication to erasure coding or vice versa), this mode allows one to choose erasure coding even if their cluster is smaller than recommended. Once the cluster has grown, more beneficial redundancy modes can be chosen.

Note: After upgrading a node in a mixed cluster, you cannot migrate VEs (virtual machines and containers) created in datastores with encoding EC 3+2, 5+2, 7+2, or 17+3 from VHS 7.5 Update 4 to VHS 7.5 Update 3. However, the migration of VEs created in local datastores and datastores with a 3-replica and 2-replica data redundancy mode is available. A mixed cluster is not supported and exists during the upgrade only.

The general recommendation is to always have at least one more node in a cluster than required by the chosen redundancy scheme. For example, a cluster using replication with 3 replicas should have four nodes, and a cluster that works in the 7+2 erasure coding mode should have ten nodes. Such a cluster configuration has the following advantages:

- The cluster will not be exposed to additional failures when in the degraded state. With one node down, the cluster may not survive another even single-disk failure without data loss.
- You will be able to perform maintenance on cluster nodes that may be needed to recover a failed node

(for example, for installing software updates).

- In most cases, the cluster will have enough nodes to rebuild itself. In a cluster without a spare node, each replica of user data is distributed to each cluster node for redundancy. If one or two nodes go down, the user data will not be lost, but the cluster will become degraded and will only start self-healing after the failed nodes are back online. During its rebuilding process, the cluster may be exposed to additional failures until all of its nodes are healthy again.
- You can replace and upgrade a cluster node without adding a new node to the cluster. A graceful release of a storage node is only possible if the remaining nodes in the cluster can comply with the configured redundancy scheme. You can, however, release a node forcibly without data migration, but it will make the cluster degraded and trigger the cluster self-healing.

You choose a data redundancy mode when configuring storage access points and their volumes. In particular, when creating LUNs for iSCSI storage access points and creating S3 clusters.


No matter what redundancy mode you choose, it is highly recommended is to be protected against a simultaneous failure of two nodes as that happens often in real-life scenarios.

Note: All redundancy modes allow write operations when one storage node is inaccessible. If two storage nodes are inaccessible, write operations may be frozen until the cluster heals itself.

1.2.4.1 Redundancy by Replication

With replication, Virtuozzo Storage breaks the incoming data stream into 256MB chunks. Each chunk is replicated and replicas are stored on different storage nodes, so that each node has only one replica of a given chunk.

The following diagram illustrates the 2 replicas redundancy mode.



Replication in Virtuozzo Storage is similar to the RAID rebuild process but has two key differences:

- Replication in Virtuozzo Storage is much faster than that of a typical online RAID 1/5/10 rebuild. The reason is that Virtuozzo Storage replicates chunks in parallel, to multiple storage nodes.
- The more storage nodes are in a cluster, the faster the cluster will recover from a disk or node failure.

High replication performance minimizes the periods of reduced redundancy for the cluster. Replication performance is affected by:

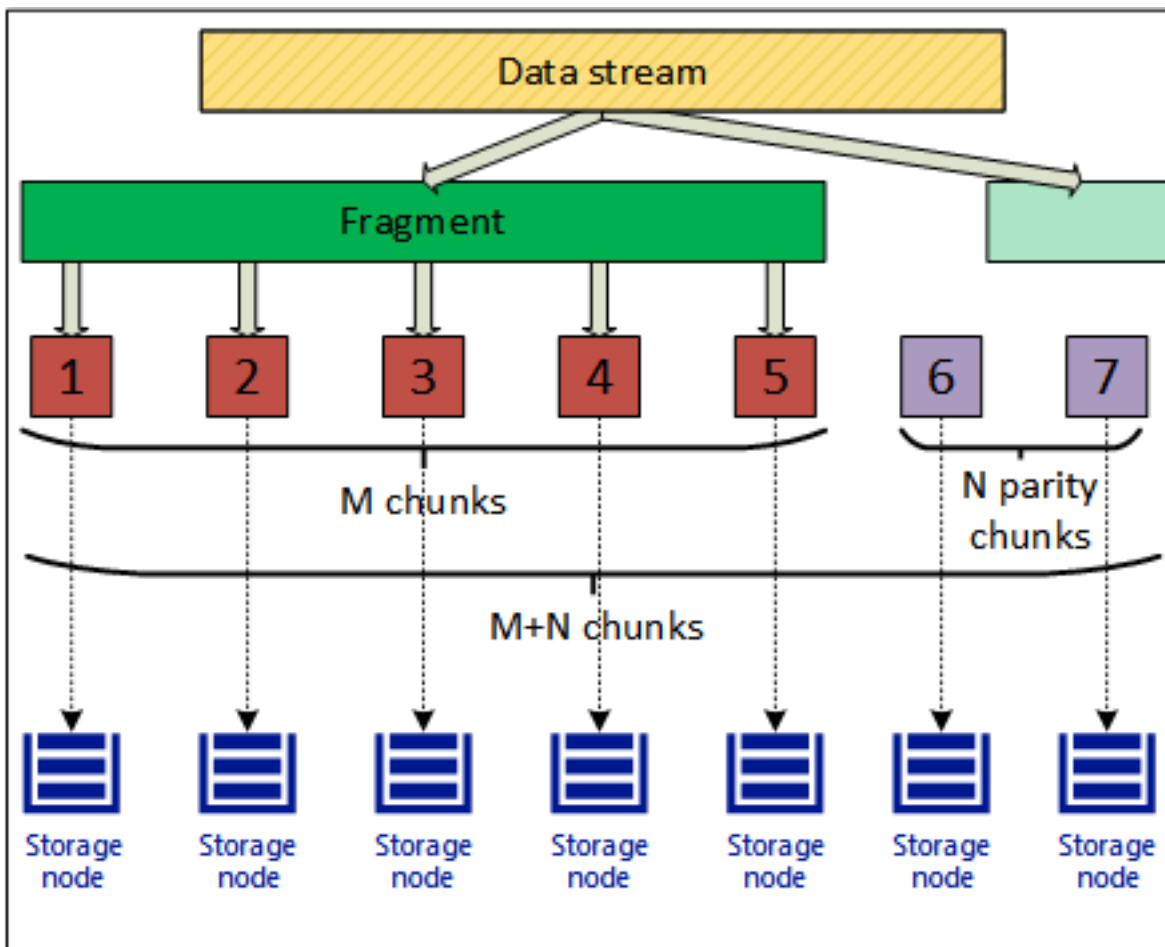
- The number of available storage nodes. As replication runs in parallel, the more available replication sources and destinations there are, the faster it is.
- Performance of storage node disks.
- Network performance. All replicas are transferred between storage nodes over network. For example, 1 Gbps throughput can be a bottleneck (see *Per-Node Network Requirements* on page 17).
- Distribution of data in the cluster. Some storage nodes may have much more data to replicate than other and may become overloaded during replication.
- I/O activity in the cluster during replication.

1.2.4.2 Redundancy by Erasure Coding

With erasure coding, Virtuozzo Storage breaks the incoming data stream into fragments of certain size, then splits each fragment into a certain number (M) of 1-megabyte pieces and creates a certain number (N) of parity pieces for redundancy. All pieces are distributed among M+N storage nodes, that is, one piece per node. On storage nodes, pieces are stored in regular chunks of 256MB but such chunks are not replicated as redundancy is already achieved. The cluster can survive failure of any N storage nodes without data loss.

The values of M and N are indicated in the names of erasure coding redundancy modes. For example, in the 5+2 mode, the incoming data is broken into 5MB fragments, each fragment is split into five 1MB pieces and two more 1MB parity pieces are added for redundancy. In addition, if N is 2, the data is encoded using the RAID6 scheme, and if N is greater than 2, erasure codes are used.

The diagram below illustrates the 5+2 mode.



1.2.4.3 No Redundancy

Warning: Danger of data loss!

Without redundancy, singular chunks are stored on storage nodes, one per node. If the node fails, the data may be lost. Having no redundancy is highly not recommended no matter the scenario, unless you only want to evaluate Virtuozzo Storage on a single server.

1.2.5 Understanding Failure Domains

A failure domain is a set of services which can fail in a correlated manner. To provide high availability of data, Virtuozzo Storage spreads data replicas evenly across failure domains, according to a replica placement policy.

The following policies are available:

- Host as a failure domain (default). If a single host running multiple CS services fails (e.g., due to a power outage or network disconnect), all CS services on it become unavailable at once. To protect against data loss under this policy, Virtuozzo Storage never places more than one data replica per host. This policy is highly recommended for clusters of five nodes and more.
- Disk, the smallest possible failure domain. Under this policy, Virtuozzo Storage never places more than one data replica per disk or CS. While protecting against disk failure, this option may still result in data loss if data replicas happen to be on different disks of the same host and it fails. This policy can be used with small clusters of up to five nodes (down to a single node).

1.2.6 Understanding Storage Tiers

In Virtuozzo Storage terminology, tiers are disk groups that allow you to organize storage workloads based on your criteria. For example, you can use tiers to separate workloads produced by different tenants. Or you can have a tier of fast SSDs for service or virtual environment workloads and a tier of high-capacity HDDs for backup storage.

When assigning disks to tiers (which you can do at any time), have in mind that faster storage drives should be assigned to higher tiers. For example, you can use tier 0 for backups and other cold data (CS without SSD

cache), tier 1 for virtual environments—a lot of cold data but fast random writes (CS with SSD cache), tier 2 for hot data (CS on SSD), caches, specific disks, and such.

This recommendation is related to how Virtuozzo Storage works with storage space. If a storage tier runs out of free space, Virtuozzo Storage will attempt to temporarily use lower tiers down to the lowest. If the lowest tier also becomes full, Virtuozzo Storage will attempt to use a higher one. If you add more storage to the original tier later, the data, temporarily stored elsewhere, will be moved to the tier where it should have been stored originally. For example, if you try to write data to the tier 2 and it is full, Virtuozzo Storage will attempt to write that data to tier 1, then to tier 0. If you add more storage to tier 2 later, the aforementioned data, now stored on the tier 1 or 0, will be moved back to the tier 2 where it was meant to be stored originally.

Inter-tier data allocation as well as the transfer of data to the original tier occurs in the background. You can disable such migration and keep tiers strict as described in [Disabling Inter-Tier Data Allocation](#).

Note: With the exception of out-of-space situations, automatic migration of data between tiers is not supported.


1.3 Planning Infrastructure for Virtuozzo Storage with CLI Management

To plan your infrastructure for Virtuozzo Storage managed by command-line tools, you will need to decide on the hardware configuration of each storage node and plan the storage networks.

Information in this section is meant to help you complete these tasks.

1.3.1 Understanding Virtuozzo Storage Architecture

Before starting the deployment process, you should have a clear idea of the Virtuozzo Storage infrastructure. A typical Virtuozzo Storage is shown below.



The basic component of Virtuozzo Storage is a cluster. The cluster is a group of physical computers connected to the same Ethernet network and performing the following roles:

- Chunk servers (CS)
- Metadata servers (MDS)
- Client computers (or clients)

All data in a Virtuozzo Storage cluster, including virtual machine and container disk images, is stored in the form of fixed-size chunks on chunk servers. The cluster automatically replicates the chunks and distributes them across the available chunk servers to provide high availability of data.

To keep track of data chunks and their replicas, the cluster stores metadata about them on metadata (MDS) servers. The central MDS server, called the master MDS server, monitors all cluster activity and keeps metadata current.

Clients manipulate data stored in the cluster by sending different types of file requests, such as modifying an existing file or creating a new one.

- **Chunk servers (CS).** Chunk servers store all the data, including the contents of virtual machines and

containers, in the form of fixed-size chunks and provide access to these chunks. All data chunks are replicated and the replicas are kept on different chunk servers to achieve high availability. If one of the chunk servers goes down, the other chunk servers will continue providing the data chunks that were stored on the failed server.

- **Metadata servers (MDS).** Metadata servers store metadata about chunk servers and control how files keeping the contents of virtual machines and containers are split into chunks and where these chunks are located. MDS servers also ensure that a cluster has enough chunk replicas and store a global log of all important events that happen in the cluster.

To provide high availability for a Virtuozzo Storage cluster, you need to set up several MDS servers in the cluster. In this case, if one MDS server goes offline, another MDS server will continue keeping control over the cluster.

Note: MDS servers deal with processing metadata only and do not normally participate in any read/write operations related to data chunks.

- **Clients.** Clients are computers with Virtuozzo Hybrid Server 7 from where you run virtual machines and containers stored in a Virtuozzo Storage cluster.

Note the following:

- You can set up any computer in the cluster to perform the role of a metadata server, chunk server, or client. You can also assign two or all three roles to one and the same computer. For example, you can configure a computer to act as a client by installing Virtuozzo Hybrid Server 7 on it and running virtual machines and containers from the computer. At the same time, if you want this computer to allocate its local disk space to the cluster, you can set it up as a chunk server.
- Though Virtuozzo Storage can be mounted as a file system, it is not a POSIX-compliant file system and lacks some POSIX features like ACL, user and group credentials, hardlinks, and some other.

1.3.2 Virtuozzo Storage Configurations

This section provides information on two Virtuozzo Storage configurations:

- *Minimum Configuration* on page 33. You can create the minimum configuration for evaluating the Virtuozzo Storage functionality. This configuration, however, is not recommended for use in production environments.


- *Recommended Configuration* on page 34. You can use the recommended Virtuozzo Storage configuration in a production environment “as is” or adapt it to your needs.

1.3.2.1 Minimum Configuration

The minimum hardware configuration for deploying a Virtuozzo Storage cluster is given below:

Server Role	Number of Servers
Metadata Server	1 (can be shared with chunk servers and clients)
Chunk Server	1 (can be shared with metadata servers and clients)
Client	1 (can be shared with chunk and metadata servers)
Total number of servers	with role sharing: 1 without role sharing: 3

Graphically, the minimum configuration can be represented as follows:



For a Virtuozzo Storage cluster to function, it must have at least one MDS server, one chunk server, and one client. The minimum configuration has two main limitations:

1. The cluster has one metadata server, which presents a single point of failure. If the metadata server fails, the entire Virtuozzo Storage cluster will become non-operational.
2. The cluster has one chunk server that can store only one chunk replica. If the chunk server fails, the

cluster will suspend all operations with chunks until a new chunk server is added to the cluster.

1.3.2.2 Recommended Configuration

The table below lists two of the recommended configurations for deploying Virtuozzo Storage clusters.

Metadata Server	Chunk Server	Total Number of Servers
3 (can be shared with chunk servers and clients)	5-9 (can be shared with metadata servers and clients)	5 or more (depending on the number of clients and chunk servers and whether they share roles)
5 (can be shared with chunk servers and clients)	10 or more (can be shared with metadata servers and clients)	5 or more (depending on the number of clients and chunk servers and whether they share roles)
Clients	1 or more You can include any number of clients in the cluster. For example, if you have 5 servers with Virtuozzo Hybrid Server, you can configure them all to act as clients. You can share servers acting as clients with chunk and metadata servers. For example, you can have 5 physical servers and configure each of them to simultaneously act as an MDS server, a chunk server, and a client.	

Even though new clusters are configured to have 1 replica for each data chunk by default, you need to configure each data chunk to have at least 3 replicas to provide high availability for your data.

In total, at least 9 machines running Virtuozzo Storage are recommended per cluster. Smaller clusters will work as fine but will not provide the significant performance advantages over direct-attached storage (DAS) or improved recovery times.

Note the following:

- For large clusters, it is critically important to configure proper failure domains to improve data availability. For more information, see [Configuring Failure Domains](#).
- In small and medium clusters, MDS servers consume little resources and do not require being set up on dedicated Hardware Nodes.

- A small cluster is 3 to 5 machines, a medium cluster is 6 to 15-20 machines, and a large cluster is 15-20 machines and more.
- Time should be synchronized on all servers in the cluster via NTP. Doing so will make it easier for the support department to understand cluster logs (migrations, failovers, etc.). By default, Virtuozzo Hybrid Server 7 uses the `chronyd` service for time synchronization. If you want to use `ntpd` or `ntpd`, stop and disable `chronyd` first.

1.3.3 Hardware Requirements

Before setting up a Virtuozzo Storage cluster, make sure you have all the necessary equipment at hand.

You are also recommended to:

- Consult [Using SSD Drives](#) to learn how you can increase cluster performance by using solid-state drives for write journaling and data caching, and how many SSDs you may need depending on the number of HDDs in your cluster.
- Check [Appendix A - Troubleshooting](#) for common hardware issues and misconfigurations that may affect your cluster performance and lead to data inconsistency and corruption.

General

- Each service (be it MDS, CS or client) requires 1.5 GB of free space on root partition for logs. For example, to run 1 metadata server, 1 client, and 12 chunk servers on a host, you will need 21 GB of free space on the root partition.

Metadata Servers

- 1 CPU core
- 1 GB of RAM per 100 TB of data in the cluster
- 3 GB of disk space per 100 TB of data in the cluster
- 1 or more Ethernet adapters (1 Gbit/s or faster)

Note: It is recommended to place the MDS journal on SSD, either dedicated or shared with CS and client caches, or at least on a dedicated HDD which has no CS services on it.

Chunk Servers

- 1/8 of a CPU core (e.g., 1 CPU core per 8 CS)
- 1 GB of RAM
- 100 GB or more of free disk space
- 1 or more Ethernet adapters (1 Gbit/s or faster)

Note the following:

- On using local RAID with Virtuozzo Storage, consult [Exploring Possible Disk Drive Configurations](#).
- Using a shared JBOD array across multiple nodes running CS services may introduce a single point of failure and make the cluster unavailable if all data replicas happen to be allocated and stored on the failed JBOD. For more information, see [Configuring Failure Domains](#).
- For large clusters, it is critically important to configure proper failure domains to improve data availability. For more information, see [Configuring Failure Domains](#).
- Do not place chunk servers on disks already used in other I/O workloads, e.g., system or swap. Sharing disks between CS and other sources of I/O will result in severe performance loss and high I/O latencies.

Clients

- 1 CPU core per 30,000 IOPS
- 1 GB of RAM
- 1 or more Ethernet adapters (1 Gbit/s or faster)

The following table lists the maximum network performance a Virtuozzo Storage client can get with the specified network interface. The recommendation for clients is to use 10Gbps network hardware between any two cluster nodes and minimize network latencies, especially if SSD disks are used.

Storage network interface	1Gbps	2 x 1Gbps	3 x 1Gbps	10Gbps	2 x 10Gbps
Entire node maximum I/O throughput	100MB/s	~175MB/s	~250MB/s	1GB/s	1.75GB/s
Single VM maximum I/O throughput (replication)	100MB/s	100MB/s	100MB/s	1GB/s	1GB/s
Single VM maximum I/O throughput (erasure coding)	70MB/s	~130MB/s	~180MB/s	700MB/s	1.3GB/s

Note: For hard disk requirements and the recommended partitioning scheme for servers that run Virtuozzo Hybrid Server and participate in clusters, see the *Partitioning the Hard Drives* on page 78.


1.3.4 Network Requirements

When planning your network, make sure that it

- Operates at 1 Gbit/s or faster (for more details, see [Using 1 GbE and 10 GbE Networks](#))
- Has non-blocking Ethernet switches

You should use separate networks and Ethernet adapters for user and cluster traffic. This will prevent possible I/O performance degradation in your cluster by external traffic. Besides, if a cluster is accessible from public networks (e.g., from the Internet), it may become a target of Denial-of-Service attacks, and the entire cluster I/O subsystem may get stuck.

The figure below shows a sample network configuration for Virtuozzo Storage.



In this network configuration:

- BackNet is a private network used solely for interconnection and intercommunication of servers in the cluster and is not available from the public network. All servers in the cluster are connected to this network via one of their network cards.
- FrontNet is a public network customers use to access their virtual machines and containers in the Virtuozzo Storage cluster.

Note the following:

- Network switches are a very common point of failure, so it is critically important to configure proper failure domains to improve data availability. For more information, see [Configuring Failure Domains](#).
- To learn more about Virtuozzo Storage networks (in particular, how to bind chunk servers to specific IP addresses), see [Securing Server Communication in the Cluster](#).

1.4 Obtaining the Distribution Image

To obtain the Virtuozzo Hybrid Server distribution ISO image, do as follows:

1. Visit the [product page](#).
2. Click **GET IN TOUCH** upper right.
3. Fill out the form and submit the request.
4. In an email that you will get shortly afterwards, click the link to download the Virtuozzo Hybrid Server distribution ISO image. Also in this email, you will find a license key number to activate a Virtuozzo Hybrid Server license (see [Activating a License](#) on page 58).

1.5 Preparing for Installation from USB Storage Drives

To install Virtuozzo Hybrid Server from a USB storage drive, you will need a 2 GB or higher-capacity USB drive and the Virtuozzo Hybrid Server 7 distribution ISO image.

Make a bootable USB drive by transferring the distribution image to it.

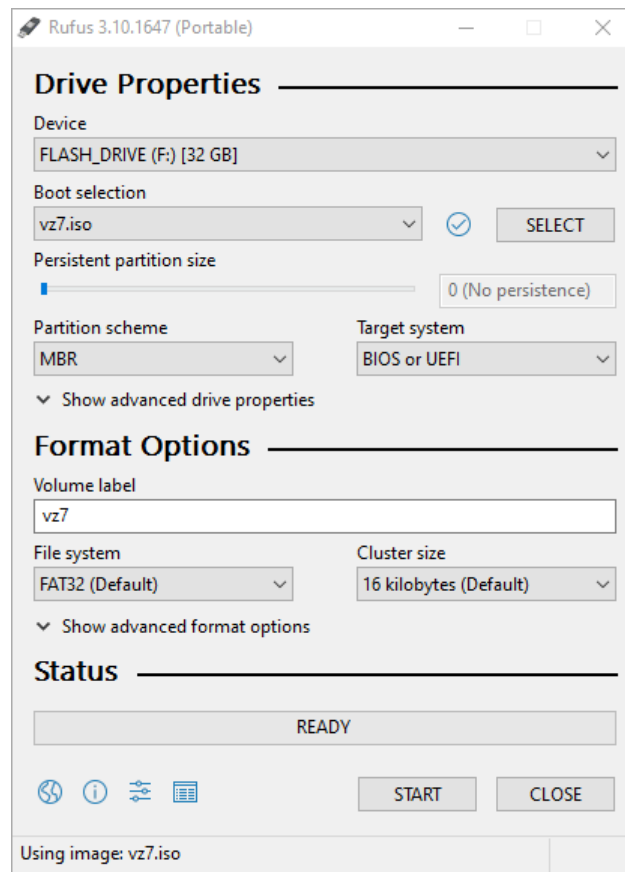
Important: Be careful to specify the correct drive to transfer the image to.

On Linux, you can use dd:


```
# dd if=<VZ7_ISO_image> of=/dev/sdb
```

On Windows, you can use Rufus:

1. In the **Drive Properties** section, select your flash drive from the **Device** drop-down menu and click **SELECT**. Then, select the distribution image from your local machine.
2. Click **START**.



3. In the pop-up window, select **Write in DD Image mode** and click **OK**.



CHAPTER 2

Installing Virtuozzo Hybrid Server

This chapter explains how to install Virtuozzo Hybrid Server.

2.1 Starting Installation

Virtuozzo Hybrid Server can be installed from:

- IPMI virtual drives
- USB drives (see *Preparing for Installation from USB Storage Drives* on page 38)
- PXE servers (see *Installation Using PXE* for information on installing Virtuozzo Hybrid Server over the network)

Note: In Virtuozzo Hybrid Server 7, time synchronization via NTP is enabled by default using the `chronyd` service. If you want to use `ntpd` or `ntpd`, stop and disable `chronyd` first.

To start the installation, do the following:

1. Configure the server to boot from the chosen media.
2. Boot the server and wait for the welcome screen.

2.1.1 Choosing Installation Type

You can install Virtuozzo Hybrid Server 7 in one of the following modes:

- Graphics. In this mode, you can install Virtuozzo Hybrid Server components with one of the following management interfaces:
 - Web-based user interface (default, recommended). The interface consists of Virtuozzo Storage management panel and Virtuozzo Automator management panel (Control Center), see *Installing with GUI Management* on page 44.
 - Command-line interface, see *Installing with CLI Management* on page 61.
- Basic graphics (in case of issues with video card drivers), see *Installing in the Basic Graphics Mode* on page 85.
- Graphics via VNC, see *Installing via VNC* on page 85.
- Text (Virtuozzo Storage cannot be installed in this mode), see *Installing in the Text Mode* on page 85.

Your further installation steps will differ depending on the chosen mode.

2.1.2 Enabling Forced Detection of SSDs

Certain solid-state drives (SSDs) may not be autodetectable by the installer. This may result in issues when you create or join Virtuozzo Storage clusters. To avoid this problem, you can force the installer to identify the required drives as SSDs by doing the following:

1. Select the required installation option and press **E** to edit it.
2. Add `ssd_hack=sd<N>[, ...]` at the end of the line starting with `linux /images/pxeboot/vmlinuz`. For example:

```
linux /images/pxeboot/vmlinuz inst.stage2=hd:LABEL=<VZ7_ISO_image> quiet ip=dhcp \  
ssd_hack=sdb,sdc
```

3. Press **Ctrl-X** to load the selected installation option.

The installer should identify the specified drives as SSDs.

2.1.3 Changing the Swap Size Calculation Method

Virtuozzo Hybrid Server requires that a swap partition be created on the system disk during installation. The size of this partition depends on the RAM size. If automatic partitioning is used, the swap size is calculated as described in *Partitioning the Hard Drives* on page 78:

- If the RAM size is 4-8 GB, the swap size is equal to the RAM size.
- Otherwise the swap size is half the RAM size.

Such method is beneficial for memory overcommit scenarios. If, however, overcommit is not required, the old method of calculating the swap size can be used:

- If the RAM size is 4-8 GB, the swap size is equal to the RAM size.
- If the RAM size is 8-64 GB, the swap size is half the RAM size.
- Otherwise the swap size is 32 GB.

To have the installer use the old method while it partitions the disks automatically, boot to the welcome screen and do the following:

1. Select the required installation option and press **E** to edit it.
2. Add `smallswap` at the end of the line starting with `linux /images/pxeboot/vmlinuz`. For example:

```
linux /images/pxeboot/vmlinuz inst.stage2=hd:LABEL=<VZ7_ISO_image> quiet ip=dhcp smallswap
```

3. Press **CTRL-X** to start booting the chosen installation option.

2.1.4 Installing on USB Drive

Aside from regular disks, you can also install Virtuozzo Hybrid Server on a 64GB+ USB drive. In this case, however, you will still need to place the swap partition on a regular HDD/SSD drive. It is also recommended to store logs on a syslog server.

To have a USB drive recognized as an installation destination, configure the installer as follows:

1. Select the required installation option and press **E** to edit it.
2. Add `allow_usb_hdd` at the end of the line starting with `linux /images/pxeboot/vmlinuz`. For example:

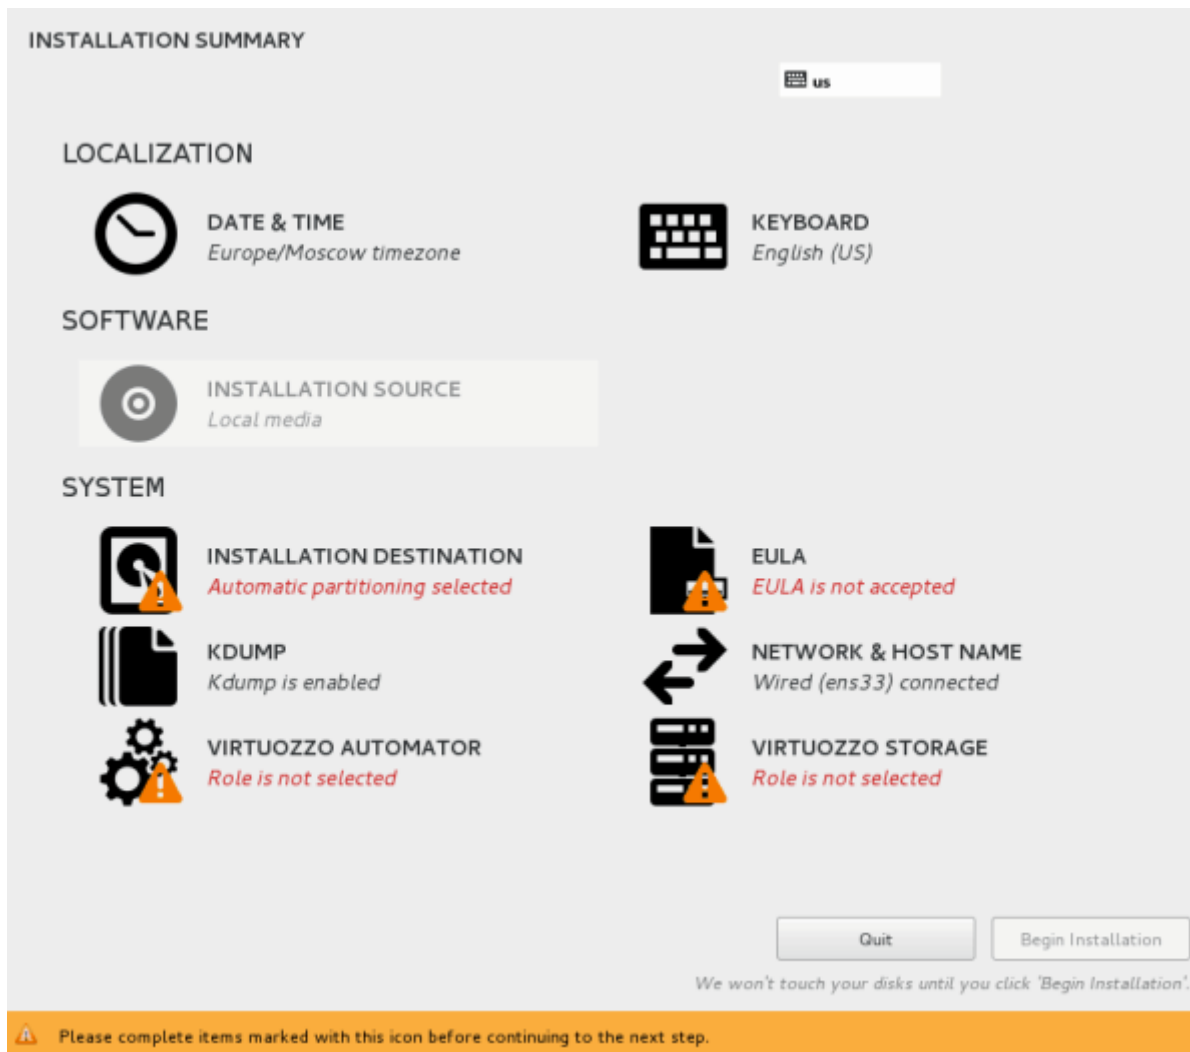
```
linux /images/pxeboot/vmlinuz inst.stage2=hd:LABEL=<VZ7_ISO_image> quiet ip=dhcp \  
allow_usb_hdd
```

Note: If you chose to install Virtuozzo Hybrid Server with GUI, also add `expert` to this line.

3. Press **Ctrl-X** to load the selected installation option.
4. On installer's **INSTALLATION DESTINATION** screen, select the USB drive as system, select **I will configure partitioning** in the bottom, and click **Done** and then **OK** to proceed to the **MANUAL PARTITIONING** screen.
5. On the **MANUAL PARTITIONING** screen, select the swap partition in the list, click **Modify** on the right, and select an HDD or SSD drive to place the swap on.
6. Click **Done**, then **Accept Changes** to return to the main installer screen and configure other options as usual.

2.2 Installing with GUI Management

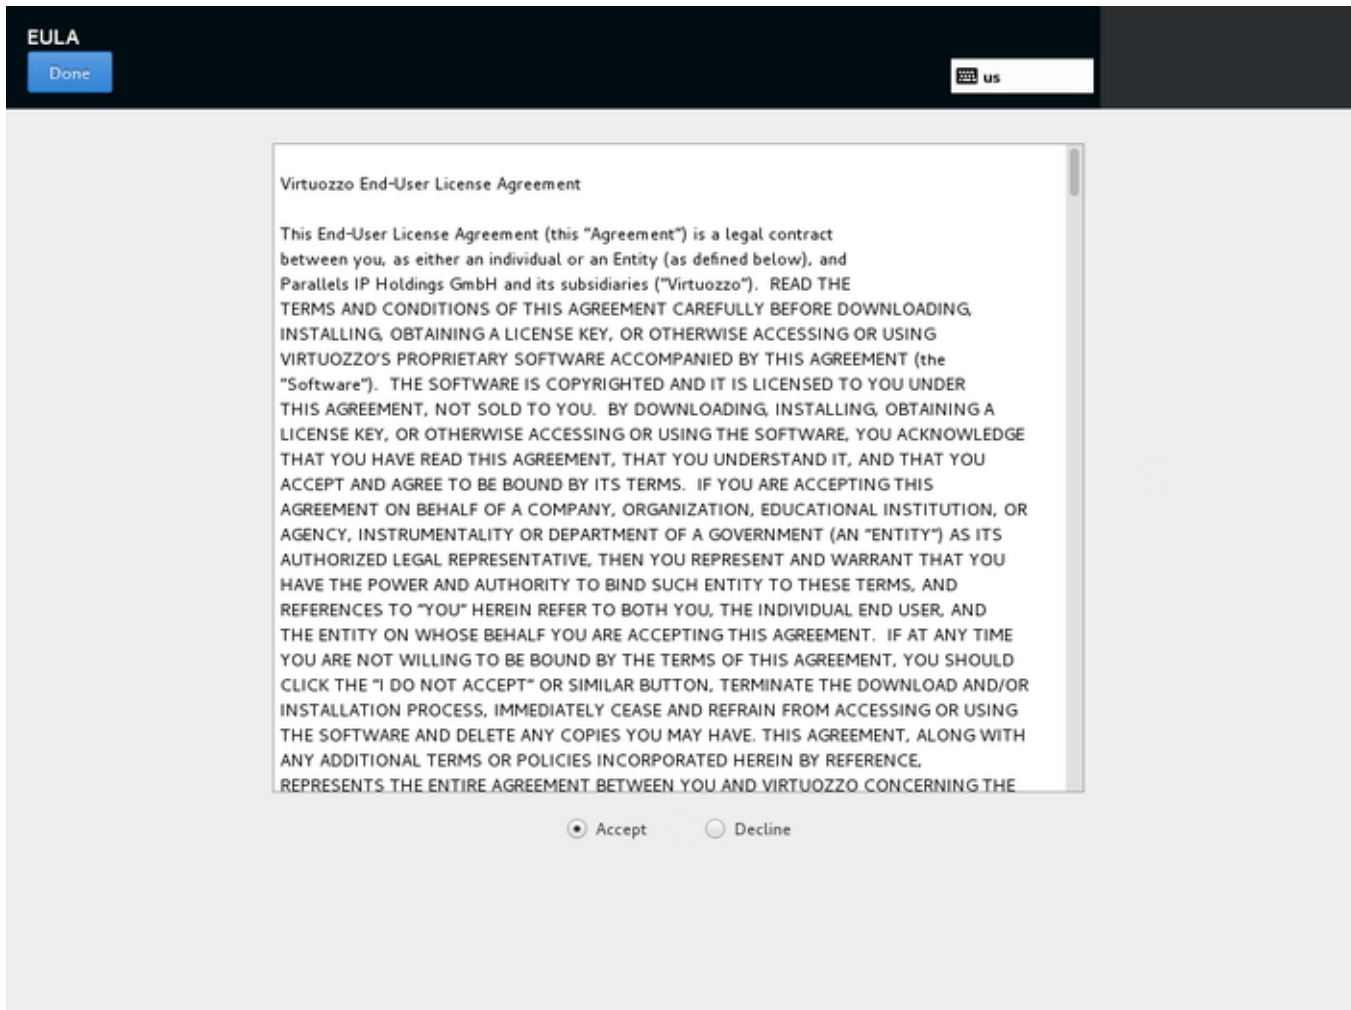
To install Virtuozzo Hybrid Server with GUI management in graphics mode, choose **Install with GUI management** on the welcome screen. After the installation program loads, you will see the **Installation Summary** screen. On this screen, you need to specify a number of parameters required to install Virtuozzo Hybrid Server.



2.2.1 Accepting EULA

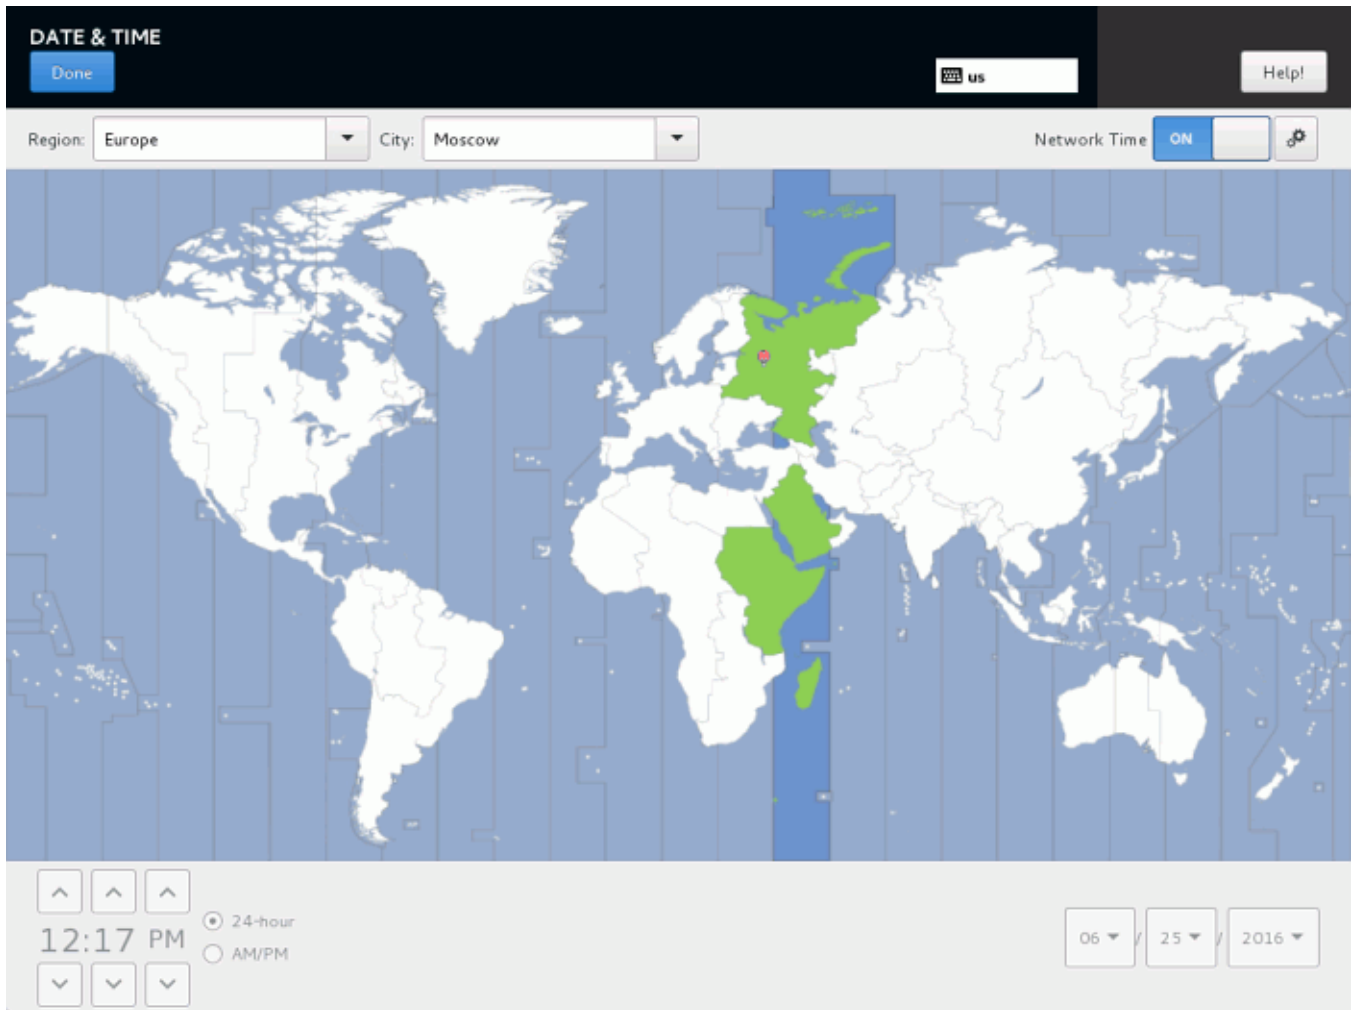
You need to accept the Virtuozzo Hybrid Server End-User License Agreement to install Virtuozzo Hybrid Server.

To accept the Virtuozzo Hybrid Server EULA, open the **EULA** screen, select **Accept**, and click **Done**.



2.2.2 Setting Date and Time

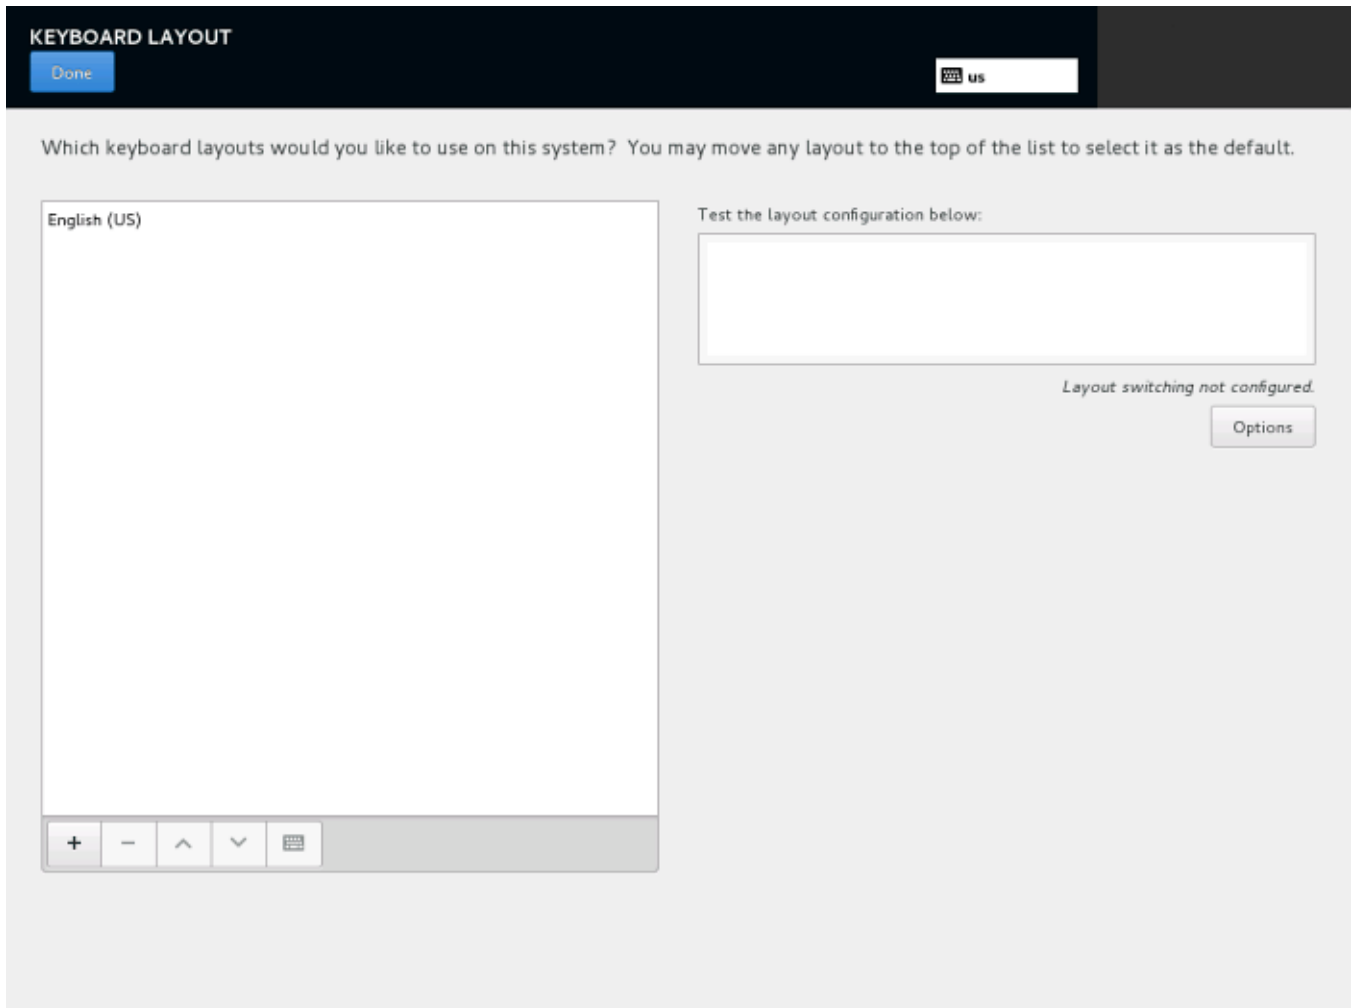
If you need to set the date and time for your Virtuozzo Hybrid Server installation, open the **DATE & TIME** screen and make the necessary changes.



2.2.3 Selecting the Keyboard Layout

The selected keyboard layout can be used during installation and, once the installation is complete, in the console (e.g., for entering localized descriptions, configuration file comments, and such).

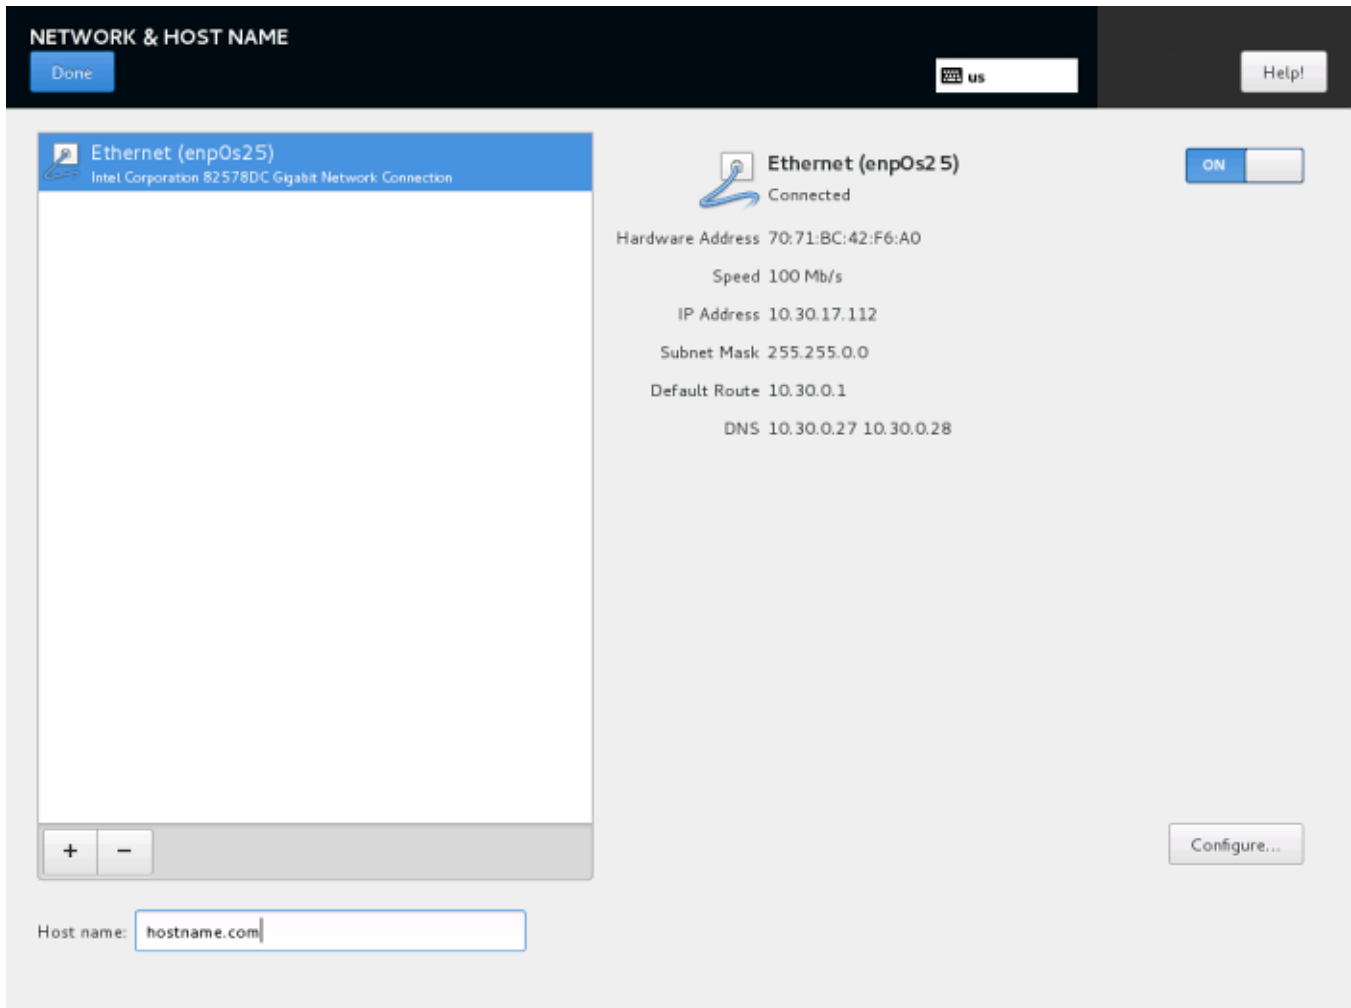
If you need to change the default English (US) keyboard layout, open the **KEYBOARD** screen, click the plus sign to add a layout, and click **Options** to choose a key combination for switching layouts.



2.2.4 Configuring Network

Usually network is configured automatically by the installation program. If you need to modify network settings, you can do so on the **NETWORK & HOST NAME** screen.

To install Virtuozzo Hybrid Server, you will need to have at least one network card configured and you will also need to provide a hostname: either a fully qualified domain name (<hostname>.<domainname>) or a short name (<hostname>).




2.2.4.1 Creating Bonded and Teamed Connections

Bonded and teamed connections offer increased throughput beyond the capabilities of a single network card as well as improved redundancy.

During installation, you can configure bonding on the **NETWORK & HOSTNAME** screen as described below. Teaming can be configured in a similar way after choosing **Team** on step 1.

1. To add a new bonded connection, click the plus button in the bottom, select **Bond** from the drop-down list, and click **Add**.

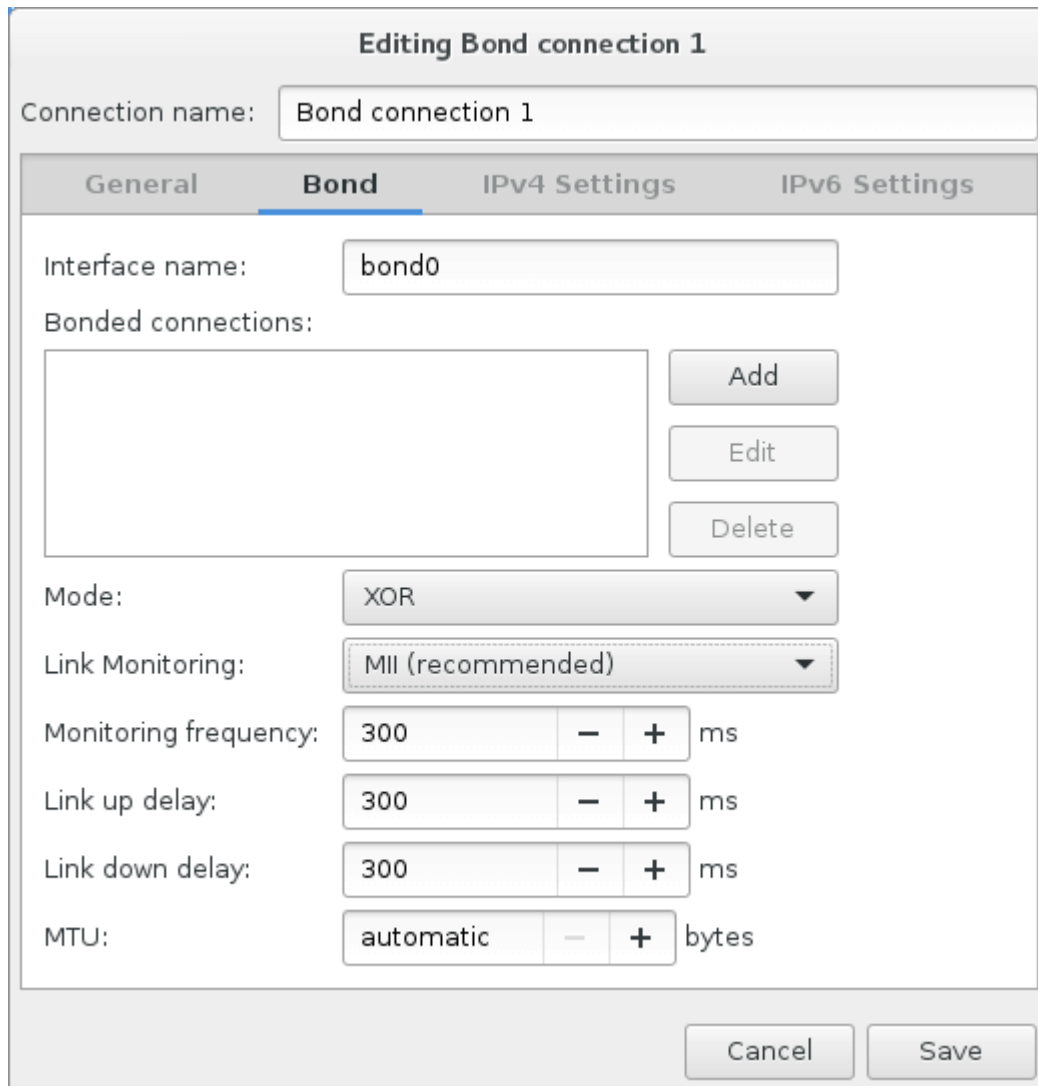


2. In the **Editing Bond connection...** window, set the following parameters:

2.1. **Mode** to XOR.

2.2. **Link Monitoring** to MII (recommended).

2.3. **Monitoring frequency**, **Link up delay**, and **Link down delay** to 300.



Note: It is also recommended to manually set `xmit_hash_policy` to `layer3+4` after the installation as detailed in [Setting Up Network Bonding](#). For more information on network bonding, see [Red Hat Enterprise Linux Deployment Guide](#) and [Linux Ethernet Bonding Driver HOWTO](#).

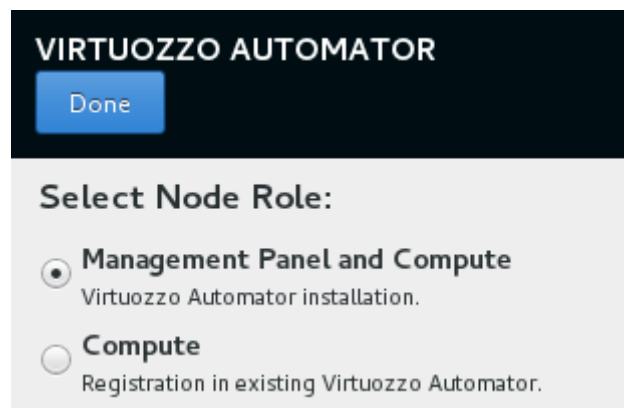
- In the **Bonded connections** section on the **Bond** tab, click **Add**.
- In the **Choose a Connection Type** window, select **Ethernet** from the in the drop-down list, and click **Create**.



- In the **Editing bond slave...** window, select a network interface to bond from the **Device** drop-down list.

2.2.5 Installing Virtuozzo Automator Components

To install Virtuozzo Automator along with Virtuozzo Hybrid Server, you need to choose the respective component to install on the **VIRTUOZZO AUTOMATOR** screen.



The following options are available:

- **Management Panel and Compute.** Install the management panel and agent components of Virtuozzo Automator. The former is installed into a container on the node and allows you to manage registered nodes and virtual environments on them. The latter is installed on the node itself and lets the node be registered in the management panel.
- **Compute.** Install the agent on the node so the node can be registered in the management panel.

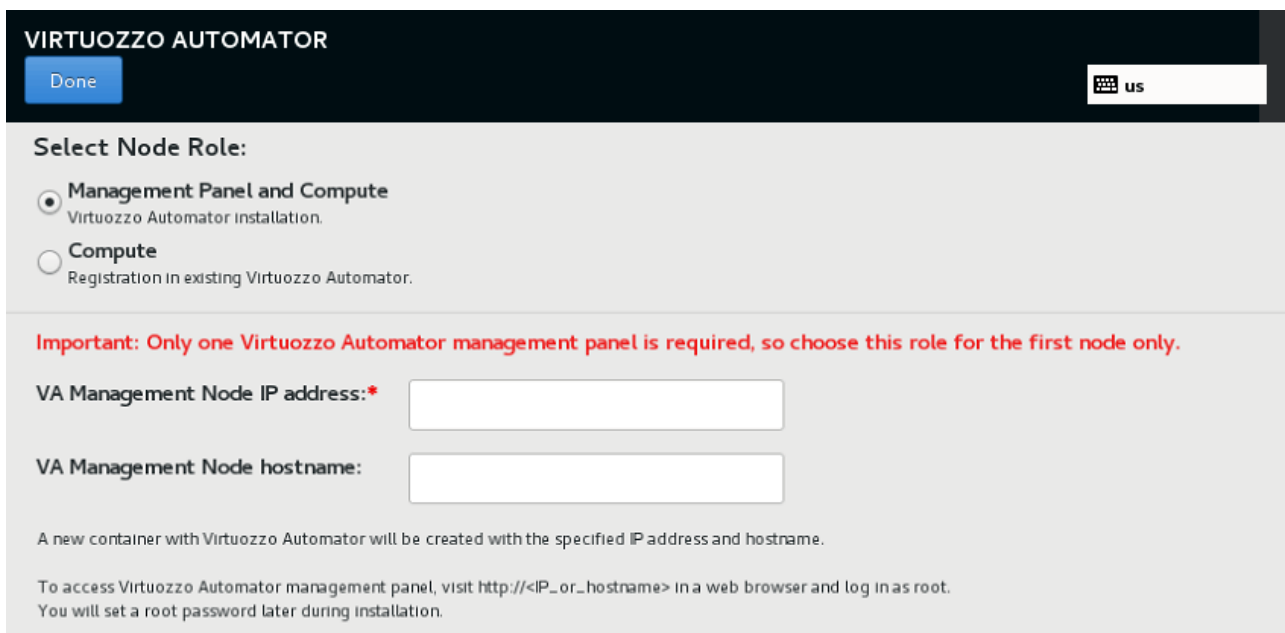
Important: Only one VA Management Node is needed, so choose **Management Panel and Compute** only for the first node in your Virtuozzo Automator infrastructure. For the second and other nodes choose **Compute**. The detailed instructions are provided in the following subsections.

2.2.5.1 Installing Management Panel and Compute on the First Server

On the first server in your Virtuozzo Automator infrastructure, you will need to install both the management panel and compute components.

Do the following on the **VIRTUOZZO AUTOMATOR** screen:

1. Choose **Management Panel and Compute**.



The screenshot shows the VIRTUOZZO AUTOMATOR installation interface. At the top, there is a dark header with the title 'VIRTUOZZO AUTOMATOR', a 'Done' button, and a language selector set to 'us'. Below the header, the 'Select Node Role:' section contains two radio button options: 'Management Panel and Compute' (selected) and 'Compute'. A red warning message states: 'Important: Only one Virtuozzo Automator management panel is required, so choose this role for the first node only.' Below this, there are two input fields: 'VA Management Node IP address: *' and 'VA Management Node hostname:'. At the bottom, there is explanatory text: 'A new container with Virtuozzo Automator will be created with the specified IP address and hostname. To access Virtuozzo Automator management panel, visit http://<IP_or_hostname> in a web browser and log in as root. You will set a root password later during installation.'

2. Specify an **VA Management Node IP address** and, optionally, a **VA Management Node hostname**.

Virtuozzo Hybrid Server will create a container with the specified IP and install the management panel in it. The management panel will be available at the specified IP address (or hostname) on default port 4648.

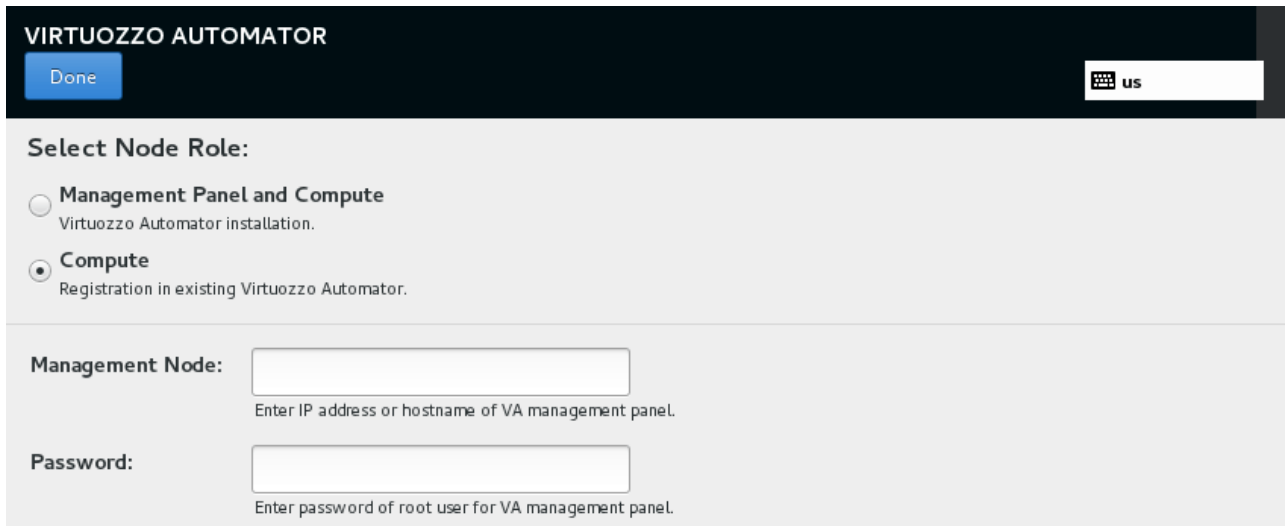
3. Click **Done** and proceed to configure Virtuozzo Storage as described in *Installing Virtuozzo Storage Components* on page 54.

2.2.5.2 Installing Compute on the Second and Other Servers

On the second and other servers in your Virtuozzo Automator infrastructure, you will need to install the compute component only.

Do the following on the **VIRTUOZZO AUTOMATOR** screen:

1. Choose **Compute**.



The screenshot shows the 'VIRTUOZZO AUTOMATOR' interface. At the top left is a 'Done' button. At the top right is a language selector showing 'us'. The main section is titled 'Select Node Role:' and contains two radio button options: 'Management Panel and Compute' (unselected) and 'Compute' (selected). Below these are two input fields: 'Management Node:' and 'Password:'. The 'Management Node' field has a placeholder text 'Enter IP address or hostname of VA management panel.' and the 'Password' field has a placeholder text 'Enter password of root user for VA management panel.'

2. To register the node in the management panel, specify the following:
 - 2.1. in the **Management node** field, the IP address or hostname of the container with the management panel;
 - 2.2. in the **Password** field, the password of the management panel administrator.

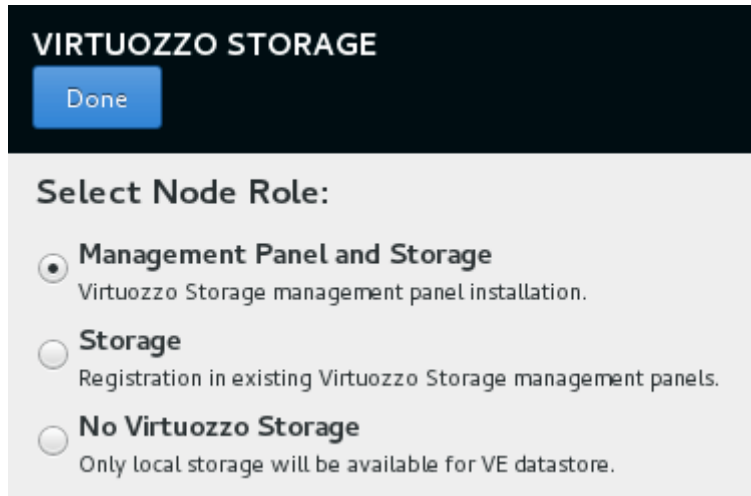
Note: Alternatively, you can register the node in Virtuozzo Automator manually later (see [Registering a Physical Server in |vauto|](#)).

3. Click **Done** and proceed to configure Virtuozzo Storage as described in *Installing Virtuozzo Storage*

Components on page 54.

2.2.6 Installing Virtuozzo Storage Components

To install Virtuozzo Storage along with Virtuozzo Hybrid Server, you need to choose the respective component to install on the **VIRTUOZZO STORAGE** screen.



The following options are available:

- **Management Panel and Storage.** Install the Virtuozzo Storage management panel and prepare the node for data storage. The management panel is installed into a container on the node. Storage nodes run data storage services as well as store and provide access to user data.
- **Storage.** Prepare the node for data storage. Storage nodes run data storage services as well as store and provide access to user data.
- **No |vstorage|.** Install Virtuozzo Hybrid Server on local hard drives. If you select this option, click **Done**, and proceed to set the installation destination as described in *Choosing the Installation Destination* on page 57.

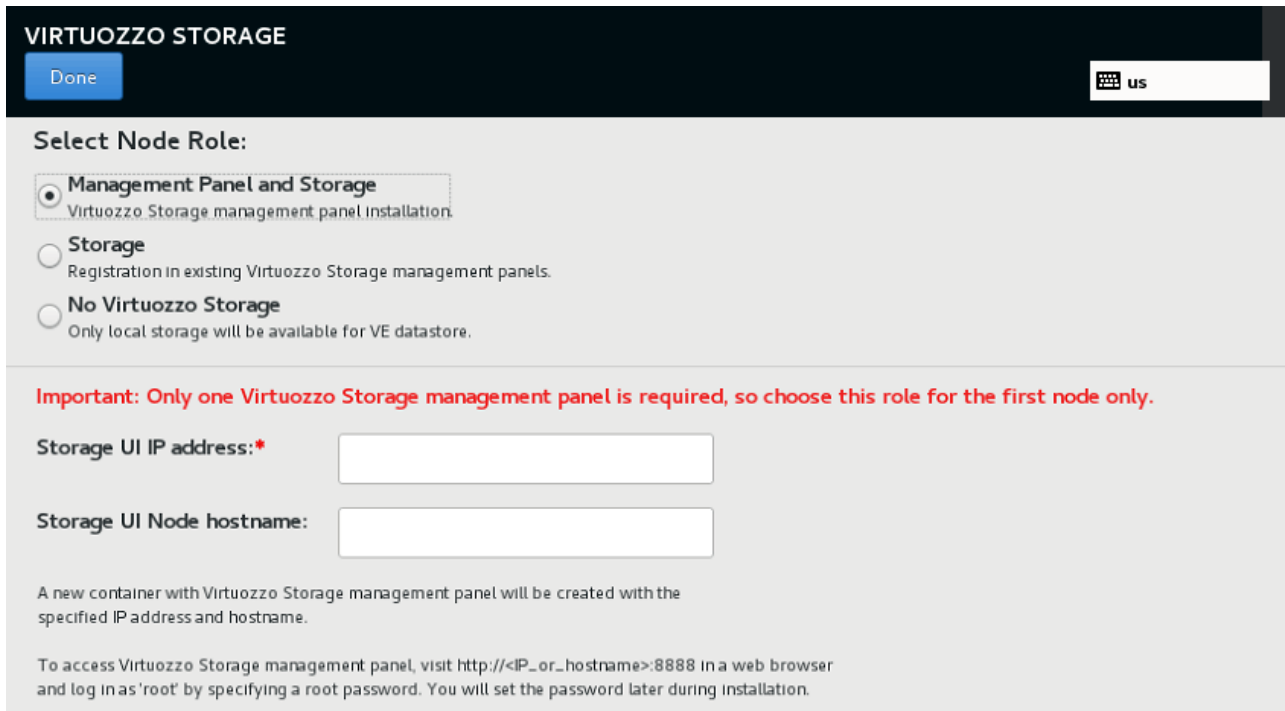
Important: Only one management panel is needed, so choose **Management Panel and Storage** only for the first node in your Virtuozzo Storage infrastructure. For the second and other nodes choose **Storage**. The detailed instructions are provided in the following subsections.

2.2.6.1 Installing Management Panel and Storage on the First Server

On the first server in your Virtuozzo Storage infrastructure, you will need to install both the management panel and storage components. This node will also become a storage node. It will be added to the Virtuozzo Storage infrastructure and registered in the management panel during installation.

Do the following on the **VIRTUOZZO STORAGE** screen:

1. Choose **Management Panel and Storage**.



VIRTUOZZO STORAGE

Done

Select Node Role:

Management Panel and Storage
Virtuozzo Storage management panel installation.

Storage
Registration in existing Virtuozzo Storage management panels.

No Virtuozzo Storage
Only local storage will be available for VE datastore.

Important: Only one Virtuozzo Storage management panel is required, so choose this role for the first node only.

Storage UI IP address:*

Storage UI Node hostname:

A new container with Virtuozzo Storage management panel will be created with the specified IP address and hostname.

To access Virtuozzo Storage management panel, visit http://<IP_or_hostname>:8888 in a web browser and log in as 'root' by specifying a root password. You will set the password later during installation.

2. Specify an **Storage UI IP address** and, optionally, a **Storage UI Node hostname**.

Virtuozzo Hybrid Server will create a new container with the specified IP and install the Virtuozzo Storage management panel in it. The management panel will be available at the specified IP address (or hostname) on default port 8888.

3. Click **Done** and proceed to set installation destination as described in *Choosing the Installation Destination* on page 57.

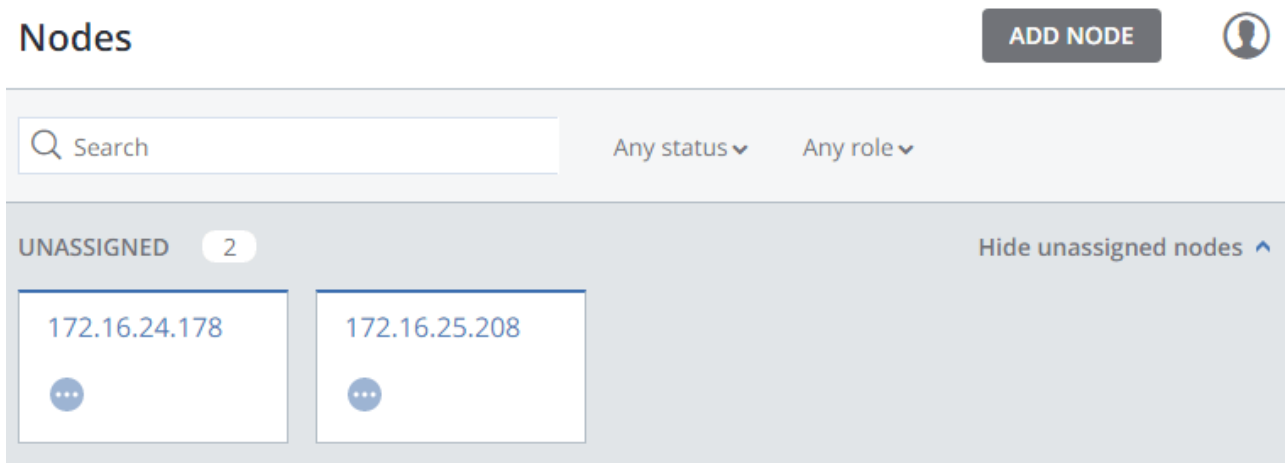
2.2.6.2 Installing Storage on the Second and Other Servers

On the second and other servers in your Virtuozzo Storage infrastructure, you will need to install the **Storage** component only. Storage nodes will be added to the Virtuozzo Storage infrastructure and registered in the management panel during installation.

For security reasons, you will need to provide a token that can only be obtained from the management panel you have installed on the first server. A single token can be used to install the storage component on multiple servers in parallel.

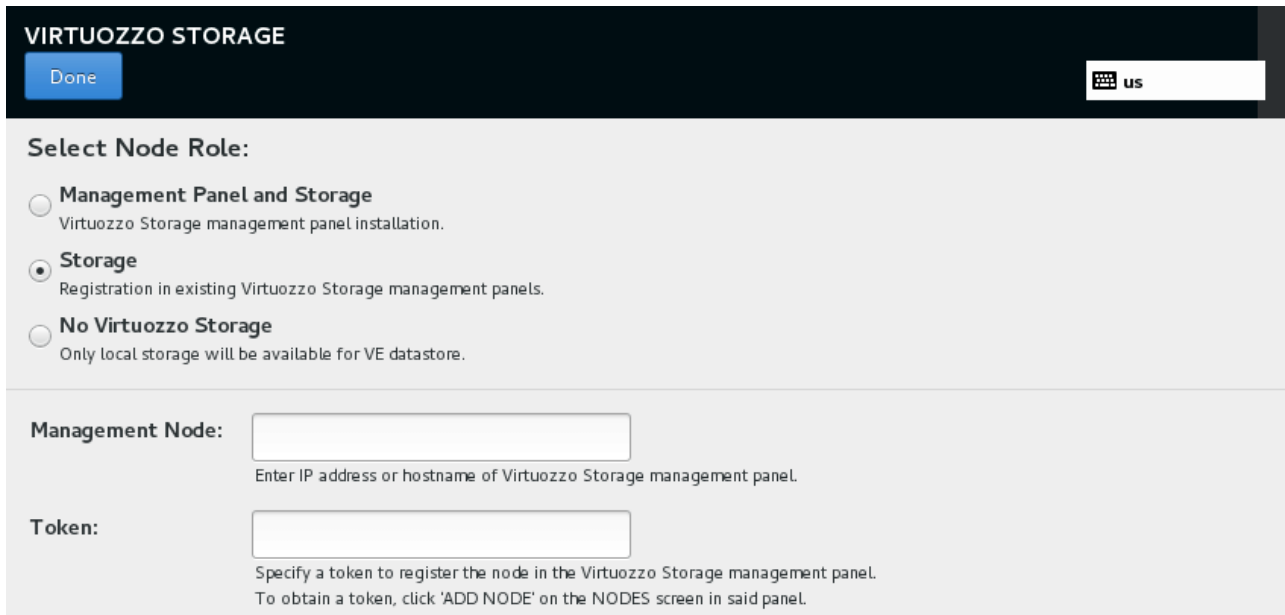
To obtain a token:

1. Log in to the Virtuozzo Storage management panel with the password created while deploying the first server. To do this, on any computer with network access to the container with the management panel, open a web browser and visit the container's IP address or hostname on port 8888.
2. In the management panel, click **ADD NODE** to open a screen with the token (you can generate a new one if needed; generating a new token invalidates the old one).



Having obtained the token, do the following on the **VIRTUOZZO STORAGE** screen:

1. Choose **Storage**.



VIRTUOZZO STORAGE

Done us

Select Node Role:

Management Panel and Storage
Virtuozzo Storage management panel installation.

Storage
Registration in existing Virtuozzo Storage management panels.

No Virtuozzo Storage
Only local storage will be available for VE datastore.

Management Node:
Enter IP address or hostname of Virtuozzo Storage management panel.

Token:
Specify a token to register the node in the Virtuozzo Storage management panel.
To obtain a token, click 'ADD NODE' on the NODES screen in said panel.

2. Register the node in the Virtuozzo Storage management panel. To do this, specify the following:

2.1. in the **Management node** field, the IP address or hostname of the container with the management panel;

2.2. in the **Token** field, the acquired token.

Note: Alternatively, you can register the node in the management panel manually later by logging in to the node via SSH and running `/usr/libexec/vstorage-ui-agent/bin/register-storage-node.sh -m <MP_addr> -t <token>` in the node's console, where `<MP_addr>` is the IP address of the management panel container and `<token>` is the token obtained in the management panel.

3. Click **Done** and proceed to set installation destination as described in [Choosing the Installation Destination](#) on page 57.

2.2.7 Choosing the Installation Destination


Having chosen the storage type, open the **INSTALLATION DESTINATION** screen. If you have multiple disks, select which one will be system. Other disks will be used according to the chosen storage type.

Warning: All disks recognized by the installation program will be cleaned from partitions once you click **Begin Installation**.

Click **Done** and proceed to finish the installation.

2.2.8 Finishing the Installation

Having configured everything necessary on the **INSTALLATION SUMMARY** screen, click **Begin Installation**.



While Virtuozzo Hybrid Server is installing, click **ROOT PASSWORD** to create a password for the root account. Installation will not finish until the password is created.

 A screenshot of the 'ROOT PASSWORD' screen. At the top left is a blue 'Done' button. At the top right is a user indicator showing 'us'. The main area contains the text: "The root account is used for administering the system. Enter a password for the root user." Below this are two input fields: "Root Password:" and "Confirm:". The "Root Password:" field has a strength indicator below it showing a green bar and the word "Strong".

While Virtuozzo Hybrid Server is installing, you can activate a license and disable the automatic updating of the Virtuozzo ReadyKernel patch if needed. If you choose not to activate a license at this time, you may need to take additional steps to switch the memory management policy. For more details, see [Configuring Automatic Memory Management Policies](#).

2.2.8.1 Activating a License

To activate a Virtuozzo Hybrid Server license, click **LICENSE KEY**, enter a key, and click **Done**.

VIRTUOZZO HYBRID SERVER
LICENSE KEY

INSTALLATION

us Done

Please enter a license key. If you do not have a license key, you can get one from the Virtuozzo website (<https://www.virtuozzo.com/hybrid-server/trial-request/>). You can also install a license later using the 'sudo vzlicload -p KEY' command.

License Key:

If you do not have a license, you can fill in the [Free Trial Request](#) form.

Alternatively, you can activate a license with the `vzlicload` command after the installation is complete. For more details, see [Understanding Licensing](#).

Note: If the license is not activated, containers and virtual machines will be deactivated (suspended without a possibility to resume) after a grace period. All the newly created containers and virtual machines will run for 5 minutes. Containers with Virtuozzo Storage and Virtuozzo Automator management panels (if installed) will run for 24 hours.

2.2.8.2 Configuring ReadyKernel

With automatic updating enabled by default, ReadyKernel will check for new patches daily at 12:00 server time. If a patch is available, ReadyKernel will download, install, and load it for the current kernel. For more details, see [Updating the Kernel with ReadyKernel](#).

In previous Virtuozzo Hybrid Server versions, you can turn off the option of automatic update of Virtuozzo ReadyKernel patches while installing Virtuozzo Hybrid Server. In the latest versions, you can switch it off after installing Virtuozzo Hybrid Server. To do this, click **READYKERNEL**, clear the **ReadyKernel automatic update** checkbox, and click **Done**.


READYKERNEL

Done us

ReadyKernel automatic update

2.2.8.3 Completing Installation

Once the installation has been complete, click **Reboot** to restart the server.



Note: If you are installing Virtuozzo Hybrid Server from a USB drive, remove the drive before restarting the server.

After restart, you can continue building your Virtuozzo Hybrid Server infrastructure as described in *Building the Infrastructure* on page 60.

2.2.8.4 Building the Infrastructure

Your next steps will depend on which Virtuozzo Storage and Virtuozzo Automator components you have already installed and which you still need to install.

- If you installed Virtuozzo Storage and Virtuozzo Automator management panels on the first server, proceed to install storage and compute components, respectively, on the second and other servers.
- If you installed storage and compute components on a server and need to install it on more servers, repeat the performed installation steps. When on the **VIRTUOZZO STORAGE** or **VIRTUOZZO AUTOMATOR** screen, follow the instructions in *Installing Storage on the Second and Other Servers* on page 56 or *Installing Compute on the Second and Other Servers* on page 53.
- If you installed storage and compute components on the last server, do the following:
 - Log in to the Virtuozzo Storage management panel and make sure that all storage nodes are present in the **UNASSIGNED** list on the **NODES** screen.
 - Log in to the Virtuozzo Automator management panel and make sure that all compute nodes are present on the **Infrastructure > Hardware Nodes** tab.

With everything ready, you can start managing your infrastructure as described in the *Storage Administrator's Guide* and *|vauto| Administrator's Guide*.

2.3 Installing with CLI Management

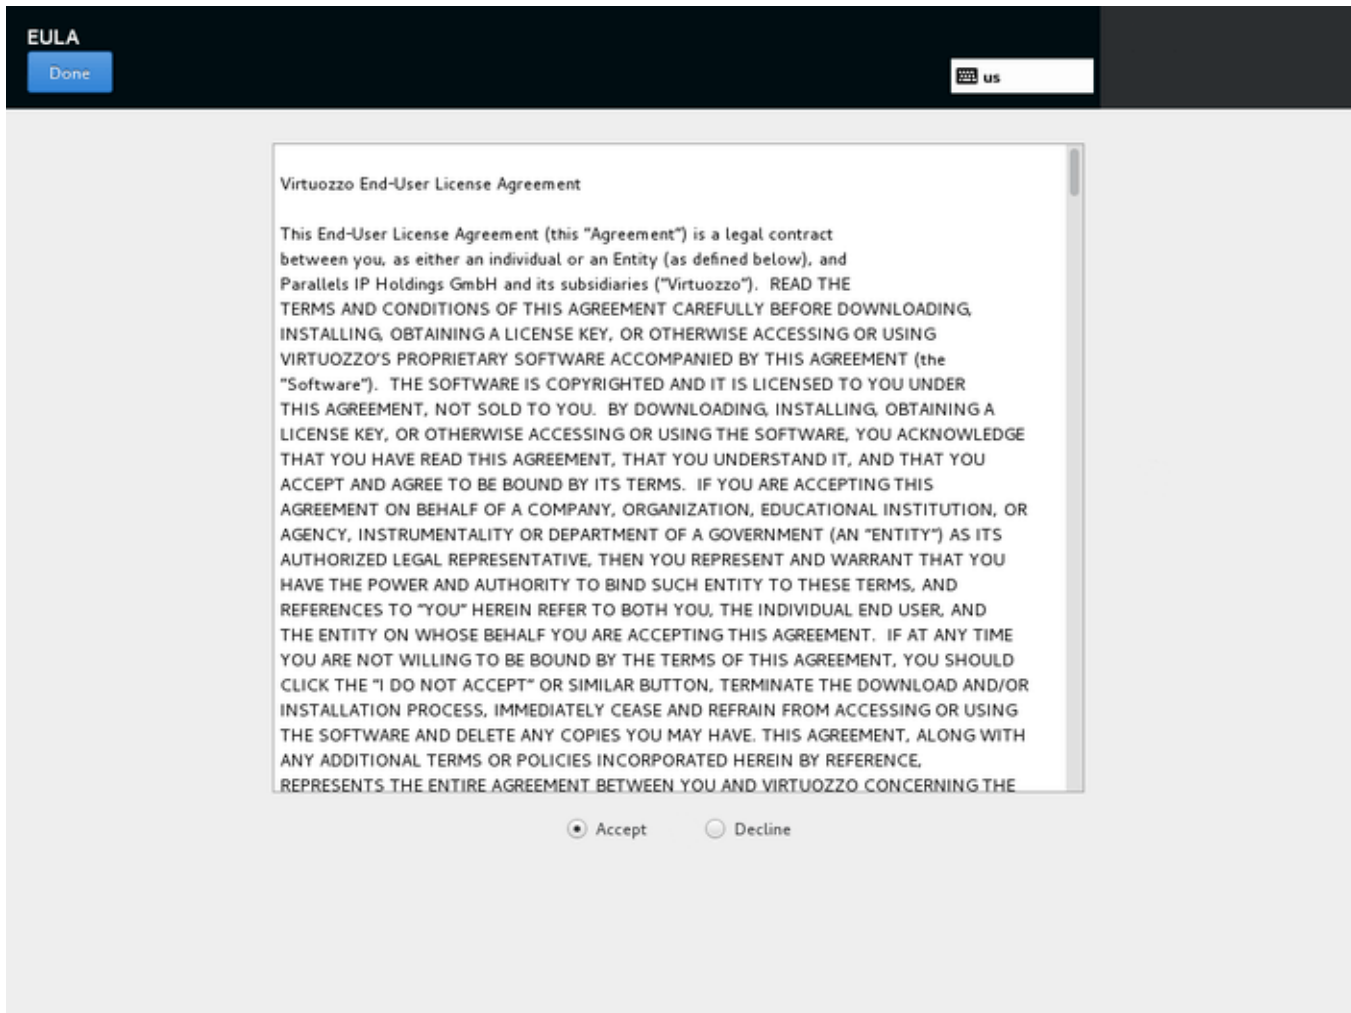
To install Virtuozzo Hybrid Server with CLI management in graphics mode, choose **Install with CLI management** on the welcome screen. After the installation program loads, you will see the **Installation Summary** screen. On this screen, you need to specify a number of parameters required to install Virtuozzo Hybrid Server.



2.3.1 Accepting EULA

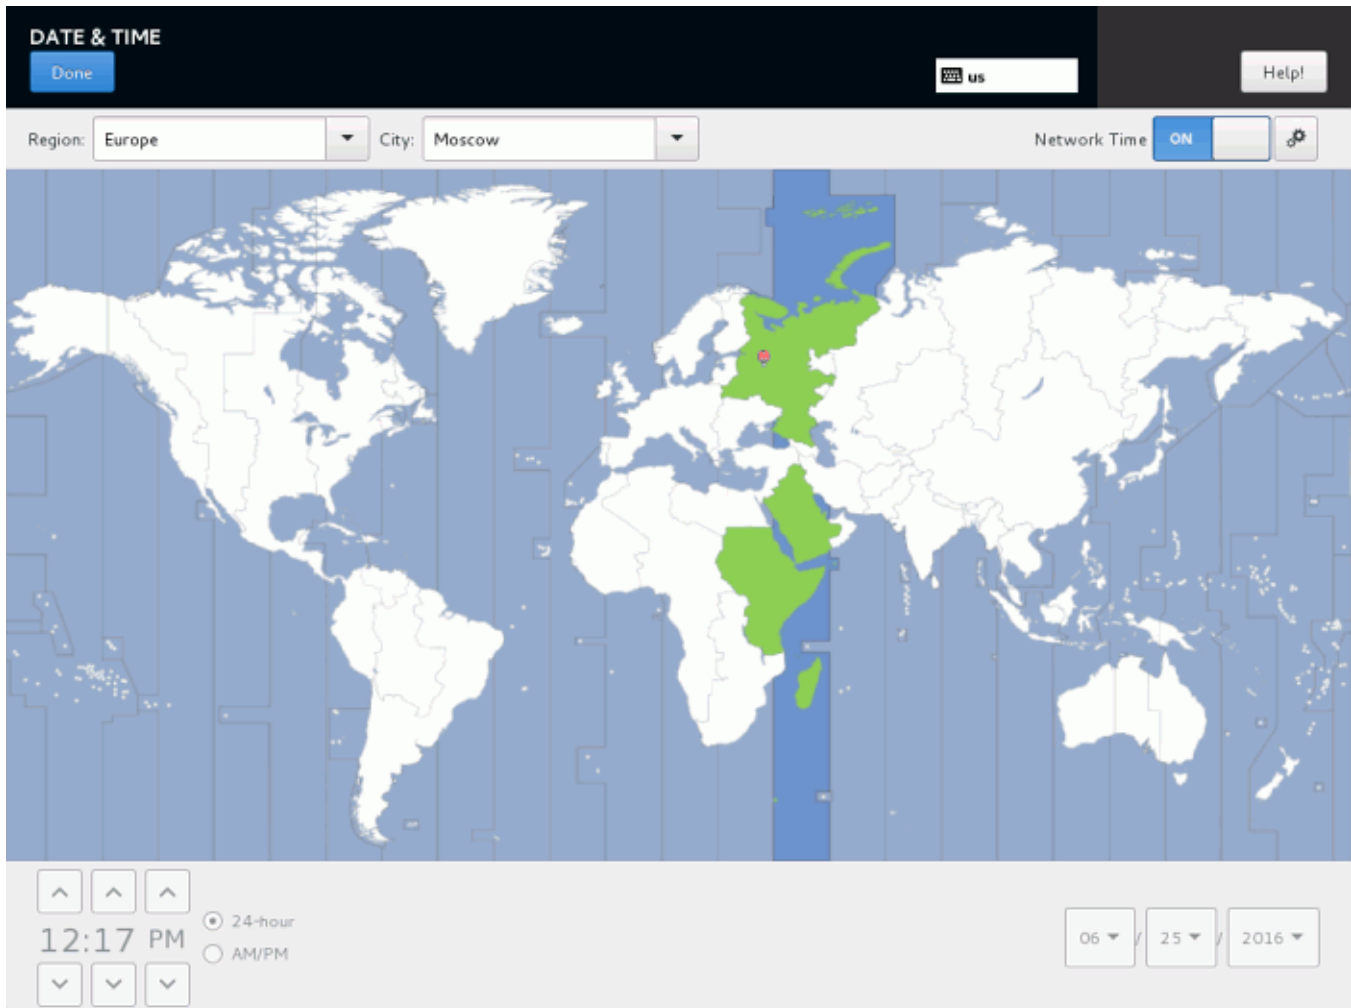
You need to accept the Virtuozzo Hybrid Server End-User License Agreement to install Virtuozzo Hybrid Server.

To accept the Virtuozzo Hybrid Server EULA, open the **EULA** screen, select **Accept**, and click **Done**.



2.3.2 Setting Date and Time

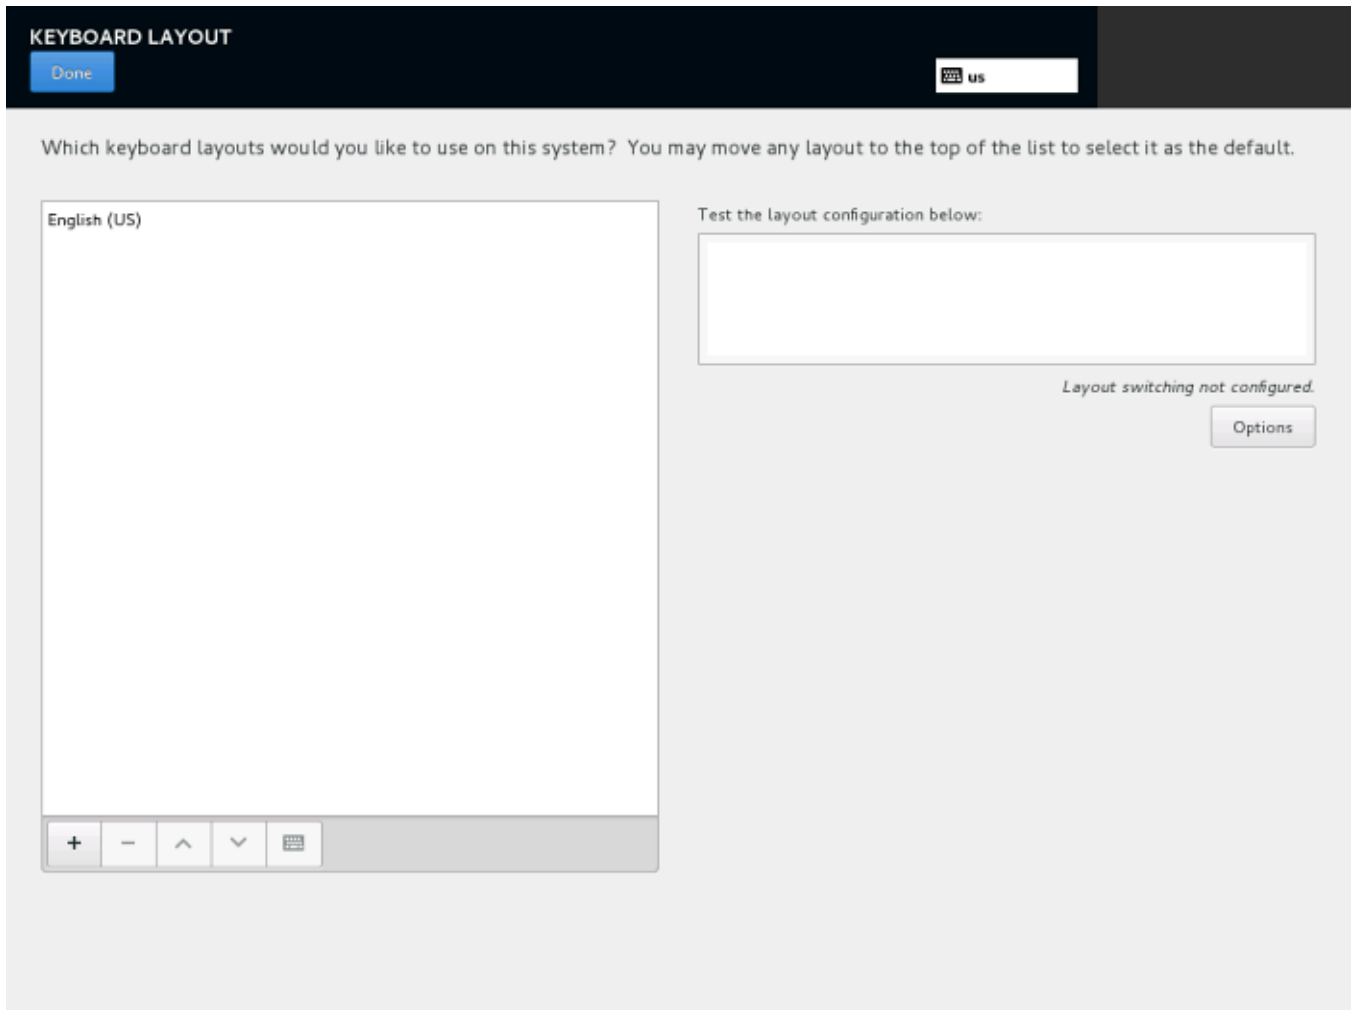
If you need to set the date and time for your Virtuozzo Hybrid Server installation, open the **DATE & TIME** screen and make the necessary changes.



2.3.3 Selecting the Keyboard Layout

The selected keyboard layout can be used during installation and, once the installation is complete, in the console (e.g., for entering localized descriptions, configuration file comments, and such).

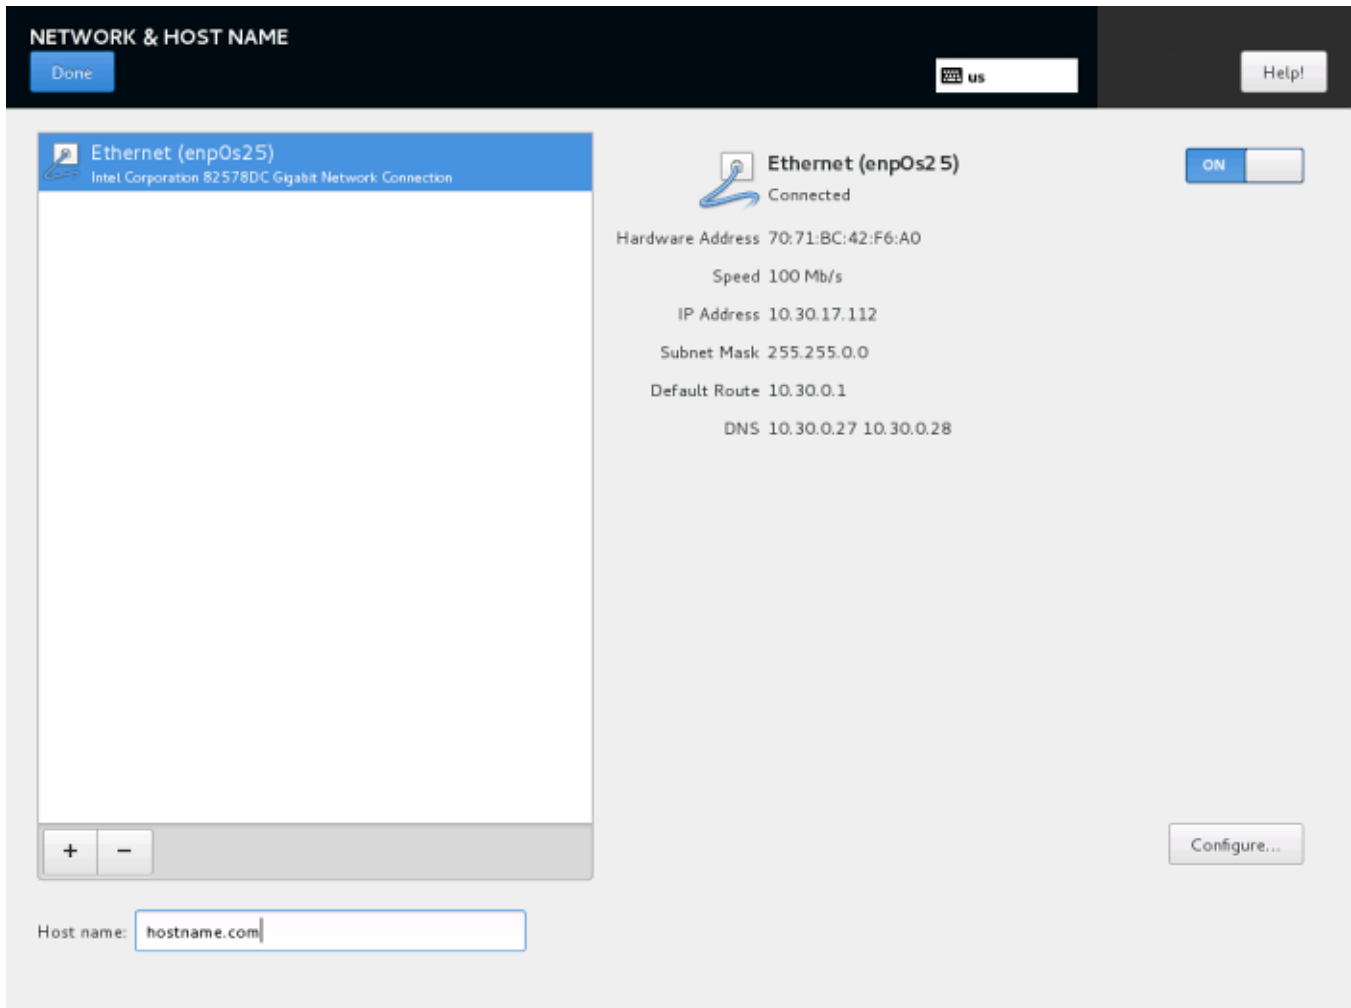
If you need to change the default English (US) keyboard layout, open the **KEYBOARD** screen, click the plus sign to add a layout, and click **Options** to choose a key combination for switching layouts.



2.3.4 Configuring Network

Usually network is configured automatically by the installation program. If you need to modify network settings, you can do so on the **NETWORK & HOST NAME** screen.

To install Virtuozzo Hybrid Server, you will need to have at least one network card configured and you will also need to provide a hostname: either a fully qualified domain name (<hostname>.<domainname>) or a short name (<hostname>).




2.3.4.1 Creating Bonded and Teamed Connections

Bonded and teamed connections offer increased throughput beyond the capabilities of a single network card as well as improved redundancy.

During installation, you can configure bonding on the **NETWORK & HOSTNAME** screen as described below. Teaming can be configured in a similar way after choosing **Team** on step 1.

1. To add a new bonded connection, click the plus button in the bottom, select **Bond** from the drop-down list, and click **Add**.

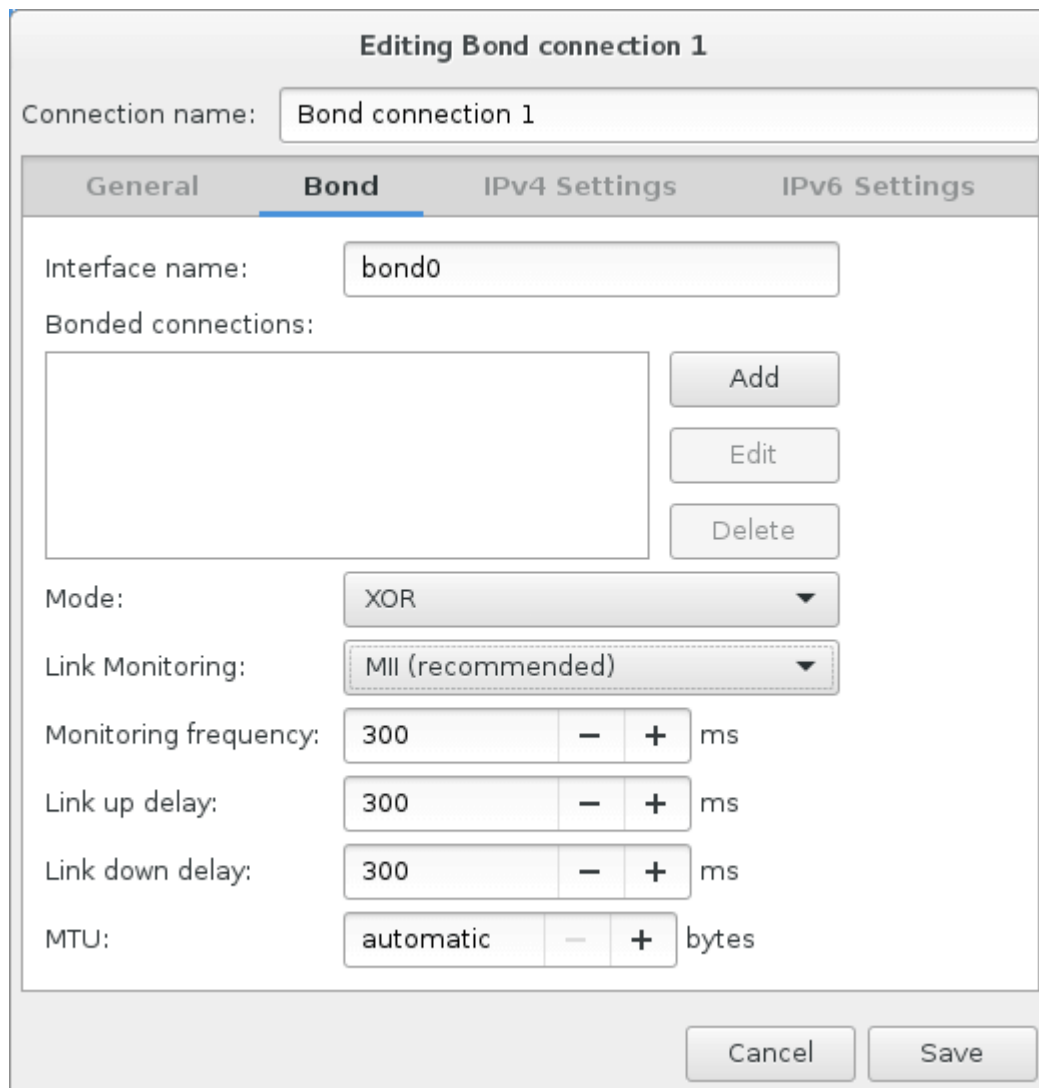


2. In the **Editing Bond connection...** window, set the following parameters:

2.1. **Mode** to XOR.

2.2. **Link Monitoring** to MII (recommended).

2.3. **Monitoring frequency**, **Link up delay**, and **Link down delay** to 300.



Note: It is also recommended to manually set `xmit_hash_policy` to `layer3+4` after the installation as detailed in [Setting Up Network Bonding](#). For more information on network bonding, see [Red Hat Enterprise Linux Deployment Guide](#) and [Linux Ethernet Bonding Driver HOWTO](#).

3. In the **Bonded connections** section on the **Bond** tab, click **Add**.
4. In the **Choose a Connection Type** window, select **Ethernet** from the in the drop-down list, and click **Create**.



5. In the **Editing bond slave...** window, select a network interface to bond from the **Device** drop-down list.

Editing bond0 slave 1

Connection name:

General	Ethernet	802.1X Security	DCB
Device: <input style="width: 90%;" type="text" value="ens37 (00:0C:29:ED:32:0E)"/>			
Cloned MAC address: <input style="width: 90%;" type="text"/>			
MTU: <input style="width: 60%;" type="text" value="automatic"/> <input type="button" value="-"/> <input type="button" value="+"/> bytes			
Wake on LAN: <input checked="" type="checkbox"/> Default <input type="checkbox"/> Phy <input type="checkbox"/> Unicast <input type="checkbox"/> Multicast <input type="checkbox"/> Ignore <input type="checkbox"/> Broadcast <input type="checkbox"/> Arp <input type="checkbox"/> Magic			
Wake on LAN password: <input style="width: 90%;" type="text"/>			

6. Configure other parameters if required and click **Save**.
7. Repeat steps 3 to 7 for each network interface you need to add to the bonded connection.
8. Configure other parameters if required and click **Save**.


The connection will appear in the list on the **NETWORK & HOSTNAME** screen.

The screenshot shows the Network & Hostname configuration window. On the left, a list of network interfaces is shown, with 'Bond (bond0)' selected. Below it, two Ethernet interfaces are listed: 'Ethernet (ens33)' and 'Ethernet (ens34)'. On the right, the configuration details for 'Bond (bond0)' are displayed. The status is 'Connected' with an 'ON' toggle switch. The configuration includes: Slaves: ens34, ens33; Hardware Address: 00:0C:29:17:90:E7; Speed: (blank); IP Address: 172.16.24.223; Subnet Mask: 255.255.254.0; Default Route: 172.16.24.1; DNS: 10.30.0.27 10.30.0.28.

2.3.4.2 Creating VLAN Adapters

While installing Virtuozzo Hybrid Server, you can also create virtual local area network (VLAN) adapters on the basis of physical adapters on the **NETWORK & HOSTNAME** screen as described below.

1. To add a new VLAN adapter, click the plus button in the bottom, select **VLAN** from the drop-down list, and click **Add**.



2. In the **Editing VLAN connection...** window:
 - 2.1. Select a physical adapter the VLAN adapter will be based on from the **Parent interface** drop-down list.
 - 2.2. Specify a VLAN adapter identifier in the **VLAN ID** field. The value must be in the 1-4094 range.

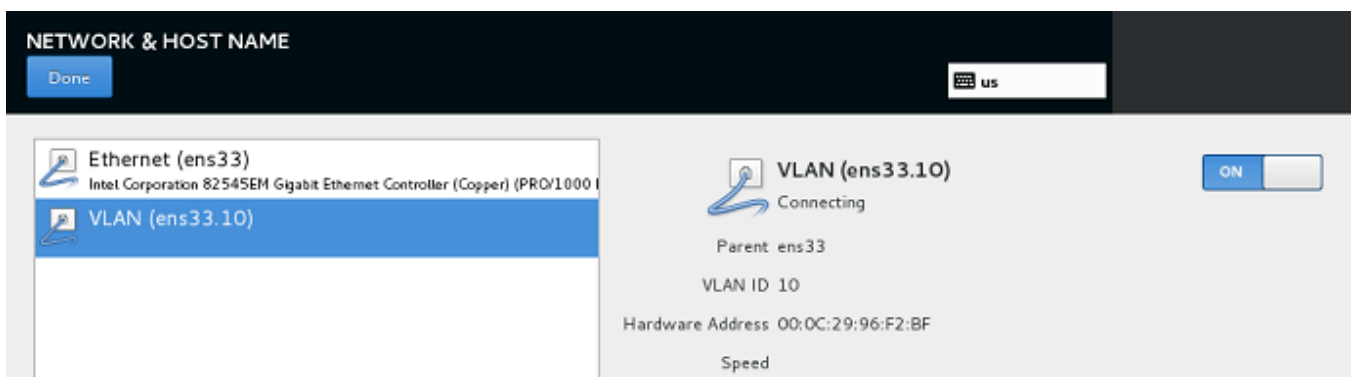
Editing VLAN connection 1

Connection name:

General	VLAN	IPv4 Settings	IPv6 Settings
Parent interface:	<input type="text" value="ens33 (00:0C:29:96:F2:BF)"/>		
VLAN id:	<input type="text" value="10"/>	<input type="button" value="-"/>	<input type="button" value="+"/>
VLAN interface name:	<input type="text"/>		
Cloned MAC address:	<input type="text"/>		
MTU:	<input type="text" value="automatic"/>	<input type="button" value="-"/>	<input type="button" value="+"/> bytes
Flags:	<input checked="" type="checkbox"/> Reorder headers <input type="checkbox"/> GVRP <input type="checkbox"/> Loose binding <input type="checkbox"/> MVRP		

3. Configure other parameters if required and click **Save**.

The VLAN adapter will appear in the list on the **NETWORK & HOSTNAME** screen.



2.3.5 Choosing the Storage Type

To choose the storage type, open the **SELECT STORAGE TYPE** screen:



Virtuozzo Hybrid Server with CLI management can be installed on two types of storage:

- Basic storage, i.e. local hard drive(s). This option is chosen by default and requires no configuration.
- Virtuozzo Storage, as a part of a new or existing Virtuozzo Storage cluster. If you choose this option, you will need to set additional options described in *Setting Virtuozzo Storage Installation Options* on page 71.

Note: Virtuozzo Storage is a solution that transforms local hard drives into a highly protected enterprise-level storage (like SAN or NAS) with data replication, high-availability, and self-healing features. Using Virtuozzo Storage, you can safely store and run virtual machines and containers, migrate them with zero downtime, provide high availability for your Virtuozzo Hybrid Server installations, and much more. For more information on Virtuozzo Storage, see the [Storage Administrator's Guide](#).

2.3.6 Setting Virtuozzo Storage Installation Options

If you choose to install on Virtuozzo Storage, you need to choose one of these options:

- Join the server to an existing Virtuozzo Storage cluster (see *Joining an Existing Virtuozzo Storage Cluster* on page 74), or

- Create a new Virtuozzo Storage cluster (see *Creating a New Virtuozzo Storage Cluster* on page 72)

Note: For detailed information on working with Virtuozzo Storage clusters, consult the [Storage Administrator's Guide](#).

2.3.6.1 Creating a New Virtuozzo Storage Cluster

If you choose to create a new Virtuozzo Storage cluster, you will need to provide the cluster name and password and select one or more roles for the server.

The screenshot shows a configuration window titled "SELECT STORAGE TYPE" with a dark header and a blue "Done" button. The main content area is light gray and contains the following sections:

- Select type of a Storage:**
 - Basic Storage
Installs or upgrades to typical types of storage devices.
 - Virtuozzo Storage
Installs or upgrades to a Virtuozzo Storage.
- Create or join Virtuozzo Storage:**
 - Join
You can join an existing Virtuozzo Storage in your network, if you have already set up one.
 - Create
You can create a new Virtuozzo Storage. Choose this option if you have no Virtuozzo Storage in your network
- Create a name for your Virtuozzo Storage cluster:**
 - Name:
 - This name will uniquely identify the Storage in your network.

First, in the **Name** field, specify a name for the cluster that will uniquely identify it among other clusters in your network. A cluster name must consist of the characters a-z, A-Z, 0-9, minus (-), underscore (_), and must not be longer than 63 characters.



Security Settings

Virtuozzo Storage Security Settings

Virtuozzo Storage clusters are password-protected. Create a password for the new cluster:

Password:

Confirm:

Show Password

Cancel OK

Next, click **Configure security** next to the **Name** field and specify a password for your cluster. The password must be at least 8 characters long. It is encrypted and saved to the file `/etc/vstorage/clusters/<cluster_name>/auth_digest.key` on the server.

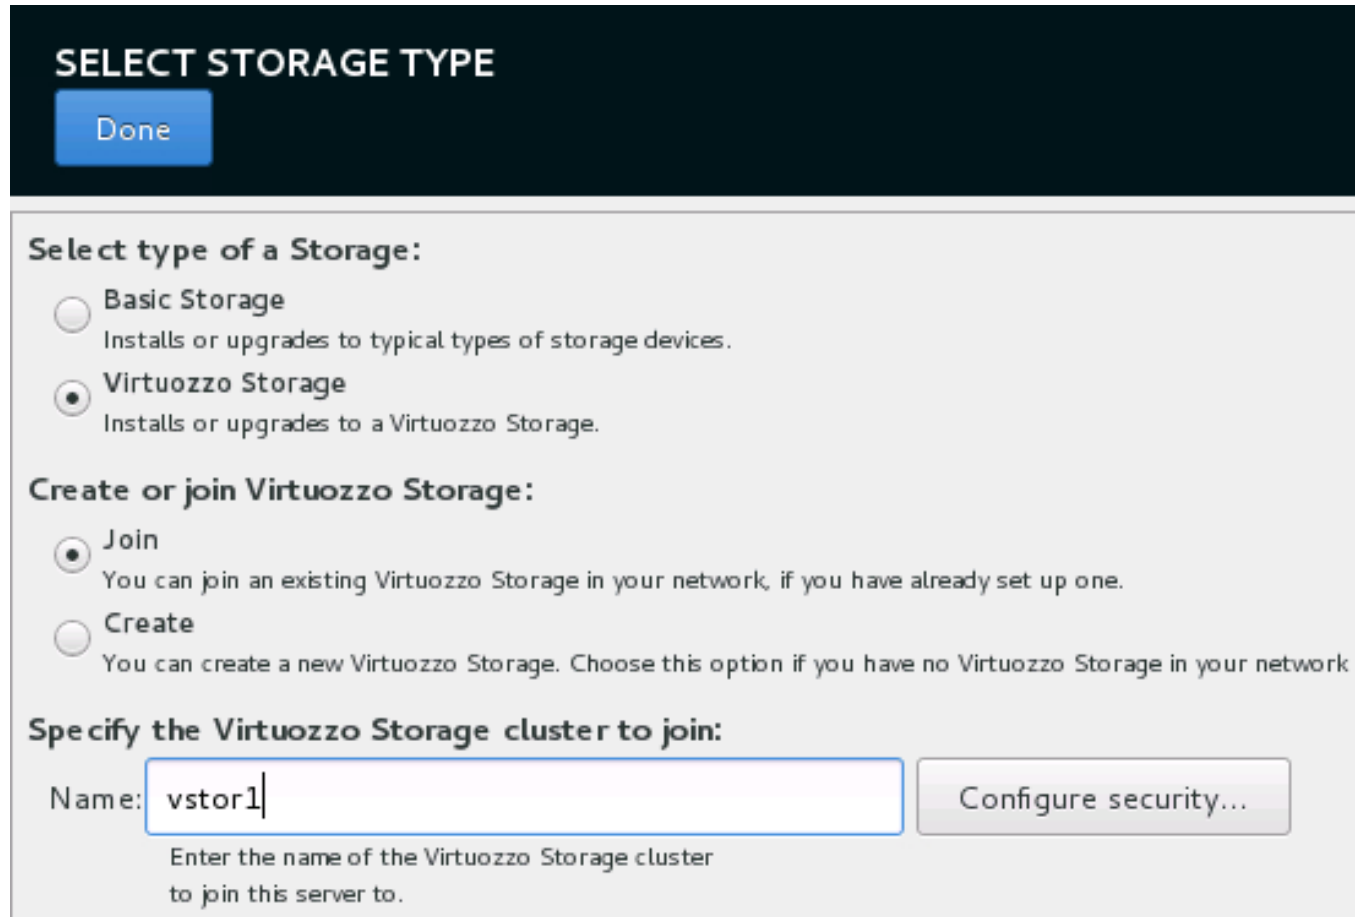
Note: A server needs to be authenticated only once. After that, you can configure it as a metadata server, chunk server, or a client. If later you decide to configure the server where you are setting the first MDS as a chunk server, no additional authentication will be required.

Next, choose the cluster role(s) for your server. See *Virtuozzo Storage Server Roles* on page 75. When you create a new Virtuozzo Storage cluster, the **Metadata Server Role** option is selected by default.

Finally, select an IP address through which Virtuozzo Storage services on this node will communicate with similar services on other cluster nodes. The network interface that has the IP address will be added to firewall's trusted zone and all ports will be open for it. It is recommended to keep this interface's network internal and separated from other networks. If the node has a single network interface, avoid opening all ports on it by creating a VLAN for internal traffic.

2.3.6.2 Joining an Existing Virtuozzo Storage Cluster

If you choose to join an existing Virtuozzo Storage cluster, you will need to provide the cluster name and password and select one or more roles for your server.



SELECT STORAGE TYPE

Done

Select type of a Storage:

- Basic Storage
Installs or upgrades to typical types of storage devices.
- Virtuozzo Storage
Installs or upgrades to a Virtuozzo Storage.

Create or join Virtuozzo Storage:

- Join
You can join an existing Virtuozzo Storage in your network, if you have already set up one.
- Create
You can create a new Virtuozzo Storage. Choose this option if you have no Virtuozzo Storage in your network

Specify the Virtuozzo Storage cluster to join:

Name:

Enter the name of the Virtuozzo Storage cluster to join this server to.

First, in the **Name** field, specify a name of the cluster to join.

Next, click **Configure security** next to the **Name** field and specify a password for your cluster.

Next, choose the role(s) for your server. See *Virtuozzo Storage Server Roles* on page 75.

Finally, select an IP address through which Virtuozzo Storage services on this node will communicate with similar services on other cluster nodes. The network interface that has the IP address will be added to firewall's trusted zone and all ports will be open for it. It is recommended to keep this interface's network internal and separated from other networks. If the node has a single network interface, avoid opening all ports on it by creating a VLAN for internal traffic.

2.3.6.3 Virtuozzo Storage Server Roles


Each Virtuozzo Hybrid Server server can have one, some, or all of the following roles.

- **Metadata Server Role.** MDS servers are an essential part of any Virtuozzo Storage cluster. They store metadata about chunk servers and control how files keeping the contents of virtual machines and containers are split into chunks and where these chunks are located.

As a static IP address is required for an MDS server, the IP address detected by the installation program is specified in the corresponding field by default. If multiple IP addresses are available, you need to choose one to assign to the MDS server. In some cases, you may need to enter a valid IP address manually.

Note: MDS servers require static IP addresses. If you are using DHCP, map an IP address to the MAC address of the MDS server.

- **Chunk Server Role.** Chunk servers (CS) store the contents of virtual machines and containers in the form of fixed-size chunks and provide access to these chunks. All data chunks are replicated and the replicas are kept on different chunk servers for high data availability. If one of the chunk servers goes down, other chunk servers continue providing the replicas of data chunks that were stored on the failed server.

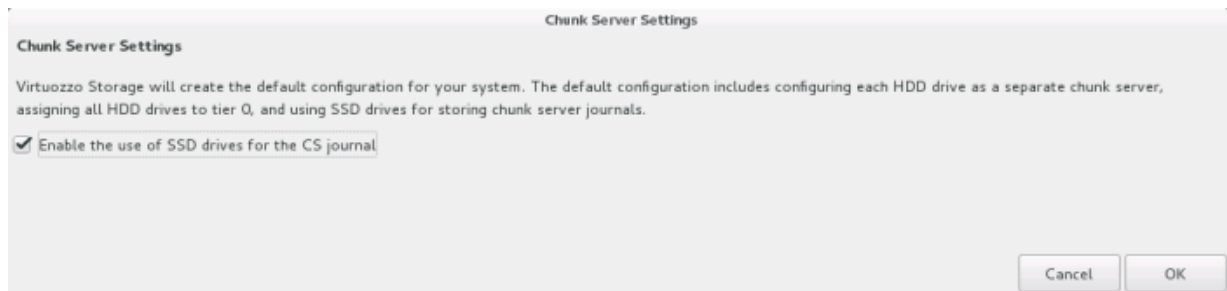


Warning: Virtuozzo Storage has redundancy built in, so you should avoid running Virtuozzo Storage on redundant types of RAID like 1, 5, or 6 over local storage. In this case, a single write operation may affect a significant number of HDDs resulting in very poor performance. For example, for 3 Virtuozzo Storage replicas and RAID5 on servers with 5 HDDs each, a single write operation may result in 15 I/O operations. For recommendations on optimal local storage configurations, consult the [Storage Administrator's Guide](#).

By default, the installer does the following:


- If your server has several disk drives, the installer will automatically configure each disk drive except system to act as a separate chunk server.
- If one or more SSD drives are available on the server, they will be set up to store chunk server write journals (one journal per chunk server). By using SSD drives for write journaling, you can boost the performance of write operations in the cluster by up to two or more times. For more information on using SSD drives, consult [Assigning Disk Roles Manually](#). If you need to disable this setting, click

Configure under the chunk server role checkbox and clear the checkbox **Enable the use of SSD drives for the CS journal**.



Note: If one or more SSDs are not detected automatically, find out their device names (for example, invoke the console by pressing **Ctrl+Alt+F2** and run `lsblk -d -o name,rota, zeroes` in the **ROTA** column will indicate non-rotational drives, i.e. SSDs), reboot to the installer welcome screen, and follow instructions in *Enabling Forced Detection of SSDs* on page 42.

- Client Server Role. Clients are computers with Virtuozzo Hybrid Server 7 from where you run virtual machines and containers stored in your Virtuozzo Storage cluster.



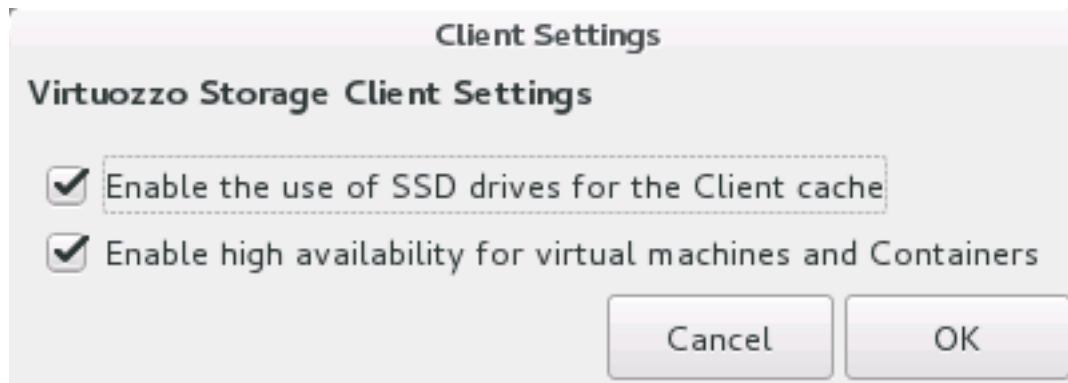
By default, the installer does the following:

- Enables high availability support for the client and for all virtual machines and containers you will create on it. With high availability turned on, if the client fails, all virtual machines and containers hosted on it will be automatically moved to a healthy server. For detailed information on high availability, consult [Managing High Availability Clusters](#).
- If one or more SSD drives are available on the server, the installer configures them to store a local cache of frequently accessed data. By having a local cache on an SSD drive, you can increase the overall cluster performance by up to 10 and more times. For more information on using SSD drives, consult [Assigning Disk Roles Manually](#).

Note: If one or more SSDs are not detected automatically, find out their drive letters (for example, invoke the console by pressing **Ctrl+Alt+F2** and run `lsblk -d -o name,rota, zeroes` in the **ROTA** column will indicate non-rotational drives, i.e. SSDs), reboot to the installer welcome screen, and

follow instructions in *Enabling Forced Detection of SSDs* on page 42.

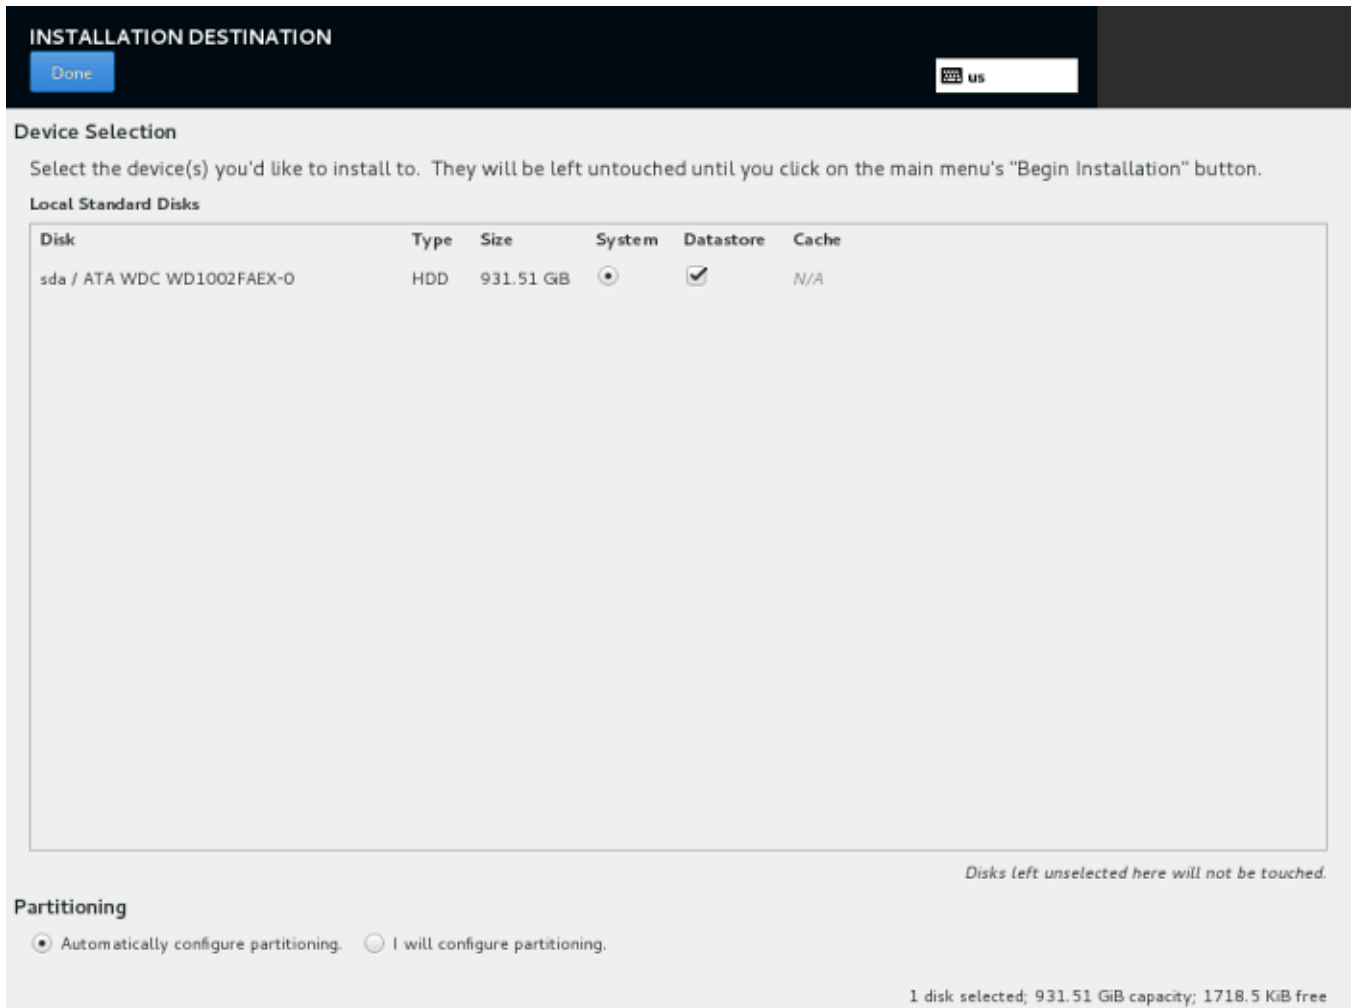
To change either of these settings, click **Configure** under the client role checkbox and set the corresponding checkboxes in the client settings window.



2.3.7 Partitioning the Hard Drives

Having chosen the storage type, you need to choose partitioning options on the **INSTALLATION DESTINATION** screen.

Warning: All the existing partitions on all disks will be deleted.



Firstly, you will need to choose which of the disks are marked as **System**, **Datastore**, and **Cache**:

- Use the **System** radio button to select a disk where the root partition with Virtuozzo Hybrid Server system files (mounted to /) will be kept.
- Use **Datastore** checkboxes to mark disks where virtual machines and containers will be kept. All such disks will be organized into a single volume group and mounted to the /vz mount point. At least one disk need to be marked as a data store. If Virtuozzo Storage chunk server role is selected, a chunk server will be created on each disk marked as a data store (in this case, the system disk cannot be marked as a data store).
- Use **Cache** checkboxes to mark SSD drives where journals and cache will be kept. This option is only applicable for SSD drives to be used in Virtuozzo Storage clusters.

Secondly, in the bottom of the screen, you will need to choose:

- **Automatically configure partitioning** and click **Done** to have the installation program create the default layout on the server.
- **I will configure partitioning** and click **Done** to manually partition your disk(s).

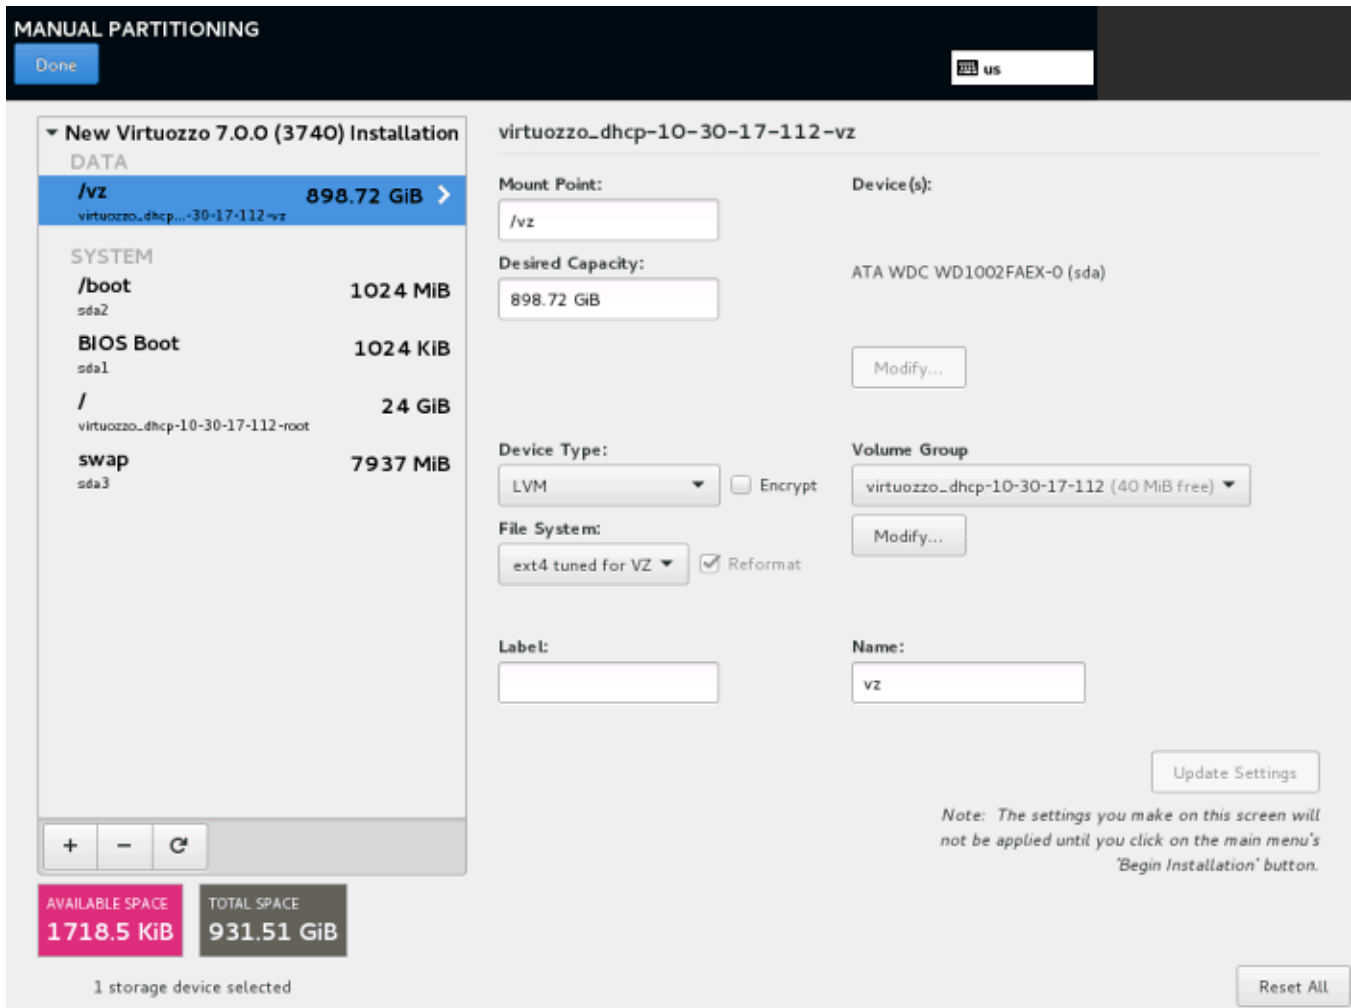
When partitioning the disks, keep in mind that Virtuozzo Hybrid Server requires these partitions:

- **Boot:** mount point `/boot`, 1 GB, a boot partition with Virtuozzo Hybrid Server boot files, created on each HDD.
- **Root:** mount point `/`, 12-24 GB, a root partition with Virtuozzo Hybrid Server files, created on the HDD marked **System**.
- **Swap:** a paging partition with the `swap` file system, created on the HDD marked **System**. The size depends on RAM size:
 - If RAM size is 4-8 GB, swap size should be equal to the RAM size.
 - Otherwise swap size should be half the RAM size.

Such method is beneficial for memory overcommit scenarios. If, however, overcommit is not required, you can determine the swap size using the old method:

- If the RAM size is 4-8 GB, the swap size should be equal to the RAM size.
 - If the RAM size is 8-64 GB, the swap size should be half the RAM size.
 - Otherwise the swap size should be 32 GB.
- **Data storage,** depends on the chosen storage type:
 - **Local:** mount point `/vz`, at least 30 GB, storage for virtual machines, containers, and OS and application templates. An LVM volume group that spans all HDDs marked **Datastore**. The `/vz` partition must be formatted to `ext4 tuned for VZ` (which is `ext4` with additional mount options that improve Virtuozzo Hybrid Server performance).
 - **Virtuozzo Storage chunk server:** mount point `<cluster_name>-cs<N>`, at least 100 GB, only required if the chunk server role is chosen. Chunk servers can only be created on partitions formatted to `ext4 tuned for CS` (which is `ext4` with additional mount options that improve Virtuozzo Storage performance).
 - In addition, either a 1 MB partition with the BIOS boot file system or a 200 MB partition with the EFI boot file system is required depending on your server configuration.

A typical partition layout for Virtuozzo Hybrid Server on basic storage may look like this:



A typical partition layout for Virtuozzo Hybrid Server on Virtuozzo Storage may look like this:

MANUAL PARTITIONING

Done en_US

▼ New Installation

DATA

/vz 30 GiB >
virtuozzo_dhcp-10-30-17-112-vz

/vstorage/vstor1-cs1 868.7 GiB
virtuozzo_dhcp-10...storage_vstor1-cs1

SYSTEM

/boot 1024 MiB
sda2

BIOS Boot 1024 KiB
sda1

/ 24 GiB
virtuozzo_dhcp-10-30-17-112-root

swap 7937 MiB
sda3

+ - ↻

AVAILABLE SPACE: 1718.5 KiB

TOTAL SPACE: 931.51 GiB

1 storage device selected

virtuozzo_dhcp-10-30-17-112-vz

Mount Point: /vz

Device (s): ATA WDC WD1002FAEX-0 (sda)

Desired Capacity: 30 GiB

Modify...

Device Type: LVM Encrypt

Volume Group: virtuozzo_dhcp-10-30-17-112 (60 MiB free)

File System: ext4 tuned for VZ Reformat

Modify...

Label:

Name: vz


Update Settings

Note: The settings you make on this screen will not be applied until you click on the main menu's 'Begin Installation' button.

Reset All

2.3.8 Finishing the Installation

Having configured everything necessary on the **INSTALLATION SUMMARY** screen, click **Begin Installation**.



While Virtuozzo Hybrid Server is installing, click **ROOT PASSWORD** to create a password for the root account. Installation will not finish until the password is created.

While Virtuozzo Hybrid Server is installing, you can activate a license and disable the automatic updating of the Virtuozzo ReadyKernel patch if needed. If you choose not to activate a license at this time, you may need to take additional steps to switch the memory management policy. For more details, see [Configuring Automatic Memory Management Policies](#).

2.3.8.1 Activating a License

To activate a Virtuozzo Hybrid Server license, click **LICENSE KEY**, enter a key, and click **Done**.

If you do not have a license, you can fill in the [Free Trial Request](#) form.


Alternatively, you can activate a license with the `vzlicload` command after the installation is complete. For more details, see [Understanding Licensing](#).

Note: If the license is not activated, containers and virtual machines will be deactivated (suspended without a possibility to resume) after a grace period. All the newly created containers and virtual machines will run for 5 minutes. Containers with Virtuozzo Storage and Virtuozzo Automator management panels (if installed) will run for 24 hours.

2.3.8.2 Configuring ReadyKernel


With automatic updating enabled by default, ReadyKernel will check for new patches daily at 12:00 server time. If a patch is available, ReadyKernel will download, install, and load it for the current kernel. For more details, see [Updating the Kernel with ReadyKernel](#).

In previous Virtuozzo Hybrid Server versions, you can turn off the option of automatic update of Virtuozzo ReadyKernel patches while installing Virtuozzo Hybrid Server. In the latest versions, you can switch it off after installing Virtuozzo Hybrid Server. To do this, click **READYKERNEL**, clear the **ReadyKernel automatic update** checkbox, and click **Done**.



2.3.8.3 Completing Installation

Once the installation is complete, click **Reboot** to restart the server.



Note: If you are installing Virtuozzo Hybrid Server from a USB drive, remove the drive before restarting the server.

After restart, you will see the login prompt as well as the server IP address and hostname that you can use to connect to the server remotely.

To manage virtual machines and containers on the Virtuozzo Hybrid Server server, you will need to log in as the root user. After you do so, you will see a shell prompt and can start creating and managing your virtual machines and containers. For quick-start instructions, run `man afterboot`. More detailed information is provided in the [Virtuozzo Hybrid Server 7 User's Guide](#).

2.4 Installing in the Basic Graphics Mode

If the installer cannot load the correct driver for your video card, you can try to install Virtuozzo Hybrid Server in the basic graphics mode. To select this mode, on the welcome screen, choose **Troubleshooting->**, then **Install in the basic graphics mode**. The installation process itself is the same as that in the default graphics mode (see *Installing with GUI Management* on page 44).

2.5 Installing via VNC

To install Virtuozzo Hybrid Server via VNC, boot to the welcome screen and do the following:

1. Select the main installation option **Install Virtuozzo Hybrid Server** and press **E** to start editing it.
2. Add text at the end of the line starting with `linux /images/pxeboot/vmlinuz`. For example:

```
linux /images/pxeboot/vmlinuz inst.stage2=hd:LABEL=<VZ7_ISO_image> quiet ip=dhcp text
```

3. Press **CTRL-X** to start booting the chosen installation option.
4. When presented with a choice of starting VNC or proceeding to the text mode, choose **1** for VNC.
5. When offered, enter a VNC password.
6. In the output that follows look up the hostname or IP address and VNC port to connect to, e.g.,
192.168.0.10:1.
7. Connect to said address in a VNC client. You will see the usual **Installation Summary** screen.

The installation process itself is the same as that in the default graphics mode (see *Installing Virtuozzo Hybrid Server* on page 41).

2.6 Installing in the Text Mode

To install Virtuozzo Hybrid Server in the text mode, boot to the welcome screen and do the following:

1. Select the main installation option **Install Virtuozzo Hybrid Server** and press **E** to start editing it.
2. Add text at the end of the line starting with `linux /images/pxeboot/vmlinuz`. For example:

```
linux /images/pxeboot/vmlinuz inst.stage2=hd:LABEL=<VZ7_ISO_image> quiet ip=dhcp text
```

3. Press **CTRL-X** to start booting the chosen installation option.
4. When presented with a choice of starting VNC or proceeding to the text mode, choose **2** for text mode.
5. In the installation menu that is shown, at least do the following: set or confirm the installation source (press **3**), the installation destination (press **6**), and the root password (press **11**) and accept the EULA (press **5**).
6. Press **b** to begin installation.
7. When installation ends, press **Enter** to reboot.

2.7 Configuring Server Ports

Virtuozzo Hybrid Server enables Linux kernel firewall during installation. This section lists ports opened by default. The set of ports differs depending on your system configuration:

- If the server does not participate in a Virtuozzo Storage cluster, see *Opened Ports on Standalone Servers* on page 86 for information on ports used by Virtuozzo Hybrid Server.
- If the server is part of a Virtuozzo Storage cluster, in addition, see *Opened Ports on Servers in Virtuozzo Storage Clusters* on page 87 for information on ports used by the cluster.

2.7.1 Opened Ports on Standalone Servers

The table below lists the ports for servers that do not participate in Virtuozzo Storage clusters. I in the **Description** column signals that the port should be open for incoming traffic and O means the same for outgoing traffic.

Port	Description
22	(IO) Used for secure logins via SSH.
80	(IO) Used for HTTP connections, e.g., to download Virtuozzo Hybrid Server updates and EZ templates from remote repositories.
21	(O) Used to connect to the Debian repository to cache Debian EZ templates.
443	(O) Used to send problem reports to the support team.
5224	(O) Used to connect to the Key Administrator server to update Virtuozzo Hybrid Server lease licenses.

Continued on next page

Table 2.7.1.1 -- continued from previous page

Port	Description
64000	(IO) Used to connect SDK with the dispatcher running on the remote server, and for communication between the dispatchers on different servers.
67	Used to support host-only adapters in virtual machines. Virtuozzo Hybrid Server does not use port 67 for any external connections.
<RPC ports>	Used by various RPC services (e.g., to support NFS shares). Port numbers may differ from system to system. To learn what RPC services are registered on your server and what ports they are using, run <code># rpcinfo -p localhost</code>
647, 847	Reserved by the Linux <code>portreserve</code> program for the DHCP server, if you use one.
5700-6900	Range of ports used for VNC connections.

You may also need to additionally open ports used to connect to remote yum repositories. Though most of the repositories can be accessed via HTTP, some may require access via HTTPS or FTP. You can check what repositories are currently configured for your system and what protocols are used to connect to them as follows:

```
# yum repolist -v | egrep -e 'baseurl|mirrors'
Repo-mirrors : <VZ7_mirrorlist_URL>
Repo-baseurl : <VZ7_base_URL>
# curl <VZ7_mirrorlist_URL>
```

2.7.2 Opened Ports on Servers in Virtuozzo Storage Clusters

A Virtuozzo Storage cluster requires ports listed below to be opened in addition to those on standalone servers. If you use Virtuozzo Storage management panel to create clusters, all the necessary ports are opened automatically. Otherwise, open these ports manually on each node participating in the cluster.

Port	Description
MDS Servers	
2510	(IO) Used for communication between MDS servers.
2511	(IO) Used for communication with chunks servers and clients.
Chunk Servers	
2511	(O) Used for communication with MDS servers.

Continued on next page

Table 2.7.2.1 -- continued from previous page

Port	Description
<random_port>	(I) Used for communication with clients. The chunk server management service automatically binds to any available port. You can also manually assign the service to a specific port.
Clients	
2511	(O) Used for communication with MDS servers.
<random_port>	(O) Used for communication with chunk servers. The client management service automatically binds to any available port. You can also manually assign the service to a specific port.

CHAPTER 3

Exploring Additional Installation Options

This chapter describes how to boot into the rescue mode, enable the expert mode, and run Virtuozzo Hybrid Server in virtual machines.

3.1 Booting into the Rescue Mode

If you experience problems with your system, you can boot into the rescue mode to troubleshoot these problems. Once you are in the rescue mode, your Virtuozzo Hybrid Server installation is mounted under `/mnt/sysimage`. You can go to this directory and make the necessary changes to your system.

To enter the rescue mode, do the following:

1. Boot your system from the chosen media.
2. On the welcome screen, click **Troubleshooting->**, then **Rescue system**.
3. Once Virtuozzo Hybrid Server boots into the emergency mode, press **Ctrl+D** to load the rescue environment.
4. In the rescue environment, you can choose one of the following options:
 - Continue (press **1**): mount the Virtuozzo Hybrid Server installation in read and write mode under `/mnt/sysimage`.
 - Read-only mount (press **2**): mount the Virtuozzo Hybrid Server installation in read-only mode under `/mnt/sysimage`.

- Skip to shell (press **3**): load shell, if your file system cannot be mounted; for example, when it is corrupted.
 - Quit (Reboot) (press **4**): reboot the server.
5. Unless you press **4**, a shell prompt will appear. In it, run `chroot /mnt/sysimage` to make the Virtuozzo Hybrid Server installation the root environment. Now you can run commands and try to fix the problems you are experiencing.
 6. After you fix the problem, run `exit` to exit the chroot environment, then `reboot` to restart the system.

3.2 Enabling the Expert Mode

The expert mode allows you to choose packages to install and manually partition disks. The latter capability is needed to install Virtuozzo Hybrid Server with GUI on a USB drive. It is recommended to only use the expert mode in this particular case.

To enable the expert mode, do the following:

1. Select the required installation option and press **E** to edit it.
2. Add `expert` at the end of the line starting with `linux /images/pxeboot/vmlinuz`. For example:

```
linux /images/pxeboot/vmlinuz inst.stage2=hd:LABEL=<VZ7_ISO_image> quiet ip=dhcp \  
expert
```

3. Press **Ctrl-X** to load the selected installation option.

3.3 Running Virtuozzo Hybrid Server in Virtual Machines

Warning: Nested virtualization is an experimental feature and tested only on Linux guests. The operation of nested virtual machines may be unstable.

Installing Virtuozzo Hybrid Server in virtual machines may prove useful if you want to evaluate Virtuozzo Hybrid Server.

To run virtual machines with Virtuozzo Hybrid Server, the physical server's processor(s) must support either of these architectures: Intel VT-x (with "unrestricted guest") and EPT. The following hypervisors are supported: Parallels Desktop for Mac, VMware Fusion, VMware Workstation, and VMware ESXi. Make sure that nested virtualization support is enabled in your hypervisor.

The following virtual hardware is recommended for virtual machines with Virtuozzo Hybrid Server:

- vCPUs: 2 or more
- RAM: 2 GB or more
- HDD: 64 GB or more

To install Virtuozzo Hybrid Server in a VM, copy the distribution ISO image to a local drive and create a new VM from it according to your virtualization software documentation. Start the VM, boot to the Virtuozzo Hybrid Server installer, and follow the instructions in *Installing Virtuozzo Hybrid Server* on page 41.

3.3.1 Restrictions and Peculiarities

When using Virtuozzo Hybrid Server in a virtualized environment, keep in mind the following restrictions and specifics:

- Running Virtuozzo Hybrid Server in a virtual machine entails a list of limitations that need to be taken into consideration:
 - Only containers are supported.
 - Performance degradation can be expected due to external limits on a given hypervisor.
 - Free memory reporting could be wrong due to overcommit management by the cloud provider.
 - SLA cannot be guaranteed if the problem is related to third-party cloud infrastructure or hypervisor. Such cases are reviewed with a cloud infrastructure/hypervisor representative in terms of the TSANET collaborative support program.
- If you change the configuration of a virtual machine where Virtuozzo Hybrid Server is installed, you may need to reactivate Virtuozzo Hybrid Server.
- When you start a virtual machine with Virtuozzo Hybrid Server, VMware Fusion may warn you that it requires full access to the network traffic. Ignore this message, and proceed with booting the virtual machine.
- To run in a virtualized Virtuozzo Hybrid Server environment, a virtual machine must have Virtuozzo

guest tools installed.

- To enable full support for virtual machines created inside Virtuozzo Hybrid Server, make sure to enable nested virtualization support for the Virtuozzo Hybrid Server VM in your virtualization software. Otherwise virtual machines created in Virtuozzo Hybrid Server will only support 32-bit operating systems and a single CPU.
- Containers created in Virtuozzo Hybrid Server running inside a virtual machine have no limitations and work as usual.