

^zVirtuozzo

Virtuozzo Storage 2.3

Installation Guide

December 14, 2017

Virtuozzo International GmbH

Vordergasse 59

8200 Schaffhausen

Switzerland

Tel: + 41 52 632 0411

Fax: + 41 52 672 2010

<https://virtuozzo.com>

Copyright ©2001-2017 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://virtuozzo.com>.

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. Introduction 1**
 - 1.1 About VirtuoZZO Storage 1
 - 1.2 VirtuoZZO Storage Deployment Overview 1

- 2. Planning VirtuoZZO Storage Infrastructure 3**
 - 2.1 Understanding VirtuoZZO Storage Architecture 3
 - 2.1.1 Storage Role 4
 - 2.1.2 Metadata Role 4
 - 2.1.3 Network Roles (Storage Access Points) 5
 - 2.1.4 Supplementary Roles 5
 - 2.2 Planning Node Hardware Configurations 6
 - 2.2.1 Hardware Requirements 6
 - 2.2.2 Hardware Recommendations 7
 - 2.2.2.1 Cluster Composition Recommendations 7
 - 2.2.2.2 General Hardware Recommendations 8
 - 2.2.2.3 Storage Hardware Recommendations 8
 - 2.2.2.4 Network Hardware Recommendations 10
 - 2.2.3 Hardware and Software Limitations 11
 - 2.2.4 Minimum Configuration 11
 - 2.2.5 Recommended Configuration 13
 - 2.2.5.1 HDD Only 13
 - 2.2.5.2 HDD + System SSD (No Cache) 14
 - 2.2.5.3 HDD + SSD 15
 - 2.2.5.4 SSD Only 15
 - 2.2.5.5 HDD + SSD (No Cache), 2 Tiers 16
 - 2.2.5.6 HDD + SSD, 3 Tiers 17

2.2.6	Raw Disk Space Considerations	18
2.3	Planning Network	18
2.3.1	General Network Requirements	19
2.3.2	Network Limitations	19
2.3.3	Per-Node Network Requirements	20
2.3.4	Network Interface Roles	24
2.3.5	Sample Network Configuration	25
2.4	Understanding Data Redundancy	27
2.4.1	Redundancy by Replication	28
2.4.2	Redundancy by Erasure Coding	29
2.4.3	No Redundancy	30
2.5	Understanding Failure Domains	31
2.6	Understanding Storage Tiers	31
2.7	Understanding Cluster Rebuilding	32
3.	Installing Virtuozzo Storage	34
3.1	Preparing for Virtuozzo Storage Installation	34
3.1.1	Preparing for Installation from USB Storage Drives	35
3.2	Starting Virtuozzo Storage Installation	35
3.3	Setting Date and Time	35
3.4	Configuring Network	36
3.4.1	Creating Bonded and Teamed Connections	36
3.5	Choosing Virtuozzo Storage Components to Install	39
3.5.1	Choosing the Components to Install on the First Server	39
3.5.2	Choosing the Components to Install on the Second and Other Servers	41
3.6	Selecting Destination Partition	44
3.7	Finishing Virtuozzo Storage Installation	44

CHAPTER 1

Introduction

To support the growing demand for both high performance and high data availability, modern data centers need a fast, flexible storage solution. Existing solutions, however, are often difficult to manage and maintain, or not flexible enough (e.g., local RAID arrays), or too expensive (e.g., storage area networks).

Virtuozzo Storage is designed to solve these issues. It can run on commodity hardware, so no significant infrastructure investments are needed. It is also easy to set up and grow on demand.

1.1 About Virtuozzo Storage

Virtuozzo Storage is a software-defined storage solution that allows you to quickly and easily transform low-cost commodity hardware and network equipment into protected enterprise-grade storage like storage area networks (SAN) or network-attached storage (NAS).

Virtuozzo Storage is optimized for storing large amounts of data and provides data redundancy (replication and erasure coding), high availability, self-healing, and storage sharing.

In Virtuozzo Storage, user data is stored on organized clusters of servers in the form of fixed-size chunks. These chunks are automatically replicated and distributed across available servers in the cluster to ensure high availability of user data.

Cluster storage space can be exported through access points like iSCSI, S3, or Acronis Backup Gateway.

1.2 Virtuozzo Storage Deployment Overview

To deploy Virtuozzo Storage for evaluation purposes or in production, you will need to do the following:

1. Plan the Virtuozzo Storage infrastructure.
2. Install and configure Virtuozzo Storage on each server in the planned infrastructure.
3. Create one or more clusters in Virtuozzo Storage.
4. Set up data export for the cluster(s).
5. Populate the cluster(s) with user data.

CHAPTER 2

Planning Virtuozzo Storage Infrastructure

To plan your Virtuozzo Storage infrastructure, you will need to decide on the hardware configuration of each server, plan the Virtuozzo 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 chapter is meant to help you complete all of these tasks.

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)
 - Acronis Backup Gateway access point service (ABGW)
 - S3 gateway (access point) service (GW)

- S3 name service (NS)
- S3 object service (OS)
- NFS service
- 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.

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 6).

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 provide system reliability, Virtuozzo Storage uses the Paxos consensus algorithm. It guarantees fault-tolerance if the majority of nodes running metadata services are healthy.

2.1. Understanding Virtuozzo Storage Architecture

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

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 and use the clusters as backend storage for Acronis Backup Cloud.

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, Dovecot, and Acronis Access. 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.
- **Acronis Backup Gateway**, allows you to connect Virtuozzo Storage to Acronis Backup Cloud via Acronis FES API.
- **NFS**, allows you to create redundant NFS exports in Virtuozzo Storage that can be mounted like regular directories shared over NFS.

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.

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 4. 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.

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.

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
System disk	See Storage disk below	250GB SATA HDD
Storage disk	Three 100GB SATA HDDs (one system, one storage, one MDS (on five nodes))	Four or more HDDs or SSDs; 1 DWPD endurance minimum, 10 DWPD recommended
Disk controller	None	HBA or RAID
Network	1Gbps or faster network interface	Two 10Gbps network interfaces; dedicated links for internal and public networks
SSD	None	One or more recommended enterprise-grade SSDs with power loss protection 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.

2.2. Planning Node Hardware Configurations

Type	Minimal	Recommended
Sample configuration		Intel Xeon E5-2620V2, 32GB, 2xST1000NM0033, 32xST6000NM0024, 2xMegaRAID SAS 9271/9201, Intel X540-T2, Intel P3700 800GB

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:

2.2.2.1 Cluster Composition Recommendations

Designing an efficient cluster means finding a compromise between performance and cost that suits your purposes. When planning, keep in mind that a cluster with many nodes and few disks per node offers higher performance while a cluster with the minimal number of nodes (5) and a lot of disks per node is cheaper. See the following table for more details.

Design considerations	Minimum nodes (5), many disks per node	Many nodes, few disks per node
Optimization	Lower cost.	Higher performance.
Free disk space to reserve	More space to reserve for cluster rebuilding as fewer healthy nodes will have to store the data from a failed node.	Less space to reserve for cluster rebuilding as more healthy nodes will have to store the data from a failed node.
Redundancy	Fewer erasure coding choices.	More erasure coding choices.
Cluster balance and rebuilding performance	Worse balance and slower rebuilding.	Better balance and faster rebuilding.
Network capacity	More network bandwidth required to maintain cluster performance during rebuilding.	Less network bandwidth required to maintain cluster performance during rebuilding.
Favorable data type	Cold data (e.g., backups).	Hot data (e.g., virtual environments).
Sample server configuration	Supermicro SSG-6047R-E1R36L (Intel Xeon E5-2620 v4 CPU, 32GB RAM, 36 x 12TB HDDs, 1 x 500GB system disk).	Supermicro SYS-2028TP-HC0R-SIOM (4 x Intel E5-2620 v4 CPUs, 4 x 16GB RAM, 24 x 1.9TB Samsung SM863a SSDs).

Note:

1. These considerations only apply if failure domain is host.
2. The speed of rebuilding in the replication mode does not depend on the number of nodes in the cluster.
3. Virtuozzo Storage supports hundreds of disks per node. If you plan to use more than 36 disks per node, contact sales engineers who will help you design a more efficient cluster.

2.2.2.2 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 in virtual machines or on top of SAN/NAS hardware that has its own redundancy mechanisms. Doing so may negatively affect performance and data availability.
- To achieve best performance, keep at least 20% of cluster capacity free.
- During disaster recovery, Virtuozzo Storage may need additional disk space for replication. Make sure to reserve at least as much space as any single storage node has.

2.2.2.3 Storage Hardware Recommendations

- It is possible to use disks of different size in the same cluster. However, keep in mind that, given the same IOPS, smaller disks will offer higher performance per terabyte of data compared to bigger disks. It is recommended to group disks with the same IOPS per terabyte in the same tier.
- Using the recommended SSD models may help you avoid loss of data. Not all SSD drives can withstand

2.2. Planning Node Hardware Configurations

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. We recommend the following enterprise-grade SSD drives which are the best in terms of performance, endurance, and investments: Intel S3710, Intel P3700, Huawei ES3000 V2, Samsung SM1635, and Sandisk Lightning.
- 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.

2.2.2.4 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

2.2. Planning Node Hardware Configurations

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.

2.2.3 Hardware and Software Limitations

Hardware limitations:

- Each physical server must have at least 3 disks: for the operation system, metadata, and storage. Servers with fewer disks cannot be added to clusters.
- Five servers are required to test all the features of the product.
- The system disk must have at least 100 GBs of space.

Software limitations:

- 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 19.

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, Object Storage private, S3 public, NFS, ABGW

Node #	1st disk role	2nd disk role	3rd and other disk roles	Access points
2	System	Metadata	Storage	iSCSI, Object Storage private, S3 public, NFS, ABGW
3	System	Metadata	Storage	iSCSI, Object Storage private, S3 public, NFS, ABGW
4	System	Metadata	Storage	iSCSI, Object Storage private, ABGW
5	System	Metadata	Storage	iSCSI, Object Storage private, ABGW
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, an 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.

Important: If you deploy Virtuozzo Storage on a single node, you must take care of making its storage persistent and redundant to avoid data loss. If the node is physical, it must have multiple disks so you can replicate the data among them. If the node is a virtual machine, make sure that this VM is made highly available by the solution it runs on.

Acronis Backup Gateway works with the local object storage in the staging mode. It means that the data to be replicated, migrated, or uploaded to a public cloud is first stored locally and only then sent to the destination. It is vital that the local object storage is persistent and redundant so the local data does not get lost. There are multiple ways to ensure the persistence and redundancy of the local storage. You can deploy your Acronis Backup Gateway on multiple nodes and select a good redundancy mode. If your gateway is deployed on a single node in Virtuozzo Storage, you can make its storage redundant by replicating it among multiple local disks. If your entire Virtuozzo Storage installation is deployed in a single virtual machine with

2.2. Planning Node Hardware Configurations

the sole purpose of creating a gateway, make sure this VM is made highly available by the solution it runs on.

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, Object Storage private, S3 public, ABGW
Nodes 6+	System	SSD; cache	Storage	iSCSI, Object Storage private, ABGW
5 or more nodes in total		5 MDSs in total	5 or more CSs in total	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.

2.2.5.1 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

Disk No.	Disk Type	Disk Role(s)
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

2.2.5.2 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

2.2. Planning Node Hardware Configurations

2.2.5.3 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

2.2.5.4 SSD Only

This configuration does not require SSDs for cache.

When choosing hardware for this configuration, have in mind the following:

- Each Virtuozzo Storage client will be able to obtain up to about 40K sustainable IOPS (read + write) from the cluster.
- If you use the erasure coding redundancy scheme, each erasure coding file, e.g., a single VM's or container's HDD disk, will get up to 2K sustainable IOPS. That is, a user working inside a VM or container will have up to 2K sustainable IOPS per virtual HDD at their disposal. Multiple VMs and containers on a node can utilize more IOPS, up to the client's limit.
- In this configuration, network latency defines more than half of overall performance, so make sure that the network latency is minimal. One recommendation is to have one 10Gbps switch between any two

nodes in the cluster.

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
...		
N	SSD	CS

2.2.5.5 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)

2.2. Planning Node Hardware Configurations

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

2.2.5.6 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	
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

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.

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.

2.3. Planning Network



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.

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.

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 8888 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.




2.3. Planning Network

- 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 management node must have the port 8888 open by default to allow access to the management panel from the public network and to the cluster node from the internal network.

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



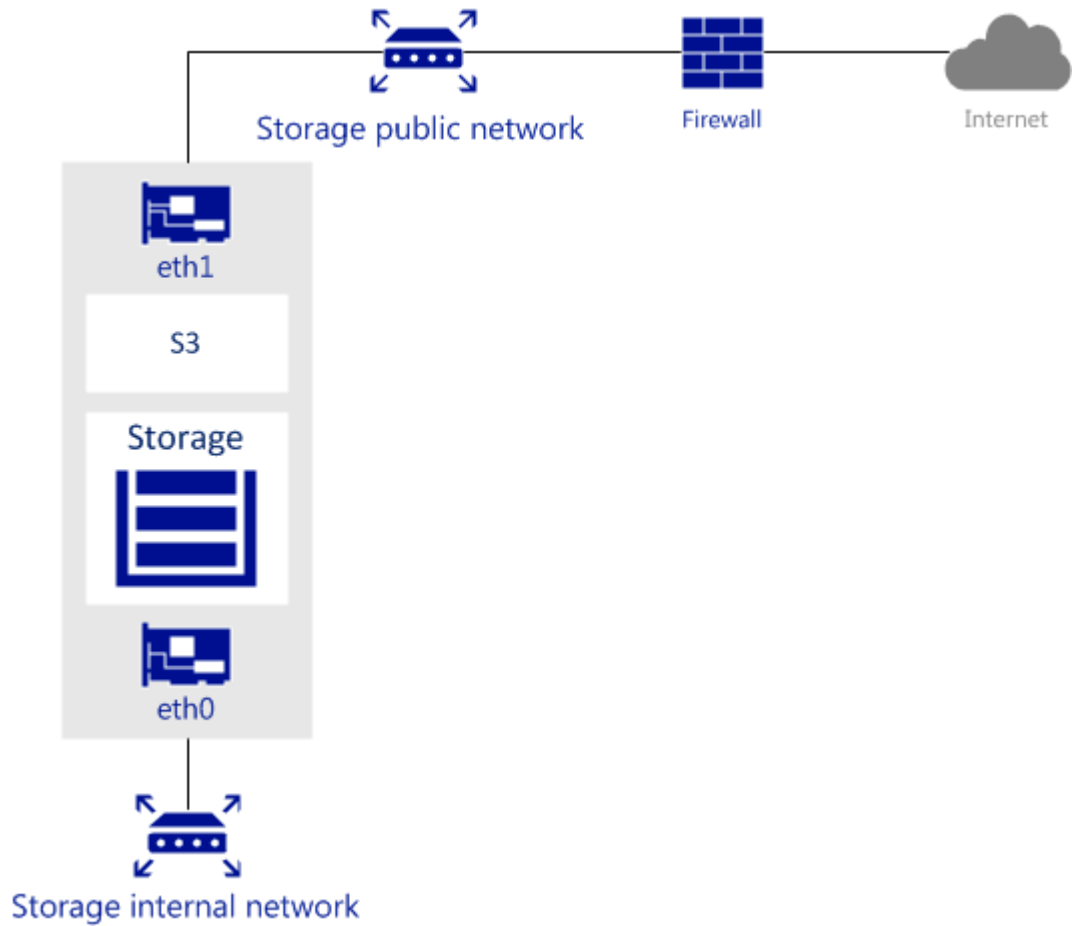
- A node that runs one or more storage access point services must have a network interface for the internal network traffic and a network interface for the public network traffic.

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.

2.3. Planning Network



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

The next figure shows a sample network configuration for a node with an Acronis Backup Gateway storage access point. Acronis Backup Gateway access points use port 44445 for incoming connections from both internal and public networks and ports 443 and 8443 for outgoing connections to the public network.



2.3.4 Network Interface Roles

For an 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 private.** The network interface will be used by the S3 and NFS storage access

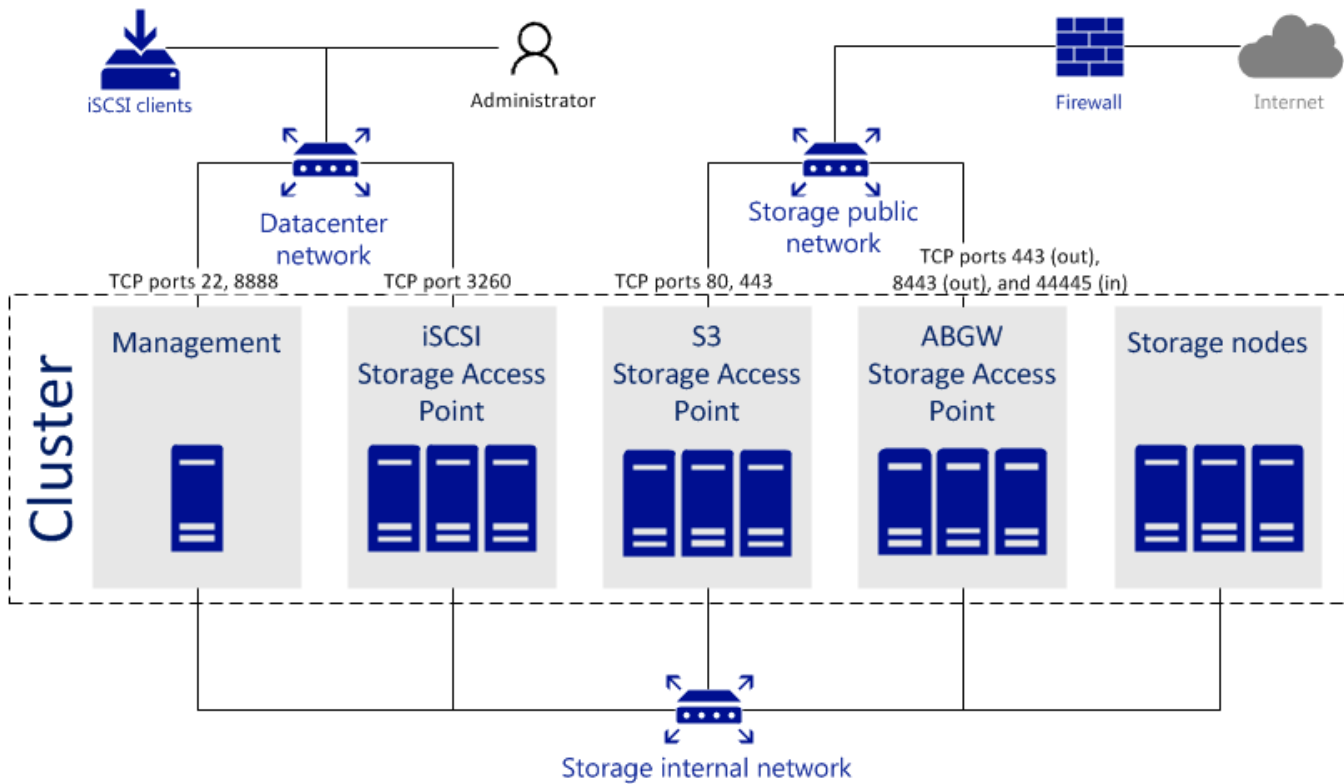
2.3. Planning Network

points. 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 or NFS storage access point service.

- **ABGW private.** The network interface will be used by the Acronis Backup gateway 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 Acronis Backup gateway 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.
 - **ABGW public.** The network interface will be used by the Acronis Backup gateway 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 Acronis Backup Cloud agents.
 - **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.
 - **NFS.** The network interface will be used by the NFS 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 NFS clients.
- **Custom.** These roles allow you to open specific ports on public network interfaces.

2.3.5 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.

2.4. Understanding Data Redundancy

- An external public network can be used to export the storage space to the outside services through S3 and Acronis Backup Gateway storage access points, that is, for the S3 (public) and Acronis Backup Gateway roles.

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. In general, replication offers better performance while erasure coding leaves more storage space available for data (see table).

Virtuozzo Storage supports a number of modes for each redundancy method. The following table illustrates data overhead of various redundancy modes. The first three lines are replication and the rest are erasure coding.

Redundancy mode	Minimum number of nodes required	How many nodes 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	2	200	300GB
Encoding 3+2	5	2	67	167GB

Redundancy mode	Minimum number of nodes required	How many nodes can fail without data loss	Storage overhead, %	Raw space required to store 100GB of data
Encoding 5+2	7	2	40	140GB
Encoding 7+2	9	2	29	129GB
Encoding 17+3	20	3	18	118GB

Note: 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.

You choose a data redundancy mode when configuring storage access points and their volumes. In particular, when:

- creating LUNs for iSCSI storage access points,
- creating S3 clusters,
- configuring Acronis Backup Gateway storage access points.

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.

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.

2.4. Understanding Data Redundancy



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 20).
- 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.

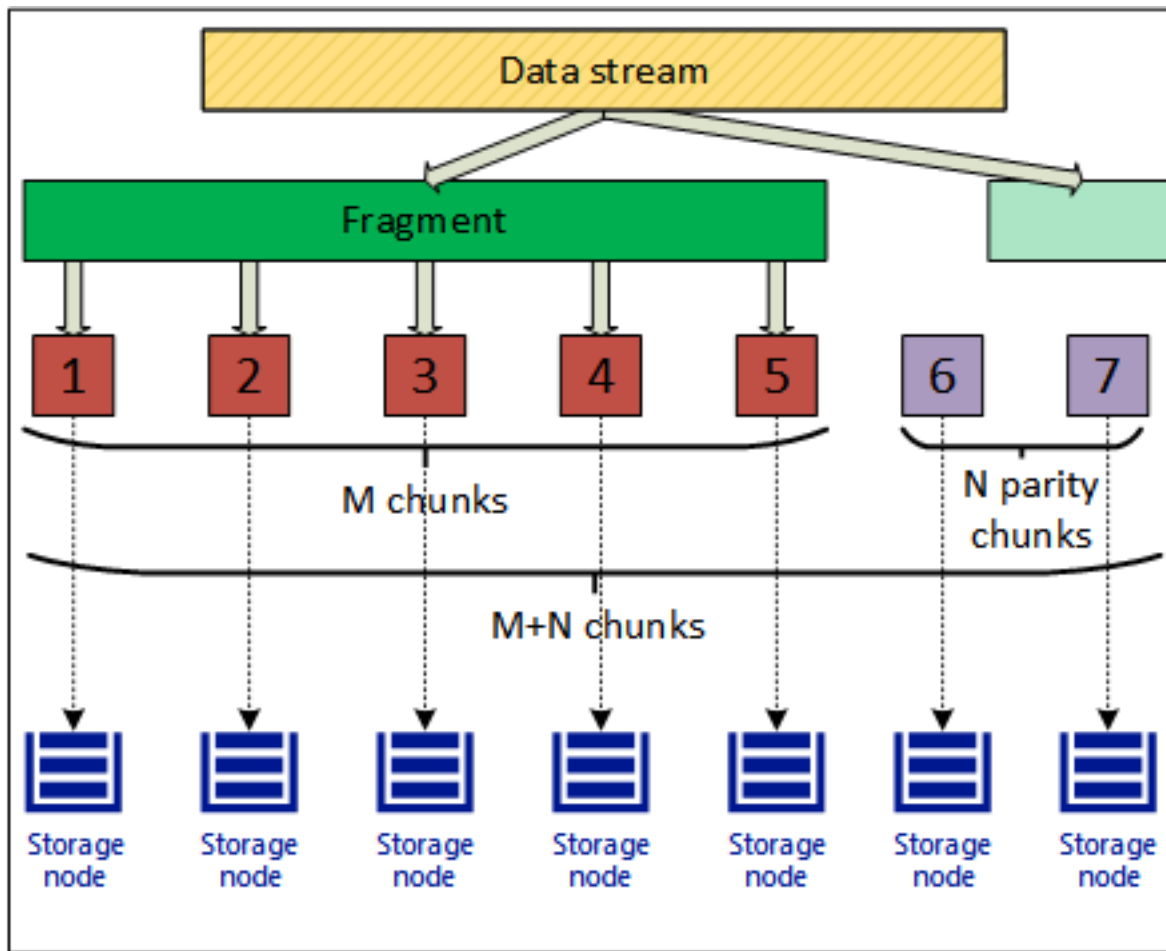
2.4.2 Redundancy by Erasure Coding

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.



2.4.3 No Redundancy

Warning: Danger of data loss!

2.5. Understanding Failure Domains

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.

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).

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 the *Virtuozzo Storage Administrator's Command Line Guide*.

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

2.7 Understanding Cluster Rebuilding

Virtuozzo Storage clusters are self-healing. If a node or disk fails, a cluster will automatically try to restore the lost data, i.e. rebuild itself.

The rebuild process involves the following steps. Every CS sends a heartbeat message to an MDS every 5 seconds. If a heartbeat is not sent, the CS is considered *inactive* and the MDS informs all cluster components that they stop requesting operations on its data. If no heartbeats are received from a CS for 15 minutes, the MDS considers that CS *offline* and starts cluster rebuilding (if prerequisites below are met). In the process, the MDS finds CSs that do not have pieces (replicas) of the lost data and restores the data—one piece (replica) at a time—as follows:

- If replication is used, the existing replicas of a degraded chunk are locked (to make sure all replicas remain identical) and one is copied to the new CS. If at this time a client needs to read some data that has not been rebuilt yet, it reads any remaining replica of that data.
- If erasure coding is used, the new CS requests almost all the remaining data pieces to rebuild the missing ones. If at this time a client needs to read some data that has not been rebuilt yet, that data is rebuilt out of turn and then read.

Self-healing requires more network traffic and CPU resources if replication is used. On the other hand, rebuilding with erasure coding is slower.

For a cluster to be able to rebuild itself, it must have at least:

2.7. Understanding Cluster Rebuilding

1. as many healthy nodes as required by the redundancy mode;
2. enough free space to accommodate as much data as any one node can store.

The first prerequisite can be explained on the following example. In a cluster that works in the 5+2 erasure coding mode and has seven nodes (i.e. the minimum), each piece of user data is distributed to 5+2 nodes for redundancy, i.e. each node is used. If one or two nodes fail, the user data will not be lost, but the cluster will become degraded and will not be able to rebuild itself until at least seven nodes are healthy again (that is, until you add the missing nodes). For comparison, in a cluster that works in the 5+2 erasure coding mode and has ten nodes, each piece of user data is distributed to the random 5+2 nodes out of ten to even out the load on CSs. If up to three nodes fail, such a cluster will still have enough nodes to rebuild itself.

The second prerequisite can be explained on the following example. In a cluster that has ten 10 TB nodes, at least 1 TB on each node should be kept free, so if a node fails, its 9 TB of data can be rebuilt on the remaining nine nodes. If, however, a cluster has ten 10 TB nodes and one 20 TB node, each smaller node should have at least 2 TB free in case the largest node fails (while the largest node should have 1 TB free).

Two recommendations that help smooth out rebuilding overhead:

- To simplify rebuilding, keep uniform disk counts and capacity sizes on all nodes.
- Rebuilding places additional load on the network and increases the latency of read and write operations. The more network bandwidth the cluster has, the faster rebuilding will be completed and bandwidth freed up.

CHAPTER 3

Installing Virtuozzo Storage

After creating a plan of your Virtuozzo Storage infrastructure, proceed to install Virtuozzo Storage on each server included in the plan.

Virtuozzo Storage is installed in a similar way on all required servers. One exception is the first server where you must also install the management panel (only one is allowed per cluster).

Note: On all nodes in the same cluster, time needs to be synchronized via NTP. Make sure the nodes can access the NTP server.

3.1 Preparing for Virtuozzo Storage Installation

Virtuozzo Storage can be installed from

- DVD discs (burn the distribution ISO image onto a DVD disc),
- PXE servers (see the *Installation via PXE Server* guide for information on installing Virtuozzo Storage over the network),

Note: Time synchronization via NTP is enabled by default.

- USB drives.

3.2. Starting Virtuozzo Storage Installation

3.1.1 Preparing for Installation from USB Storage Drives

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

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

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

For example, on Linux:

```
# dd if=storage-image.iso of=/dev/sdb
```

And on Windows (with `dd` for Windows):

```
C:\>dd if=storage-image.iso of=\\?\Device\Harddisk1\Partition0
```

3.2 Starting Virtuozzo Storage Installation

To start the installation, do the following:

1. Configure the server to boot from a DVD or USB drive.
2. Boot the server from the chosen media and wait for the welcome screen.
3. On the welcome screen, choose **Install Virtuozzo Storage**. After the installation program loads, you will see the **Installation Summary** screen. On it, you need to specify a number of parameters required to install Virtuozzo Storage.

3.3 Setting Date and Time

If you need to set the date and time for your Virtuozzo Storage installation, open the **DATE & TIME** screen and make the necessary changes. Make sure that NTP is enabled to synchronize time on each node.

3.4 Configuring Network

Virtuozzo Storage requires at least one network interface per server for management (that is, the management role). You will specify this network interface on the **Component Installation** screen and will not be able to remove the management role from it later (you will, however, be able to add more network roles to it).

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


If manual configuration is required, specify the necessary parameters for at least one network card and provide a hostname: either a fully qualified domain name (hostname, domainname) or a short name (hostname).

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.

3.4. Configuring Network

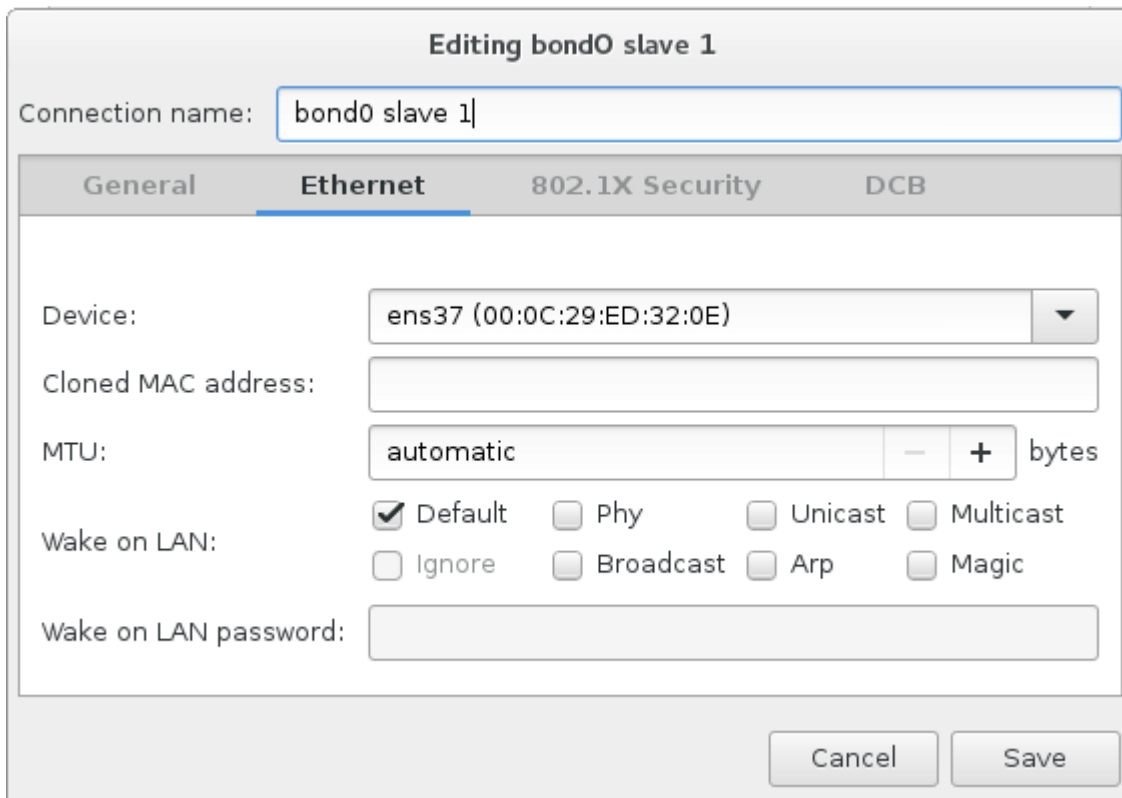
The screenshot shows a dialog box titled "Editing Bond connection 1". At the top, the "Connection name" is "Bond connection 1". Below this are four tabs: "General", "Bond", "IPv4 Settings", and "IPv6 Settings". The "Bond" tab is selected. Under "Interface name", the value is "bond0". The "Bonded connections" section is empty, with "Add", "Edit", and "Delete" buttons to its right. The "Mode" is set to "XOR" in a dropdown menu. "Link Monitoring" is set to "MII (recommended)" in a dropdown menu. "Monitoring frequency" is set to "300" with minus and plus buttons and "ms" to the right. "Link up delay" is set to "300" with minus and plus buttons and "ms" to the right. "Link down delay" is set to "300" with minus and plus buttons and "ms" to the right. "MTU" is set to "automatic" with minus and plus buttons and "bytes" to the right. At the bottom right are "Cancel" and "Save" buttons.

Note: It is also recommended to manually set `xmit_hash_policy` to `layer3+4` after the installation. 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**.



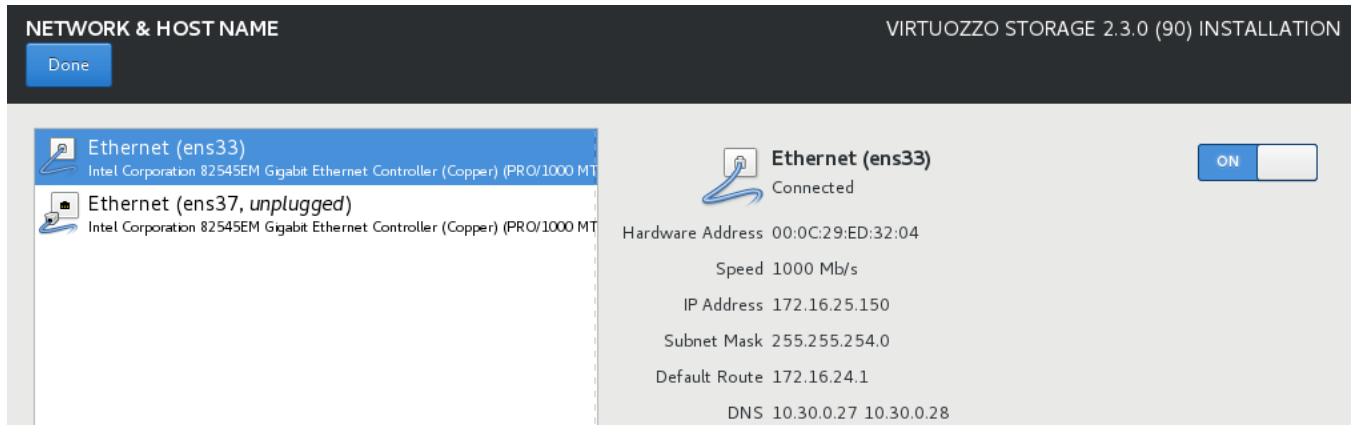
5. In the **Editing bond slave...** window, select a network interface to bond from the **Device** drop-down list.



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.

3.5. Choosing Virtuozzo Storage Components to Install



3.5 Choosing Virtuozzo Storage Components to Install

To install Virtuozzo Storage on a server, you need to choose a component to install on the **Virtuozzo Storage** screen.

The following options are available:

- **Management Panel.** Install the web-based user interface for managing Virtuozzo Storage clusters.
- **Storage.** Turn the server into a node ready to run Virtuozzo Storage services related to data storage.
- **Management Panel and Storage.** Install both above components at once.

You will need to install **Management Panel** or **Management Panel and Storage** on the first server and **Storage** on all other servers. The detailed instructions are provided in the following sections.

Note: The management panel will be installed on the system disk.

3.5.1 Choosing the Components to Install on the First Server

On the first server, you will need to install the management panel (with or without storage, as per your plan).

Do the following on the **Virtuozzo Storage** screen:

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

2. In the **Management Panel network** drop-down list, select a network interface that will provide access to the management panel.
3. In the **Management network** drop-down list, select a network interface and specify a port for internal management and configuration purposes (the port 8888 is used by default).
4. Create a password for the `admin` account of the management panel and confirm it in the corresponding fields.
5. Click **Done**.

Component Installation

- Management Panel.** The web user interface for adding and managing storage nodes.
- Storage.** Choose this option only if the Management Panel is already installed.
- Management Panel and Storage.** Both components at once.

Important: Only one management panel is required, so choose this option for the first node only!

Management network

eno16777736 - 10.30.26.129 ▼

The Management Network is used by the management node to configure and manage storage nodes. It can also be used by storage administrators for accessing storage nodes directly via SSH. This network should be protected and inaccessible over WAN. It can be the same as the private Storage Network used for communication between storage nodes.

Management Panel network

eno16777736 - 10.30.26.129 ▼

The Management Control Panel Network is used by storage administrators to access the web control panel of Storage. In most cases, it can be the same as the Management Network. If, however, the Management Network is only accessible by storage nodes, choose another network for the control panel, one that can be accessed by storage administrators. For security reasons, the web control panel should not be accessible from public/WAN networks.

Create a password for the Management Panel

Confirm the password

After completing the steps above, proceed to *Selecting Destination Partition* on page 44.

3.5. Choosing Virtuozzo Storage Components to Install

3.5.2 Choosing the Components to Install on the Second and Other Servers

On the second and other servers, you will need to install the **Storage** component only. Such servers will run services related to data storage and will be added to the Virtuozzo Storage infrastructure 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. On any computer with access to the management panel network, open a web browser and visit the management node IP address on port 8888: https://<management_node_IP_address>:8888. If prompted, add the security certificate to browser's exceptions.
2. In the management panel:
 - If you only installed the management component on the first server, you will see the welcome screen where a token will be shown (you can generate a new one if needed; generating a new token invalidates the old one).

Add storage nodes to Virtuozzo Storage

On this screen, you can get instructions on how to add storage nodes to your Virtuozzo Storage infrastructure. To add storage nodes to Virtuozzo Storage, do the following:

- 1 On this screen, generate a token required for adding the storage nodes.

b2843c5d	
Valid till Apr 25, 2017, 4:52 pm	Generate new token

The token can be used to add multiple storage nodes. Generating a new token invalidates the previous token, so you will need to provide the new token when adding nodes.

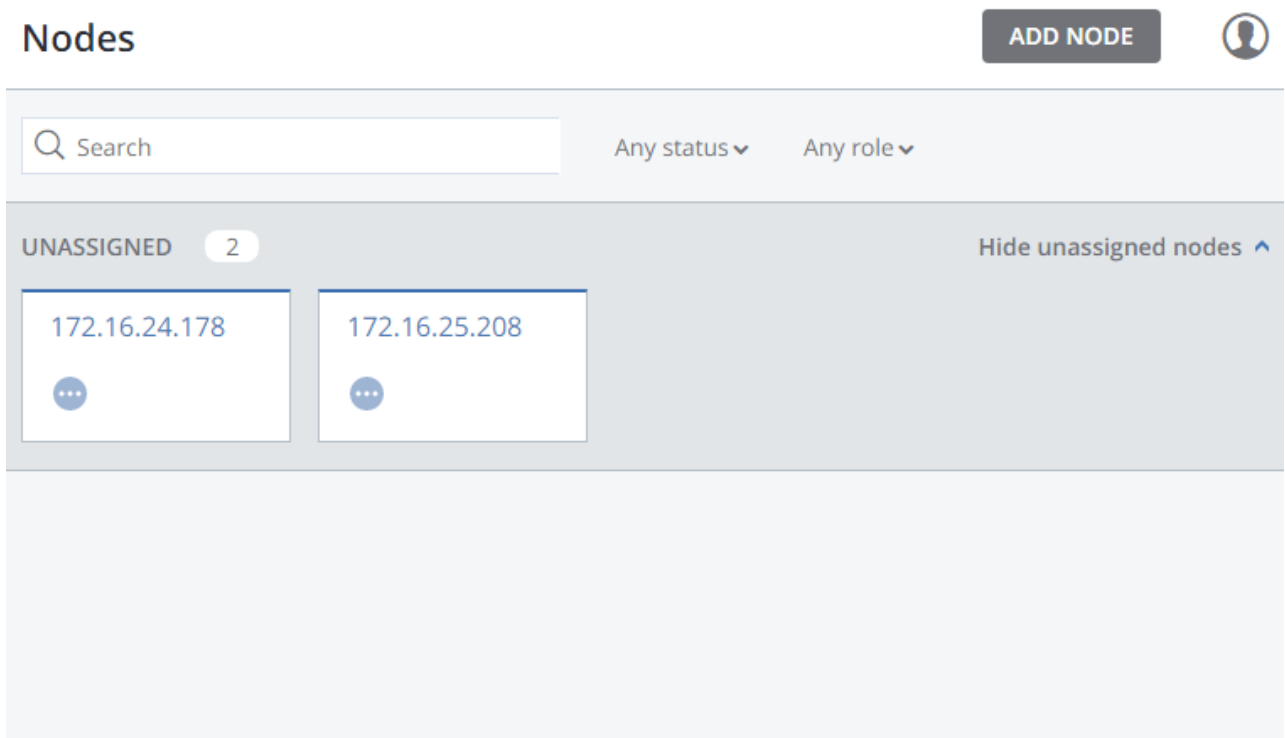
- 2 Boot the node from the ISO image.
- 3 Choose the Storage option on the VIRTUOZZO STORAGE screen.
- 4 Specify the IP address or hostname of the management panel in the Management panel field.
- 5 Specify the token shown on this screen in the Token field.
- 6 Complete other installation steps if required and finish the installation.
- 7 Repeat the steps from 2 to 6 for each storage node you need to add to Virtuozzo Storage.



- If you installed both the **Management** and **Storage** components on the first server, you will see the **NODES** screen where the only node will be shown in the **UNASSIGNED** list.

Click **ADD NODE** to open a screen similar to the welcome one. On it, a token will be shown (you can generate a new one if needed; generating a new token invalidates the old one).

3.5. Choosing Virtuozzo Storage Components to Install



Having obtained the token, do the following on the **Virtuozzo Storage** screen:

1. Choose **Storage**.

Component Installation

Management Panel. The web user interface for adding and managing storage nodes.

Storage. Choose this option only if the Management Panel is already installed.

Management Panel and Storage. Both components at once.

Management Node

Enter the IP address or hostname of the node with the Management Panel

Token

Enter a token for the new storage node.
To obtain a token, click "ADD NODE" on the "Nodes" screen in the Management Panel.

2. In the **Management node** field, specify the IP address of the node with the management panel.
3. In the **Token** field, specify the acquired token.
4. Click **Done** and proceed to *Selecting Destination Partition* on page 44.

3.6 Selecting Destination Partition

You need to choose on which server disk the operating system will be installed. This disk will have the system supplementary role and will not be used for data storage. To choose a system disk, open the **INSTALLATION DESTINATION** screen and select a device in the **Device Selection** section. Configure other options if required.

3.7 Finishing Virtuozzo Storage Installation

On the **ROOT PASSWORD** screen, create a password for the root account. Installation cannot be started until the password is created.

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

Once the installation is complete, the server will automatically reboot.

Your next steps depend on which server you installed Virtuozzo Storage:

- If you installed the management component on the first server (with or without the storage component), proceed to install the storage component on the second and other servers.
- If you installed the storage component on a server and need to install it on more servers, repeat the installation steps. When on the **Virtuozzo Storage** screen, follow the instructions in *Choosing the Components to Install on the Second and Other Servers* on page 41.
- If you installed the storage component on the last server, log in to the management panel and make sure that all the storage nodes are present in the **UNASSIGNED** list on the **NODES** screen.

With the management panel ready and with all the nodes present in the **UNASSIGNED** list, you can start managing your Virtuozzo Storage infrastructure as described in the *Virtuozzo Storage 2 Administrator's Guide*.