

---

Parallels

# Deploying RHCS Clusters in Parallels Virtuozzo-Based Systems



*ISBN: N/A*  
*Parallels Holdings Ltd.*  
*13755 Sunrise Valley Drive*  
*Suite 600*  
*Herndon, VA 20171*  
*USA*  
*Tel: +1 (703) 815 5670*  
*Fax: +1 (703) 815 5675*

*© 1999-2008 Parallels Holdings Ltd. All rights reserved.*  
*Distribution of this work or derivative of this work in any form is prohibited unless prior written permission is obtained from the copyright holder.*

---

# Contents

<b>Introduction</b>	<b>4</b>
<hr/>	
<b>Deploying Virtuozzo Failover Cluster</b>	<b>5</b>
<hr/>	
Cluster Configuration Options .....	5
Setting Up Cluster Servers.....	6
Installing Host OS .....	6
Installing and Configuring Cluster Software .....	7
Configuring Shared Storage .....	7
Installing Parallels Containers 4.0 .....	8
Setting Up Clustered Services .....	9
Configuring Virtuozzo Failover Cluster .....	9
Testing Virtuozzo Cluster.....	12
<hr/>	
<b>Deploying Data Sharing Cluster</b>	<b>13</b>
<hr/>	
Cluster Creation Process Overview .....	14
Setting Up Hardware For Cluster .....	14
Installing and Configuring Host OS .....	15
Installing Clustering Software .....	15
Setting Up Red Hat Cluster .....	16
Configuring Cluster Storage .....	18
Installing and Configuring Virtuozzo Containers 4.0 .....	19
Creating and Configuring Cluster Services .....	20
<hr/>	
<b>Index</b>	<b>23</b>
<hr/>	

---

## CHAPTER 1

# Introduction

Virtuozzo Containers 4.0 for Linux allows you to use Red Hat Cluster Suite (RHCS) on Hardware Nodes running the Red Hat Enterprise Server 5 (RHEL 5) operating system to deploy the following types of clusters:

- A Virtuozzo failover cluster used to ensure high-availability for your Virtuozzo Containers 4.0 installations. In this type of cluster, each Virtuozzo critical service and Container is treated by the RHCS clustering software as a clustered service and, if necessary, failed over from the active cluster node to a standby one.
- A GFS(Global File System) Virtuozzo cluster. This type of cluster provides a consistent file system image across all 'RHEL 5'-based Hardware Nodes in the cluster and allows the Nodes to simultaneously read and write to a single shared file system on a SAN data storage.

The procedure of deploying both types of RHCS clusters is described in the following sections in detail.

---

**Note:** The document provides the description of the Virtuozzo-specific cluster configuration only. For general information regarding Red Hat Cluster Suite, please turn to the **Red Hat Cluster Suite: Configuring and Managing a Cluster** guide (available at <http://www.redhat.com/docs/manuals/csgfs>).

---

## CHAPTER 2

# Deploying Virtuozzo Failover Cluster

This section familiarizes you with the way to use the Red Hat Cluster Suite (RHCS) software to provide greater availability for Virtuozzo mission-critical services and Containers.

## In This Chapter

Cluster Configuration Options .....	5
Setting Up Cluster Servers .....	6
Setting Up Clustered Services.....	9

---

## Cluster Configuration Options

Virtuozzo Containers 4.0 allows you to create a cluster of two or more Hardware Nodes running one and the same version of the Virtuozzo Containers software to provide a high degree of availability for your Virtuozzo Containers installation. The implementation of the Virtuozzo Containers software Failover Cluster is based on the following statements:

- Red Hat Cluster Suite (RHCS) from the RHEL 5 Linux distribution is used as the clustering software.
- In the given Failover Clustering scenario, the Virtuozzo service acts as a clustered service, which means the following: if the Hardware Node becomes inaccessible in the case of a hardware or software failure or the Virtuozzo service is stopped accidentally, the clustering software will try to restart the service and all the Containers on the problem Node and, if the latter is not possible, on another (passive) Node in the cluster.
- The cluster configuration is as follows:
  - A shared SCSI storage device (iSCSI or Fibre Channel) is used as a cluster storage.
  - A single LUN (Logical Unit Number) is created per clustered service (in our case each Virtuozzo Containers installation is regarded as a clustered service).
  - The /vz partition is located on the shared SCSI storage device and formatted as ext3.
  - The clustered file system (GFS) is NOT used.
- A typical cluster configuration consists of 'X' + 'Y' servers where 'X' and 'Y' denote the number of active and passive Hardware Nodes in the cluster, respectively. It means that you can run 'X' Virtuozzo Containers installations and recover from 'Y' Hardware Node failures. For example, if your cluster contains 3 active nodes and 1 passive node, only one active node will be able to fail its resources (the Virtuozzo service and all Containers on that node) to the passive node.
- The minimal hardware requirements include:
  - 'X' + 'Y' physical servers;
  - all physical servers should have one and the same system architecture (e.g. IA32 or EM64T) to guarantee that Containers can be started on any server in the cluster;

- 'X' shared LUNs;
- 'X' + 'Y' ports on fencing devices.

Cluster configuration examples: 4 active nodes and 1 stand-by node or 10 active nodes and 2 standby nodes.

- One and the same version of the Virtuozzo Containers software should be installed on all servers in the cluster. However, only one `/vz` partition should be created per clustered service.
- The following cluster resources should be present in the cluster: the `/vz` file system, the `/etc/init.d/vz-cluster` script, and the 'IP address' resource.

---

## Setting Up Cluster Servers

The process of setting up a cluster server includes the following steps:

- 1 Installing a Host operating system on the server.
- 2 Installing and configuring the clustering software on the server.
- 3 Configuring your shared cluster storage to communicate with the server.
- 4 Installing and configuring the Parallels Virtuozzo 4.0 software on the server.

## Installing Host OS

In the case of a clustered Virtuozzo environment, the server configuration and partitioning requirements slightly differ from those used for a standalone installation. Please follow the rules below when installing a Host OS (Red Hat Enterprise or CentOS) on each of the servers to be included in the cluster:

- If your shared storage is iSCSI, consider using Gigabit Ethernet adapters and the corresponding network infrastructure. In most cases, a 100 Mb network is very likely to become a bottleneck.
- When installing the Host OS on a local storage, make sure that at least 8 GB of disk space is allocated to the root partition; bigger partitions (10 GB or more) are preferred, however, not required.
- When partitioning your server, do not create the `/vz` partition. The servers in the cluster will not need it - the `/vz` partition will be located on a remote SCSI storage.
- As a rule, the “Server” configuration in the installer is sufficient for your needs. However, the cluster administration GUI tool (*system-config-cluster*) requires X-Window to be installed on the server. So, if you plan to use this tool for administering your cluster, install the X-Window component on one of your servers.
- When configuring your Host OS installation, disable the RHEL or CentOS standard firewall since both Virtuozzo Containers 4.0 and the data sharing cluster use a multitude of TCP/IP ports. To prevent unsafe network packets from accessing your cluster servers, consider using an external firewall.

## Installing and Configuring Cluster Software

After the Host OS installation, you should update your server software to its latest version (e.g. by means of the `up2date` utility) and then complete the following tasks on each of the cluster nodes:

- 1 Install the clustering software. You can obtain the needed packages using the Red Hat `up2date` utility. The following packages from Red Hat Cluster Suite are required:
  - `perl-Net-Telnet`;
  - `cman`;
  - `system-config-cluster`;
  - `rgmanager`;
  - `openais`.

Your typical packages installation session may look as follows:

```
# rpm -ihv openais-0.80.3-7.el5.i386
cman-2.0.73-1.el5_1.1.i386 rgmanager-2.0.31-1.el5.i386.rpm
system-config-cluster-1.0.50-1.3.noarch.rpm
perl-Net-Telnet-3.03-5.noarch
```

**Note:** The `dlm` lock manager has been tested and recommended for the production usage; other locking schemes (`quorum disk`, etc.) are not supported.

- 2 Edit the `/usr/share/cluster/fs.sh` script by locating the following string in this script

```
<parameter name="mountpoint" unique="1" required="1">
```

and changing it as follows:

```
<parameter name="mountpoint" unique="0" required="1">
```

## Configuring Shared Storage

After installing the clustering software, you need to configure the shared storage which is to be used as the `/vz` partition. The shared storage can be configured from a single server or from several servers (it does not really matter). The space requirements for this partition should be calculated on the basis of the space demands for your Containers and Virtuozzo templates. The number of partitions should correspond to the number of clustered services and be less than the number of physical servers in your cluster.

Please keep in mind that all your servers should be able to access your shared LUNs. The cluster will decide where to start the clustered service by itself; so, it must be able to mount the partition with all the required Containers and templates.

Using volume labels for your partitions will greatly simplify the further management of your cluster. We recommend that you have a match between a clustered service and a partition name. For example, the `Virtuozzo-1` service can use a partition with the `vz-1` label. Also note that these volume labels do not correspond to the mount points, which are always `/vz` on all the servers in the cluster.

After the corresponding LUNs have been successfully created, make the partitions on them using the `fdisk` utility or other disk partitioning software and then create the file system. For example:

```
# mke2fs -j -J size=100 -L vz-1 -b 4096 -i 4096 /dev/sdc1
```

Please do not forget to replace `/dev/sdc1` in the example above with your real disk device and `vz-1` with a label set for a particular clustered service. Now check that the device can be successfully mounted:

```
# mount LABEL=vz-1 /vz
```

After that, unmount the device:

```
# umount /vz
```

## Installing Parallels Containers 4.0

To install Parallels Containers 4.0, complete the following cluster-related configuration tasks on each server in the cluster:

- 1 Make sure that you have the `/vz` partition mounted on the server. If you have a number of shared `/vz` partitions, and all of them have already been mounted on other servers, you should not mount any of them on the server where you wish to install Virtuozzo Containers 4.0. In this case the server may have no `/vz` partition mounted. For example, if your cluster configuration consists of 3 servers and 2 shared `/vz` partitions, you should mount the first `/vz` partition on the first server in the cluster and the second `/vz` partition - on the second server; the third server should not have any `/vz` partitions mounted.

- 2 Install the Parallels Virtuozzo software on the server:

- If your server has a `/vz` partition mounted, add the `/vz` entry to the `/etc/fstab` file and run the Parallels Virtuozzo installer by issuing the following command:

```
# ./install
```

- If your server does not have a `/vz` partition mounted, make sure that there is no `/vz` entry in the `/etc/fstab` file and run the Parallels Virtuozzo installer by issuing the following command:

```
# ./install --clustermode
```

- 3 Review and modify, if needed, your boot loader configuration (`/etc/grub.conf` by default) to make sure a new kernel will be loaded, and then restart your server.
- 4 Wait for the Virtuozzo post-install procedures to complete and log in again.
- 5 Update your Virtuozzo Containers software with `vzup2date` (one more reboot may be needed).
- 6 If your server has the `/vz` partition mounted, do the following:

- remove the `/vz` entry from the `/etc/fstab` file;
- disable the `vz` service:

```
# chkconfig vz off
```

- stop the `vz` service:

```
# /etc/init.d/vz stop
```

# Setting Up Clustered Services

## Configuring Virtuozzo Failover Cluster

The configuration procedures described below supposes the usage of the `system-config-cluster` utility. With this utility, you should do the following:

- Use an existing cluster or create a new one (the creation of a new cluster is recommended).
- Add your servers to the cluster under the **Cluster Nodes** tree.
- If the same cluster manages other services, you need to create a new failover domain (e.g. the `VZ-servers` name can be used) and add all your Virtuozzo cluster nodes to this domain.
- Add and configure fencing devices (note that using fencing devices is mandatory). Manual fencing is supported for testing purposes only and should not be used in production.
- Create resources:
  - File system: one resource per each shared SCSI partition. Make sure that the partition is referred to by the volume label instead of by the device name. Also, do not forget to use the `noatime` option. Please see the picture below:

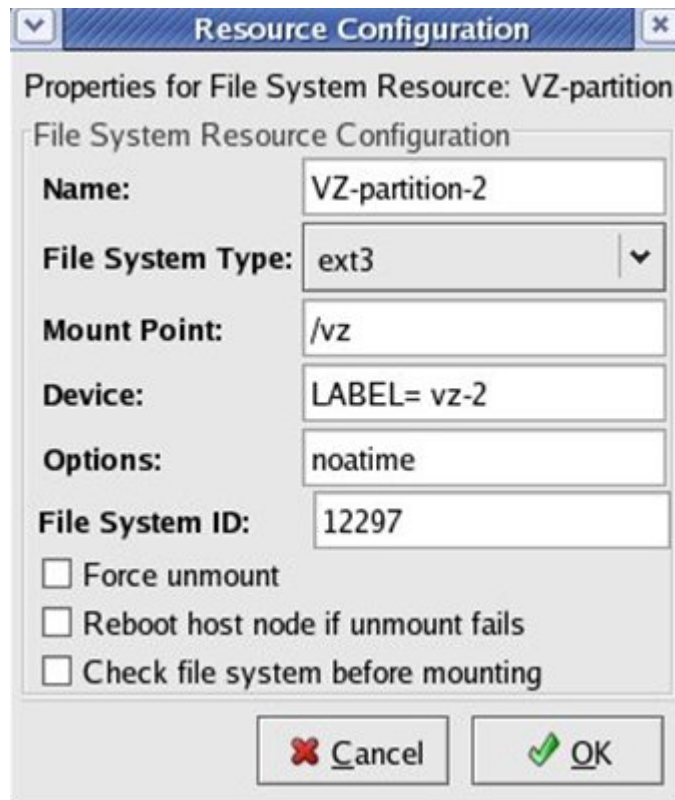


Figure 1: Cluster Resource Configuration - Configuring File System

- Script: one script per cluster:

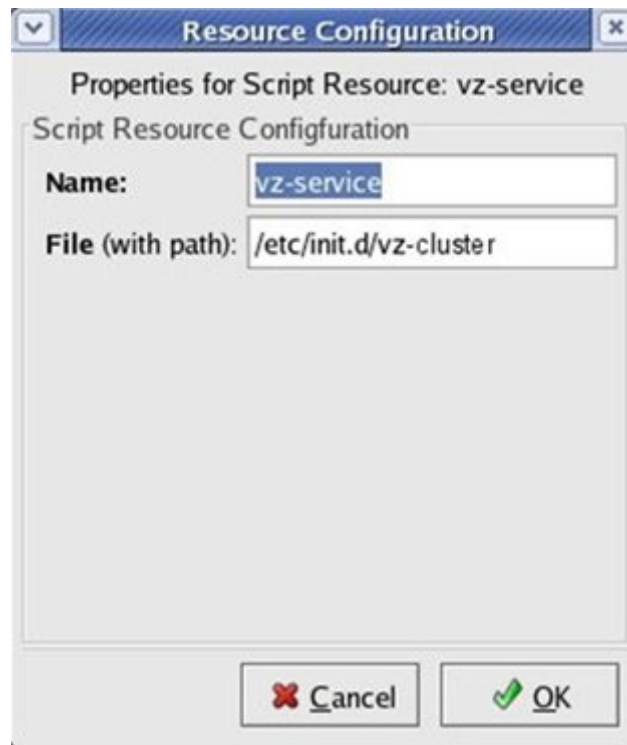


Figure 2: Cluster Resource Configuration - Defining Scripts

- IP address: an IP address is needed for each Virtuozzo service (it will be used to connect to the Service Container and for a direct SSH connection to the host). Note that the IP address will be managed by the cluster; so, it should not be already in use and assigned to the nodes directly:




Figure 3: Cluster Resource Configuration - Assigning IP Address

- Create the corresponding clustered services. The number of services corresponds to the number of active servers and shared partitions; also make sure that:
  - The Service Autostart is enabled.
  - The Service is configured to run exclusively; thus, the cluster will not attempt to run more than one Virtuozzo service on the same physical server (which is not currently supported).
  - The Service recovery policy is 'Relocate' or 'Restart'; in the latter case, if the Virtuozzo service is stopped for some reason, the cluster will attempt to restart Virtuozzo Containers 4.0 on the same server before relocating it to another one.
  - If Virtuozzo Containers 4.0 shares the cluster with other clustered applications, make sure the proper failover domain is specified (in the picture below, the failover domain is not set).
  - The 'Script' Resource (vz-service in our example) is attached to the 'File System' resource (VZ-partition-2). This ensures the proper order of the services startup:

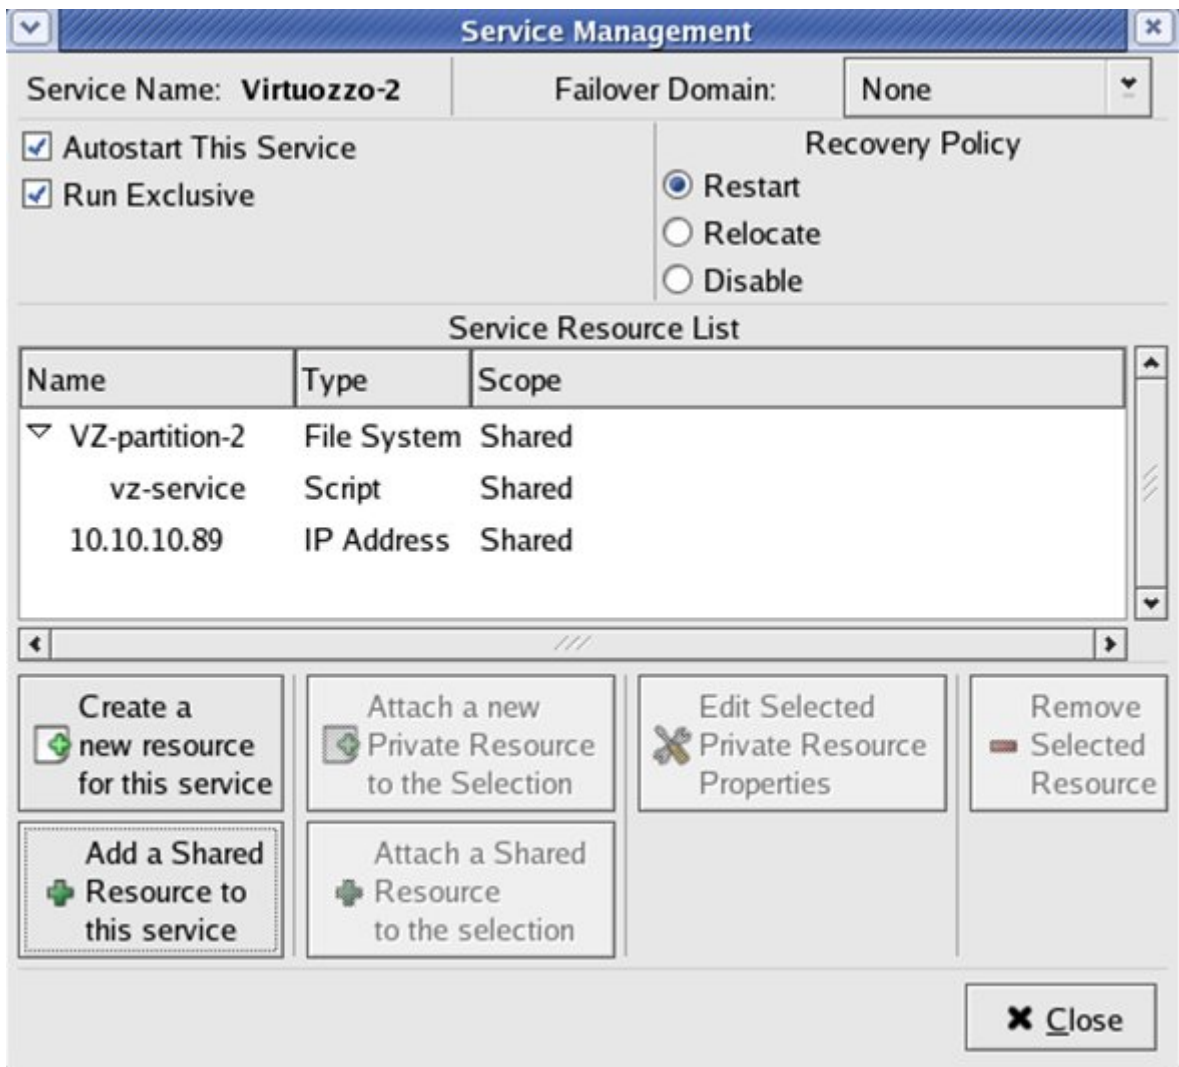


Figure 4: Cluster Configuration - Viewing Cluster Service Parameters

- Distribute the configuration file (`/etc/cluster/cluster.conf`) to all the servers and start the clustering service:

```
# service cman start
# service rgmanager start
```

- Configure the clustering service on each node in the cluster to start in the default runlevel. For example, if your system default runlevel is set to 3, you can enable the service by executing the following commands on each of the cluster nodes:

```
# chkconfig --level 3 cman on
# chkconfig --level 3 rgmanager on
```

- After starting the service, you may further synchronize your configuration using the **Send to Cluster** button in the cluster configuration utility.
- Locate the **Cluster Management** section of the GUI tool or run the `clustat` command line utility. Make sure that all the services have been successfully started. If they have not, investigate the cluster logs stored in `/var/log/messages` by default. Please keep in mind that the information you are looking for may be placed on different servers in the cluster.

## Testing Virtuozzo Cluster

To test the created cluster, you can shut down or unplug a power cord on one of the servers in the cluster which currently runs the Virtuozzo Containers software (if you are using manual fencing, you may need to use the `fence_ack_manual` tool to inform the cluster of the fenced server). Use any cluster monitoring tool to ascertain that, in less than a minute, the Virtuozzo service is automatically relocated from the failed server to a standby one.

## CHAPTER 3

# Deploying Data Sharing Cluster

Data sharing clusters are server farms that share storage devices on a storage area network (SAN) and share data on those storage devices. In a data sharing cluster, data can be written or read by any server to or from any file on a common shared storage device. Virtuozzo Containers 4.0 allows you to deploy data sharing clusters in your working environments by combining the functionality provided by Red Hat Global File System (GFS) with that of Red Hat Cluster Suite. By implementing a data sharing cluster made up of Virtuozzo Hardware Nodes, you can achieve the following goals:

- Simplify the process of managing your Containers and Virtuozzo Templates since all Containers and templates are residing on a single SAN storage shared by all servers in the cluster.
- Greatly speed up the process of migrating running Containers between the cluster nodes. In fact, the migration is almost imperceptible to users since all Container data in the cluster is stored on a shared SAN storage and there is no need to move this data between the Nodes during the Container migration.
- Provide failover support for Virtuozzo vital services and Containers. Each server in the cluster is running the clustering software responsible for monitoring the health of Virtuozzo Containers installations and failing over the services and Containers from a failed node to the standby node.

The following subsections provide detailed information on how to deploy a data sharing cluster based on Red Hat GFS and Red Hat Cluster Suite.

## In This Chapter

Cluster Creation Process Overview .....	14
Setting Up Hardware For Cluster.....	14
Installing and Configuring Host OS.....	15
Installing Clustering Software.....	15
Setting Up Red Hat Cluster.....	16
Configuring Cluster Storage .....	18
Installing and Configuring Virtuozzo Containers 4.0 .....	19
Creating and Configuring Cluster Services.....	20

---

## Cluster Creation Process Overview

To set up a data sharing cluster, you should complete the following steps:

- 1 Set up hardware for the cluster.
- 2 On each node in the cluster, install the Host operating system. Currently, GFS clusters can be deployed on servers running RHEL 5 and CentOS 5 only.
- 3 On each node in the cluster, install and configure the Red Hat clustering software, including the necessary GFS user-level tools.
- 4 Create and configure the Red Hat cluster.
- 5 Set up and configure the shared cluster storage device in the cluster.
- 6 On each node in the cluster, install and configure the Virtuozzo Containers 4.0 software.
- 7 Create and configure one or more cluster services that can be failed over in the event of a hardware or software failure.

All these operations are described in the following subsections in detail.

---

## Setting Up Hardware For Cluster

On the first step, you should set up hardware for your cluster. This procedure consists of the following main steps:

- Deciding on the hardware to be included in the cluster. The amount and type of hardware may vary depending on the purpose and availability requirements of your cluster. However, the following type of hardware is always present in any cluster:
  - *cluster nodes* represented by two or more physical servers capable of running RHEL 5 or CentOS 5;
  - *storage* usually represented by a common disk array on a SAN and used to store all Virtuozzo-related data (Container data, templates, backups, etc.);
  - *switches* (Fibre Channel or iSCSI) providing client access to the cluster and enabling the communication between each cluster node and the shared cluster storage.

---

**Note:** If you are planning to use an iSCSI storage device in your cluster, you are highly recommended using Gigabit network adapters on your servers. This will allow you to prevent possible bottlenecks which are likely to arise in a 100-Mb network.

---

- Uniting all hardware components into one subnet and ensuring that each hardware component can access all the other components in the subnet.

For more information about installing and configuring cluster hardware, please refer to the [Configuring and Managing a Red Hat Cluster](https://www.redhat.com/docs/manuals/enterprise/) document (available at <https://www.redhat.com/docs/manuals/enterprise/>).

---

## Installing and Configuring Host OS

Next, you should install and configure the Host operating system on each server (node) that will participate in the data sharing cluster. When installing the Host OS (either RHEL 5 or CentOS 5) on your cluster nodes, follow these guidelines:

- When partitioning the server hard drives:
  - Make sure that at least 8 GB of disk space is allocated to the root (/) partition.
  - Do not create the /vz partition. All nodes in the cluster will use the /vz partition located on a shared data storage (Fibre Channel or iSCSI) and mounted to them.
- When deciding on the packages to be installed on your nodes, keep in mind that the cluster administration GUI tool - *system-config-cluster* - requires the 'X Window System' component for its functioning. So, if you plan to use this tool for configuring and administering your cluster, install this component on at least one of the cluster nodes.
- When configuring your Host OS installation, disable the RHEL or CentOS standard firewall since both Virtuozzo Containers 4.0 and the data sharing cluster use a multitude of TCP/IP ports. To prevent unsafe network packets from accessing your cluster servers, consider using an external firewall.

Upon the installation completion, you are highly recommended to update your Host OS to its latest version (e.g. using the `up2date` utility).

---

## Installing Clustering Software

On the next step, you should install the necessary clustering software on each node in the cluster. The following packages from Red Hat Cluster Suite are required:

- `perl-Net-Telnet;`
- `cman;`
- `system-config-cluster;`
- `rgmanager;`
- `openais;`
- `gfs-utils.`

You can use the `up2date` utility to download the latest versions of the aforementioned packages.

To install the obtained packages, use the `rpm -i` command. For example:

```
# rpm -ihv openais-0.80.3-7.e15.i386 cman-2.0.73-1.e15_1.1.i386
rgmanager-2.0.31-1.e15.i386.rpm
system-config-cluster-1.0.50-1.3.noarch.rpm
perl-Net-Telnet-3.03-5.noarch
gfs-utils-0.1.11-1.e15.i386.rpm
```

## Setting Up Red Hat Cluster

After you have successfully installed the clustering software on each cluster node, you should set up a Red Hat cluster that will ensure the high-availability of your Virtuozzo Containers 4.0 installations. You can use one of the following tools to set up a new cluster:

- `system-config-cluster`: this GUI tool can be used to easily install, configure, and manage Red Hat clusters.
- `Conga`: this comprehensive user interface can be used for installing, configuring, and managing Red Hat clusters, servers, and storage devices attached to clusters and servers.
- `command line tools`: this is a set of command line tools for configuring and managing Red Hat clusters.

The example below demonstrates how to set up a new Red Hat cluster using the `system-config-cluster` cluster configuration GUI tool:

- 1 Log in to any of you cluster nodes (this node should have the 'X-Window' component installed) and launch the Cluster Configuration Tool as follows:

```
# system-config-cluster
```

- 2 Your further steps will be different depending on whether you are starting the tool for the first time or have already launched it before:
  - click **Create New Configuration** in the displayed dialog window if you are running the Cluster Configuration Tool for the first time;
  - otherwise, click **File -> New**.
- 3 In the **New Configuration** window, enter a descriptive name for the cluster in the field provided and click **OK**. Please keep in mind that you will need to specify this name when creating a GFS file system with the `gfs_mkfs` utility. Please see **Step 4** in the **Configuring Shared Storage** subsection (p. 18) for more information.
- 4 Set one or more fence devices for the cluster:
  - Select the **Fence Devices** item and click **Add a Fence Device** at the bottom of the Cluster Configuration window:



Figure 5: Setting Up Data Sharing Cluster - Defining Fencing Devices

- In the Fence Device Configuration window, select the type of fence device to configure and provide the necessary information in the appropriate fields.

---


**Note:** Please refer to the *Configuring and Managing a Red Hat Cluster* document for detailed information on available fence devices and their parameters.

---

- Click OK.

**5** Add all your Virtuozzo servers to the cluster:

- Select the Cluster Nodes item and click **Add a Cluster Node** at the bottom of the Cluster Configuration window.
- In the Node Properties window, type the name of the server you wish to add to the cluster (or its IP address on the cluster subnet) and click **OK**. For example:



*Figure 6: Setting Up Data Sharing Cluster - Adding Node to Cluster*

- Configure fencing for the added node using the **Manage Fencing For This Node** button at the bottom of the Cluster Configuration right frame.
  - Repeat the aforementioned operations for each Virtuozzo server to be included in the cluster.
- 6** Add one or more failover domains to the cluster. This step is optional and should be performed if you wish to specify one or more subsets of cluster nodes that will be eligible to run the corresponding cluster services. For more information on cluster services, please refer to the *Creating and Configuring Cluster Services* subsection (p. 20).
- Select the Failover Domains item and click **Create a Failover Domain** at the bottom of the Cluster Configuration window:



*Figure 7: Setting Up Data Sharing Cluster - Adding Failover Domain*

- Type an arbitrary name for the failover domain (it will be used to distinguish the given failover domain among the other domains in the cluster) and click **OK**.
- In the **Failover Domain Configuration** window, specify the cluster nodes to be included in the failover domain, configure the domain properties (e.g. set the cluster nodes prioritization), and click **Close**.

---

**Note:** Each failover domain must contain at least one standby node where all Virtuozzo mission-critical services and Containers will be able to fail over in the event of an active node failure.

---

- 7 Save the current cluster configuration by clicking **File --> Save** on the main menu of the Cluster Configuration Tool. By default, you are offered to save the cluster configuration into the `cluster.conf` file in the `/etc/cluster` directory. Click **OK** to confirm the offered file name and location or change them to your preference.
- 8 Distribute the cluster configuration file (`/etc/cluster/cluster.conf` by default) to all nodes in the cluster. For example, you can do it using the `scp` utility.
- 9 Start the Red Hat clustering service on each server in the cluster by executing the following command:

```
# service cman start
```

---

## Configuring Cluster Storage

After you have set up the Red Hat cluster and started the clustering software on each cluster node, you should configure data storage (iSCSI or Fibre Channel) for the cluster. This data storage will be used for hosting the `/vz` partition with all Container data and Virtuozzo Templates from all Virtuozzo servers in the cluster and shared by all cluster nodes. To set up a shared data storage, you should perform the following operations:

- 1 Log in to any of your cluster nodes.
- 2 Use standard Linux tools (e.g. Logical Volume Manager) to set up a logical volume on your data storage (e.g. `/dev/vg01/lvol1`). This logical volume will keep the `/vz` partition. Please notice that one logical volume is required for each Red Hat GFS file system.

---

**Note:** If you are going to use Logical Volume Manager (LVM) for creating logical volumes, make sure that it is configured with the clustered locking support; otherwise, the LVM metadata may become corrupted. For detailed information on LVM and its configuration settings, please turn to the LVM documentation and `lvm.conf` man pages.

---

- 3 Ascertain that the created logical volumes can be accessed by all servers in the cluster. This will ensure that the clustering software will be able to mount the `/vz` partition that you will create on the logical volume in the next step to any of your cluster nodes.
- 4 Create the `/vz` partition on the logical volume (e.g. using the `fdisk` utility or any other disk partitioning software).
- 5 Create a GFS file system on the logical volume using the `gfs_mkfs` utility. For example, you can issue the following command for this purpose:

```
# gfs_mkfs -p lock_dlm -t gfs_cluster_1:gfs_1 -j 4 /vg01/lvol1
```

In this example:

- `-p lock_dlm` denotes the name of the locking protocol that will be used by the GFS file system. The currently recognized cluster-locking protocols include `lock_dlm` and `lock_nolock`.
- `-t gfs_cluster_1:gfs_1` denotes:
  - a** the name of the cluster (`gfs_cluster_1`) for which the GFS file system is created;
  - b** the name that will be assigned to the GFS file system (`gfs_1`).

- `-j 4` is the number of journals that will be created by the `gfs_mkfs` utility. When deciding on the number of journals, keep in mind that one journal is required for each cluster node which is to mount the GFS file system. You can also make additional journals at the time of the GFS file system creation to reserve them for future use.
- `/vg01/lv011` denotes the logical volume where the GFS file system is to be located.

So, as a result of the aforementioned command a GFS file system with the `gfs_1` name for the `gfs_cluster_1` cluster will be created. The file system will use the `lock_dlm` protocol, contain 4 journals, and reside on the `/vg01/lv011` volume.

- 6 Tell each node in the cluster to automatically mount the `/vz` partition on the node boot. To this effect, add the `/vz` entry to the `/etc/fstab` file on each node in the cluster. Assuming that your GFS file system resides on the `/vg01/lv011` logical volume, you can add the following entry to the `fstab` file:

```
/dev/vg01/lv011 /vz gfs defaults,noatime 0 0
```

- 7 Configure the `gfs` service on each node in the cluster to start in the default runlevel. For example, if your system default runlevel is set to 3, you can enable the `gfs` service by executing the following command on each of the cluster nodes:

```
# chkconfig --level 3 gfs on
```

---

## Installing and Configuring Virtuozzo Containers 4.0

Now that you have set up the shared data storage, you should install Virtuozzo Containers 4.0 on each node in the data sharing cluster. To this effect, perform the following operations on each cluster node:

- 1 Make sure that the node where you wish to install Virtuozzo Containers 4.0 has the shared `/vz` partition mounted. If it does not, mount the `/vz` partition as follows:

```
# mount /vz
```

- 2 Start the Virtuozzo Containers installation by executing the following command on the node:

```
# ./install
```

- 3 Follow the instructions of the Parallels Virtuozzo Containers Installation wizard to install the Virtuozzo Containers software on the server.

---

**Note:** Detailed information on the Virtuozzo Containers installation procedure is provided in the [Parallels Virtuozzo Containers Installation Guide](#).

---

- 4 Upon the Virtuozzo Containers installation completion, reboot the Hardware Node and wait for the necessary Virtuozzo post-install operations to finish.
- 5 Repeat Steps 1-4 for all the remaining nodes in the cluster.

---

## Creating and Configuring Cluster Services

Finally, you should create and configure one or more cluster services that will be able to fail over from a problem node in the cluster to a healthy one. In our data sharing cluster, each Virtuozzo Containers installation is regarded as a separate cluster service; so, the number of cluster services to make should correspond to the number of your active Virtuozzo servers in the cluster. To make a cluster service, perform the following operations:

- 1 Log in to any of your cluster nodes (this node should have the 'X-Window' component installed) and launch the Cluster Configuration Tool as follows:

```
# system-config-cluster
```

- 2 Add the `vz-cluster` script resource to the cluster. `vz-cluster` is a special script located in the `/etc/init.d` directory on each cluster node and used to perform cluster-specific tasks in Virtuozzo-based systems. You will need to specify the path to this script when configuring your cluster services (see Step 5). To add the script:
  - Select the Resources item and click Create a Resource at the bottom of the Cluster Configuration window.
  - Under Select a Resource Type, select Script.
  - Type an arbitrary name for the script in the Name file and specify the full path to the script (`/etc/init.d/vz-cluster`) in the File field. For example:

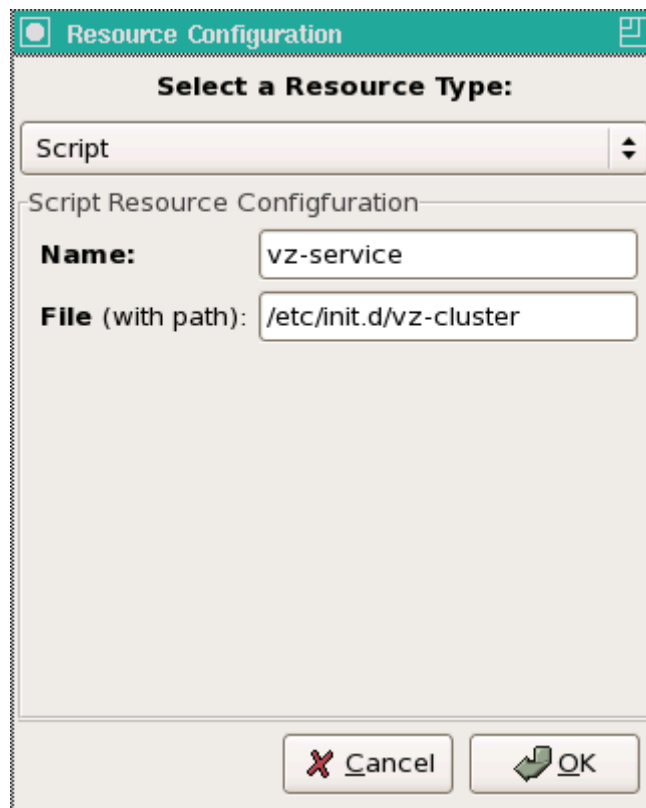


Figure 8: Setting Up Data Sharing Cluster - Creating Script Resource

- Click OK.
- 3 Select the **Services** item and click **Create a Service** at the bottom of the **Cluster Configuration** window:



*Figure 9: Setting Up Data Sharing Cluster - Adding New Clustered Service*

- 4 In the displayed window, specify a descriptive name for the cluster service (this name will be used to distinguish this service from the other services in the cluster) and click **OK**.
- 5 In the **Service Management** window, configure the service as follows:
- Make sure the **Autostart This Service**, **Run Exclusive**, and **Relocate** check boxes are selected.
  - In the **Failover Domain** field, select the name of the failover domain to be assigned to the service if you wish to restrict the members on which this cluster service will be able to run.
  - Assign the `vz-cluster` script resource to the cluster service by clicking the **Add a Shared Resource to this service** button, selecting the script name in the displayed window, and clicking **OK**.
  - Create a new IP address resource for the service by clicking the **Create a new resource for this service** button, selecting **IP address** under **Select a Resource Type**, typing a valid IP address in the field provided, and clicking **OK**. A separate IP address is required for each cluster service. When deciding on an IP address, make sure that it is not already in use (e.g. not assigned to any node in the cluster).

After performing the aforementioned operations, your **Service Management** window may look like the following:

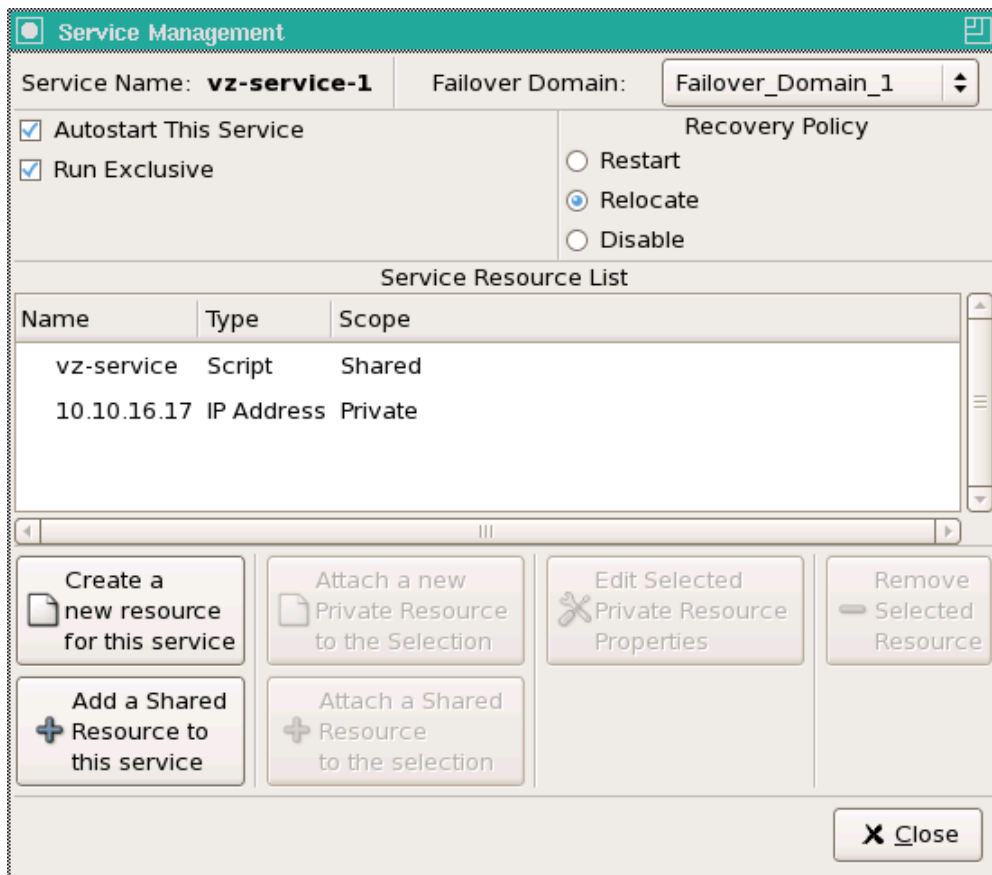


Figure 10: Setting Up Data Sharing Cluster - Configuring Clustered Service

- Click Close to create the cluster service.
- 6 Perform Steps 2-4 for all cluster services you are going to have (please remember that the number of cluster services should correspond to the number of your active Virtuozzo servers in the cluster).
  - 7 Save the current cluster configuration by clicking File --> Save on the main menu of the Cluster Configuration Tool.
  - 8 Distribute the updated cluster configuration file (`/etc/cluster/cluster.conf` by default) to all nodes in the cluster. For example, you can do it by executing the following command on the node where you are currently logged in:

```
# ccs_tool update /etc/cluster/cluster.conf
```

- 9 Reboot all nodes in the cluster.

You can check that all clustered services have been successfully started by using the Cluster Management tab of the Cluster Configuration Tool or by running the `clustat` command line utility. If some service has failed to start, investigate the cluster logs stored in `/var/log/messages` by default. Please keep in mind that the information you are looking for may be placed on different servers in the cluster.

---

# Index

## C

- Cluster Configuration Options • 5
- Cluster Creation Process Overview • 16
- Configuring Cluster Storage • 21
- Configuring Shared Storage • 8
- Configuring Virtuozzo Failover Cluster • 10
- Creating and Configuring Cluster Services • 23

## D

- Deploying Data Sharing Cluster • 15
- Deploying Virtuozzo Failover Cluster • 5

## I

- Installing and Configuring Cluster Software • 7
- Installing and Configuring Host OS • 17
- Installing and Configuring Virtuozzo Containers 4.0 • 22
- Installing Clustering Software • 17
- Installing Host OS • 6
- Installing Parallels Containers 4.0 • 8
- Introduction • 4

## S

- Setting Up Cluster Servers • 6
- Setting Up Clustered Services • 10
- Setting Up Hardware For Cluster • 16
- Setting Up Red Hat Cluster • 18

## T

- Testing Virtuozzo Cluster • 14