



OpenVZ Readme

December 20, 2016

Parallels IP Holdings GmbH
Vordergasse 59
8200 Schaffhausen
Switzerland
Tel: + 41 52 632 0411
Fax: + 41 52 672 2010
<http://www.virtuozzo.com>

Copyright © 1999-2016 Parallels IP Holdings GmbH and its affiliates. 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 <http://www.virtuozzo.com/legal/>.

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.

Table of Contents

1. OpenVZ Overview	4
2. What's New	5
3. Known Issues and Restrictions	7
4. Dropped Functionality	8

Chapter 1. OpenVZ Overview

OpenVZ is a bare-metal virtualization solution that includes container virtualization, KVM-based virtual machines, software-defined storage along with enterprise features and production support. It runs on top of VzLinux, a RHEL-based Linux distribution.

Chapter 2. What's New

The key changes in OpenVZ are:

- OpenVZ becomes a complete Linux distribution based on our own VzLinux.
- The main difference between the Virtuozzo (commercial) and OpenVZ (free) versions are the EULA, packages with paid features, and Anaconda installer.
- EZ templates can be used instead of tarballs with template caches.
- RHEL7 (3.10+) kernel.
- KVM/QEMU hypervisor.
- Guest tools for virtual machines that currently allow the following: to execute commands in VMs from the host, to set user passwords, to set and obtain network settings, to change SIDs, to enter VMs.
- Unified management of containers and KVM virtual machines with the `prctl` tool and SDK. You get a single universal toolset for all your CT/VM management needs.
- UUIDs are used to identify both virtual machines and containers. With containers, `prctl` treats the former VEID parameter as name.
- Virtual machine HDD images are stored in the QCOW2 format.
- Ability to manage containers and VMs with `libvirt` and `virt-manager` or `virsh` via a single driver for containers and virtual machines. `Libvirt` is an open-source API, daemon, and management tool for managing virtualization platforms. The API is widely used in the orchestration layer of hypervisors for cloud-based solutions. OpenVZ considers `libvirt` as the standard API for managing both virtual machines and containers. `Libvirt` provides storage management on the physical host through storage pools and volumes which can be used in OpenVZ containers.
- Memory guarantees. A memory guarantee is a percentage of container's or virtual machine's RAM that said container or VM is guaranteed to have.
- Memory hotplugging for containers and VMs that allows both increasing and reducing CT/VM memory size on the fly, without the need to reboot. Your customers can now scale their workloads without any downtime. This feature also enables you to make PAYG offerings, allowing customers to change VM resources depending on workload and potentially pay less.
- Kernel same-page merging. To optimize memory usage by virtual machines, OpenVZ uses a Linux feature called Kernel Same-Page Merging (KSM). The KSM daemon `ksmd` periodically scans memory for pages with identical content and merges those into a single page.
- VCMMD, the fourth-generation unified memory manager, and `vcmmmd`, a single daemon for managing memory of both virtual machines and containers. OpenVZ uses `memcg`. Balancing and configuring `memcg` limits enables getting the exact OpenVZ parameters like `overcommit`, `shadow gangs`, `swap`, `page cache overuse`.
- Container live migration via `CRIU` and `P.Haul`. In the previous versions of OpenVZ, most operations performed during migration were done in the kernel space. As a result, the migration process imposed

a lot of restrictions. To improve upon migration, Virtuozzo launched the CRIU project aiming to move most of the migration code to the user space, make the migration process reliable, and remove excessive restrictions.

- Containers use cgroups and namespaces that limit, account for, and isolate resource usage as isolated namespaces of a collection of processes. The beancounters interface remains in place for backward compatibility and, at the same time, acts as a proxy for actual cgroups and namespaces implementation.
- SimFS remains in OpenVZ, however, the support is limited and we don't have plans to improve it in future.

Chapter 3. Known Issues and Restrictions

- ploops over NFS are not supported. (PSBM-20108)
- Linux virtual machines with EFI firmware cannot be migrated to OpenVZ. (PSBM-32920)
- Attaching backups as block devices to VMs and containers is not supported. (PSBM-40870)
- In Windows VMs, each new SCSI HDD may be added as offline. As a workaround, you can either bring each disk back online manually. Alternatively, before the first boot with a new disk, run `san policy=onlineall` in DISKPART to set the SAN policy to `OnlineAll`. (PSBM-41663)
- Migration of containers with NFS client inside is not supported. (PSBM-47068) Alternatively, before the first boot with a new disk, run `san policy=onlineall` in DISKPART to set the SAN policy to `OnlineAll`. (PSBM-48700)
- Secondary disks of containers restored from OpenVZ backups may contain stale data. (PSBM-49091)
- Direct and offset autofs mounts served by automount are in inactive state after container migration. To work around the issue, kill the automount process with SIGKILL signal and restart the automount daemon. (PSBM-49104)
- VMs with Ubuntu 14.04 guests may hang on backup. (PSBM-49106)
- Changes to boot order are only applied after a complete VM shutdown. (PSBM-49119)
- Migration of containers with secondary disks between OpenVZ servers is not supported. (PSBM-49179)
- Limited simfs support (feature provided as is). (OVZ-6613) Unlike OpenVZ, the simfs layout is based on bindmounts in OpenVZ. When a simfs-based container is started, its private area is bindmounted to the root container area. To create a simfs container: 1) Set `VEFSTYPE=simfs` in `/etc/vz/vz.conf`, 2) Run `vzctl create <CT_name>`. The simfs limitations in OpenVZ are: 1) No support for first- or second-level quotas, 2) No support for online migration of simfs-based containers.
- Private networks are not supported.
- Basic firewall is not supported.

Chapter 4. Dropped Functionality

- VZFS
- Delayed `/vz` mounting
- Commands `prlctl --memquota`, `prlsrvctl --mem-limit`.
- The `vzbackup`, `vzabackup`, and `pbackup` utilities. For more information, see <http://kb.virtuozzo.com/119017>.
- Migration of physical servers to containers and migration of Xen virtual machines to virtual machines. For more information, see <http://kb.virtuozzo.com/119016>.
- Start-As-User.
- Shared smartcard (CCID) support.
- The global iptables mask configured in `/etc/sysconfig/vz`.
- All UBC resources except `physpages` and `swappages`.
- The `vziptables` parameter configured in the kickstart file used for unattended installations of Virtuozzo.