



Virtuozzo 7 Readme

July 25, 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. Virtuozzo Overview 4
- 2. What's New 5
- 3. Known Issues and Restrictions 7
- 4. Dropped Functionality 9

Chapter 1. Virtuozzo Overview

Virtuozzo 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 Virtuozzo Linux, a RHEL-based Linux distribution.

Chapter 2. What's New

The key changes in Virtuozzo 7 are:

- Virtuozzo 7 is based on RHEL 7 and Kernel 3.10+.
- Virtuozzo 7 uses the KVM/QEMU hypervisor and enables customers to manage virtual machines via the libvirt API and VirtManager.
- 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. At the same time, it acts as a proxy for actual cgroups and namespaces implementation.
- UUIDs are used to identify both virtual machines and containers. With containers, `prlctl` treats the former VEID parameter as name.
- VCMMD, the fourth-generation unified memory manager, and `vcmmmd`, a single daemon for managing memory of both virtual machines and containers. Virtuozzo 7 uses memcg. Balancing and configuring memcg limits enables getting the exact overcommit, shadow gangs, swap, page cache overuse Virtuozzo parameters.
- Virtual machine HDD images are stored in the QCOW2 format.
- Backups are stored now in QCOW2 format.
- ReadyKernel (commercial feature). ReadyKernel is a new way of applying kernel updates without the need to reboot the physical server. For comparison, the previous versions of Virtuozzo support rebootless kernel upgrade that delivers a complete new kernel to the server, freezes all running instances, and tries to swap the kernel. While simple, this approach has a significant drawback of having to freeze running containers for a couple of minutes. In turn, ReadyKernel is based on a new upstream technology called kpatch. It does not require freezing running processes as it swaps calls one by one using ftrace. The overall patching process may take longer but end users would see no downtime. In addition, each ReadyKernel update is prepared and tested by kernel engineers.
- Container live migration via CRIU. In the previous versions of Virtuozzo, most operations performed during migration were done in the kernel space. As a result, the migration process imposed a lot of restrictions. Transfers of running instances were only possible between similar servers with the same kernel version. Any errors during the process could lead to a kernel and physical server crash. To improve upon migration, Virtuozzo launched a new project aiming to move most of the migration code to the user space, make the migration process reliable, and remove excessive restrictions. CRIU (which stands for Checkpoint/Restore in Userspace) became a huge success and is being widely adopted by the open-source community as well as commercial companies. As of today, the most active contributors to the project are engineers working at Ubuntu, Google, and Docker.
- Libvirt support. 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. Virtuozzo considers libvirt as a standard API for managing both virtual machines and containers. For demonstration or proof-of-concept purposes, you can use the Virtual Machine Manager to manage Virtuozzo hosts.
- Any backup (full or incremental) can now be deleted from the backup chain.

- Virtuozone provides 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.
- Automatic memory management (technical preview, commercial feature). Virtuozone 7 features a technical preview of automatic memory management for virtual machines and containers. In older versions of Virtuozone, overcommitted memory was reclaimed by means of swapping. Now several policies for the `vmmmd` daemon are available, enabling optimized balancing of VM and container memory in various environments. The policies involve the use of cgroups limits, swap, ballooning for virtual machines, and kernel same-page merging.
- Memory guarantees. A memory guarantee is a percentage of container or virtual machine's RAM that said container or VM is guaranteed to have. For virtual machines, the memory guarantee value is set to 40% by default; For containers, the default memory guarantee value is 0%. To change the default value, use the new `prlctl set --memguarantee` command.
- Memory hotplugging. The new and improved memory hotplugging technique allows both increasing and reducing virtual machine RAM size on the fly, without the need to reboot the VM. Configured by the command `prlctl --mem-hotplug`, memory hotplugging is implemented as a combination of ballooning and addition of virtual DIMM slots. When a command to increase VM memory size is run, the memory is first expanded by deflating the VM's balloon. If fully deflating the balloon is not enough to get the requested memory size, virtual DIMM slots are added. When a command to reduce VM memory size is run, the memory is shrunk by inflating the VM's balloon. This feature is disabled by default and only supported for virtual machines with at least 1GB of RAM.
- Kernel same-page merging. To optimize memory usage by virtual machines, Virtuozone uses a feature of Linux called Kernel Same-Page Merging (KSM). The KSM daemon `ksmmd` periodically scans memory for pages with identical content and merges those into a single page. Said page is marked as copy-on-write (COW), so when its contents are changed by a virtual machine, the kernel creates a new copy for that virtual machine. KSM enables the host to: avoid swapping due to merging of identical pages, run more virtual machines, overcommit virtual machine memory, and speed up RAM and hence certain applications and guest operating systems.

Note: For information on how to purchase a Virtuozone license and enable commercial features, please visit <https://virtuozone.com/>.

Chapter 3. Known Issues and Restrictions

- ploops over NFS are not supported. (PSBM-20108)
- In-place upgrading from Virtuozzo 6 to Virtuozzo 7 is not available. (PSBM-29221)
- Migration of Virtuozzo 6 VMs to Virtuozzo 7 is not supported for Ubuntu and Debian guests. (PSBM-31022)
- Linux virtual machines with EFI firmware cannot be migrated to Virtuozzo 7. (PSBM-32920)
- Virtuozzo Image Catalog is not available. (PSBM-39127)
- Hybrid clusters made of mixed Virtuozzo 6 and Virtuozzo 7 servers are not supported. (PSBM-40861)
- 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)
- Virtuozzo Automator and Power Panel are not supported (PSBM-41740)
- Restoring VMs from Virtuozzo 6 backups is not supported. (PSBM-45736)
- Migration of VMs on NFS or Virtuozzo Storage is not supported. (PSBM-46623, PSBM-40856)
- Migration of containers with NFS client inside is not supported. (PSBM-47068)
- Windows-based VMs could freeze on Virtuozzo Storage after multiple cluster failures. The workaround is to remove USB devices from problematic VMs.(PSBM-47925)
- Secondary SCSI disks can be offline after v2v migration from Virtuozzo 6 to Virtuozzo 7. As a workaround, each disk can be brought 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-48700)
- Secondary disks of containers restored from Virtuozzo 7 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 Virtuozzo 7 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 Virtuozzo 7. 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 Virtuozzo 7 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.