



Virtuozzo Hybrid Server 7

Readme

December 20, 2022

Virtuozzo International GmbH

Vordergasse 59

8200 Schaffhausen

Switzerland

Tel: + 41 52 632 0411

Fax: + 41 52 672 2010

<https://virtuozzo.com>

Copyright ©2016-2022 Virtuozzo International GmbH. All rights reserved.

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

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

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

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

Contents

- 1. **Overview** 1
- 2. **What's New** 2
- 3. **Known Issues and Restrictions** 5
- 4. **Dropped Functionality** 7
- 5. **Deprecated Features** 8
- 6. **Changelog** 9

CHAPTER 1

Overview

Virtuozzo Hybrid Server 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 Virtuozzo Hybrid Server 7 are:

- Virtuozzo Hybrid Server 7 is based on RHEL 7 and Kernel 3.10+.
- Virtuozzo Hybrid Server 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, `pr1ct1` 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 Hybrid Server 7 uses memcg. Balancing and configuring memcg limits enables getting the exact overcommit, shadow gangs, swap, page cache overuse Virtuozzo Hybrid Server parameters.
- Virtual machine HDD images are stored in the QCOW2 format.
- Backups are stored now in QCOW2 format.
- ReadyKernel (starting from Virtuozzo Hybrid Server 7.0 RTM Hotfix 1 and kernel 3.10.0-327.18.2.vz7.15.2). ReadyKernel is a new way of applying kernel updates without the need to reboot the physical server. For comparison, the previous versions of Virtuozzo Hybrid Server 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 Hybrid Server, 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 Hybrid Server hosts.
- Any backup (full or incremental) can now be deleted from the backup chain.
- Virtuozzo 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). Virtuozzo Hybrid Server 7 features a technical preview of automatic memory management for virtual machines and containers. In older versions of Virtuozzo Hybrid Server, overcommitted memory was reclaimed by means of swapping. Now several policies for the `vcmmnd` 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 80% 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, Virtuozzo Hybrid Server uses a feature of Linux 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. 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.

CHAPTER 3

Known Issues and Restrictions

- ploops over NFS are not supported. (PSBM-20108)
- In-place upgrading from Virtuozzo 6 to Virtuozzo Hybrid Server 7 is available for nodes with containers only. (PSBM-29221)
- 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)
- Secondary SCSI disks can be offline after v2v migration from Virtuozzo 6 to Virtuozzo Hybrid Server 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)
- 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)
- Changes to boot order are only applied after a complete VM shutdown. (PSBM-49119)
- Limited simfs support (feature provided as is). Unlike OpenVZ, the simfs layout is based on bindmounts in Virtuozzo Hybrid Server 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 `prlctl create <CT_name>`. The simfs limitations in Virtuozzo Hybrid Server 7 are: 1) No support for first- or second-level quotas, 2) No support for online migration of simfs-based containers. (OVZ-6613)
- Virtuozzo 6 OS and application EZ templates are not supported. (PSBM-54009)
- Disk hotplugging is not supported in Virtuozzo Storage. (PSBM-57929)
- No support for connecting physical floppy and CD/DVD drives to VMs. (PSBM-54477)

- Different UUIDs may be reported on host and in VM for a physical iSCSI HDD assigned to VM. VM restart fixes the issue. (PSBM-57700)
- Cloning of running virtual machines is not supported.
- Migration of virtual machines with snapshots is not supported. (PSBM-45517)
- Live migration and snapshots creation are not supported for containers with the enabled NFS server feature. (PSBM-67698)
- Live migration of Virtuozzo 6 containers to Virtuozzo Hybrid Server 7 servers is not supported. (PSBM-18457)
- Virtuozzo 6 Windows VMs without installed Virtuozzo guest tools cannot be migrated to Virtuozzo Hybrid Server 7, restored from Virtuozzo 6 or backed up to Virtuozzo Hybrid Server 7.
- CPU hotplugging is not available and results in crash in virtual machines with Windows Server 2008 R2 Service Pack 1. (PSBM-62201)
- Windows Server 2016 guests do not see hotplugged CPUs and do not display existing CPUs in device manager's device tree. This is a known issue caused by a conflict with one of the default Microsoft drivers shipped with Windows Server 2016. A workaround is available [here](#). If the workaround is applied, Windows Server 2016 guests may still behave erratically (crash, lock-up, etc.) on attempts to hot-add new vCPUs. This is a known guest problem which affects other hypervisors as well. (PSBM-65677)
- Network settings configured from inside Windows VMs currently cannot be preserved during migration of such VMs. (PSBM-49112)
- Containers with `ipt_owner` iptables rules cannot be migrated online. (PSBM-69409)

CHAPTER 4

Dropped Functionality

- Legacy VZFS container file system.
- Delayed /vz mounting and the corresponding VZMOUNTS parameter in the /etc/vz/vz.conf file.
- The 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.
- The `vziptables` parameter configured in the kickstart file used for unattended installations of Virtuozzo Hybrid Server.
- The `vzquota` utility.

CHAPTER 5

Deprecated Features

Features considered deprecated will be removed in the next major release of Virtuozzo Hybrid Server. They will, however, work and remain supported in Virtuozzo Hybrid Server 7. The following features are deprecated:

- The `ip_contrack_disable_ve0` kernel module option in the `/etc/modprobe.d/vz.conf` file.

CHAPTER 6

Changelog

The list of release notes for Virtuozzo Hybrid Server is available [here](#).