



OpenVZ Command Line Reference

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. Introduction	6
1.1. About OpenVZ	6
1.2. About This Guide	7
2. Managing OpenVZ	8
2.1. OpenVZ Configuration Files	8
2.1.1. Global OpenVZ Configuration File	8
2.1.2. Container Configuration File	11
2.1.2.1. Miscellaneous Parameters	12
2.1.2.2. Resource Management Parameters	13
2.1.2.3. Networking Parameters	16
2.1.3. Linux Distribution Configuration Files	17
2.1.4. Memory and IOPS Deduplication Configuration File	19
2.1.5. Network Classes Definition File	20
2.1.6. Kernel Parameters	20
2.1.7. Offline Management Configuration Files	21
2.1.8. vztt Configuration File	22
2.1.9. pcompact.conf	22
2.2. OpenVZ Utilities	23
2.2.1. prlsrvctl	23
2.2.1.1. prlsrvctl info	23
2.2.1.2. prlsrvctl net	24
2.2.1.3. prlsrvctl problem-report	27
2.2.1.4. prlsrvctl set	27
2.2.1.5. prlsrvctl shutdown	29
2.2.1.6. prlsrvctl usb	30
2.2.1.7. prlsrvctl user list	31
2.2.1.8. prlsrvctl user set	31
2.2.1.9. prlsrvctl privnet	32
2.3. OpenVZ Updates	33
3. Managing Containers	34
3.1. Matrix of OpenVZ Command-Line Utilities	34
3.2. prlctl	34
3.2.1. prlctl console	35
3.2.2. prlctl create	36
3.2.3. prlctl delete	36
3.2.4. prlctl exec, enter	37
3.2.5. prlctl migrate	37
3.2.6. prlctl mount, umount	38
3.2.7. prlctl move	38
3.2.8. prlctl problem-report	39
3.2.9. prlctl register, unregister	39
3.2.10. prlctl reinstall	40
3.2.11. prlctl set	41
3.2.11.1. General Options	41
3.2.11.2. Resource Management Options	42
3.2.11.3. Network Options	44
3.2.11.4. Hard Disk Drive Management Options	46

3.2.12. prctl snapshot, snapshot-list, snapshot-switch, snapshot-delete	48
3.2.13. prctl start, stop, restart, status	48
3.2.14. prctl suspend, resume	50
3.2.15. prctl list	50
3.2.15.1. prctl list Output Parameters	51
3.3. Migration Utilities	52
3.3.1. prctl clone	52
3.4. EZ Template Management Utilities	52
3.4.1. vzpkg	52
3.4.2. vzpkg install template	53
3.4.3. vzpkg update template	54
3.4.4. vzpkg remove template	54
3.4.5. vzpkg list	55
3.4.6. vzpkg info	56
3.4.7. vzpkg status	59
3.4.8. vzpkg install	59
3.4.9. vzpkg update	60
3.4.10. vzpkg remove	61
3.4.11. vzpkg create cache	62
3.4.12. vzpkg update cache	63
3.4.13. vzpkg remove cache	63
3.4.14. vzpkg create appcache	64
3.4.15. vzpkg update appcache	64
3.4.16. vzpkg remove appcache	65
3.4.17. vzpkg localinstall	65
3.4.18. vzpkg localupdate	66
3.4.19. vzpkg upgrade	67
3.4.20. vzpkg fetch	68
3.4.21. vzpkg clean	69
3.4.22. vzpkg update metadata	69
3.5. Supplementary Tools	70
3.5.1. pcompact	70
3.5.2. prl_disk_tool	70
3.5.2.1. prl_disk_tool compact	71
3.5.2.2. prl_disk_tool merge	71
3.5.2.3. prl_disk_tool resize	71
3.5.3. vzpid	72
3.5.4. vzps, vztop	72
3.5.5. vzsplite	72
4. Managing Virtual Machines	74
4.1. prctl	74
4.1.1. General Syntax	74
4.1.2. prctl clone	74
4.1.3. prctl create	75
4.1.4. prctl delete	76
4.1.5. prctl enter	76
4.1.6. prctl exec	76
4.1.7. prctl list	77
4.1.7.1. prctl list Output Parameters	77
4.1.8. prctl statistics	78

4.1.9. prctl migrate	80
4.1.10. prctl mount, umount	81
4.1.11. prctl move	81
4.1.12. prctl pause, suspend, resume	81
4.1.13. prctl problem-report	82
4.1.14. prctl register, unregister	82
4.1.15. prctl reset-uptime	83
4.1.16. prctl set	83
4.1.16.1. Modifying Virtual Machine Configuration	83
4.1.16.2. Managing Virtual Devices	88
4.1.17. prctl snapshot, snapshot-list, snapshot-switch, snapshot-delete	94
4.1.18. prctl start, stop, restart, reset, status	95

Chapter 1. Introduction

OpenVZ is a virtualization solution that allows you to run multiple virtual machines and containers on a single physical server.


This chapter provides general information about OpenVZ and this guide.

1.1. About OpenVZ

OpenVZ is a virtualization solution that allows you to simultaneously run multiple OpenVZ virtual machines and containers on a single physical server. With OpenVZ, you can efficiently share your server's hardware resources among virtual machines and containers.

OpenVZ is installed directly on the server hardware and does not need any operating system to function. Once it is installed, OpenVZ allows you to create virtual machines and containers and manage them using the OpenVZ command-line interface (CLI). The command-line interface comprises a set of OpenVZ command-line utilities that you can use to manage virtual machines and containers, both locally and remotely.

Graphically, a server with the OpenVZ software installed can be represented as follows:



1.2. About This Guide

This guide is a reference of OpenVZ configuration files and command-line utilities. It familiarizes you with the way to configure OpenVZ to meet your requirements and to perform various tasks by using the corresponding command-line utilities.

The primary audience for this guide is anyone who is looking for an explanation of a particular configuration option, needs help for a particular command, or is seeking for a command to perform a certain task.

Chapter 2. Managing OpenVZ

This chapter provides instructions on configuration files, scripts, and command-line utilities that can be used to configure the settings related to the OpenVZ software and the hardware node.

2.1. OpenVZ Configuration Files

The table below lists the configuration files available in OpenVZ. Most files are located in the `/etc` directory on a hardware node. If a configuration file is stored in a place other than the hardware node, its exact location is specified.

Name	Description
<code>/etc/vz/vz.conf</code>	Global configuration file. This file keeps system-wide settings, such as the default location of templates and global network settings.
<code>/etc/vz/conf/<CT_name>.conf</code>	Private configuration file of a container with the name <code><CT_name></code> . This file keeps container-specific settings: resource management parameters, the location of its private area, IP address, and so on.
<code>/etc/vz/conf/ve-<name>.conf-sample</code>	Sample files containing a number of default container configurations. Some pre-created samples file are shipped with OpenVZ (e.g., <code>basic</code> and <code>confixx</code>), but you can also create your own samples to meet your demands.
<code>/usr/libexec/libvzctl/dists/<distribution_name>.conf</code>	Linux distribution configuration files. These files define what scripts should be run when you perform specific operations with containers (e.g., when you set a new IP address for a container). The scripts differ from OpenVZ action scripts and depend on the Linux version a particular container is running.
<code>/etc/vz/pfcache.conf</code>	Configuration file used by the <code>pfcache</code> utility to manage memory and IOPS deduplication.
<code>/etc/vz/oom-groups.conf</code>	OOM killer configuration file with task badness adjustments.
<code>/etc/vz/conf/networks_classes</code>	Configuration file defining the network classes for traffic shaping and bandwidth management.
<code>/etc/sysctl.conf</code>	Kernel parameters. OpenVZ adjusts a number of kernel <code>sysctl</code> parameters and modifies the default <code>/etc/sysctl.conf</code> file.
<code>/etc/vztt/vztt.conf</code>	Configuration file used by the <code>vzpkg</code> utility to manage OS and application EZ templates.

2.1.1. Global OpenVZ Configuration File

OpenVZ keeps its system wide configuration parameters in the `/etc/vz/vz.conf` configuration file. This file is in shell format. Keep in mind that OpenVZ scripts source this file - thus, shell commands in

this file will cause system to execute them under root account. Parameters in this file are presented in the form `PARAMETER="value"`. Logically all the parameters belong to the following groups: global parameters, logging, disk quotas, template, network traffic, containers, validation and overcommitment, supplementary parameters, and name-based hosting parameters. Below is the description of all the parameters defined in this version of OpenVZ.

Name	Description	Default Value
VIRTUOZZO	This can be either yes or no. OpenVZ System V startup script checks this parameter. If set to no, then OpenVZ modules are not loaded. You might set it to "no" if you want to perform system maintenance and do not want to bring up all containers on the server.	yes
HTTP_PROXY	Specifies either the hostname or the IP address of the HTTP proxy server. After setting this parameter and in case you use an HTTP proxy server for handling all HTTP requests, the OpenVZ utilities communicating with the outer world through HTTP will use this server for managing all your HTTP messages.	n/a
ACTIONLOGDIR	This is the directory where <code>prlctl</code> keeps a log of its actions in the format suitable for OpenVZ statistics daemon <code>hwcoll</code> .	/vz/ actionlog
LOCKDIR	Actions on a container should be serialized, since two simultaneous operations on the same container may break its consistency. OpenVZ keeps lock files in this directory in order to serialize access to one container.	/vz/lock
VEFSTYPE	File system to use when caching OS templates: • ext4, • simfs.	ext4
IPV6	Defines whether the IPv6 support is enabled on the hardware node.	yes
GOLDEN_IMAGE	Enables (yes) or disables (no) embedding application templates into OS EZ template cache prior to creating containers based on this cache.	yes
PFCACHE	Path to the memory and IOPS deduplication cache with common container files.	/vz/pfcache
PFCACHE_IMAGE	Path to the private area of the memory and IOPS deduplication cache.	/vz/ pfcache.hdd
PFCACHE_IMAGE_SIZE	Image size (in 1KB blocks) of the memory and IOPS deduplication cache.	10485760
PFCACHE_INCLUDES	Directories for which memory and IOPS deduplication is enabled by default.	bin lib lib64 opt sbin usr

Name	Description	Default Value
VZ_TOOLS_BCID	Enables limits for the backup, restore, and migration operations.	
VZ_TOOLS_IOLIMIT	Sets the disk I/O limit for the backup, restore, and migration operations, in bytes per second. Not set by default.	

Table 2.1. Logging parameters

Name	Description	Default Value
LOGGING	This parameter defines whether <code>prlctl</code> should log its actions.	yes
LOGFILE	File where <code>libvzctl</code> logs the actions of programs linked to this library.	/var/log/vzctl.log
LOG_LEVEL	Logging verbosity, from 0 to 10 (higher is more verbose).	0

Table 2.2. Disk quota parameters

Name	Description	Default Value
DISK_QUOTA	Enables or disables disk quotas for containers. If set to no then disk space accounting will be disabled.	yes

Table 2.3. Network traffic parameters

Name	Description	Default Value
TRAFFIC_SHAPING	Traffic shaping allows you to limit the bandwidth consumed by containers for outgoing traffic. If it is set to "yes", then limitations will be turned on. If you want to use this feature, <code>TRAFFIC_ACCOUNTING</code> should be set to yes as well.	no
BANDWIDTH	This is the list of network interfaces on which we want to shape the traffic and their speed in the form of "dev:rate". The rate is measured in Kbps. If you want to shape traffic on more than one interface, set this parameter to <code>dev1:rate1 dev2:rate2</code> . For example, for two 100 Mbps Ethernet cards, set it to <code>eth0:102400 eth1:102400</code> .	eth0:102400
TOTALRATE	This parameter sets the size of the bandwidth pool for all containers. It is the upper limit for the bandwidth available to all your containers and is specified in the form of "dev:class:rate". The rate is measured in Kbps. Containers can consume bandwidth up to this limit in addition to the limit specified by the <code>RATE</code> parameter. Default value corresponds to 4 Mbps limit for the Class 1 containers.	eth0:1:4096
RATE	This parameter is the default bandwidth guaranteed to a container for outgoing traffic if the container configuration file does not explicitly specify a different value. This value is in the same format as <code>TOTALRATE</code> and its default value is "eth0:1:8". The rate is measured in Kbps. Note that 8 Kbps, offered by the default configuration, is the guarantee and the	eth0:1:8

Name	Description	Default Value
	container cannot consume less than this value and more than the sum of this value and TOTALRATE.	
RATEMPU	This optional parameter (where MPU stands for "minimum packet unit") limits the packet rate by making packets smaller than MPU in size consume HTB tokens. With it, small packets can be accounted as larger ones and limited by TOTALRATE and RATE parameters. Approximately, the maximum packets per second rate can be calculated as TOTALRATE / RATEMPU.	*:1:1000

Table 2.4. Template parameters

Name	Description	Default Value
TEMPLATE	This is the directory where to find templates. It is not recommended to redefine this option since all OpenVZ templates use the default directory.	/vz/template

Table 2.5. Container default parameters

Name	Description	Default Value
VE_ROOT	The mount point for container's root. Must contain the literal string \$VEID that will be substituted with the actual container UUID.	/vz/root/\$VEID
VE_PRIVATE	The directory where all the files and directories specific to the container are stored. Must contain the literal string \$VEID that will be substituted with the actual container UUID.	/vz/private/\$VEID
CONFIGFILE	The default configuration file sample to be used for the container creation; it may be overridden with the --config option of the vzctl create command.	basic
DEF_OSTEMPLATE	The default OS template to be used for the container creation.	centos-7
VE_ENVIRONMENT	Additional environment variables to be passed to the container init process. Should be provided as any number of <name>=<value> pairs separated by spaces.	

2.1.2. Container Configuration File

Each container has its own configuration file, which is stored in the /etc/vz/conf directory and has a name like <CT_name>.conf. This file has the same format as the global configuration file. The settings specified in this file can be subdivided into the following categories:

- miscellaneous,
- resource management parameters,

- networking.

2.1.2.1. Miscellaneous Parameters

The table below list the miscellaneous parameters you can set in the configuration file of a container:

Name	Description
VERSION	Specifies the OpenVZ version the configuration file applies to. 2 relates to OpenVZ version 4 and later.
ONBOOT	Specifies whether the container should be started automatically on system startup. OpenVZ automatically starts all containers that have this parameter set to "yes" upon startup.
ALLOWREBOOT	Specifies whether the container may be restarted with the <code>reboot</code> command run from inside. If omitted or set to yes, restarting is allowed.
OSTEMPLATE	The name of the OS template that was used for creating the container. You do not have to change this parameter; <code>prlctl</code> will set it for you upon calling the <code>prlctl create</code> command (or using the defaults from the global configuration file). The <code>.</code> symbol before the OS template name, if specified, indicates that this is an EZ OS template.
TEMPLATES	In a configuration file of an existing container, this parameter lists application templates installed with the <code>prlctl create</code> or <code>vzpkg install</code> commands. In this case you should not modify it, because it is used by template management utilities to track installation history. This parameter is omitted if no templates have been installed to the container.
VE_ROOT	Overrides the <code>VE_ROOT</code> parameter from the global configuration file.
VE_PRIVATE	Overrides the <code>VE_PRIVATE</code> parameter from the global configuration file.
VE_ENVIRONMENT	Overrides the <code>VE_ENVIRONMENT</code> parameter from the global configuration file.
TECHNOLOGIES	Determines a set of technologies which should be provided by the OpenVZ kernel for container operation. Currently, this parameter can contain the information about the following technologies: • The system architecture of the container (x86, x86_64, or i64). • Whether the container is based on the OS template supporting the Native POSIX Thread Library (NPTL). In this case, the <code>nptl</code> entry is specified as the value of this parameter. • Whether the OS EZ template the container is based on requires the <code>sysfs</code> filesystem support (e.g., the OS EZ template for SUSE Linux Enterprise 10).
DISABLED	If set to yes, disables the container making it impossible to start the container once it was stopped. You can start the disabled container after setting the value of this parameter to no.
DESCRIPTION	Sets the description for the container. Note: You are allowed to use only symbols in the A-z and 0-9 ranges in your descriptions.

Name	Description
NAME	Container name that can be used to refer to said container in commands. Names must be alphanumeric and may contain the characters \, -, _. Names with white spaces must be enclosed in quotation marks.
ORIGIN_SAMPLE	The configuration sample the container was based on when created.
CONFIG_CUSTOMIZED	Indicates whether any of the container configuration parameters have been modified as regards its original configuration sample. If this parameter is omitted, its value is considered as no.
UUID	The container unique identifier. This identifier is used by certain OpenVZ utilities during their execution.

2.1.2.2. Resource Management Parameters

All resource management parameters can be subdivided into the CPU, disk, system, and VSwap categories for your convenience. Any parameter can be set with the `prlctl set` command and the corresponding option name (in the lower case, e.g., `--cpuunits` for CPUUNITS, etc.). See [Chapter 3, Managing Containers](#) on page 34 for more details. The **Typical value** column, if present, specifies a range of reasonable parameter values for different applications, from light to huge heavy loaded containers. If the barrier and limit fields are in use, ranges for both thresholds are given.

Table 2.6. CPU Parameters

Parameter	Description	Typical value
CPUUNITS	CPU weight. This is a positive integer number that defines how much CPU time the container can get as compared to the other virtual machines and containers running on the server. The larger the number, the more CPU time the container can receive. Possible values range from 8 to 500000. If this parameter is not set, the default value of 1000 is used.	250 . . . 1000
CPULIMIT, CPULIMIT_MHZ	CPU limit, in per cent (CPULIMIT) or megahertz (CPULIMIT_MHZ), the container is not allowed to exceed. The parameter is not set for newly created containers; so they can consume all free CPU power of the server. When setting this parameter in per cent, keep in mind that one CPU core makes up 100%. So if the server has 4 CPU cores, the total CPU power will equal 400%.	
CPUS	Number of CPU cores defining the CPU limit for a container. The limit is calculated by multiplying the power of one CPU core by the number of the specified CPU cores. This option also defines the number of CPUs shown to users from inside a container. This parameter is not set for newly created containers; so they can consume all free CPU power of the server.	

Parameter	Description	Typical value
CPUMASK	The CPU affinity mask defining which CPUs on the Node can be used to handle the processes running in the container. The CPU mask can be specified as both separate CPU index numbers (1,2,3) and CPU ranges (2-4,5-7).	
NODEMASK	The NUMA node mask defining a NUMA node to bind the container to. Once you set the mask, the processes running in the container will be executed only on the CPUs that belong to the specified NUMA node.	

Table 2.7. Disk Parameters

Parameter	Description	Typical value
DISKSPACE	Total size of disk space that can be consumed by the container, in 1 KB blocks.	
QUOTAUGIDLIMIT	This parameter enables (if set to a value other than 0) or disables (if set to 0) per-user and per-group quotas for further management with the standard Linux quota utility. Enabling per-user and per-group quotas for a container requires restarting the container.	0 . . . N
IOPRIO	The container priority for disk I/O operations. The higher the priority, the more time the container has for writing to and reading from the disk. The default container priority is 4.	0 - 7
IOPSLIMIT	The maximum number of disk input and output operations per second a container is allowed to perform. By default, any newly created container does not have the IOPS limit set and can perform so many disk I/O operations per second as necessary.	
IOLIMIT	The bandwidth a container is allowed to use for its disk input and output (I/O) operations. By default, the limit is set in megabytes per second. However, you can use the following suffixes to use other measurement units: • G - sets the limit in gigabytes per second. • K - sets the limit in kilobytes per second. • B - sets the limit in bytes per second. In the current version of OpenVZ, the maximum I/O bandwidth limit you can set for a container is 2 GB per second.	

Parameter	Description	Typical value
	The default I/O bandwidth limit for all newly created containers is set to 0, which means that no limits are applied to any containers.	

Table 2.8. System Parameters

Parameter	Description	Typical value
NUMPROC	Number of processes and threads allowed. Upon hitting this limit, container will not be able to start a new process or thread.	40 . . . 400
AVNUMPROC	Number of processes expected to run in the container on average. This is informational parameter used to ensure configuration correctness.	0 . . . NUMPROC
VMGUARPAGES	Memory allocation guarantee, in pages. Applications are guaranteed to be able to allocate memory while the amount of memory accounted as <code>privvmpages</code> does not exceed the configured barrier of the <code>vmguarpages</code> parameter. Above the barrier, memory allocation is not guaranteed and may fail in case of overall memory shortage.	1725 . . . 107520
LOCKEDPAGES	Memory not allowed to be swapped out (locked with the <code>mlock()</code> system call), in pages (one page is 4 KB).	4 . . . 4096
SHMPAGES	Total size of shared memory (including IPC, shared anonymous mappings and <code>tmpfs</code> objects), allocated by processes of a particular container, in pages.	512 . . . 16384
PRIVVMPAGES	Size of private (or potentially private) memory, allocated by an application. Memory that is always shared among different applications is not included in this resource parameter.	
NUMFILE	Number of files opened by all container processes.	512 . . . 8192
NUMFLOCK	Number of file locks created by all container processes.	50 . . . 200 - 60 . . . 220
NUMPTY	Number of pseudo-terminals. For example, the <code>ssh</code> session, <code>screen</code> , the <code>xterm</code> application consumes pseudo-terminal resources.	4 . . . 64
NUMSIGINFO	Number of <code>siginfo</code> structures (essentially this parameter limits the size of signal delivery queue).	256 . . . 512
PHYSPAGES	Total size of RAM used by processes. This parameter is used for accounting purposes only. It shows the usage of RAM by the container. For memory pages used by several different containers (mappings of shared libraries, for example), only	Not limited

Parameter	Description	Typical value
	a fraction of a page is charged to each container. The sum of the physpages for all containers corresponds to the total number of pages used in the system by all accounted users.	
NUMIPTENT	The number of IP packet filtering entries.	12 . . . 128

Table 2.9. VSwap Parameters

Parameter	Description	Typical value
PHYS_PAGES	Amount of RAM that can be used by the processes of a container, in 4KB pages.	
SWAP	Amount of swap space that can be used by the container for swapping out memory once the RAM is exceeded, in 4KB pages.	
VM_OVERCOMMIT	Memory overcommit factor that defines the memory allocation limit for a container. The limit is calculated as (PHYS_PAGES + SWAP) * factor.	Not limited

2.1.2.3. Networking Parameters

Network-related parameters allow you to set bandwidth management parameters, hostname and IP addresses that a container can use, and other parameters.

Name	Description
HOSTNAME	If this parameter is specified, then <code>pr1ct1</code> will set the hostname to its value upon the next container start. This parameter can be omitted. In this case, the container administrator should configure the hostname manually.
IP_ADDRESS	This is the list of IP addresses, which can be used on container network interfaces. This list is an argument of the container start call and it is impossible to assign IP address from inside the container if the address is not on the list. Any IP address assigned from within the container will be visible only within the container.
NAMESERVER	The IP address of the DNS server the container is supposed to use. More than one server can be specified in the space-separated format.
SEARCHDOMAIN	DNS search domains for the container. More than one domain can be specified.
NETDEV	The names of physical network adapters that have been moved from the server to the given container.
NETFILTER	Indicates which <code>iptables</code> modules are allowed for the container. If some of the allowed modules are not loaded on the destination Hardware Node after migration or restoration from backup, they will be automatically loaded on the migrated or restored container start. The following modes are available: - disabled: none. - stateless: (default) all modules except <code>conntrack</code> and NAT-related.

Name	Description
	• <code>stateful</code>: all modules except NAT-related. • <code>full</code>: all modules.
NETIF	Specifies a number of parameters for the virtual network adapters existing inside the container. These parameters include: • <code>ifname</code>: the name of the veth virtual Ethernet interface inside the container. • <code>mac</code>: the MAC address assigned to the veth virtual Ethernet interface inside the container. • <code>host_mac</code>: the MAC address assigned to the veth virtual Ethernet interface on the server. • <code>network</code>: the name of the virtual network where the veth virtual network adapter is included. • <code>ip</code>: the IP address(es) assigned to the veth virtual network adapter.
RATE	If traffic shaping is turned on, then this parameter specifies bandwidth guarantee, in Kbps, for the container. The parameters should be set in the form of <code>eth0:1:8</code> .
RATEBOUND	If set to yes, the bandwidth guarantee is also the limit for the container, and the container cannot borrow the bandwidth from the TOTALRATE bandwidth pool.

2.1.3. Linux Distribution Configuration Files

Some OpenVZ tools (e.g., `prlctl`) need to run special scripts inside a container to perform certain operations on it. However, carrying out one and the same operation inside containers running different Linux versions may require execution of different actions. This may be caused by the fact that different Linux distributions store files in different locations, use different commands to complete one and the same task, and so on. To distinguish between containers running different Linux versions and to determine what scripts should be executed while performing the relevant container-related operations, OpenVZ uses special distribution configuration files located in the `/usr/libexec/libvzctl/dists` directory on the server.

There are a number of distribution configuration files shipped with OpenVZ by default (`centos.conf`, `fedora-core.conf`, `gentoo.conf`, etc.). To view all configuration files available on your OpenVZ, you can go to the `/usr/libexec/libvzctl/dists` directory and issue the `ls` command. The distribution configuration files will be displayed in the form of `<Linux_distribution>-<version>.conf` where `<Linux_distribution>` and `<version>` denote the name of the Linux distribution and its version, respectively (e.g., `centos-7.conf`).

Any distribution configuration file consists of a number of entries in the form of `<parameter_name>=<script_name>` where `<parameter_name>` denotes the name of the parameter defining the operation when the script in the right part of the entry is to be executed and `<script_name>` is the name of the script to be run on performing the operation defined by the parameter in the left part of the entry. In the current version of OpenVZ, the following parameters are used to define what scripts should be executed for the corresponding Linux version a container is running:

- `ADD_IP`: the script specified as the value of this parameter has the default name of `<distribution_name>-add_ip.sh` and is used to configure the network settings during the container

startup and the IP address(es) assignment. The script is launched inside the container on executing the following commands:

```
prlctl start <CT_name>
prlctl set <CT_name> --ipadd <IP_address>
prlctl set <CT_name> --ipadd <IP_address> --ipdel all
```

- DEL_IP: the script specified as the value of this parameter has the default name of `<distribution_name>-del_ip.sh` and is used to delete an existing IP address from the container. The script is launched inside the container on executing the following commands:

```
prlctl set <CT_name> --ipdel <IP_address>
prlctl set <CT_name> --ipdel all
```

- SET_HOSTNAME: the script specified as the value of this parameter has the default name of `<distribution_name>-set_hostname.sh` and is used to configure the hostname of the container. The script is launched inside the container on executing the following command:

```
prlctl set <CT_name> --hostname <name>
```

- SET_DNS: the script specified as the value of this parameter has the default name of `<distribution_name>-set_dns.sh` and is used to configure DNS parameters in the `/etc/resolv.conf` file. The script is launched inside the container on executing the following command:

```
prlctl set <CT_name> --searchdomain <domain> --nameserver <IP_address>
```

- SET_USERPASS: the script specified as the value of this parameter has the default name of `<distribution_name>-set_userpass.sh` and is used to add a new user or change the current password. The script is launched inside the container on executing the following command:

```
prlctl set <CT_name> --userpasswd <user>:<passwd>
```

- SET_UGID_QUOTA: the script specified as the value of this parameter has the default name of `<distribution_name>-set_ugid_quota.sh` and is used to set up per-user/group quota. The script is launched inside the container on executing the following command:

```
prlctl set <CT_name> --quotaugidlimit <num>
```

- POST_CREATE: the script specified as the value of this parameter has the default name of `<distribution_name>-postcreate.sh` and is used to perform certain tasks (e.g., to modify the crontab files) after the container creation. This script is launched on the server on executing the following command:

```
prlctl create <CT_name>
```

- POST_MIGRATE: the script specified as the value of this parameter has the default name of `<distribution_name>-post_migrate.sh` and is used to perform certain operations on the container where the physical server has been successfully migrated. This script is launched inside the container on executing the following command:

```
vzpv2v [<options>] --ctid <CT_name>
```

The scripts specified in distribution configuration files are located in the `/usr/libexec/libvzctl/dists/scripts` directory on the server and executed on performing the aforementioned operations on the containers. After an operation has been initiated, the `prlctl` or `vzp2v` utility turns to the corresponding container configuration file, looks for the value of the `DISTRIBUTION` variable or, if the latter is not present, of the `OSTEMPLATE` variable in this file, and defines on their basis what Linux version the given container is running. After that, `prlctl` reads the corresponding configuration file for the determined Linux version from the `/usr/libexec/libvzctl/dists/` directory and executes the scripts specified in this file.

Note: If no distribution is specified as the value of the `DISTRIBUTION` and `OSTEMPLATE` variables in the container configuration file or no configuration file for the given Linux version was found in the `/usr/libexec/libvzctl/dists` directory, the default file from this directory is used.

2.1.4. Memory and IOPS Deduplication Configuration File

Contained in the `/etc/vz/pfcache.conf` file, memory and IOPS deduplication parameters allow you to tailor cache behavior and performance to your needs.

Table 2.10. Parameters

Name	Description	Default Value
COUNT	The minimum number of file copies required for the file to become cacheable. Copies may exist in the same container or different containers.	2
MINSIZE	Minimal cacheable file size, bytes. Files smaller than this value will not be cached.	0
MAXSIZE	Maximal cacheable file size, bytes. Files larger than this value will not be cached.	2147483648
TIMEOUT	Time between caching attempts, seconds.	5
PFCACHE_IOLIMIT	Memory and IOPS deduplication cache IO bandwidth limit, bps. Unlimited by default.	
PFCACHE_IOPSLIMIT	Memory and IOPS deduplication cache IOPS limit. Unlimited by default.	
LOGLEVEL	Logging verbosity. Messages are logged in the system log file <code>/var/log/messages</code> .	1
PAGEMIN	The total number of memory pages used in containers: • 0 - Cache even files without memory pages. • 1 - Cache only files in use. • N - Cache only when the total number of memory pages in containers reaches N.	1
PURGEAHEAD	Extra cache space to free up in addition to the requested space. In per cent of the requested space. Used with the <code>pfcache purge --size</code> command.	20%

2.1.5. Network Classes Definition File

In OpenVZ, both traffic accounting and bandwidth management are based on network classes. The network classes' definition file (/etc/vz/conf/networks_classes) describes network classes that OpenVZ recognizes. Currently, there can be up to 15 classes defined.

The lines in this file have the following format:

```
<class_ID> <IP_address>/<prefix_length>
```

where <class_ID> defines the network class identifier, <IP_address> defines the starting IP address, and <prefix_length> defines the subnet mask. In pair <IP_address> and <prefix_length> define the range of IP addresses for this class. There may be several lines for each class. Classes should be defined after Class 1 and represent exceptions from the "matching-everything" rule of Class 1. Class 0 has a special meaning and defines the IP ranges for which no accounting is done (this server container addresses).

The definition of class 1 is required; any class except class 1 can be omitted. However, it is recommended to define class 0 correctly as it will improve performance. For example:

```
# HW node VPS's networks
0 10.10.10.0/24
0 10.10.15.0/24
# all IP("local" traffic)
1 0.0.0.0/0
# class 2 - "foreign" traffic
#2 10.0.0.0/8
#2 11.0.0.0/8
# inside "foreign" network there
# is a hole with "local" traffic
#1 10.10.16.0/24
```

2.1.6. Kernel Parameters

There is a number of kernel limits that should be set for the OpenVZ software to work correctly. OpenVZ is shipped with a tuned /etc/sysctl.conf file. Understanding what parameters were changed is essential for running the required number of containers. Below is the contents of the /etc/sysctl.conf file as shipped with OpenVZ:

```
# Controls IP packet forwarding
net.ipv4.ip_forward = 1
# Controls source route verification
net.ipv4.conf.default.rp_filter = 1
# Do not accept source routing
net.ipv4.conf.default.accept_source_route = 0
# Controls the System Request debugging functionality of the kernel
kernel.sysrq = 1
# Controls whether core dumps will append the PID to the core filename.
# Useful for debugging multi-threaded applications.
kernel.core_uses_pid = 1
```

```
# Controls the use of TCP syncookies
net.ipv4.tcp_syncookies = 1
# Disable netfilter on bridges.
net.bridge.bridge-nf-call-ip6tables = 1
net.bridge.bridge-nf-call-iptables = 1
net.bridge.bridge-nf-call-arptables = 0
# Controls the default maximum size of a message queue
kernel.msgmnb = 65536
# Controls the maximum size of a message, in bytes
kernel.msgmax = 65536
# Controls the maximum shared segment size, in bytes
kernel.shmmax = 68719476736
# Controls the maximum number of shared memory segments, in pages
kernel.shmall = 4294967296
net.ipv6.conf.all.proxy_ndp=1
net.ipv4.conf.default.proxy_arp = 0
net.ipv4.conf.all.rp_filter = 0
fs.super-max = 2560
fs.file-max = 262144
kernel.fairsched-nodes-max = 1538
net.ipv4.neigh.default.gc_thresh2 = 2048
net.ipv4.neigh.default.gc_thresh3 = 4096
net.ipv4.conf.default.send_redirects = 0
net.ipv4.conf.all.send_redirects = 0
net.ipv6.neigh.default.gc_thresh2 = 2048
net.ipv6.neigh.default.gc_thresh3 = 4096
net.nf_conntrack_max = 500000
fs.aio-max-nr = 1048576
```

Notice that some parameters of the kernel configuration depends on the maximum number of containers you plan to run. In the default configuration file, these numbers were calculated under the assumption the maximum container number is 512. If you plan to run another number of containers, it is recommended to recalculate `net.ipv4.neigh.default.gc_thresh2` and `net.ipv4.neigh.default.gc_thresh3` parameters as three per container plus 128...512. Keep the second parameter twice as great as the first one.

To apply the changes issue the following command:

```
# sysctl -p
```

Besides, it makes sense to set `net.ipv4.tcp_use_sg` to 0, since the corresponding "Scatter/gather IO" feature is not supported by the `venet` device, used in OpenVZ networking.

It is also worth mentioning that normally you should have forwarding turned on since the server forwards packets destined to or originated from containers.

2.1.7. Offline Management Configuration Files

The offline management configuration files located in the `/etc/vzredirect.d` directory define various modes of container offline management by container administrators. One configuration file describes one offline management mode. In the current OpenVZ version, two files are accessible: `vzpp.conf` and `vzpp-plesk.conf`. The first file defines the container offline management by means of Power Panel, and the second one by means of the same Power Panel with an integrated Plesk control panel.

There are two parameters in each of the files.

Format

```
<parameter_name>=<parameter_value>
```

Table 2.11. Parameters

Name	Description	Example
PORT	This port must be entered in the address line of an Internet browser after the container IP address when managing the container by means of Power Panel or the Plesk control panel.	PORT=8443
DST_VEID	The UUID of the container where the requests coming to the specified port will be redirected.	DST_VEID=1

2.1.8. vztt Configuration File

This file (/etc/vztt/vztt.conf) is the configuration file used by the vzpkg utility when managing OS and application EZ templates.

Format

```
<parameter_name>=<parameter_value>
```

Table 2.12. Parameters

Name	Description
VZTT_PROXY	The IP address or hostname of the caching proxy server to be used by the vzpkg tool for managing OS and application EZ templates.
HTTP_PROXY	The IP address or hostname of the HTTP proxy server address, if you use this server.
HTTP_PROXY_USER	The user name used by the HTTP proxy server for your authentication.
HTTP_PROXY_PASSWORD	The password of the user specified in the HTTP_PROXY_USER parameter and used for your authentication by the HTTP proxy server.
METADATA_EXPIRE	Defines the period of time, in seconds, in the course of which the downloaded software packages in the vzpkg cache are regarded as <i>not obsolete</i> . During this time, the vzpkg utility searches for the EZ template packages in the local cache only (without checking the remote repositories set for EZ templates). By default, this period is set to 86400 seconds (24 hours).
EXCLUDE	List of comma-separated packages that are not to be installed or updated during the vzpkg execution. The package names should correspond to the name of real packages in the repository and can contain file globs (e.g., * and ?).

2.1.9. pcompact.conf

The /etc/vz/pcompact.conf file is used by the pcompact utility to compact virtual disks in containers.

Format

```
<parameter_name>=<parameter_value>
```

Table 2.13. Parameters

Name	Description
THRESHOLD=<number>	Compact the virtual disk if unused space on it exceeds THRESHOLD percent of the ploop size.
DELTA=<number>	Reduce disk space to be compacted by DELTA percent of the ploop size.
DEFRAG=<yes no>	Perform or skip file system defragmentation.

2.2. OpenVZ Utilities

This section provides information on utilities that can be used to manage OpenVZ parameters.

2.2.1. prlsrvctl

The `prlsrvctl` command-line utility is used to perform management tasks on the hardware node and OpenVZ. The tasks include getting the OpenVZ information, modifying its preferences, installing a license, obtaining statistics and problem reports, and some others.

Syntax

```
prlsrvctl [<command> [<options>]
          [-l, --login [<user>[:<passwd>]@]<server>[:<port>]]]
```

Table 2.14. Options

Name	Description
<command>	The command to execute.
<options>	Command options. See individual commands for available options.
-l, --login	Connect to the remote hardware node and execute a command there. If this parameter is omitted, the command will be executed on the local server.
<user>	The name of the user used to log in to the remote server.
<passwd>	The user password. If the password is omitted, you will be prompted to enter it.
<server>:<port>	The remote server IP address or hostname and port number, If port number is omitted, the default port will be used.

Note: To display help, enter `prlsrvctl` on the command-line without any options.

2.2.1.1. prlsrvctl info

Displays the hardware node and OpenVZ configuration information.

Syntax

```
prlsrvctl info
```

The information returned by the `info` command includes the following:

- Server ID and hostname.
- OpenVZ version number.
- Default directory for storing virtual machine files.
- OpenVZ memory limits.
- OpenVZ minimum allowable security level.
- Default directory for storing virtual machine backups.
- OpenVZ license information.
- Server hardware configuration information.
- Other miscellaneous info.

2.2.1.2. prlsrvctl net

The `prlsrvctl net` command is used to create and configure virtual networks.

Table 2.15. Subcommands

Name	Description
<code>net add</code>	Creates a new virtual network
<code>net set</code>	Configures the parameters of an existing virtual network.
<code>net del</code>	Removes an existing virtual network.
<code>net list</code>	List the available virtual networks.

net add

Creates a new virtual network.

Syntax

```
prlsrvctl net add <vnetwork_ID> [-i, --ifname <if>] [-m, --mac <mac_address>]
[-t, --type <bridged|host-only>] [-d, --description <desc>]
[--ip <IP_address>[/<mask>]] [--dhcp-server <on|off>]
[--dhcp-ip <IP_address>] [--ip-scope-start <IP_address>]
[--ip-scope-end <IP_address>] [--ip6 <IP_address>[/<mask>_]]
[--dhcp6-server <on|off>] [--dhcp-ip6 <IP_address>]
[--ip6-scope-start <IP_address>] [--ip6-scope-end <IP_address>]
```

Table 2.16. Options

Name	Description
<code><vnetwork_ID></code>	A user-defined name that will identify the new virtual network.
<code>-i, --ifname <if></code>	The name of a physical network adapter on the hardware node to which this virtual network should be bound.

Name	Description
<code>-m, --mac <mac_address></code>	The MAC address of a virtual network adapter on the hardware node to which this virtual network should be bound.
<code>-t, --type <bridged host-only></code>	The type of the virtual network to create. Possible values are: <code>bridged</code>. A virtual machine and container connected to this type of virtual network appears as an independent computer on the network. <code>host_only</code> (default). A virtual machine and container connected to this type of virtual network can access only the hardware node and the virtual machines and containers connected to the same virtual network.
<code>-d, --description <desc></code>	A user-defined description of the virtual network. Descriptions with white spaces must be enclosed in quotation marks.
<code>--ip <IP_address>[/<mask>]</code> <code>--ip6 <IP_address>[/<mask>]</code>	Set an IPv4/IPv6 address and subnet mask for the OpenVZ virtual adapter.
<code>--dhcp-server <on off></code> <code>--dhcp6-server <on off></code>	Enable or disable the OpenVZ virtual DHCPv4/DHCPv6 server.
<code>--dhcp-ip <IP_address></code> <code>--dhcp-ip6 <IP_address></code>	Set an IPv4/IPv6 address for the OpenVZ virtual DHCPv4/DHCPv6 server.
<code>--ip-scope-start <IP_address></code> <code>--ip-scope-end <IP_address></code> <code>--ip6-scope-start <IP_address></code> <code>--ip6-scope-end <IP_address></code>	Set the starting and ending IPv4/IPv6 addresses for the DHCPv4/DHCPv6 pool. The virtual machines and containers connected to the network you are creating will automatically receive their IPv4/IPv6 addresses from the respective DHCPv4/DHCPv6 pool.

net set

Configures the settings of an existing virtual network.

Syntax

```
prlsrvctl net set <vnetwork_ID> [-i, --ifname <if>] [-m, --mac <mac_address>]
[-t, --type <bridged|host-only>] [-d, --description <desc>]
[--ip <IP_address>[/<mask>]] [--dhcp-server <on|off>]
[--dhcp-ip <IP_address>] [--ip-scope-start <IP_address>]
[--ip-scope-end <IP_address>] [--ip6 <IP_address>[/<mask>_]]
[--dhcp6-server <on|off>] [--dhcp-ip6 <IP_address>]
[--ip6-scope-start <IP_address>] [--ip6-scope-end <IP_address>]
```

Table 2.17. Options

Name	Description
<code><vnetwork_ID></code>	The name of the virtual network to modify.
<code>-i, --ifname <if></code>	The name of a physical network adapter on the hardware node to which this virtual network should be bound.

Name	Description
<code>-m, --mac <mac_address></code>	The MAC address of a virtual network adapter on the hardware node to which this virtual network should be bound.
<code>-t, --type <bridged host-only></code>	The type of the virtual network to modify. Possible values are: • <code>bridged</code>. A virtual machine and container connected to this type of virtual network appears as an independent computer on the network. • <code>host_only</code> (default). A virtual machine and container connected to this type of virtual network can access only the hardware node and the virtual machines and containers connected to the same virtual network.
<code>-d, --description <desc></code>	A user-defined description of the virtual network. Descriptions with white spaces must be enclosed in quotation marks.
<code>--ip <IP_address>[/<mask>]</code> <code>--ip6 <IP_address>[/<mask>]</code>	Set an IPv4/IPv6 address and subnet mask for the OpenVZ virtual adapter.
<code>--dhcp-server <on off></code> <code>--dhcp6-server <on off></code>	Enable or disable the OpenVZ virtual DHCPv4/DHCPv6 server.
<code>--dhcp-ip <IP_address></code> <code>--dhcp-ip6 <IP_address></code>	Set an IPv4/IPv6 address for the OpenVZ virtual DHCPv4/DHCPv6 server.
<code>--ip-scope-start <IP_address></code> <code>--ip-scope-end <IP_address></code> <code>--ip6-scope-start <IP_address></code> <code>--ip6-scope-end <IP_address></code>	Set the starting and ending IPv4/IPv6 addresses for the DHCPv4/DHCPv6 pool. The virtual machines and containers connected to the network you are creating will automatically receive their IPv4/IPv6 addresses from the respective DHCPv4/DHCPv6 pool.

net del

Deletes an existing virtual network.

Syntax

```
prlsrvctl net del <vnetwork_ID>
```

Table 2.18. Options

Name	Description
<code><vnetwork_ID></code>	The name of the virtual network to delete.

net list

Lists the existing virtual networks.

Syntax

```
prlsrvctl net list
```

2.2.1.3. prlsrvctl problem-report

Generates and displays problem reports.

Syntax

```
prlsrvctl problem-report
```

The command collects technical data about OpenVZ and the hardware node and displays the report on screen (the output can also be piped to a file). The report can then be directed to the OpenVZ technical support team for analysis.

2.2.1.4. prlsrvctl set

Configures OpenVZ preferences.

Syntax

```
prlsrvctl set [--mem-limit <auto>|<size>]
               [-s, --min-security-level <low|normal|high>]
               [-c, --cep <on|off>] [--mng-settings <allow|deny>]
               [--device <device> --assignment <host|vm>_]
               [--backup-storage [<user>[:<passwd>]@]<server>[:<port>]]
               [--backup-tmpdir <tmp_dir>] [--backup-path <path>]
               [--idle-connection-timeout <timeout>]
               [--verbose-log <on|off>] [--cluster-mode <on|off>]
               [--cpu-features-mask <{+|-}feature1,feature2=value[,...]>]
               [--vm-cpulimit-type <full|guest>]
```

Table 2.19. Options

Name	Description
--mem-limit {auto <size>}	Sets the upper limit of the memory size that can be reserved for use by virtual machines. The following options are available: • auto - if this option is used, the memory size will be calculated automatically. • size - user-defined memory size, in megabytes.
-s, --min-security-level <low normal high>	The lowest allowable security level that can be used to connect to the hardware node. The following options are available: • low - plain TCP/IP (no encryption). • normal - most important data is sent and received using SSL over TCP/IP (user credentials during login, guest OS clipboard, etc.) Other data is sent and received using plain TCP/IP with no encryption.

Name	Description
	• high - all of the data is sent and received using SSL.
-c, --cep <on off>	Enables/disables the participation in the Customer Experience Program (CEP). The following options are available: • on - enables CEP. • off - disables CEP.
--mng-settings <allow deny>	Allows to grant or deny permission to new users to modify OpenVZ preferences. By default, only administrators of the host OS can modify OpenVZ preferences. When a new user profile is created (this happens when a user logs in to the hardware node for the first time), he/she will be granted or denied this privilege based on the default setting. This parameter allows you to set that default setting. Please note that this parameter only affects new users (the users that will be created in the future). The profiles of the existing users will not be modified.
--device <device> --assignment <host vm>	Allows to set the assignment mode for the specified VTd device. The following options are available: • host - assign the device to the hardware node. • vm - assign the device to virtual machines.
--backup-storage [<user>[:<passwd>]@]<server>[:<port>]	The default backup server where to store virtual machine backups.
--backup-path <path>	The name and path of the default directory on the backup server where to store virtual machine backups.
--verbose-log <on off>	Turns the verbose output for the command on or off.
--cluster-mode <on off>	Turns the cluster mode on or off.
--idle-connection-timeout <timeout>	Sets a timeout interval in seconds after which, if no data has been received from the storage server or backup client, the process of backup/restore is terminated.
--backup-tmpdir <tmp_dir>	Specifies a temporary directory where special snapshots created during virtual machine backup will be stored. This may be necessary so as not to run out of storage space on physical servers where most of the storage space is allocated to virtual machines and very little is left for the server itself.

Name	Description
<pre>--cpu-features-mask <{+ -}feature1,feature2=value[,...]></pre>	Changes CPU features mask on the host. To mask/unmask features, use the +feature/-feature syntax respectively. Omitting the sign is equivalent to unmasking. Features that require specific value can be set using the feature=value syntax. To view a full list of host CPU features which are supported, unmaskable and already masked, run the <code>prlsrvctl info --full</code> command. Note: All virtual machines and containers on the host must be stopped.
<pre>--vm-cpulimit-type <full guest></pre>	Specifies the type of virtual machine threads to be affected by the CPU limit: • <code>full</code> (default) - both hardware emulation and guest OS threads are limited. • <code>guest</code> - only guest OS threads are limited. With the <code>guest</code> option, the guest OS is guaranteed to have all the resources implied by the VM configuration. At the same time, the VM's hardware emulation threads spend additional resources of the host. For example, for a VM with two 2.8 GHz vCPUs, switching to <code>guest</code> means that VM's guest applications will have all the resources of two 2.8 GHz vCPUs at their disposal. Notes: 1. Some types of guest applications, like voice-over-IP software, significantly increase expenses on hardware emulation threads. 2. After changing this parameter, restart running virtual machines for the changes to take effect.

2.2.1.5. prlsrvctl shutdown

Shuts down the OpenVZ component responsible for managing virtual machines and containers. No operations on virtual machines and containers are possible.

Syntax

```
prlsrvctl shutdown [-f, --force]
```

Table 2.20. Options

Name	Description
-f, --force	Specifies whether the shutdown operation should be forced. If one or more virtual machines and containers are running, clients are connected, or some tasks are currently in progress, then forcing the shutdown will stop all processes automatically and will shut down the OpenVZ component.

2.2.1.6. prlsrvctl usb

The `prlsrvctl usb` command is used to permanently assign a USB device to a specific virtual machine. A permanently assigned USB device will be connected to the virtual machine automatically on server restart. This functionality works only with virtual machines (not containers).

Table 2.21. Subcommands

Name	Description
usb list	Lists USB devices connected to the server together with the information about their virtual machine assignments for the current user.
usb set	Permanently assigns a USB device to the specified virtual machine.
usb del	Removes a previously created USB device assignment.

usb list

Lists the USB devices connected to the physical server.

Syntax

```
prlsrvctl usb list
```

Returns a list of USB devices in tabular format with the following columns:

- **Name** - the USB device name.
- **ID** - a string that uniquely identifies the USB devices on the physical server. The ID never changes even if the device is disconnected from the server and then reconnected again. Please note that if a device ID is listed in quotes, they are a part of the ID and must be included in other calls that use it as an input parameter.
- **VM UUID** - a universally unique ID of the virtual machine to which this USB device is permanently assigned. If a USB device is not assigned to any virtual machine, this column will be empty.

usb set

Permanently assigns a USB device to the specified virtual machine. A permanently assigned USB device will be connected to the virtual machine automatically on server restart. The USB device assignment is performed for the current user only. Other users may create their own USB device assignments. This functionality works only with virtual machines (not containers).

Syntax

```
prlsrvctl usb set <usb\_dev\_ID> <VM\_name>
```

Table 2.22. Options

Name	Description
<usb_dev_ID>	The USB device ID. To obtain the list of USB devices connected to the server use the <code>usb list</code> command.
<VM_name>	The name of the virtual machine to which to assign the USB device.

usb del

Deletes a USB device assignment previously created with the `usb set` command. The USB device assignment is performed on the user level, so if you remove an assignment, it will only be removed for the current user. Other users may have their own USB devices assignments, which will not be affected.

Syntax

```
prlsrvctl usb del <usb_dev_ID>
```

Table 2.23. Options

Name	Description
<usb_dev_ID>	The USB device ID. To see the current USB device assignments for the current user use the <code>usb list</code> command.

2.2.1.7. prlsrvctl user list

Displays the list of OpenVZ users. Only those users are displayed who has created at least one virtual machine and container.

Syntax

```
prlsrvctl user list [-o, --output <name|mng_settings|def_vm_home>]
```

Table 2.24. Options

Name	Description
-o, --output <name mng_settings def_vm_home>	Fields to include in the output. The following fields are available: • <code>name</code> - User name. • <code>mng_settings</code> - Indicates whether the user is allowed to modify OpenVZ preferences. • <code>def_vm_home</code> - The user default virtual machine folder. The fields must be specified in lowercase.

2.2.1.8. prlsrvctl user set

Configures the profile of the user currently logged in to the OpenVZ server.

Syntax

```
prlsrvctl user set [--def-vm-home <path>]
```

Table 2.25. Options

Name	Description
--def-vm-home <path>	The default virtual machine and container directory name and path.

2.2.1.9. prlsrvctl privnet

The `prlsrvctl privnet` command is used to manage private networks on physical servers.

Table 2.26. Subcommands

Name	Description
add	Creates a new private network.
set	Configures the parameters of an existing private network.
del	Removes an existing private network.
list	Lists the available private networks.

add

Creates a new private network.

Syntax

```
prlsrvctl privnet add <private_network_ID> [-a, --ipadd <addr>[/<mask>]]
                        [--global <yes|no>]
```

Table 2.27. Options

Name	Description
<private_network_ID>	The private network ID.
-a, --ipadd <addr>[/<mask>]	Add a subnet to the private network. The network can have multiple subnets.
--global <yes no>	Make the private network weak, allowing access to and from external networks.

set

Configures an existing private network.

Syntax

```
prlsrvctl privnet set <private_network_ID> [-a, --ipadd <addr>[/<mask>]]
                        [-d, --ipdel <addr>[/<mask>]] [--global <yes|no>]
```

Table 2.28. Options

Name	Description
<private_network_ID>	The private network ID.

Name	Description
a, --ipadd <addr>[/<mask>]	Add a subnet to the private network. The network can have multiple subnets.
-d, --ipdel <addr>[/<mask>]	Delete the specified subnet from the private network.
--global <yes no>	Make the specified private network weak, allowing access to and from external networks.

del

Deletes a private network from the physical server.

Syntax

```
prlsrvctl privnet del <private_network_ID>
```

Table 2.29. Options

Name	Description
<private_network_ID>	The ID of the private network to delete.

list

Lists the private networks that exist on the physical server.

Syntax

```
prlsrvctl privnet list
```

2.3. OpenVZ Updates

OpenVZ allows quick and easy updates with the yum utility standard for RPM-compatible Linux operating systems. For more information on yum, See **Updating OpenVZ** in the *OpenVZ User's Guide* and the yum manual page.

Chapter 3. Managing Containers

OpenVZ containers can be managed using the `prlctl` command-line utility. The utility is installed on the hardware node during the product installation.

3.1. Matrix of OpenVZ Command-Line Utilities

The table below contains the full list of OpenVZ command-line utilities and command you can use for managing containers.

Table 3.1. General Utilities

Name	Description
<code>prlctl</code>	Utility to control containers.
<code>prlctl list</code>	Utility to view a list of containers existing on the server with additional information.

Table 3.2. Container Migration Utilities

Name	Description
<code>prlctl clone</code>	Command for the local cloning of containers.

Table 3.3. Template Management Utilities

Name	Description
<code>vzpkg</code>	Utility to manage OS and application EZ templates either inside your containers or on the server itself.

Table 3.4. Supplementary Utilities

Name	Description
<code>vzps</code> , <code>vztop</code>	Utilities working as the standard <code>ps</code> and <code>htop</code> utilities, with container-related functionality added.
<code>vzpid</code>	Utility that prints container UUID the process belongs to.
<code>vzsplit</code>	Utility to generate container configuration file sample, "splitting" the server into equal parts.
<code>pfcache</code>	Memory and IOPS deduplication management utility.
<code>pcompact</code>	Utility to compact containers by removing unused blocks from their virtual disks.

3.2. `prlctl`

`prlctl` is the primary tool for container management. To use it, you have to log in to the server as the root user.

Syntax

```
prlctl <subcommand> <CT_name>
prlctl --version
prlctl --help
```

Table 3.5. Subcommands

Name	Description
create	Creates a new container.
delete	Deletes a container.
mount	Mounts the container private area and executes the container mount script.
umount	Unmounts the container private area and executes the unmount script.
start	Starts a container.
stop	Stops a container.
restart	Restarts a container.
status	Displays the container status.
set	Sets container parameters: resource control settings, hostname, IP addresses, and so on.
enter	Logs in to a container without knowing its root password.
exec	Runs arbitrary commands in a container without logging in to it.
suspend	Saves the state of a running container in a dump file.
resume	Restores a container from its dump file.
snapshot snapshot-list snapshot-switch snapshot-delete	Creates and manages container snapshots.
console	Creates a command prompt channel to a container.

Table 3.6. Options

Name	Description
--version	Displays the prlctl package version currently installed on the server.
--help	Displays the usage information about prlctl.

3.2.1. prlctl console

Creates a command prompt channel to a container. Allows to log in to and execute commands in running containers; provides container startup/shutdown information that may be used for troubleshooting purposes. Logging in to containers requires a virtual terminal (e.g., mingetty) to be installed in the container.

Note: To exit the console, press **Esc** and then . (period).

Syntax

```
prlctl console <CT_name>
```

Table 3.7. Options

Name	Description
<CT_name>	Container name.

3.2.2. prlctl create

This command is used to create new containers.

Syntax

```
prlctl create <CT_name> --vmtype ct [<options>]
```

With this command, you can create regular containers. A unique container name is required for this command.

Table 3.8. Options

Name	Description
<CT_name>	An arbitrary name to assign to the new container.
--vmtype ct	Tells the prlctl create command to make a container. If the option is omitted, a virtual machine is created instead.
--ostemplate <name>	OS EZ template to use for creating the container. If omitted, this value is taken from the DEF_OSTEMPLATE parameter in the global OpenVZ configuration file.
--config <name>	Container sample configuration file to use for creating the container. Sample configuration files are located in /etc/vz/conf and have names in the format ve-<name>.conf-sample. The sample configuration files usually have a number of resource control limits for the container and some application templates to be added to the container immediately upon its creation. If you skip this option and the default configuration file name is not specified in the global OpenVZ configuration file, you will have to set resource control parameters for the container using the prlctl set command.
--uuid <uuid>	A custom UUID to assign to the container.

3.2.3. prlctl delete

Deletes a container from the server.

Syntax

```
prlctl delete <CT_name>
```

Table 3.9. Options

Name	Description
<CT_name>	Container name.

When executed, `prctl delete` physically removes all the files located in the container private area (specified as the `VE_PRIVATE` variable in the container configuration file) and renames the container configuration file in `/etc/vz/conf` from `<CT_name>.conf` to `<CT_name>.conf.destroyed`. It also renames container action scripts, if any, in a similar manner.

Note: A container must be stopped before its private area can be unmounted.

3.2.4. `prctl exec, enter`

Allow running arbitrary commands in a container.

Syntax

```
prctl exec <CT_name> [--without-shell] <command>
prctl enter <CT_name>
```

where `<command>` is a string to be executed in the container. If `<command>` is specified as `-`, then the commands for execution will be read from the standard input until the end of file or `exit` is encountered.

Table 3.10. Options

Name	Description
<CT_name>	Container name.
--without-shell	Run commands directly without bash or cmd shell.

When using `prctl exec`, remember that the shell parses the command-line and, if your command has shell metacharacters in it, you should escape or quote them.

The `prctl enter` command is similar to `prctl exec /bin/bash`. The difference between the two is that `prctl enter` makes the shell interpreter believe that it is connected to a terminal. As such, you receive a shell prompt and are able to execute multiple commands as if you were logged in to the container.

3.2.5. `prctl migrate`

Migrates a container from one server to another.

Syntax

```
prctl migrate [<source_server>/]<CT_name> <destination_server>[/<CT_name>]
               [--dst=<path>] [--keep-src|--remove-src] [--no-compression] [--ssh <options>]
```

Table 3.11. Options

Name	Description
<CT_name>	The source container name.
<source_server>	The source server information. Use the following format to specify this info: [<user>[:<password>]@]<server_IP_address_or_hostname>[:<port>].
<destination_server>	The destination server information. If omitted, the migration will be performed locally. Use the following format to specify this info: [<user>[:<password>]@]<server_IP_address_or_hostname>[:<port>].
--dst=<path>	Name and path of the directory on the destination server where the container files should be stored.
--keep-src	Do not remove the original container from the source server. Cannot be used together with --remove-src.
--remove-src	Remove the original container from the source server. Cannot be used together with --keep-src.
--no-compression	Disable data compression during migration.
--ssh	Additional options to pass to ssh to connect to the destination server. All standard ssh options are supported. Note: Do not specify the destination server hostname or IP address as an ssh option.

3.2.6. prlctl mount, umount

The `prlctl mount` command mounts the container private area to the container root directory (`/vz/root/<CT_name>` on the server) without starting it. Normally, you do not have to use this command as the `prlctl start` command mounts the container private area automatically.

The `prlctl umount` command unmounts the container private area. Usually, there is no need in using this command either because `prlctl stop` unmounts the container private area automatically.

Syntax

```
prlctl mount <CT_name>
prlctl umount <CT_name>
```

Table 3.12. Options

Name	Description
<CT_name>	Container name.

3.2.7. prlctl move

Moves container's private area to a new location on the same server. The container can be stopped, suspended or running.

Syntax

```
prlctl move <CT_name> --dst=<path>
```

Table 3.13. Options

Name	Description
<CT_name>	Container name.
--dst=<path>	Path to container's new private area.

3.2.8. prlctl problem-report

Generates a problem report for the specified container and either sends it to the OpenVZ technical support team or displays it on the screen.

Syntax

```
prlctl problem-report <CT_name> <-d, --dump|-s, --send [--proxy [<user> \
[:<passwd>]@<proxyhost>[:<port>]]] [--no-proxy]>
```

Table 3.14. Options

Name	Description
<CT_name>	The name of the container for which to generate the problem report.
-d, --dump	Collect technical data about the specified container and display it on the screen. You can also pipe the output to a file and then send it to the OpenVZ technical support team to analyze your problem.
-s, --send	Send the generated problem report to the OpenVZ technical support team.
--proxy [<user>[:<passwd>]@\<proxyhost>[:<port>]]	Use the specified information to send the generated report through a proxy server, if you use one to connect to the Internet.
--no-proxy	Do not use a proxy server to send the generated report. This is the default behavior, so you can omit this parameter.

3.2.9. prlctl register, unregister

The register command is used to register a container with OpenVZ.

The unregister command removes a container from the OpenVZ registry.

Syntax

```
prlctl register <path> [--preserve-uuid <yes|no>]
```

```
prlctl unregister <CT_name>
```

Table 3.15. Options

Name	Description
<path>	Full path to the container directory.
<CT_name>	The name of the container to remove from the OpenVZ registry.
--preserve-uuid <yes no>	Specifies what to do with the container UUID (universally unique identifier). If you specify yes, the UUID is preserved. If you specify no, the UUID is regenerated.
Note: By default, UUIDs are regenerated.	

- Use the register command when you have a container on the server that does not show up in the list of the containers registered with the OpenVZ. This can be a container that was previously removed from the registry or that was copied from another location.
- The unregister command removes a container from the OpenVZ registry, but does not delete the container files from the server. You can re-register the container later using the register command.

3.2.10. prlctl reinstall

Recreates a container from scratch according to its configuration file. Copies old private area content to the /vz/root/<CT_name>/old directory.

Syntax

```
prlctl reinstall <CT_name> [--skipbackup] [--resetpwdb] [--scripts <script> [...]]
prlctl reinstall <CT_name> [--listscripts] [--desc]
```

Table 3.16. Options

Name	Description
<CT_name>	Container name.
--resetpwdb	Removes container's user database and creates a clean database as for any new installation.
--skipbackup	Does not save the old private area contents to the /old directory.
--scripts <script>[...]	Specifies the scripts to be executed during reinstallation. These scripts are used to customize application templates in the new container and bring them to the same state as in the old container. By default, all available scripts are executed.
--listscripts	Lists the scripts to be executed during container reinstallation.
--desc	Displays the description of the scripts to be executed during container reinstallation. Used together with the --listscripts option.

Note: Currently, the reinstall command may not be supported by the prlctl utility. Use vzctl instead.

3.2.11. prlctl set

This command is used for setting container parameters.

Syntax

```
prlctl set <CT_name> <option> <value>
```

where <CT_name> is container name.

The command options specified in this file can be subdivided into the following categories:

- miscellaneous
- networking
- resource management
- hard disk drive management

3.2.11.1. General Options

The table below lists the general options you can use with `prlctl set`.

Name	Description
--onboot <yes no>	Set to yes to have OpenVZ automatically start this container on next system startup.
--offline_management <yes no>	Enabling/disabling the direct managing of the container through a common Internet browser by means of Power Panels and the Plesk control panel (as defined by the OFFLINE_SERVICE parameter in the global or container configuration file).
--offline_service <service_name>	Defines whether the container can be managed by means of Power Panel or Plesk or both. Valid only if the OFFLINE_MANAGEMENT parameter is set to yes. The names of the available services can be taken from the file names (excluding the .conf extension) in the /etc/vzredirect.d directory on the server.
--userpasswd <user>:<password>	This setting creates a new user with the specified password in the container, or changes the password of an already existing user. This command modifies not the container configuration file, but the /etc/passwd and /etc/shadow files inside the container. In case the container root is not mounted, it is automatically mounted to apply the changes and then unmounted.
--crypted	Used with --userpasswd. Indicates that the specified password is already a hash.
--features {<name>:on off}	Enables/disables the support for the following functionality inside the container: • nfs: mounting NFS shares • ipip: creating IPIP tunnels • sit: using the Simple Internet Transition (SIT) mechanisms • ipgre: creating IP-GRE tunnels • bridge: using bridges to connect virtual Ethernet devices

Name	Description
	• <code>nfsd</code>: running an NFS-kernel-space server
<code>--name <new_name></code>	Changes the container name. You can change the names of both stopped and running containers.
<code>--description <desc></code>	Custom container description. Descriptions must be alphanumeric. Descriptions with white spaces must be enclosed in quotation marks.
<code>--vnc-mode <auto manual off></code>	Enables or disables access to the container via the VNC protocol.
<code>--vnc-port <port></code>	Sets the VNC port number for the container. Used with <code>--vnc-mode manual</code> .
<code>--vnc-passwd <passwd> --vnc-nopasswd</code>	Sets the VNC password for the container or specifies that no password is needed for VNC connections. Either of these options is mandatory for any VNC setup.
<code>--autocompact <on off></code>	Enables or disables compaction for all disks in the container. Note: For details on how to enable or disable compaction for a specific disk in the container, see Section 3.2.11.4, “Hard Disk Drive Management Options” on page 46.

3.2.11.2. Resource Management Options

Resource management options control the amount of resources a container may consume. If the setting has `bar:lim` after it than this setting requires specifying both barrier and limit values separated by colons.

Name	Description
<code>--applyconfig <name></code>	This option lets you set the resource parameters for the container not one by one, but by reading them from the container sample configuration file. All container sample configuration files are located in the <code>/etc/vz/conf</code> directory and are named according to the following pattern: <code>ve-<name>.conf-sample</code> , so you should specify only the <code><name></code> part of the corresponding sample name after the <code>--applyconfig</code> option. Note that the names of sample configuration files cannot contain spaces. The <code>--applyconfig</code> option applies all the parameters from the specified sample file to the given container, except for the <code>OSTEMPLATE</code> , <code>TEMPLATES</code> , <code>VE_ROOT</code> , <code>VE_PRIVATE</code> , <code>HOSTNAME</code> , <code>IP_ADDRESS</code> , <code>TEMPLATE</code> , <code>NETIF</code> parameters (if they exist in the configuration sample file).
<code>--cpuunits <units></code>	CPU weight. This is a positive integer number that defines how much CPU time the container can get as compared to the other virtual machines and containers running on the server. The larger the number, the more CPU time the container can receive. Possible values range from 8 to 500000. If this parameter is not set, the default value of 1000 is used.
<code>--cpulimit {<percent> <megahertz>}</code>	CPU limit, in per cent or megahertz (MHz), the container is not allowed to exceed. This parameter is not set for newly created containers; so they can consume all free CPU power of the server. By default, the limit is set in percent. To set the limit in MHz, specify <code>m</code> after the value. When setting this

Name	Description
	parameter in per cent, keep in mind that one CPU core makes up 100%. So if the server has 4 CPU cores, the total CPU power will equal 400%.
<code>--cpus <num></code>	Number of CPU cores defining the CPU limit for a container. The limit is calculated by multiplying the power of one CPU core by the number of the specified CPU cores. This option also defines the number of CPUs shown to container users. This parameter is not set for newly created containers; so they can consume all free CPU power of the server.
<code>--cpumask <num></code>	CPU affinity mask. This mask defines the CPUs on the server that can be used to handle the processes running in the container. The CPU mask can be specified as both separate CPU index numbers (1,2,3) and CPU ranges (2-4,5-7).
<code>--nodemask {<num> all}</code>	The NUMA node mask defining a NUMA node to bind the container to. Once you set the mask, the processes running in the container will be executed only on the CPUs that belong to the specified NUMA node.
<code>--diskspace <amount></code>	Total disk space consumed by the container, in megabytes. You can use the following suffixes to specify measurement units: • G for gigabytes • M for megabytes • K for kilobytes • B for bytes
<code>--quotauidlimit {<0 <N>}</code>	Enables (if set to a value other than 0) or disables (if set to 0) per-user/group quotas for further management with the standard Linux quota utility. Keep in mind the following: • Enabling per-user and per-group quotas for a container requires restarting the container. • If you delete a registered user but some files with their ID continue residing inside your container, the current number of UGIDs (user and group identities) inside the container will not decrease. • If you copy an archive containing files with user and group IDs not registered inside your container, the number of UGIDs inside the container will increase by the number of these new IDs.
<code>--ioprio <num></code>	The container priority for disk I/O operations. The allowed range of values is 0-7. The greater the priority, the more time the container has for writing to and reading from the disk. The default container priority is 4.
<code>--iolimit <num></code>	The bandwidth a container is allowed to use for its disk input and output (I/O) operations. By default, the limit is set in megabytes per second. You can use the following suffixes to specify measurement units: • G: sets the limit in gigabytes per second • M: sets the limit in megabytes per second • K: sets the limit in kilobytes per second • B: sets the limit in bytes per second In the current version of OpenVZ, the maximum I/O bandwidth limit you can set for a container is 2 GB per second. The default I/O bandwidth limit for all

Name	Description
	newly created containers is set to 0, which means that no limits are applied to any containers.
<code>--iopslimit <num></code>	The maximum number of disk input and output operations per second a container is allowed to perform. By default, any newly created container does not have the IOPS limit set and can perform so many disk I/O operations per second as necessary.
<code>--rate <class>:_<Kbits>_</code>	If traffic shaping is turned on, then this parameter specifies bandwidth guarantee for the container. The format is <code><class>:_<Kbits>_</code> where <code><class></code> is the network class (group of IP addresses) and <code><Kbits></code> is the traffic bandwidth.
<code>--ratebound <yes no></code>	If set to "yes", the bandwidth guarantee is also the limit for the container and the container cannot borrow the bandwidth from the TOTALRATE bandwidth pool.
<code>--memsize <size></code>	The amount of RAM that can be used by the processes of a container, in megabytes. You can use the following suffixes to specify measurement units: • G for gigabytes • M for megabytes • K for kilobytes • B for bytes
<code>--memguarantee <size></code>	Sets a percentage of container's RAM that said container is guaranteed to have. By default, set to 0%.
<code>--swappages <pages></code>	The amount of swap space that can be used by the container for swapping out memory once the RAM is exceeded, in 4KB pages. You can use the following suffixes to specify measurement units: • G for gigabytes • M for megabytes • K for kilobytes • B for bytes
<code>--swap <size></code>	The amount of swap space that can be used by the container for swapping out memory once the RAM is exceeded, in bytes. You can use the following suffixes to specify measurement units: • G for gigabytes • M for megabytes • K for kilobytes • B for bytes

3.2.11.3. Network Options

Network-related options allow you to set the hostname, the domain to search when a not fully qualified domain name is used, the DNS server address and the IP addresses (both IPv4 and IPv6) that container can use, and other parameters.

Name	Description
<code>--hostname <name></code>	Sets the hostname to the specified name.
<code>--ipadd <addr></code>	Adds an IP address to a list of IP addresses the container can use and brings up the network interface with this address inside the container. If used with the <code>--ifname</code> option, adds an IP address to the specified container virtual network adapter.
<code>--ipadd <addr>/<net_mask></code>	Assigns the IP address and network mask to the container. Note: You can assign network masks to containers operating in the <code>venet0</code> networking mode only if the <code>USE_VENET_MASK</code> parameter in the OpenVZ containers configuration file is set to yes.
<code>--ipdel {<addr> all}</code>	Allows you to revoke IP address from the container. If "all" is used instead of IP address than all IP addresses will be revoked. If used with the <code>--ifname</code> option, deletes an IP address from the specified container virtual network adapter.
<code>--nameserver <addr></code>	The DNS server IP address for the container. If used with the <code>--ifname</code> option, sets the DNS server for the specified container virtual network adapter.
<code>--searchdomain <domain></code>	The DNS search domain for the container. More than one domain may be specified.
<code>--netfilter <disabled stateless stateful full></code>	Indicates which iptables modules are allowed for the container. If some of the allowed modules are not loaded on the destination Hardware Node after migration or restoration from backup, they will be automatically loaded on the migrated or restored container start. The following modes are available: • disabled: none. • stateless: (default) all modules except conntrack and NAT-related. • stateful: all modules except NAT-related. • full: all modules.
<code>--netif_add <name>[, <MAC>, <host_MAC>]</code>	Creates a new veth virtual network adapter and assigns the name of <code><name></code> to the Ethernet interface inside the container. Along with the Ethernet interface name inside the container, you can set the following parameters when creating the veth adapter: • <code><MAC></code>: the MAC address to be assigned to the veth Ethernet interface inside the container. • <code><host_MAC></code>: the MAC address to be assigned to the veth Ethernet interface on the server.

Name	Description
	Only the Ethernet interface name (<i><name></i>) is mandatory; all the other parameters, if not specified, are automatically generated by OpenVZ during the veth adapter creation.
<code>--netif_del <name></code>	Removes the veth virtual network adapter with the specified name from the container.
<code>--ifname <name></code>	Specifies the name of the veth virtual network adapter whose settings are to be configured. This option can be used along with one of the following options: <code>--ipadd</code> , <code>--ipdel</code> , <code>--nameserver</code> , <code>--gw</code> , <code>--network</code> , <code>--dhcp</code> , <code>--mac</code> , <code>--host_mac</code> .
<code>--mac <addr></code>	The MAC address to be assigned to the veth virtual Ethernet interface inside the container. Should be used along with the <code>--ifname</code> option.
<code>--host_mac <addr></code>	The MAC address to be assigned to the veth virtual Ethernet interface on the server. Should be used along with the <code>--ifname</code> option.
<code>--host_ifname <name></code>	The name to be assigned to the veth virtual Ethernet interface on the server. Should be used along with the <code>--ifname</code> option.
<code>--network <network_ID></code>	Connects the veth virtual network adapter to the bridge associated with the specified network ID. Should be used along with the <code>--ifname</code> option. You can also use this option to disconnect the veth virtual network adapter from the bridge. To this effect, you should specify "" after the option.
<code>--dhcp <yes no></code>	Specifies whether the virtual network adapter should obtain the IPv4 settings through a DHCP server.
<code>--dhcp6 <yes no></code>	Specifies whether the virtual network adapter should obtain the IPv6 settings through a DHCP server.
<code>--gw <gw></code>	The default gateway to be used by the container.
<code>--gw6 <gw></code>	The default IPv6 gateway to be used by the container.
<code>--apply-iponly <yes no></code>	If set to yes, the hostname, nameserver, and search domain settings from the container configuration file are ignored.
<code>--configure <yes no></code>	If set to yes, the settings above are applied to the virtual network adapter instead of its original settings. Configuring any of the settings above automatically sets this option to yes.

3.2.11.4. Hard Disk Drive Management Options

This group of options is used to manage virtual hard disks in a container.

Syntax

```
prlctl set <CT_name> --device-add hdd [--image <file>] [--size <size>]
      [--mnt <path>] [--iface <ide|scsi|virtio>] [--position <pos>]
```

```
prlctl set <CT_name> --device-set hdd<N> [--image <file>] [--size +<size>]
      [--mnt <path>] [--iface <ide|scsi|virtio>] [--position <pos>]
      [--autocompact <on|off>]
prlctl set <CT_name> --backup-add <backup_ID> [--disk <disk_name>]
prlctl set <CT_name> --device-del hdd<N> [--detach-only|--destroy-image]
prlctl set <CT_name> --backup-del {<backup_ID>|all}
```

Table 3.17. Options

Name	Description
<CT_name>	Container name.
--device-add hdd	Adds a virtual hard disk to the container. If no other options are specified, the command creates a new unmounted disk with the following parameters: - name: hdd<N> where <N> is the next available disk index. - size: 65536 MB - image location: /vz/private/<CT_UUID>/disk-<ID>.hdd.
--device-set hdd<N>	Modifies the parameters of the virtual hard disk hdd<N>. Note: For the list of disks, use the <code>prlctl list -i</code> command.
--image <file>	Specifies an image file that will be used to emulate the virtual disk. - If the specified image does not exist, it is created and used to emulate the virtual hard disk. - If the specified image exists, it is used to emulate the virtual hard disk.
--size <size>	Specifies the size of the virtual hard disk, in megabytes.
--mnt <path>	Specifies the mount point of the virtual hard disk inside the container. A corresponding entry is also added to container's <code>/etc/fstab</code> file, so the disk is mounted automatically on container start.
--autocompact <on off>	Enables or disables compaction for the specified disk in the container. Note: For details on how to enable or disable compaction for all disks in the container, see Section 3.2.11.1, “General Options” on page 41.
--backup-add <backup_ID>	Attach the backup with the identifier <backup_ID> to the virtual machine as a virtual hard disk. To obtain the backup ID, use the <code>prlctl backup-list -f</code> command.
--disk <disk_name>	Used with <code>--backup-add</code> . The name of the disk in the backup to attach. If a disk is not specified, all disks contained in the

Name	Description
	backup will be attached. To obtain the disk name, use the <code>prlctl backup-list -f</code> command.
<code>--device-del hdd<N></code>	Deletes a virtual hard disk from the stopped container.
<code>--detach-only</code>	Removes the virtual disk from the container configuration but leaves its image file intact.
<code>--destroy-image</code>	Removes the virtual disk from the container configuration and deletes its image file.
<code>--backup-del {<backup_ID> all}</code>	Detach either the backup with the identifier <code><backup_ID></code> or detach all backups from the virtual machine.

3.2.12. prlctl snapshot, snapshot-list, snapshot-switch, snapshot-delete

Takes, displays, reverts to, and deletes container snapshots.

Syntax

```
prlctl snapshot <CT_name> [-n, --name <name>] [-d, --description <desc>]
prlctl snapshot-list <CT_name> [-t, --tree] [-i, --id <snapshot_ID>]
prlctl snapshot-switch <CT_name> -i, --id <snapshot_ID>
prlctl snapshot-delete <CT_name> -i, --id <snapshot_ID>
```

Table 3.18. Options

Name	Description
<code><CT_name></code>	Container name.
<code>-n, --name <name></code>	User-defined snapshot name. Names with white spaces must be enclosed in quotation marks.
<code>-d, --description <desc></code>	User-defined snapshot description. Descriptions with white spaces must be enclosed in quotation marks.
<code>-t, --tree</code>	Displays the snapshot list as a tree. The default display format is tabular with Parent Snapshot ID and Snapshot ID as columns.
<code>-i, --id <snapshot_ID></code>	- Use with <code>prlctl snapshot-list</code> to specify the ID of the snapshot to use as the root. If this parameter is omitted, the entire snapshot tree will be displayed. - Use with <code>prlctl snapshot-switch</code> to specify the ID of the snapshot to revert to. - Use with <code>prlctl snapshot-delete</code> to specify the ID of the snapshot to delete.

Note: If the snapshot you want to delete has children snapshots derived from it, they will not be deleted.

3.2.13. prlctl start, stop, restart, status

These commands start, stop, restart, and show the current state of containers, respectively.

Syntax

```
prlctl start <CT_name> [--wait]
prlctl stop <CT_name> [--fast]
prlctl restart <CT_name>
prlctl status <CT_name>
```

Table 3.19. Options

Name	Description
<CT_name>	Container name.

The first command is used to start a container. It will set up all network interfaces inside the container, initialize the container quota, if needed, start the `init` process inside the container, and exit. You can also make the `prlctl start` command wait for all the necessary startup processes to complete and the container to boot into the default runlevel by passing the `--wait` option to this command.

`prlctl stop` shuts the container down. If the container is not down after a two-minute timeout due to an error in an application, for example, `prlctl` will forcibly kill all the processes inside the container. To avoid waiting for two minutes in case of a corrupted container, you may use the `--fast` option with this command.

When starting or stopping a container, `prlctl` executes a number of helper scripts located in the `/vz/private/<CT_UUID>/scripts` and `/usr/libexec/libvzctl/scripts/` directories.

Table 3.20. Scripts:

Name	Description
<code>vz-create_prvt</code>	This script is expected to create the container private area from a private area template.
<code>vz-net_add</code>	This script sets up the necessary routing entries for container IP addresses and adds public ARP records on all interfaces.
<code>vz-net_del</code>	This script deletes routing entries and ARP records for container IP addresses from all interfaces.
<code>vz-pci_configure</code>	This configuration script is executed after a PCI device is added to or removed from a container.
<code>vz-setrate</code>	This script configures the network traffic shaping for a container.
<code>vz-start</code>	This script is called just before a container is started and used to perform any additional setup of the container, such as network setup.
<code>vz-stop</code>	This script is called just after a container is shut down and can be used to perform any additional cleanup of the container, such as network cleanup.

Use action scripts `vz-start/vz-stop` to perform actions during container startup/shutdown. To add more functionality to other shipped scripts, create a new script and add a call to it to a shipped script. Doing so will facilitate upgrades to future versions of OpenVZ.

The `prlctl restart <CT_name>` command consecutively performs the stopping and starting of the corresponding container.

The `prlctl status` command shows the current container state. It outputs the following information: whether the container private area exists, whether it is mounted, and whether the container is running.

3.2.14. prlctl suspend, resume

The `prlctl suspend` command is used to save the state of a running container.

Syntax

```
prlctl suspend <CT_name>
```

Table 3.21. Options

Name	Description
<CT_name>	Container name.

During the `prlctl suspend` execution, the current container state is saved to a special dump file and the container itself is stopped. The created dump file is saved to the Dump file in the `/vz/private/<CT_UUID>/dump` directory on the server.

The `prlctl resume` command is used to restore the container from its dump file created with the `prlctl suspend` command.

Syntax

```
prlctl resume <CT_name>
```

When executed, `prlctl resume` searches for the Dump file in the `/vz/private/<CT_UUID>/dump` directory on the server and restores the container from this file. You can restore the container dump file on the Source Server, i.e. on the server where this container was running before its dumping, or transfer the dump file to another server and restore it there.

Note: Before restoring a container from its dump file, make sure that the file system on the Destination Server is identical to that at the moment of the container dumping. Otherwise, the container restoration may fail.

3.2.15. prlctl list

Displays a list of containers on the Hardware Node. Displays information on containers on the Hardware Node.

Syntax

```
prlctl list --vmtypes ct [-a, --all] [-o, --output <field>[,...]]
                        [-s, --sort <field>|-<field>] [-t, --template] [-j, --json]
prlctl list -i, --info --vmtypes ct [<CT_name>] [-f, --full] [-t, --template]
                                   [-j, --json]
```

Table 3.22. Options

Name	Description
-a, --all	List all running, stopped, suspended, and paused containers. If this and the rest of the parameters are omitted, only the running containers will be displayed.
-t, --template	List available container templates instead of actual containers.
-o, --output <field>[,...]	Display only the specified fields. Type field names in lower case. Separate multiple fields with commas. For the list of fields, see Section 3.2.15.1, “prctl list Output Parameters” on page 51.
-s, --sort {<field> -<field>}	Sort containers by the specified field in either ascending or descending order.
-i, --info	Display detailed information about the specified container.
-f, --full	Display detailed information about network cards in containers. Used with the --info option.
<CT_name>	The name of the container for which to display the detailed information. If not specified, the information will be displayed for all registered containers.
-j, --json	Produce machine-readable output in the JSON format.

3.2.15.1. prctl list Output Parameters

Listed below are the parameters that can be specified after the -o switch.

Name	Output Column	Description
uuid	UUID	Container UUID.
hostname	HOSTNAME	Container hostname.
name	NAME	Container name.
description	DESCRIPTION	Container description.
ostemplate	OSTEMPLATE	Specifies the name of the OS template your container is based on (e.g., centos-6-x86_64).
ip	IP_ADDR	Container IP address.
status	STATUS	Container status (e.g., running or stopped).
numproc	NPROC	The number of processes and threads allowed.
mac	MAC	Network device's MAC address.
netif	NETIF	Network devices in the container.
iolimit	IOLIMIT	The bandwidth a container is allowed to use for its disk input and output (I/O) operation, in bytes per second.
ha_enable	HA_ENABLE	Indicates whether the container is joined to the High Availability Cluster.
ha_prio	HA_PRIO	Container priority in the High Availability Cluster (0 is the lowest). Higher-priority virtual environments are restarted first in case of failures.

3.3. Migration Utilities

This section describes the utilities you can use to move/clone containers within the same server.

3.3.1. prctl clone

Creates an exact copy of the specified container.

Syntax

```
prctl clone <CT_name> --name <new_name> [--template] [--dst=<path>]
```

Table 3.23. Options

Name	Description
<CT_name>	Name of the container to clone.
--name <new_name>	Name to be assigned to the new container.
--template	Create a container template instead of a clone. Template cannot be started.
--dst=<path>	Full path to the directory for storing the contents of the cloned container. If this parameter is omitted, the clone is created in the default directory.

3.4. EZ Template Management Utilities

This section described the utilities you can use to manage OS and application templates.

3.4.1. vzpkg

The vzpkg utility is used to manage OS and application EZ templates either inside your containers or on the server itself. This tool can also be used to manage standard software packages (e.g., `mysql.rpm`) inside containers.

Syntax

```
vzpkg <command> [<options>] {<CT_name>|<object>}
vzpkg --help
```

Table 3.24. Commands

Name	Description
install template	Installs OS and application EZ templates on the server.
update template	Updates OS and application EZ templates installed on the server.
remove template	Removes OS and application EZ templates from the server.
list	Outputs a list of EZ templates, OS template caches with preinstalled application templates, or software packages either on the server or inside a particular container.

Name	Description
info	Outputs information on any EZ templates or software packages available on the server or inside the container.
status	Outputs information on updates for the packages installed inside a container.
install	Adds application EZ templates to or to install software packages inside the container.
update	Updates application EZ templates and software packages inside the container.
remove	Removes application EZ templates or software packages from the container.
create cache	Creates a tarball (cache) for the given OS EZ template.
update cache	Updates the existing tarball (cache) for the given OS EZ template.
remove cache	Removes a tarball (cache) for the given OS EZ template.
create appcache	Creates a cache of an OS EZ template with preinstalled application templates.
update appcache	Updates or recreates a cache of an OS EZ template with preinstalled application templates.
remove appcache	Removes a cache of an OS EZ template with preinstalled application templates.
localinstall	Installs a software package inside a container from the corresponding file on the server.
localupdate	Updates the software packages installed inside your container(s) by means of the <code>vzpkg install</code> or <code>vzpkg localinstall</code> commands.
upgrade	Upgrades an OS EZ template the container is based on to a newer version.
fetch	Downloads packages included in EZ templates to the server and to store them in the <code>vzpkg</code> local cache.
clean	Removes all locally cached data from the template directories on the server.
update metadata	Updates the local metadata on the server.

3.4.2. vzpkg install template

This command is used to install an OS or application EZ template on the server from an RPM package or OpenVZ repositories.

Syntax

```
vzpkg install template [<options>] <object> [...]
```

where *<object>* is a path to an RPM package or an EZ template name.

Table 3.25. Options

Name	Description
-q, --quiet	Disables logging to the screen and to the log file.

Name	Description
-f, --force	Forces installation of the EZ template on the server.

Note: To install multiple EZ templates, specify multiple RPM package or EZ template names separated by white spaces.

3.4.3. vzpkg update template

This command is used to update an OS or application EZ template on the server from an RPM package or OpenVZ repositories.

Syntax

```
vzpkg update template [<options>] <object> [...]
```

where *<object>* is a path to an RPM package or an EZ template name.

Table 3.26. Options

Name	Description
-q, --quiet	Disables logging to the screen and to the log file.
-f, --force	Forces update of the EZ template.

Note: To update multiple EZ templates, specify multiple RPM package or EZ template names separated by white spaces.

3.4.4. vzpkg remove template

This command removes an OS or application EZ template from the server.

Syntax

```
vzpkg remove template [<options>] <template_name> [...]
```

Table 3.27. Options

Name	Description
-F, --for-os <OS_template>	Specifies the OS EZ template to delete the application EZ template from.
-q, --quiet	Disables logging to screen and file.
-f, --force	Forces deletion of the EZ template.

When executed, the `vzpkg remove template` command removes the specified OS EZ template from the server. To delete an application EZ template, additionally specify the name of the OS EZ template (*<OS_template>*) under which this application template is to be run.

3.4.5. vzpkg list

The `vzpkg list` command is used to list

- EZ templates installed on the server, in a container, or available in remote EZ template repositories
- YUM software groups or individual packages installed in a container

Syntax

```
vzpkg list [<options>] [<OS_template>|<CT_name> [...]]
```

If you indicate a container name, the command will list all EZ templates applied to the specified container. If you indicate an OS EZ template, `vzpkg list` will display a list of application EZ templates available for this OS EZ template. Without any options, the utility lists all EZ templates installed on the server.

Table 3.28. Options

Name	Description
<code>-p, --package</code>	Lists the software packages installed in the container or included in the OS EZ template.
<code>-g, --groups</code>	Lists the YUM software groups installed in the container or available for the OS EZ template. The <code>-g</code> option works only for containers running RPM-based Linux distributions.
<code>-O, --os</code>	Displays the OS EZ template the container is based on.
<code>-A, --app</code>	Displays the application EZ templates installed in the container or included in the OS EZ template.
<code>-C, --cache</code>	Lists the packages included in the specified EZ template or applied to the specified container from the local <code>vzpkg</code> cache. You can omit this parameter if the elapsed time from the last <code>vzpkg</code> cache update does not exceed the value of the <code>METADATA_EXPIRE</code> parameter specified in the <code>/etc/vztt/vztt.conf</code> file. Should be used along with the <code>-p</code> option.
<code>-r, --remote</code>	If the elapsed time from the last <code>vzpkg</code> cache update does not exceed the value of the <code>METADATA_EXPIRE</code> parameter specified in the <code>/etc/vztt/vztt.conf</code> file, you should use this option to make <code>vzpkg list</code> list the packages included in the specified EZ template or applied to the specified container in the remote repositories. Should be used along with the <code>-p</code> option.
<code>-u, --custom-pkg</code>	Displays a list of packages that are applied to the specified container but absent from the repository set to handle the EZ template(s) where these packages are included.
<code>-i, --pkgid</code>	Displays the ID assigned to the EZ template instead of its name; these IDs are unique within the given system. If the <code><CT_name></code> argument is given, the command shows the IDs of the EZ templates available inside the container. If the <code><OS_template></code> argument is given, the command displays the IDs of the OS EZ template specified and all its EZ application templates.
<code>-S, --with-summary</code>	In addition to listing the EZ templates available either in the container (if the <code><CT_name></code> argument is given) or installed on the server (if the <code><CT_name></code>

Name	Description
	argument is omitted), this option makes <code>vzpkg list</code> display the summary information on the corresponding EZ templates/packages.
<code>-c, --cached</code>	This option has no effect if the <code><CT_name></code> argument is given. If used for listing the EZ templates available on the server, it makes <code>vzpkg list</code> omit all application and OS EZ templates for which the cache has not been created (by running the <code>vzpkg create cache</code> command). In other words, with this option on, <code>vzpkg list</code> will list only the OS EZ templates ready to be used for the container creation.
<code>appcache</code>	Outputs a list of OS EZ template caches with preinstalled applications.
<code>-d, --debug <num></code>	Sets the debugging level to one of the specified values (from 0 to 10). 10 is the highest debug level and 0 sets the debug level to its minimal value.
<code>-q, --quiet</code>	Disables logging to the screen and to the log file.

3.4.6. vzpkg info

This command displays information about EZ templates, YUM software packages, and individual software packages.

Syntax

```
vzpkg info [-F {<OS_template>|<CT_name>} -q|-d <app_template> [<parameters> ...]
vzpkg info -p|-g [-C|-r] [-F {<OS_template>|<CT_name>} -q|-d {<package_name>|<yum_package_group>}]
```

Table 3.29. Options

Name	Description
<code><CT_name></code>	Container name.
<code><OS_template></code>	OS EZ template.
<code><app_template></code>	Application EZ template.
<code><package_name></code>	Software package name.
<code><yum_package_group></code>	YUM software group name.
<code>-F, --for-os {<OS_template> <CT_name>_}</code>	Displays information on the application EZ template or the software package (if the <code>-p</code> option is specified) included in the specified OS EZ template or applied to the indicated container.
<code>-p, --package</code>	Displays information about the specified software package. Must be used with the <code>-F</code> option.
<code>-g, --groups</code>	Displays information about the packages included in the specified YUM software group.
<code>-C, --cache</code>	Displays the information on the specified package from the local <code>vzpkg</code> cache. You can omit this parameter if the elapsed time from the last <code>vzpkg</code> cache update does not exceed the value of the <code>METADATA_EXPIRE</code> parameter specified in the <code>/etc/vztt/vztt.conf</code> file.

Name	Description
-r, --remote	If the elapsed time from the last vzpkg cache update does not exceed the value of the METADATA_EXPIRE parameter specified in the /etc/vztt/vztt.conf file, you should use this option to make vzpkg info get the information on the specified package from the remote repositories set for handling the EZ template where this package is included.
-d, --debug <num>	Sets the debugging level to one of the specified values (from 0 to 10). 10 is the highest debug level and 0 sets the debug level to its minimal value.
-q, --quiet	Disables logging to the screen and to the log file.

While executed, vzpkg info parses the subdirectories and files located in the /vz/template/<os_name>/<os_version>/<arch>/config directory and containing the EZ template meta data. To run the command, you should specify either the OS EZ template or the container name. In either case, detailed information on the corresponding OS EZ template is displayed. You can also use the -F option to get the necessary information on any application EZ template included into the OS EZ template or applied to the container.

By default, vzpkg info displays all meta data on the EZ template/package specified. However, you can reduce the amount of the output information by using special parameters (<parameters>) listed in the table below.

Table 3.30. Output Parameters

Name	Description
name	The name of the EZ template/package.
packages	The packages included in the EZ template. For EZ templates only.
repositories	The repository where the packages comprising the EZ template are stored. For EZ templates only.
mirrorlist	The URL to the file containing a list of repositories from where the packages comprising the EZ template are to be downloaded. For EZ templates only.
distribution	The Linux distribution on the basis the OS EZ template has been created or under which the application EZ template is to be run. For EZ templates only.
summary	Brief information on the EZ template/package.
description	Detailed information on the EZ template/package. As distinct from summary, it can contain additional data on the EZ template/package.
technologies	For EZ templates only. Displays the following information: • The microprocessor architecture where the EZ template is to be used (x86, x86_64); • Specifies whether the EZ template can be used only on the servers with the Native POSIX Thread Library (NPTL) support. In this case the nptl entry is displayed after the vzpkg info execution.
version	The version of the software package.
release	The release of the software package.

Name	Description
arch	The system architecture where the EZ template/package is to be used. It can be one of the following: • x86 if the EZ template/package is to be used on 32-bit platforms. • x86_64 if the EZ template is to be used on 64-bit platforms (e.g., on servers with the AMD Opteron and Intel Pentium D processors installed).
config_path	Displays the path to the EZ template configuration directory containing the template meta data where the meta data for the base OS EZ template are stored (the default directory path is <code>/vz/template/<OS_name>/<OS_version>/<arch>/config/os/default</code>).
package_manager_type	For EZ templates only. The packaging system used to handle the packages included in the specified EZ template. It can be one of the following: • rpm for RPM-based Linux distributions (Fedora Core, Red Hat Enterprise Linux, etc.); • dpkg for Debian-based Linux distributions (e.g., Debian and Ubuntu).
package_manager	The package manager type for managing the packages included in the specified EZ template. It can be one of the following: x86 Linux distributions • rpm49db5x86: Fedora 17 • rpm49x86: Fedora 15 and 16 • rpm47x86: Red Hat Enterprise Linux 6 and CentOS 6 • rpm44x86: Red Hat Enterprise Linux 5 and CentOS 5 • rpm43x86: Red Hat Enterprise Linux 3 and 4, CentOS 3 and 4 • rpmzypp44x86: SUSE Linux Enterprise Server 11 with Service Pack 2 • rpm41x86: SUSE Linux Enterprise Server 10 and SUSE Linux 10.x • rpm41s9x86: SUSE Linux Enterprise Server 9 • rpmzypp49x86: openSUSE 12.1 • dpkg: Debian and Ubuntu x86-64 Linux distributions • rpm49db5x64: Fedora 17 • rpm49x64: Fedora 15 and 16 • rpm47x64: Red Hat Enterprise Linux 6 and CentOS 6 • rpm44x64: Red Hat Enterprise Linux 5 and CentOS 5 • rpm43x64: Red Hat Enterprise Linux 3 and 4, CentOS 3 and 4 • rpmzypp44x64: SUSE Linux Enterprise Server 11 with Service Pack 2 • rpm41x64: SUSE Linux Enterprise Server 10 and SUSE Linux 10.x • rpm41s9x64: SUSE Linux Enterprise Server 9 • rpmzypp49x64: openSUSE 12.1 • dpkgx64: Debian and Ubuntu

3.4.7. vzpkg status

This command is used to check the status of the packages either installed inside a container or included in an OS EZ template.

Syntax

```
vzpkg status [<options>] {<CT_name>|<OS_template>}
```

Table 3.31. Options

Name	Description
-C, --cache	Makes the <code>vzpkg status</code> command look for available updates in the local <code>vzpkg</code> cache only. You can omit this parameter if the elapsed time from the last <code>vzpkg</code> cache update does not exceed the value of the <code>METADATA_EXPIRE</code> parameter specified in the <code>/etc/vztt/vztt.conf</code> file.
-r, --remote	If the elapsed time from the last <code>vzpkg</code> cache update does not exceed the value of the <code>METADATA_EXPIRE</code> parameter specified in the <code>/etc/vztt/vztt.conf</code> file, you should use this option to make <code>vzpkg status</code> look for the package updates in the remote repositories set for handling the corresponding EZ template.
-d, --debug <num>	Sets the debugging level to one of the specified values (from 0 to 10). 10 is the highest debug level and 0 sets the debug level to its minimal value.
-q, --quiet	Disables logging to the screen and to the log file.

When executed, the command performs the following operations:

- Checks all the packages installed inside the specified container or included in the specified OS EZ template.
- Checks the repository used to install/update packages inside the container/OS EZ template.
- Compares the packages in the repository with those inside the container/OS EZ template.
- Lists the found packages updates for the container/OS EZ template, if any, or informs you that the container/OS EZ template is up-to-date.

Note: The `vzpkg status` command can be executed for running containers only.

3.4.8. vzpkg install

This command is used to install application EZ templates, YUM software groups, or individual software packages into containers.

Syntax

```
vzpkg install [<options>] <CT_name> <object> [...]
```

The `vzpkg install` command will add an `<object>` to the specified container. An object can be an application EZ template, a YUM software group, or a standard software package. You can specify several objects to install into the container by separating them by spaces.

When executed, `vzpkg install` automatically handles the interdependencies among the packages to be installed into a container and ensures that all dependencies are satisfied. If the package dependencies cannot be resolved, the installation process fails and the corresponding message is displayed.

Table 3.32. Options

Name	Description
<code>-p, --package</code>	Installs a software package instead of an EZ template.
<code>-g, --groups</code>	Installs a YUM software group instead of an EZ template. The <code>-g</code> option works only for containers running RPM-based Linux distributions.
<code>-f, --force</code>	Forces the EZ template/package installation.
<code>-C, --cache</code>	Makes the <code>vzpkg install</code> command look for the packages included in the EZ template in the local <code>vzpkg</code> cache only. If there is a package not available locally, the command will fail. You can omit this parameter if the elapsed time from the last <code>vzpkg</code> cache update does not exceed the value of the <code>METADATA_EXPIRE</code> parameter specified in the <code>/etc/vztt/vztt.conf</code> file.
<code>-r, --remote</code>	If the elapsed time from the last <code>vzpkg</code> cache update does not exceed the value of the <code>METADATA_EXPIRE</code> parameter specified in the <code>/etc/vztt/vztt.conf</code> file, you should use this option to make <code>vzpkg install</code> look for the packages in the remote repositories set for handling the corresponding EZ template.
<code>-n, --check-only</code>	Simulates the same operations as <code>vzpkg install</code> completes without specifying this option (downloads the software packages to the server, handles the package interdependencies, etc.); however, the packages themselves are not installed in the specified the container.
<code>-d, --debug <num></code>	Sets the debugging level to one of the specified values (from 0 to 10). 10 is the highest debug level and 0 sets the debug level to its minimal value.
<code>-q, --quiet</code>	Disables logging to the screen and to the log file.

By default, the specified object is treated by `vzpkg install` as an application EZ template. However, you can use the `-p` or `-g` option to explicitly specify the type of the object.

Note: A container has to be running in order to apply an application EZ template to or install a package inside this container.

3.4.9. vzpkg update

The `vzpkg update` command is used to update the following components of a container:

- OS EZ template
- application EZ templates
- YUM software groups
- individual software packages

Syntax

```
vzpkg update [<options>] <CT_name> [<object> [...]]
```

Table 3.33. Options

Name	Description
-C, --cache	Makes the <code>vzpkg update</code> command look for the package updates in the local <code>vzpkg</code> cache only. You can omit this parameter if the elapsed time from the last <code>vzpkg</code> cache update does not exceed the value of the <code>METADATA_EXPIRE</code> parameter specified in the <code>/etc/vztt/vztt.conf</code> file.
-r, --remote	If the elapsed time from the last <code>vzpkg</code> cache update does not exceed the value of the <code>METADATA_EXPIRE</code> parameter specified in the <code>/etc/vztt/vztt.conf</code> file, you should use this option to make <code>vzpkg update</code> look for the package updates in the remote repositories set for handling the corresponding EZ templates.
-p, --package	Updates the packages installed in the container by using the <code>vzpkg install</code> command.
-g, --groups	Updates the YUM software group in the container. The <code>-g</code> option works only for containers running RPM-based Linux distributions.
-f, --force	Forces the EZ template/package update procedure.
-n, --check-only	Simulates the same operations as <code>vzpkg update</code> completes without specifying this option (downloads the updated packages to the server, handles their interdependencies, etc.); however, the packages themselves are not updated.
-d, --debug <num>	Sets the debugging level to one of the specified values (from 0 to 10). 10 is the highest debug level and 0 sets the debug level to its minimal value.
-q, --quiet	Disables logging to the screen and to the log file.

Without any options specified, `vzpkg update` updates all EZ templates (including the OS EZ template) in the specified container. However, you can make the command update a particular EZ template by specifying its name as `<object>`. You can also use the `-p` or `-g` option to update YUM software groups or individual software packages in the container.

3.4.10. vzpkg remove

This command is used to remove an application EZ template, YUM software group, or a software package from a container.

Syntax

```
vzpkg remove [<options>] <CT_name> <object> [...]
```

This command will remove `<object>` from the container with the name `<CT_name>`. The `<object>` can be an application EZ template, a YUM software group, or a software package installed with the `vzpkg install` command. You may specify a number of objects for removing.

Table 3.34. Options

Name	Description
-p, --package	Removes the specified package(s) from the container.

Name	Description
-g, --groups	Removes the specified YUM software group from the container. The -g option works only for containers running RPM-based Linux distributions.
-w, --with-depends	Removes also the packages having dependencies with the object specified.
-f, --force	Forces the EZ template/package deletion.
-n, --check-only	Simulates the same operations as <code>vzpkg remove</code> completes without specifying this option (handles interdependencies of the packages to be removed from the server, etc.); however, the packages themselves are not deleted from the specified container(s).
-d, --debug <num>	Sets the debugging level to one of the specified values (from 0 to 10). 10 is the highest debug level and 0 sets the debug level to its minimal value.
-q, --quiet	Disables logging to the screen and to the log file.

By default, the specified object is treated by `vzpkg remove` as an application EZ template. However, you can use the -p or -g option to explicitly specify the type of the object.

Note: A container has to be running in order to remove an application EZ template/package from it.

3.4.11. vzpkg create cache

This command is used to create tarballs (caches) for OS EZ templates. You should execute this command before you start using a newly installed OS EZ template for creating containers.

Syntax

```
vzpkg create cache [<options>] [<OS_template> [...]]
```

Table 3.35. Options

Name	Description
-C, --cache	Makes the <code>vzpkg create cache</code> command check for the packages included in the EZ OS template in the local <code>vzpkg</code> cache only and use them for the cache creation. You can omit this parameter if the elapsed time from the last <code>vzpkg</code> cache update does not exceed the value of the <code>METADATA_EXPIRE</code> parameter specified in the <code>/etc/vztt/vztt.conf</code> file. In this case <code>vzpkg create cache</code> will also check the local <code>vzpkg</code> cache only.
-r, --remote	If the elapsed time from the last <code>vzpkg</code> cache update does not exceed the value of the <code>METADATA_EXPIRE</code> parameter specified in the <code>/etc/vztt/vztt.conf</code> file, you should use this option to make <code>vzpkg create cache</code> check for the packages included in the EZ OS template in the remote repositories set for its handling.
-d, --debug <num>	Sets the debugging level to one of the specified values (from 0 to 10). 10 is the highest debug level and 0 sets the debug level to its minimal value.
-q, --quiet	Disables logging to the screen and to the log file.
-f, --force	Forces the process of the cache creation.

`vzpkg create cache` checks the template area on the server (by default, the `/vz/template` directory is used) and if it finds an OS EZ template for which no tar archive exists, it creates a gzipped tarball for the corresponding OS EZ template and places it to the `/vz/template/cache` directory. When a container is being created, `prlctl` just unpacks the tar archive.

By default, `vzpkg create cache` checks the tar archive existence for all OS EZ templates installed on the server and creates some, if necessary. However, you can explicitly indicate what OS EZ template should be cached by specifying its name as `<OS_template>`. If the cache of the OS template specified already exists on the server, the command will fail and you will be presented with the corresponding error message.

3.4.12. vzpkg update cache

This command is used to update tarballs (caches) of the OS EZ templates installed on the server.

Syntax

```
vzpkg update cache [<options>] [<OS_template> [...]]
```

Table 3.36. Options

Name	Description
<code>-C, --cache</code>	Makes the <code>vzpkg update cache</code> command check for the packages updates in the local <code>vzpkg</code> cache only and use them for the cache creation. You can omit this parameter if the elapsed time from the last <code>vzpkg</code> cache update does not exceed the value of the <code>METADATA_EXPIRE</code> parameter specified in the <code>/etc/vztt/vztt.conf</code> file. In this case <code>vzpkg update cache</code> will also check the local <code>vzpkg</code> cache only.
<code>-r, --remote</code>	If the elapsed time from the last <code>vzpkg</code> cache update does not exceed the value of the <code>METADATA_EXPIRE</code> parameter specified in the <code>/etc/vztt/vztt.conf</code> file, you should use this option to make <code>vzpkg update cache</code> check for the packages updates in the remote repositories set for handling the given EZ OS template.

The `vzpkg update cache` command checks the cache directory in the template area (by default, the template area is located in the `/vz/template` directory on the server) and updates all existing tarballs in this directory. However, you can explicitly indicate what OS EZ template tarball is to be updated by specifying its name as `<OS_template>`. Upon the `vzpkg update cache` execution, the old tarball is renamed by receiving the `-old` suffix (e.g., `centos-6-x86_64.tar.gz-old`).

If the `vzpkg update cache` command does not find a tarball for one or more OS EZ templates installed on the server, it creates the corresponding tar archive(s) and puts them to the `/vz/template/cache` directory.

3.4.13. vzpkg remove cache

This command removes the cache for the OS EZ templates specified.

Syntax

```
vzpkg remove cache [<options>] [<OS_template> [...]]
```

Table 3.37. Options

Name	Description
-d, --debug <num>	Sets the debugging level to one of the specified values (from 0 to 10). 10 is the highest debug level and 0 sets the debug level to its minimal value.
-q, --quiet	Disables logging to the screen and to the log file.

By default, `vzpkg remove cache` deletes all caches located in the `/vz/template/cache` directory on the server. However, you can explicitly indicate what OS EZ template tar archive is to be removed by specifying its name as `<OS_template>`.

Note: The OS EZ template caches having the `-old` suffix are not removed from the `/vz/template/cache` directory. You should use the `-rm` command to delete these caches from the server.

3.4.14. vzpkg create appcache

This command combines an OS EZ template cache and one or more application EZ templates into a new OS and applications cache. If the OS EZ template cache has not been created yet, this will be done before application templates are added.

Syntax

```
vzpkg create appcache --config <config> [<options>]
```

Table 3.38. Options

Name	Description
--config <config>	Specifies the path to the configuration file with the information on what OS and application templates to use.
--ostemplate <OS_template>	Specifies the OS EZ template to use in cache creation. This option redefines the OS EZ template specified in the configuration file.
--apptemplate <app_template>	Specifies one or more application EZ templates (comma-separated) to be added to the resulting cache. This option redefines the application EZ templates specified in the configuration file.
-d, --debug <num>	Sets the debugging level (0 to 10), 10 being the highest.
-q, --quiet	Disables logging to screen and log file.
-f, --force	Forces cache creation.

3.4.15. vzpkg update appcache

This command updates an existing OS EZ template cache with preinstalled application templates if the `--update-cache` option is provided. Otherwise, the cache is created anew.

Syntax

```
vzpkg update appcache --config <config> [<options>]
```

Table 3.39. Options

Name	Description
--config <config>	Specifies the configuration file with the information on what OS and application templates to use.
--ostemplate <OS_template>	Specifies the OS EZ template, cache of which needs to be recreated or updated. This option redefines the OS EZ template specified in the configuration file.
--apptemplate <app_template>	Specifies all application EZ templates (comma-separated) preinstalled in the cache which needs to be updated. This option redefines the application EZ templates specified in the configuration file.
--update-cache	Instructs the command to check for updates for the existing OS and application cache. Otherwise, the cache is created anew.
-d, --debug <num>	Sets the debugging level (0 to 10), 10 being the highest.
-q, --quiet	Disables logging to screen and log file.
-f, --force	Forces cache creation.

3.4.16. vzpkg remove appcache

This command removes an existing OS EZ template cache with preinstalled application templates.

Syntax

```
vzpkg remove appcache --config <config> [<options>]
```

Table 3.40. Options

Name	Description
--config <config>	Specifies the configuration file with the information on what OS and application templates to use.
--ostemplate <OS_template>	Specifies the OS EZ template, cache of which needs to be removed. This option redefines the OS EZ template specified in the configuration file.
--apptemplate <app_template>	Specifies all application EZ templates (comma-separated) preinstalled in the cache which needs to be removed. This option redefines the application EZ templates specified in the configuration file.
-d, --debug <num>	Sets the debugging level (0 to 10), 10 being the highest.
-q, --quiet	Disables logging to screen and log file.
-f, --force	Forces cache creation.

3.4.17. vzpkg localinstall

The `vzpkg localinstall` command is used to install a software package inside a container from the corresponding file on the server.

Syntax

```
vzpkg localinstall [<options>] <CT_name> <rpm_file_path> [...]
```

Table 3.41. Options

Name	Description
-C, --cache	When handling the package interdependencies, makes the vzpkg localinstall command look for the needed packages in the local vzpkg cache only. If there is a package not available locally, the command will fail. You can omit this parameter if the elapsed time from the last vzpkg cache update does not exceed the value of the METADATA_EXPIRE parameter specified in the /etc/vztt/vztt.conf file.
-r, --remote	If the elapsed time from the last vzpkg local cache update does not exceed the value of the METADATA_EXPIRE parameter specified in the /etc/vztt/vztt.conf file, you should use this option to make vzpkg localinstall look for the packages in the remote repository.
-n, --check-only	Simulates the same operations as vzpkg localinstall completes without specifying this option (e.g., handles the package interdependencies); however, the package itself is not installed in the specified container.
-d, --debug <num>	Sets the debugging level to one of the specified values (from 0 to 10). 10 is the highest debug level and 0 sets the debug level to its minimal value.
-q, --quiet	Disables logging to the screen and to the log file.

When executed, the command installs the package, the full path to which is specified as *<rpm_file_path>*, inside the container with the name *<CT_name>*. You may specify multiple packages to be installed inside the container.

During its execution, vzpkg localinstall automatically handles the interdependencies among the packages to be installed inside a container and ensures that all dependencies are satisfied. If the package dependencies cannot be resolved, the installation process will fail and the corresponding message will be displayed.

3.4.18. vzpkg localupdate

The vzpkg localupdate command is used to update the software packages installed inside your container(s) by means of the vzpkg install or vzpkg localinstall commands.

Syntax

```
vzpkg localupdate [<options>] <CT_name> <rpm_file_path> [...]
```

Table 3.42. Options

Name	Description
-C, --cache	When handling the package interdependencies, makes the vzpkg localupdate command look for the needed packages in the local vzpkg cache only. If there is a package not available locally, the command will

Name	Description
	fail. You can omit this parameter if the elapsed time from the last vzpkg cache update does not exceed the value of the METADATA_EXPIRE parameter specified in the /etc/vztt/vztt.conf file.
-r, --remote	If the elapsed time from the last vzpkg local cache update does not exceed the value of the METADATA_EXPIRE parameter specified in the /etc/vztt/vztt.conf file, you should use this option to make vzpkg localupdate look for the packages in the remote repository.
-n, --check-only	Simulates the same operations as vzpkg localupdate completes without specifying this option (e.g., handles the package interdependencies); however, the package itself is not installed in the specified container.
-d, --debug <num>	Sets the debugging level to one of the specified values (from 0 to 10). 10 is the highest debug level and 0 sets the debug level to its minimal value.
-q, --quiet	Disables logging to the screen and to the log file.

When executed, vzpkg localupdate compares the file on the server the full path to which is specified as *<rpm_file_path>* with the corresponding package inside the container with the name *<CT_name>* and updates it, if necessary. You may specify a number of packages at once to be updated inside your container.

3.4.19. vzpkg upgrade

The vzpkg upgrade command is used to upgrade an OS EZ template the container is based on to a newer version.

Syntax

```
vzpkg upgrade [<options>] <CT_name>
```

Table 3.43. Options

Name	Description
-C, --cache	Makes the vzpkg upgrade command check for the packages included in the OS EZ template in the local vzpkg cache only. If any package is not available locally, the command will fail. You can omit this parameter if the elapsed time from the last vzpkg cache update does not exceed the value of the METADATA_EXPIRE parameter specified in the /etc/vztt/vztt.conf file; in this case vzpkg upgrade will also check the local vzpkg cache only.
-r, --remote	If the elapsed time from the last local vzpkg cache update does not exceed the value of the METADATA_EXPIRE parameter specified in the /etc/vztt/vztt.conf file, you should use this option to make vzpkg upgrade check for the packages in the remote repositories set for handling the given EZ OS template.
-n, --check-only	Simulates the same operations as vzpkg upgrade completes without specifying this option (downloads the packages to the server, handles their interdependencies, etc.); however, the packages themselves inside the container are not upgraded.

Name	Description
-f, --force	Forces the upgrade of the OS EZ template.
-d, --debug <num>	Sets the debugging level to one of the specified values (from 0 to 10). 10 is the highest debug level and 0 sets the debug level to its minimal value.
-q, --quiet	Disables logging to the screen and to the log file.

3.4.20. vzpkg fetch

This command is used to download packages included in the corresponding OS EZ template or their updates from the remote repository to the vzpkg local cache on the server and to prepare them for installation.

Syntax

```
vzpkg fetch [<options>] <OS_template>
```

Table 3.44. Options

Name	Description
-O, --os	Download packages/updates for the specified EZ OS template.
-A, --app	Download packages/updates for EZ application templates used with the EZ specified OS template.
-C, --cache	Makes the vzpkg fetch command look for the metadata in the vzpkg local cache only. You can omit this parameter if the elapsed time from the last vzpkg cache update does not exceed the value of the METADATA_EXPIRE parameter specified in the /etc/vztt/vztt.conf file.
-r, --remote	If the elapsed time from the last vzpkg cache update does not exceed the value of the METADATA_EXPIRE parameter specified in the /etc/vztt/vztt.conf file, you should use this option to make vzpkg fetch look for the OS EZ template metadata in the remote repositories set for handling the corresponding EZ template.
-f, --force	Forces the process of downloading packages and/or their updates to the server.
-d, --debug <num>	Sets the debugging level to one of the specified values (from 0 to 10). 10 is the highest debug level and 0 sets the debug level to its minimal value.
-q, --quiet	Disables logging to the screen and to the log file.

You can make vzpkg fetch run as a cron job (e.g., nightly) checking for available packages or packages updates for your EZ templates and keeping them in the local cache. Having all the necessary packages in the vzpkg local cache can greatly speed up the execution of the vzpkg install, vzpkg update, or vzpkg create cache commands since the packages are available locally and there is no need to check for them in the corresponding remote repositories.

3.4.21. vzpkg clean

This command is used to remove the software packages, their headers, and metadata downloaded to the server from the repository during the vzpkg execution (e.g., while caching an OS EZ template or adding an application EZ template to a container for the first time).

Syntax

```
vzpkg clean [<options>] [<OS_template> [...]]
```

Table 3.45. Options

Name	Description
-k, --clean-packages	Removes the packages, headers, and metadata of the specified EZ OS template from the local vzpkg cache. This is also the default behavior of vzpkg clean.
-t, --clean-template	Checks the template area for the specified EZ OS template (the template area has the default path of /vz/template) and removes all packages that are currently not used by any container on the server and not included in the EZ OS template cache.
-a, --clean-all	Removes both: - the packages, headers, and metadata of the specified EZ OS template from the vzpkg local cache, and - the packages that are currently not used by any container on the server and not included in the EZ OS template cache.
-f, --force	Forces the vzpkg clean execution.
-n, --check-only	Simulates the same operations as vzpkg clean completes without specifying this option; however, the packages and headers are not removed from the server.
-d, --debug <num>	Sets the debugging level to one of the specified values (from 0 to 10). 10 is the highest debug level and 0 sets the debug level to its minimal value.
-q, --quiet	Disables logging to the screen and to the log file.

3.4.22. vzpkg update metadata

This command is used to update the OS EZ template local metadata on the server.

Syntax

```
vzpkg update metadata [<options>] [<OS_template> [...]]
```

Table 3.46. Options

Name	Description
-C, --cache	Makes the vzpkg update metadata command look for available metadata updates in the local vzpkg cache only. You can omit this parameter if the

Name	Description
	elapsed time from the last vzpkg cache update does not exceed the value of the METADATA_EXPIRE parameter specified in the /etc/vztt/vztt.conf file.
-r, --remote	If the elapsed time from the last vzpkg cache update does not exceed the value of the METADATA_EXPIRE parameter specified in the /etc/vztt/vztt.conf file, you should use this option to make vzpkg update metadata look for the updated metadata in the remote repositories set for handling the corresponding OS EZ template.
-d, --debug <num>	Sets the debugging level to one of the specified values (from 0 to 10). 10 is the highest debug level and 0 sets the debug level to its minimal value.
-q, --quiet	Disables logging to the screen and to the log file.

When executed without any options, the command updates the metadata of all OS EZ templates installed on the server. If you specify one or more OS EZ templates, the command will update the metadata of the indicated OS templates only. You can run this command a cron job at regular intervals to be sure that your OS EZ templates metadata are always up-to-date.

3.5. Supplementary Tools

3.5.1. pcompact

Utility to compact containers by removing unused blocks from their virtual disks. By compacting virtual disks, you can increase free disk space on the physical server.

Syntax

```
pcompact [-v] [-n] [-s] [-t <timeout>[s|m|h]]
```

Table 3.47. Options

Name	Description
-v	Increase the command output verbosity. Multiple -v options can be specified to produce a more verbose output.
-n	Display the actions the command will execute but do not actually compact the disks.
-s	Stop the command execution after compacting the first virtual disk.
-t <timeout>[s m h]	Terminate the command after the specified timeout, in seconds (default), minutes or hours.

3.5.2. prl_disk_tool

The prl_disk_tool utility is used to manage virtual hard disk drives.

3.5.2.1. prl_disk_tool compact

Removes all empty blocks from the expanding virtual disk to reduce its size on the physical hard disk. The virtual disk must be formatted to NTFS, ext2/ext3/ext4, btrfs, or xfs.

Syntax

```
prl_disk_tool compact --hdd <disk_path> [--force]
prl_disk_tool compact -i, --info --hdd <disk_path>
```

Table 3.48. Options

Name	Description
--hdd <disk_path>	Full path to the virtual disk.
--force	Forces the compacting operation for suspended virtual disks.
-i, --info	Do not compact the virtual disk; just display the information about the size the disk will have after compacting.

3.5.2.2. prl_disk_tool merge

Merges all snapshots of the virtual hard disk.

Syntax

```
prl_disk_tool merge --hdd <disk_path>
```

Table 3.49. Options

Name	Description
--hdd <disk_path>	Full path to the virtual disk.

3.5.2.3. prl_disk_tool resize

Changes the capacity of the specified virtual disk. During resizing, all data present on the disk volumes are left intact. You can also resize the last partition using the --resize_partition option. The supported file systems are NTFS, ext2/ext3/ext4, btrfs, or xfs.

Syntax

```
prl_disk_tool resize --size <size>[K|M|G|T] [--resize_partition] --hdd <disk_path>
                        [--force]
prl_disk_tool resize -i, --info --hdd <disk_path>
```

Table 3.50. Options

Name	Description
--size	New size of the virtual disk. It can be set either in kilobytes (K), megabytes (M, default), gigabytes (G), or terabytes (T).

Name	Description
--resize_partition	Resizes the last partition of the specified virtual disk. Note: You cannot reduce XFS filesystems (the default choice for CentOS 7 and Red Hat Enterprise Linux 7).
--hdd <disk_path>	Full path to the virtual disk.
--force	Forces the resizing operation for suspended virtual disks.
-i, --info	Do not resize the virtual disk; just show the size the disk will have after resizing.

3.5.3. vzpid

This utility prints the ID of the container where the process is running.

Syntax

```
vzpid <pid> [...]
```

Multiple process IDs can be specified as arguments.

3.5.4. vzps, vztop

These two utilities can be run on the server just as the standard Linux `ps` and `htop` utilities. For information on the `ps` and `htop` utilities, consult their man pages. The `vzps` and `vztop` utilities provide certain additional functionality related to monitoring separate containers running on the server.

The `vzps` utility has the following functionality added: the `-E <CT_name>` command-line switch can be used to show only the processes running inside the container with the specified ID.

The `vztop` utility has the **CTID** column added to display the container UUID where a particular process is running (`@` stands for the server itself).

3.5.5. vzsplit

This utility is used to generate a sample container configuration file with a set of system resource control parameters.

Syntax

```
vzsplit [-n <num>] [-f <sample_name>] [-s <swap_size>]
```

This utility is used for dividing the server into equal parts. It generates a full set of containers system resource control parameters based on the total physical memory of the server it runs on and the number of containers the server shall be able to run even if the given number of containers consume all allowed resources.

Without any option the utility prompts for the desired number of containers and outputs the resulting resource control parameters to the screen.

Table 3.51. Options

Name	Description
-n <num>	Desired number of containers to be simultaneously run on the server.
-f <sample_name>	Name of the sample configuration to create.
-s <swap_size>	Size of the swap file on the server. It is recommended to specify the swap size to be taken into account when the utility generates sample configurations.

The resulting sample configuration will be created in the `/etc/vz/conf` directory. The file name will be `ve-<sample_name>.conf-sample`. Now you can pass `<sample_name>` as an argument to the `--config` option of the `prlctl create` command. If a sample with this name already exists, the utility will output an error message and will not overwrite the existing configuration.

Chapter 4. Managing Virtual Machines

4.1. prlctl

OpenVZ virtual machines can be managed using the `prlctl` command-line utility. The utility is installed on the hardware node during the product installation.

4.1.1. General Syntax

The `prlctl` utility is used to perform administration tasks on virtual machines. The utility supports a full range of tasks from creating and administering virtual machines to getting statistics and generating problem reports.

Syntax

```
prlctl <command> <VM_name> [<options>] [-v, --verbose <number>] [--timeout <sec>]
      [-l, --login [<user>[:<passwd>_]@]<server>] [-p, --read-passwd <file>]]
```

Table 4.1. Options

Name	Description
<command>	The name of the command to execute.
<VM_name>	The name of the virtual machine to perform the operation on. To obtain the list of the available virtual machines, use the <code>prlctl list</code> command.
<options>	Command options. See individual commands for available options.
-v, --verbose <number>	Enables verbose output. The greater the <number>, the higher the verbosity.
-l, --login [<user>[:<passwd>_]@]<server>	Connect to a remote <server> with the specified credentials. If this flag is omitted, the <code>prlctl</code> command is assumed to be run locally.
-p, --read-passwd <file>	Use the password from the file <file> to log in to a remote hardware node, other credentials being specified with the <code>--login</code> option. The <code>--read-passwd</code> option can be specified multiple times in order to form a password stack for operations requiring multiple passwords. Each password must be supplied in a separate file.

To display help, enter `prlctl` without any options.

4.1.2. prlctl clone

Creates an exact copy of the specified virtual machine.

Syntax

```
prlctl clone <VM_name> --name <new_name> [--template] [--dst=<path>] [--changesid]
          [--detach-external-hdd <yes|no>]
```

Table 4.2. Options

Name	Description
<VM_name>	Name of the virtual machine to clone.
--name <new_name>	Name to be assigned to the new virtual machine.
--template	Create a virtual machine template instead of a real virtual machine. Templates are used as a basis for creating new virtual machines.
--dst=<path>	Full path to the directory where the new virtual machine will be stored. If this option is omitted, the new virtual machine will be created in the default directory.
--changesid	Generate a new Windows security identifier (SID) for a Windows-based virtual machine. For this parameter to work, OpenVZ tools must be installed in the virtual machine.
--detach-external-hdd <yes no>	If set to no, hard disks located outside the source virtual machine are not removed from the configuration of the resulting clone. Setting the parameter to yes removes external hard disks from the configuration. Note: External hard disks are not copied to the cloned virtual machine.

4.1.3. prlctl create

Creates a new virtual machine. A virtual machine can be created from scratch or from a virtual machine template. When created from scratch, the target operating system type or version must be specified. To create a virtual machine from a template, the template name must be passed to the command.

Syntax

```
prlctl create <VM_name> [<options>]
```

Table 4.3. Options

Name	Description
<VM_name>	User-defined new virtual machine name. If the name consists of two or more words separated by spaces, it must be enclosed in quotes.
-d, --distribution {<name> list}	The operating system distribution the virtual machine will be optimized for. For the full list of supported distributions, refer to the prlctl man pages.
--ostemplate <template_name>	The name of the virtual machine template from which to create the new virtual machine. Use the prlctl list --template command to obtain the list of the available templates.
--dst <vm_path>	User-defined new VM home path. If this parameter is omitted, the VM will be placed in the default virtual machine directory.

Name	Description
<code>--uuid <uuid></code>	A custom UUID to assign to the virtual machine.

Note: You can use either `--distribution` or `--ostemplate`, not both options at once.

When creating a virtual machine from scratch, you may specify the operating system family or version. If an operating system version is specified using the `--distribution` parameter, the virtual machine will be configured for that operating system. If an operating system family is specified using the `--ostype` parameter, the virtual machine will be configured for the default version of this OS family. The default versions are determined internally by OpenVZ. The best way to find out the default versions used in your OpenVZ installation is by creating a sample virtual machine.

4.1.4. prlctl delete

Deletes a virtual machine from the hardware node. The command removes a virtual machine from the OpenVZ registry and permanently deletes all its files from the server. Once completed, this operation cannot be reversed.

Syntax

```
prlctl delete <VM_name>
```

Table 4.4. Options

Name	Description
<code><VM_name></code>	The name of the virtual machine to delete.

4.1.5. prlctl enter

Creates a command prompt channel to a virtual machine. By using this command, you can create a command prompt channel and execute commands in a virtual machine. OpenVZ tools must be installed in a virtual machine to use this utility.

Syntax

```
prlctl enter <VM_name>
```

Table 4.5. Options

Name	Description
<code><VM_name></code>	The name of the virtual machine.

4.1.6. prlctl exec

Executes a command inside a virtual machine. OpenVZ tools must be installed in a virtual machine to use this utility. By default, running `prlctl exec <command>` is equivalent to executing `bash -c <command>` in a Linux VM or `cmd /c <command>` in a Windows VM.

Syntax

```
prlctl exec <VM_name> [--without-shell] <command>
```

Table 4.6. Options

Name	Description
<VM_name>	The name of the virtual machine.
<command>	A command to execute.
--without-shell	Run commands directly without bash or cmd shell.

4.1.7. prlctl list

Displays a list of virtual machines on the Hardware Node. Displays information on virtual machines on the Hardware Node.

Syntax

```
prlctl list --vmtype vm [-a, --all] [-o, --output <field>[,...]]
               [-s, --sort {<field>|-<field>}] [-t, --template] [-j, --json]
prlctl list -i, --info --vmtype vm [<VM_name>] [-f, --full] [-t, --template]
                                   [-j, --json]
```

Table 4.7. Options

Name	Description
-a, --all	List all running, stopped, suspended, and paused virtual machines. If this and the rest of the parameters are omitted, only the running virtual machines will be displayed.
-t, --template	List available virtual machine templates instead of actual virtual machines.
-o, --output <field>[,...]	Display only the specified fields. Type field names in lower case. Separate multiple fields with commas. For the list of fields, see Section 3.2.15.1, “prlctl list Output Parameters” on page 51.
-s, --sort {<field> -<field>}	Sort virtual machines by the specified field in either ascending or descending order.
-i, --info	Display detailed information about the specified virtual machine.
-f, --full	Display detailed information about network cards in virtual machines. Used with the --info option.
<VM_name>	The name of the virtual machine for which to display the detailed information. If not specified, the information will be displayed for all registered virtual machines.
-j, --json	Produce machine-readable output in the JSON format.

4.1.7.1. prlctl list Output Parameters

Listed below are the parameters that can be specified after the -o switch.

Name	Output Column	Description
uuid	UUID	Virtual machine UUID.
hostname	HOSTNAME	Virtual machine hostname.
name	NAME	Virtual machine name.
description	DESCRIPTION	Virtual machine description.
ostemplate	OSTEMPLATE	Specifies the name of the OS template the virtual machine is based on (e.g., centos-6-x86_64).
ip	IP_ADDR	Virtual machine IP address.
status	STATUS	Virtual machine status (e.g., running or stopped).
numproc	NPROC	The number of processes and threads allowed.
mac	MAC	Network device's MAC address.
netif	NETIF	Network devices in the virtual machine .
iolimit	IOLIMIT	The bandwidth the virtual machine is allowed to use for its disk input and output (I/O) operations, in bytes per second.
ha_enable	HA_ENABLE	Indicates whether the virtual machine is joined to the High Availability Cluster.
ha_prio	HA_PRIORITY	Virtual machine priority in the High Availability Cluster (0 is the lowest). Higher-priority virtual environments are restarted first in case of failures.

4.1.8. prlctl statistics

Print statistics for running containers on the server.

Syntax

```
prlctl statistics {<CT_UUID_or_name>|-a, --all} [--loop] [--filter <filter>]
```

Table 4.8. Options

Name	Description
`-a, --all	Print statistics for all virtual machines and containers on the server.
--loop	Print statistics every second until the program is terminated.
--filter <pattern>;	Specifies the subset of performance statistics to collect and print. If omitted, all available statistics is shown. Asterisks (*) can be used as wildcards for any number of arbitrary characters. The available filters are listed below (<N>; is the device or filesystem index). Storage device statistics • devices.{ide
scsi	sata}<N>.read_requests - Total count of read requests to IDE, SCSI, or SATA controller * devices.{ide

Name	Description
scsi	sata}<N>.read_total - Total count of read bytes for IDE, SCSI, or SATA controller * devices.{ide
scsi	sata}<N>.write_requests - Total count of write requests to IDE, SCSI, or SATA controller * devices.{ide
scsi	sata}<N>.write_total - Total count of written bytes for IDE, SCSI, or SATA controller Network statistics • net.nic<N>.pkts_in - Total number of incoming packets for network adapter • net.nic<N>.pkts_out - Total number of outgoing packets for network adapter • net.nic<N>.bytes_in - Total number of incoming bytes for network adapter • net.nic<N>.bytes_out - Total number of outgoing bytes for network adapter Classful network statistics The result is provided in five columns: Class, Input(bytes), Input(packets), Output(bytes), Output(packets). • net.classful.traffic - Total counters for IPv4 and IPv6 traffic • net.classful.traffic.ipv4 - Counters for IPv4 traffic • net.classful.traffic.ipv6 - Counters for IPv6 traffic CPU statistics • guest.cpu.usage - Guest OS CPU usage, in percent • guest.cpu.time - Sum of guest CPU time differences since the last query for each vCPU averaged by the number of host CPUs, in microseconds • host.cpu.time - Sum of host CPU time differences since the last query for each vCPU averaged by the number of host CPUs, in microseconds • guest.vcpu<N>.time - per-vCPU statistics, in nanoseconds RAM statistics • guest.ram.usage - Guest OS used RAM, in MiB • guest.ram.cached - Guest OS cached RAM, in MiB • guest.ram.total - Guest OS total RAM, in MiB • guest.ram.swap_in - Guest OS virtual memory stats, in counts

Name	Description
	• guest.ram.swap_out - Guest OS virtual memory stats, in counts • guest.ram.minor_fault - Guest OS minor page fault count • guest.ram.major_fault - Guest OS major page fault count • guest.ram.balloon_actual - Guest OS balloon size, in MiB Mounted filesystems statistics • guest.fs<N>.name - Device name as seen from inside the guest filesystem • guest.fs<N>.total - Total size of the filesystem, in KiB • guest.fs<N>.free - Amount of free space on the filesystem, in KiB • guest.fs<N>.disk.<N> - Disk indices

4.1.9. prlctl migrate

Migrates a virtual machine from one server to another.

Syntax

```
prlctl migrate [<source_server>]/<VM_name> <destination_server>[/<VM_name>]
               [--dst=<path>] [--keep-src|--remove-src] [--changesid] [--no-compression]
               [--no-tunnel] [--ssh <options>]
```

Table 4.9. Options

Name	Description
<VM_name>	The source virtual machine name.
<source_server>	The source server information. Use the following format to specify this info: [<user>[:<password>]@]<server_IP_address_or_hostname>[:<port>].
<destination_server>	The destination server information. If omitted, the migration will be performed locally. Use the following format to specify this info: [<user>[:<password>]@]<server_IP_address_or_hostname>[:<port>].
--dst=<path>	Name and path of the directory on the destination server where the virtual machine files should be stored.
--keep-src	Do not remove the original VM from the source server. Cannot be used together with --remove-src.
--remove-src	Remove the original VM from the source server. Enabled by default. Cannot be used together with --keep-src.
--changesid	Changes the resulting virtual machine SID.
--no-compression	Disable data compression during migration.
--no-tunnel	Disables connection tunneling for migration. Connection tunneling provides secure data transmission. The option works only for VM live migration.

Name	Description
--ssh	Additional options to pass to ssh to connect to the destination server. All standard ssh options are supported.
	Note: Do not specify the destination server hostname or IP address as an ssh option.

4.1.10. prlctl mount, umount

Mounts or unmounts the hard disks of a virtual machine to the `/vz/root/<UUID>` directory on the hardware node.

Syntax

```
prlctl mount <VM_name> [-o <ro|rw> | --info]
prlctl umount <VM_name>
```

Table 4.10. Options

Name	Description
<VM_name>	Virtual machine name.
-o <ro rw>	Sets access rights: - ro - read-only, - rw - read-write.
--info	Show information about the mounted virtual disks.

4.1.11. prlctl move

Moves the files of a virtual machine to a new location on the same server. The virtual machine must be stopped or suspended.

Syntax

```
prlctl move <VM_name> --dst=<path>
```

Table 4.11. Options

Name	Description
<VM_name>	Virtual machine name.
--dst=<path>	Path to the new virtual machine files location.

4.1.12. prlctl pause, suspend, resume

Pause, suspend, and resume a virtual machine.

Syntax

```
prlctl pause <VM_name>
prlctl suspend <VM_name>
prlctl resume <VM_name>
```

Table 4.12. Options

Name	Description
<VM_name>	The name of the virtual machine to pause, suspend, or resume.

The pause command pauses a virtual machine. To continue the virtual machine operation, use the prlctl start command.

The suspend command suspends the virtual machine operation. When a running virtual machine is suspended, the state of the virtual machine processes is saved to a file on the host. After that, the machine is stopped. To resume the machine, use the resume command.

4.1.13. prlctl problem-report

Obtains a problem report for the specified virtual machine and either sends it to the OpenVZ technical support team or displays it on the screen.

Syntax

```
prlctl problem-report <VM_name> <-d, --dump|-s, --send [--proxy [<user> \
[:<passwd>]@<proxyhost>[:<port>]]] [--no-proxy]>
```

Table 4.13. Options

Name	Description
<VM_name>	The name of the virtual machine for which to obtain the problem report. If the name consists of separate words, it must be enclosed in quotes.
-d, --dump	Collect technical data about a virtual machine and display it on the screen. You can also pipe the output to a file and then send it to the OpenVZ technical support team to analyze your problem.
-s, --send	Send the generated problem report to the OpenVZ technical support team.
--proxy [<user>[:<passwd>]@\<proxyhost>[:<port>]]	Use the specified information to send the generated report through a proxy server, if you use one to connect to the Internet.
--no-proxy	Do not use a proxy server to send the generated report. This is the default behavior, so you can omit this parameter.

4.1.14. prlctl register, unregister

The register command is used to register a virtual machine with OpenVZ.

The `unregister` command removes a virtual machine from the OpenVZ registry.

Syntax

```
prctl register <path> --preserve-uuid <yes|no>
prctl unregister <VM_name>
```

Table 4.14. Options

Name	Description
<path>	An absolute path to the virtual machine directory.
<VM_name>	The name of the virtual machine to remove from the OpenVZ registry.
--preserve-uuid <yes no>	Specifies what to do with the UUID (universally unique identifier). If you specify yes, the UUID is preserved. If you specify no, the UUID is regenerated (default behavior).

Use the `register` command when you have a virtual machine on the server that does not show up in the list of the virtual machines registered with the OpenVZ. This can be a machine that was previously removed from the registry or a machine that was copied from another location.

The `unregister` command removes a virtual machine from the OpenVZ registry, but does not delete the virtual machine files from the server. You can re-register such a machine with OpenVZ later using the `register` command.

4.1.15. `prctl reset-uptime`

Resets a virtual machine uptime counter as well as count start date and time.

Syntax

```
prctl reset-uptime <VM_name>
```

Table 4.15. Options

Name	Description
<VM_name>	Virtual machine name. Names consisting of multiple words must be enclosed in quotes.

4.1.16. `prctl set`

The `prctl set` command is used to modify the configuration of a virtual machine and manage virtual machine devices. The following subsections provide technical information on how to use the command to perform these tasks.

4.1.16.1. Modifying Virtual Machine Configuration

The `prctl set` command is used to modify the virtual machine configuration parameters.

Syntax

```
prlctl set <VM_name> [--cpus <number>] [--memsize <number>]
    [--videosize <number>] [--memguarantee <percentage>]
    [--mem-hotplug <on|off>] [--distribution _<name>] [--description <desc>]
    [--autostart <on|off|auto>] [--autostart-delay <number>]
    [--autostop <stop|suspend>] [--applyconfig <conf>] [--name <new_name>]
    [--start-as-user <administrator|owner|<user>:<passwd>_]
    [--vnc-mode <auto|manual|off> {--vnc-passwd <passwd> | --vnc-nopasswd}]
    [--vnc-port <port>] [--vnc-address <address>] [--cpu-hotplug <on|off>]
    [--cpuunits <units>] [--cpulimit {<percent>|<megahertz>}]
    [--ioprio <priority>] [--iolimit <limit>] [--iopslimit <limit>]
    [--cpumask {<N>[,N,N1-N2] | all}] [--offline-management <on|off>]
    [--offline-service <service_name>] [--userpasswd <user>:<passwd>
    [--encrypted]] [--rate <rate>] [--ratebound <on|off>_]
    [--apply-iponly <yes|no>] [--efi-boot <on|off>]
    [--tools-autoupdate <on|off>]
```

Table 4.16. Options

Name	Description
<VM_name>	Target virtual machine name.
--cpus <number>	Number of virtual CPUs in the virtual machine. If the server has several CPU cores, this option also defines the number of CPUs shown to users from inside a virtual machine.
--memsize <number>	The amount of memory (RAM) available to the virtual machine, in megabytes. You can use the following suffixes to specify measurement units: • G for gigabytes • M for megabytes • K for kilobytes • B for bytes
--videosize <number>	The amount of video memory available to the virtual machine graphics card. You can use the following suffixes to specify measurement units: • G for gigabytes • M for megabytes • K for kilobytes • B for bytes
--memguarantee <size>	Sets a percentage of virtual machine's RAM that said VM is guaranteed to have. By default, set to 40%.
--mem-hotplug <on off>	Enables or disables memory (RAM) hotplug support in the virtual machine. This feature is disabled in the virtual machine by default. The guest operating system must support memory hotplug for this functionality to work.

Name	Description
<code>--distribution <name></code>	Optimize the virtual machine for use with the operating system <i><name></i> . You can get the list of available distributions using the <code>prlctl set <VM_name> -d list</code> command.
<code>--description <desc></code>	Sets virtual machine description. Descriptions with white spaces must be enclosed in quotation marks.
<code>--autostart <on off auto></code>	Defines the virtual machine start-up options: • <code>on</code> - the virtual machine is started automatically when the hardware node starts or the OpenVZ component responsible for managing virtual machines is disabled. • <code>off</code> - the autostart is off. This is the default virtual machine start-up mode. • <code>auto</code> - resume the virtual machine state prior to shutting down the hardware node or disabling the OpenVZ component responsible for managing virtual machines. If you set this option to <code>on</code> or <code>auto</code>, you must additionally specify the <code>--start-as-user</code> option.
<code>--autostart-delay <number></code>	Sets the time delay used during the virtual machine automatic startup.
<code>--autostop < suspend></code>	Sets the automatic shutdown mode for the specified virtual machine: • <code>stop</code> - the virtual machine is stopped when you shut down the hardware node or disable the OpenVZ component responsible for managing virtual machines. • <code>suspend</code> - the virtual machine is suspended when the hardware node is shut down or the OpenVZ component responsible for managing virtual machines is disabled.
<code>--applyconfig <conf></code>	Applies the resource parameter values from the specified VM sample file in <code>/etc/parallels/samples</code> to the virtual machine. The following parameters are applied: • all memory-related parameters (both RAM and video) • all CPU-related parameters • IO and IOPS limits • disk size
<code>--name <new_name></code>	Changes the virtual machine name. You can only change the names of stopped virtual machines.

Name	Description
<code>--start-as-user <administrator owner <user>:<passwd></code>	Specifies the account to use to autostart the virtual machine: • <code>administrator</code> - start the virtual machine as the administrator of the host operating system. • <code>owner</code> - start the virtual machine as the virtual machine owner. • <code><user>:<passwd></code> - start the virtual machine as the specified user.
<code>--vnc-mode <auto manual off></code>	Enables or disables access to the virtual machine via the VNC protocol.
<code>--vnc-port <port></code>	Sets the VNC port number.
<code>--vnc-passwd <passwd> --vnc-nopasswd</code>	Sets the VNC password or specifies that no password is needed for VNC connections. Either of these options is mandatory for any VNC connection.
<code>--vnc-address <address></code>	Sets the IP address to use for logging in to the virtual machine via VNC. It must be one of the IP addresses assigned to the hardware node. By default, you can use any of the IP addresses of the hardware node to log in to the virtual machine.
<code>--cpu-hotplug <on off></code>	Enables or disables CPU hotplug support in the virtual machine. This feature is disabled by default. The guest operating system must support CPU hotplug for this functionality to work.
<code>--cpuunits <units></code>	Sets the CPU weight for the virtual machine. This is a positive integer number that defines how much CPU time the virtual machine can get as compared to the other virtual machines and containers running on the server. The larger the number, the more CPU time the virtual machine can receive. Possible values range from 8 to 500000. If this parameter is not set, the default value of 1000 is used.
<code>--cpulimit {<percent> <megahertz>}</code>	CPU limit, in percent or megahertz (MHz) the virtual machine is not allowed to exceed. By default, the limit is set in percent. To set the limit in MHz, specify "m" after the value. Note: If the server has 2 processors, the total CPU time equals 200%.
<code>--ioprio <priority></code>	Disk I/O priority level from 0 to 7. The default is 4.
<code>--iolimit <limit></code>	Disk I/O bandwidth limit. The default is 0 (no limit). By default the limit is set in megabytes per second.

Name	Description
	You can use the following letters following the number to specify units of measure: • G - gigabytes per second (e.g., 1G). • K - kilobytes per second (e.g., 10K). • B - bytes per second (e.g., 100B). The default I/O bandwidth limit for all newly created virtual machines is set to 0, which means that no limits are applied to them.
<code>--iopslimit <limit></code>	Maximum number of disk input and output operations per second a virtual machine is allowed to perform. By default, any newly created container does not have the IOPS limit set and can perform so many disk I/O operations per second as necessary.
<code>--cpumask {<N>[,N,N1-N2] all}</code>	An affinity mask indicating what CPU(s) the virtual machine processes should be run on. You can specify a list of CPUs identified by their index numbers separated by commas (0, 1, 2, 3, etc.) or a range (4-6). To make all CPUs available for the virtual machine processes specify <code>--cpumask all</code> .
<code>--offline-management <on off></code>	Turns the offline management on or off.
<code>--offline-service <service_name></code>	The name of the service to use for offline management.
<code>--userpasswd <user>:<passwd></code>	Sets the password for the specified user in the virtual machine. If the user account does not exist, it will be created. OpenVZ tools must be installed in the virtual machine for the command to work.
<code>--cryptd</code>	Used with <code>--userpasswd</code> . Indicates that the specified password is already a hash.
<code>--rate <rate></code>	Sets the guaranteed outgoing traffic rate in Kbps for the virtual machine.
<code>--ratebound <on off></code>	Turns the network traffic rate limitation set by the <code>--rate</code> parameter (above) on or off. The default value is off.
<code>--apply-iponly <yes no></code>	If set to yes, the hostname, nameserver, and search domain settings from the virtual machine configuration file are ignored.
<code>--efi-boot <on off></code>	If set to on, the virtual machine will boot using the EFI firmware. If set to off (default), the virtual machine will boot using the BIOS firmware.
<code>--tools-autoupdate <on off></code>	Enables or disables automatic update of OpenVZ tools inside Windows virtual machines. If set to on, OpenVZ tools are updated automatically on user

Name	Description
	log in. A reboot is required to complete the update. If set to off, OpenVZ tools are not updated automatically, so that you can do it manually at a convenient time.

4.1.16.2. Managing Virtual Devices

The `prctl set` command allows to add, modify, and delete virtual devices of virtual machines.

Syntax

```
prctl set <VM_name> --device-add <dev_type> <options>
prctl set <VM_name> --device-set <dev_name> <options>
prctl set <VM_name> --device-del <dev_name> <options> --destroy-image-force
prctl set <VM_name> --device-connect <dev_name>
prctl set <VM_name> --device-disconnect <dev_name>
prctl set <VM_name> --device-bootorder "dev_name1 dev_name2 [...]"
```

Table 4.17. Options

Name	Description
<VM_name>	Virtual machine name.
--device-add <dev_type> <options>	Adds a virtual device of the type <dev_type> to a virtual machine. The <dev_type> parameter can be: hdd, cdrom, net, fdd, serial, usb, pci. Except for SCSI and VirtIO hard disks, devices can only be added to stopped virtual machines.
--device-set <dev_name> <options>	Modifies the configuration of the virtual device <dev_name> in a stopped virtual machine.
--device-del <dev_name> <options>	Deletes the virtual device <dev_name> from a stopped virtual machine.
--destroy-image-force	Used with the --device-del option. Deletes a virtual machine HDD even if it is used in that virtual machine's snapshots.
--device-connect <dev_name>	Connects the virtual device <dev_name> to a running virtual machine.
--device-disconnect +<dev_name>	Disconnects the virtual device <dev_name> from a running virtual machine.
--device-bootorder "dev_name1 dev_name2 [...]"	Specifies the boot order for a virtual machine.

Note: Device names can be obtained with the `prctl list -i` command.

The device-related <options> can be subdivided into the following categories:

- hard disk drives
- optical disk drives

- network cards
- floppy disk drives
- serial ports
- USB devices

Each group of options is explained in the following subsections in detail.

Hard Disk Drive Management Options

This group of options is used to add and configure virtual hard disks in a virtual machine. The first syntax uses a file to emulate a hard disk drive. The second syntax connects a physical hard disk on the host server to the virtual machine.

Syntax

```
prlctl set +<VM_name> [--device-add hdd | --device-set hdd<N>]
    [--image <file>] [--type <expanded|plain>] [--size <size>] [--split]
    [--iface <ide|scsi|virtio>] [--position <pos>] [--enable|--disable]
prlctl set <VM_name> --device-add hdd --device <dev_name> [--position <pos>]
    [--iface <ide|scsi|virtio>]
prlctl set <VM_name> --backup-add <backup_ID> [--disk <disk_name>]
    [--iface <ide|scsi|virtio>] [--position <pos>]
prlctl set <VM_name> --backup-del {<backup_ID>|all}
```

Table 4.18. Options

Name	Description
<VM_name>	Virtual machine name.
--device-add hdd	Adds a virtual hard disk to the VM. New hard disks are created in the virtual machine directory and are automatically named <code>harddisk<N>.hdd</code> , where <code><N></code> is the next available disk index. SCSI and VirtIO hard disks can be added to both running and stopped VMs, IDE disks can only be added to stopped VMs.
--device-set hdd<N>	Modifies the parameters of an existing virtual hard disk. Virtual hard disks are named using the <code>hdd<N></code> format where <code><N></code> is the drive index number starting from 0 (e.g., <code>hdd0</code> , <code>hdd1</code>). To obtain the list of disk names, use the <code>prlctl list</code> command with the <code>--info</code> option.
--image <file>	Specifies an image file that will be used to emulate the virtual disk. • If the specified image does not exist, it is created and used to emulate the virtual hard disk. • If the specified image exists, it is used to emulate the virtual hard disk.
--device <dev_name>	This option is used to connect a physical hard disk on the hardware node to the virtual machine. You can obtain the names of the existing hard disks on the server using the <code>prlsrvctl info</code> command.
--size <size>	The size of the virtual hard disk, in megabytes. The default size is 65536 MB.
--split	Splits the hard disk image file into 2 GB pieces. You should split a virtual disk if it is stored on a file system that cannot support files larger than 2 GB (e.g., FAT16).

Name	Description
--enable	Enables the specified virtual disk drive. All newly added disk drives are enabled by default (provided the --disable option is omitted).
--disable	Disables the specified virtual disk drive. The disk drive itself is not removed from the virtual machine configuration.
--backup-add <backup_ID>	Attach the backup with the identifier <backup_ID> to the virtual machine as a virtual hard disk. To obtain the backup ID, use the <code>prlctl backup-list -f</code> command.
--disk <disk_name>	Used with --backup-add. The name of the disk in the backup to attach. If a disk is not specified, all disks contained in the backup will be attached. To obtain the disk name(s), use the <code>prlctl backup-list -f</code> command.
+--backup-del {<backup_ID> all}	Detach either the backup with the identifier <backup_ID> or detach all backups from the virtual machine.
--iface <ide scsi virtio>	The disk drive Interface type. If omitted, the SCSI interface will be used.
--position <pos>	The SCSI or IDE device identifier to be used for the virtual disk.

Optical Disk Drive Management Options

This group of options is used to add and configure virtual optical disk drives, such as DVD or CD drives.

Syntax

```
prlctl set <VM_name> {--device-add cdrom | --device-set cdrom<N>}
                    {--device <dev_name> | --image <file>} [--iface <ide|scsi>]
                    [--position <pos>] [--enable|--disable] [--connect|--disconnect]
```

Table 4.19. Options

Name	Description
<VM_name>	Virtual machine name.
--device-add cdrom	Adds a DVD/CD drive to the virtual machine.
--device-set cdrom<N>	Modifies the parameters of an existing virtual optical disk. The <N> postfix indicates the drive index number. To obtain the list of the available drives, use the <code>prlctl list</code> command with the --info option.
--device <dev_name>	The name of the physical optical disk to connect to the virtual machine.
--image <file>	The name of an existing disk image file to mount in the virtual machine. Currently, the following image file formats are supported: .iso, .cue, .ccd, and .dmg. The image must not be compressed and/or encrypted.
--iface <ide scsi>	Interface type: - ide - IDE disk. - scsi - SCSI disk (default).
--position <pos>	The SCSI or IDE device identifier to be used for the DVD/CD drive. You can use one of the following formats for specifying IDs: <ID>:._<bus>_, <ID>-<bus>, <ID>. For example, if you specify 3:0 (or 3-0 or 3) as <number>

Name	Description
	for a SCSI drive, the guest OS will see the drive as having ID 3 on SCSI bus 0.
--enable	Enables the specified DVD/CD drive. All newly added drives are enabled by default (provided the --disable option is omitted).
--disable	Disables the specified optical disk drive. The disk drive itself is not removed from the virtual machine configuration.
--connect	Automatically connect the specified optical disk drive during the virtual machine startup process.
--disconnect	Do not automatically connect the specified optical disk drive during the virtual machine startup process.

Network Adapter Management Options

This group of options is used to manage virtual network adapters in a virtual machine.

Syntax

```
prlctl set <VM_name> {--device-add net | --device-set net<N>}
    [--type routed | --network <network_ID>} [--mac {<addr>|auto}]
    [{--ipadd <addr>[/<mask>] | --ipdel <addr>[/<mask>] | --dhcp <yes|no> |
    --dhcp6 <yes|no>}] [--gw <gw>] [--gw6 <gw>]
    [--nameserver <addr>] [--searchdomain <addr>]
    [--configure <yes|no>] [--ipfilter <yes|no>] [--macfilter <yes|no>]
    [--preventpromisc <yes|no>] [--enable|--disable]
    [--connect|--disconnect] [--adapter-type <e1000|rtl|virtio>]
```

Table 4.20. Options

Name	Description
<VM_name>	Virtual machine name.
--device-add net	Adds a new virtual network adapter to the virtual machine.
--device-set net<N>	Modifies an existing virtual network adapter. To obtain the list of the available adapters, use the prlctl list command with the --info option.
--type routed	Sets the networking mode for the virtual network adapter to "routed". In this mode, the network adapter is communicating with the outside world through an internal virtual network adapter.
--network <network_ID>	Sets the networking mode for the virtual network adapter to "virtual_network". In this mode the adapter is connected to a virtual network specified by <network_ID>.
--mac {<addr> auto}	Specifies the MAC address to assign to an existing network adapter. Specify a desired MAC address using the addr parameter value or use the auto option to generate the existing address automatically.
--ipadd <addr>[/<mask>]	Adds an IP address and a mask (optional) to the network adapter.
--ipdel <addr>[/<mask>]	Deletes an IP address from the network adapter.

Name	Description
--dhcp <yes no>	Specifies whether the virtual network adapter should obtain the IPv4 settings through a DHCP server.
--dhcp6 <yes no>	Specifies whether the virtual network adapter should obtain the IPv6 settings through a DHCP server .
--gw <gw>	The default gateway to be used by the virtual machine.
--gw6 <gw>	The default IPv6 gateway to be used by the virtual machine.
--nameserver <addr>	The default DNS server address to be used by the virtual machine.
--searchdomain <addr>	The default search domain to be used by the virtual machine.
--configure <yes no>	If set to yes, the settings above are applied to the virtual network adapter instead of its original settings. Configuring any of the settings above automatically sets this option to yes.
--ipfilter <yes no>	Determines if the specified network adapter is configured to filter network packages by IP address. If set to yes, the adapter is allowed to send packages only from IPs in the network adapter IP addresses list.
--macfilter <yes no>	Determines if the specified network adapter is configured to filter network packages by MAC address. If set to yes, the adapter is allowed to send packages only from its own MAC address.
--preventpromisc <yes no>	Determines if the specified network adapter should reject packages not addressed to its virtual machine. If set to yes, the adapter will drop such packages.
--enable --disable	Enables or disable the network adapter. If omitted during the adapter creation, the adapter will be enabled.
--connect --disconnect	Connects or disconnects the network adapter. When disconnected, the adapter is not removed from the virtual machine.
--adapter-type <e1000 rtl virtio>	Emulated network adapter: • e1000 - Intel 82545EM, • rtl - Realtek RTL8029, • virtio - VirtIO. Note: The adapter requires no additional configuration on supported Linux and FreeBSD guest operating systems. However, additional drivers need to be installed on Windows Server 2012 R2 guest OSes. For the drivers, visit https://alt.fedoraproject.org/pub/alt/virtio-win/latest/. The VirtIO adapter is not supported on the Windows Server 2008 guest operating systems.

Floppy Disk Drive Management Options

This group of options is used to add a floppy disk drive to a virtual machine and to modify the existing virtual floppy disk drive.

Syntax

```
prlctl set <VM_name> {--device-add fdd | --device-set fdd0}
                    {--device <dev_name> | --image <file>}
                    [--enable|--disable] [--connect|--disconnect]
```

Table 4.21. Options

Name	Description
<VM_name>	Virtual machine name.
--device-add fdd	Adds a new floppy disk drive to the virtual machine.
--device-set fdd0	Modifies the parameters of the existing virtual floppy disk drive.
--device <dev_name>	The name of the physical floppy disk drive to connect to the virtual machine. If this parameter is omitted, a floppy drive image emulating the floppy disk drive will be created.
--image <file>	The name and path of an existing floppy disk image file (usually floppy.fdd) to mount in the virtual machine.
--enable	Enables the specified floppy disk drive. All newly added floppy drives are enabled by default (provided the --disable option was omitted during the drive creation).
--disable	Disables the specified floppy disk drive. The drive itself is not removed from the virtual machine configuration.
--connect	Connect the specified floppy disk drive automatically during the virtual machine startup process.
--disconnect	Use this option if you don't want the specified floppy disk drive automatically connected to the virtual machine on its start.

Serial Port Management Options

This group of options is used to manage serial ports in a virtual machine.

Syntax

```
prlctl set <VM_name> {--device-add serial | --device-add serial<N>}
                    {--device <dev_name> | --output <file> | --socket <name>}
                    [--enable|--disable] [--connect|--disconnect]
```

Table 4.22. Options

Name	Description
<VM_name>	Virtual machine name.
--device-add serial	Adds a new serial port to the virtual machine.
--device-set serial<N>	Modifies the parameters of an existing serial port.
--device <dev_name>	The name of the physical serial port to which to connect the virtual machine.
--output <file>	The name and path of the output file to which to connect the virtual serial port.
--socket <name>	The name of the physical socket to which to connect the virtual serial port.

Name	Description
--enable --disable	Enables or disables the virtual serial port. All newly added serial ports are enabled by default (provided the --disable option is omitted).
--connect	Automatically connect the virtual serial port during the virtual machine startup process.
--disconnect	Do not automatically connect the virtual serial port during the virtual machine startup process.

USB Controller Management Options

This group of options is used to manage the USB controller in a virtual machine.

Syntax

```
prlctl set <VM_name> --device-add usb [--enable|--disable]
```

Table 4.23. Options

Name	Description
<VM_name>	Virtual machine name.
--device-add usb	The type of the virtual device to add to the virtual machine (in this instance, a USB device).
--enable --disable	Enables or disables the USB controller. The controller is enabled by default (provided the --disable option is omitted).

Removing Devices from Virtual Machines

The --device-del option is used to remove virtual devices from a virtual machine.

Syntax

```
prlctl set <VM_name> --device-del <dev_name> [--detach-only|--destroy-image]
```

Table 4.24. Options

Name	Description
<dev_name>	The name of the virtual device to delete from the virtual machine. To obtain the list of virtual devices, use the <code>prlctl list -i</code> command.
--detach-only	Deletes the information about the specified device from the virtual machine configuration.
--destroy-image	Deletes the information about the specified device from the virtual machine configuration and removes the device from the server.

4.1.17. prlctl snapshot, snapshot-list, snapshot-switch, snapshot-delete

Takes, displays, reverts to, and deletes snapshots of a running virtual machine.

Syntax

```
prlctl snapshot <VM_name> [-n, --name <name>] [-d, --description <desc>]
prlctl snapshot-list <VM_name> [-t, --tree] [-i, --id <snapshot_ID>]
prlctl snapshot-switch <VM_name> -i, --id <snapshot_ID>
prlctl snapshot-delete <VM_name> -i, --id <snapshot_ID>
```

Table 4.25. Options

Name	Description
<VM_name>	Virtual machine name.
-n, --name <name>	User-defined snapshot name. Names with white spaces must be enclosed in quotation marks.
-d, --description <desc>	User-defined snapshot description. Descriptions with white spaces must be enclosed in quotation marks.
-t, --tree	Displays the snapshot list as a tree. The default display format is tabular with Parent Snapshot ID and Snapshot ID as columns.
-i, --id <snapshot_ID>	- Use with <code>prlctl snapshot-list</code> to specify the ID of the snapshot to use as the root. If this parameter is omitted, the entire snapshot tree will be displayed. - Use with <code>prlctl snapshot-switch</code> to specify the ID of the snapshot to revert to. - Use with <code>prlctl snapshot-delete</code> to specify the ID of the snapshot to delete.

Note: If the snapshot you want to delete has child snapshots derived from it, they will not be deleted.

4.1.18. prlctl start, stop, restart, reset, status

Start, stop, reset, and check the status of a virtual machine.

Syntax

```
prlctl start <VM_name>
prlctl stop <VM_name> [--kill]
prlctl restart <VM_name>
prlctl reset <VM_name>
prlctl status <VM_name>
```

Table 4.26. Options

Name	Description
<VM_name>	The name of the virtual machine to start, stop, restart, reset, or check the status of.
--kill	Perform a hard virtual machine shutdown. If this option is omitted, an attempt to perform a graceful shutdown will be made.

The `stop` command can perform a hard or a graceful virtual machine shutdown. If the `--kill` parameter is included, the hard shutdown will be performed. If the parameter is omitted, the outcome of the graceful shutdown attempt will depend on the following:

- If OpenVZ tools are installed in a virtual machine, the graceful shutdown will be performed using its facilities.
- If OpenVZ tools are not installed, the command will try to perform a graceful shutdown using ACPI. Depending on the ACPI support availability in the guest operating system, this may work or not.

The `restart` command first gracefully shuts down a virtual machine and then starts it again.

The `reset` command resets a virtual machine without shutting it down.

Note: Resetting a VM may result in loss of unsaved data stored in that VM.

The `start` command can be used to start a stopped virtual machine or to resume a paused virtual machine).