



Parallels Cloud Server 6.0

Command Line Reference Guide

July 10, 2014

Parallels IP Holdings GmbH

Vordergasse 59

8200 Schaffhausen

Switzerland

Tel: + 41 52 632 0411

Fax: + 41 52 672 2010

www.parallels.com

Copyright © 1999-2014 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.parallels.com/trademarks>.

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

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

Linux is a registered trademark of Linus Torvalds.

All other marks and names mentioned herein may be trademarks of their respective owners.

Contents

| | |
|---|-----------|
| Introduction | 8 |
| About Parallels Cloud Server 6.0 | 8 |
| About This Guide | 9 |
| Organization of This Guide | 9 |
| Documentation Conventions | 10 |
| Getting Help..... | 11 |
| Feedback..... | 11 |
| Managing Parallels Cloud Server 6.0..... | 12 |
| Parallels Cloud Server Configuration Files..... | 12 |
| Global Parallels Cloud Server Configuration File..... | 13 |
| Container Configuration File..... | 19 |
| Linux Distribution Configuration Files | 26 |
| Memory and IOPS Deduplication Configuration File..... | 28 |
| Network Classes Definition File..... | 28 |
| vzvpn Configuration File..... | 29 |
| Kernel Parameters | 30 |
| Offline Management Configuration Files..... | 31 |
| vzlmnd Configuration File..... | 32 |
| pstat Configuration File | 33 |
| vzrmnd Configuration File | 35 |
| vzstatrep Configuration File | 37 |
| Backup Configuration File..... | 38 |
| vzpkgproxy Configuration File..... | 41 |
| vztt Configuration File | 41 |
| pcompact.conf | 42 |
| shaman.conf | 42 |
| Parallels Cloud Server Scripts | 44 |
| Overview | 44 |
| Container Action Scripts..... | 45 |
| Virtual Machine Action Scripts | 46 |
| High Availability Scripts..... | 47 |

| | |
|---|-----------|
| Parallels Cloud Server Utilities | 47 |
| prlsrvctl..... | 48 |
| vzlicload | 58 |
| vzlicupdate | 59 |
| vzlicview | 60 |
| vznetcfg..... | 60 |
| vzstatrep | 62 |
| pstat..... | 63 |
| vzreboot | 69 |
| High Availability Utilities..... | 69 |
| shaman General Syntax..... | 69 |
| shaman join | 70 |
| shaman leave | 70 |
| shaman stat | 71 |
| shaman top | 71 |
| Parallels Cloud Server Updates..... | 72 |
| Managing Containers | 73 |
| Matrix of Parallels Cloud Server Command-Line Utilities | 73 |
| prlctl | 75 |
| prlctl console | 76 |
| prlctl create | 76 |
| prlctl delete, destroy | 78 |
| prlctl exec, enter | 78 |
| prlctl mount, umount | 79 |
| prlctl move..... | 79 |
| prlctl problem-report..... | 80 |
| prlctl register, unregister | 80 |
| prlctl reinstall | 81 |
| prlctl set | 82 |
| prlctl snapshot, snapshot-list, snapshot-switch, snapshot-delete..... | 95 |
| prlctl start, stop, restart, status | 96 |
| prlctl suspend, resume | 97 |
| prlctl list | 98 |
| Migration Utilities | 99 |
| pmigrate..... | 99 |

| | |
|--|-----|
| prctl clone..... | 102 |
| Backup and Restoration Utilities | 103 |
| pbackup..... | 103 |
| prestore..... | 105 |
| prctl backup, backup-list, backup-delete, restore..... | 106 |
| EZ Template Management Utilities..... | 108 |
| vzpkg | 108 |
| vzpkg install template | 109 |
| vzpkg update template | 109 |
| vzpkg remove template | 110 |
| vzpkg list | 110 |
| vzpkg info..... | 111 |
| vzpkg status | 114 |
| vzpkg install..... | 115 |
| vzpkg update | 116 |
| vzpkg remove..... | 117 |
| vzpkg create cache | 118 |
| vzpkg update cache | 119 |
| vzpkg remove cache | 119 |
| vzpkg create appcache | 120 |
| vzpkg update appcache..... | 120 |
| vzpkg remove appcache | 121 |
| vzpkg localinstall..... | 122 |
| vzpkg localupdate | 122 |
| vzpkg upgrade | 123 |
| vzpkg fetch..... | 124 |
| vzpkg clean | 125 |
| vzpkg update metadata..... | 125 |
| vzmktmpl..... | 126 |
| vzpkgproxy..... | 130 |
| vzcreaterepo | 130 |
| vzmtemplate..... | 131 |
| Supplementary Tools..... | 132 |
| pcompact..... | 132 |
| pfcache | 133 |

| | |
|--|------------|
| pnetstat | 136 |
| prl_disk_tool | 137 |
| vzcalc | 138 |
| vzcfgscale | 138 |
| vzcfgvalidate..... | 139 |
| vzcheckovr | 140 |
| vzcpucheck | 140 |
| vzctl convert | 140 |
| vzhwcalc | 141 |
| vziotop | 142 |
| vzmemcheck | 142 |
| vzpid | 142 |
| vzps, vztop | 143 |
| vzsplit | 143 |
| Managing Virtual Machines | 145 |
| prctl | 145 |
| General Syntax | 145 |
| prctl capture | 146 |
| prctl change-passwd | 146 |
| prctl clone..... | 147 |
| prctl convert..... | 147 |
| prctl create | 148 |
| prctl delete..... | 149 |
| prctl encrypt, decrypt..... | 149 |
| prctl enter | 150 |
| prctl exec..... | 150 |
| prctl list | 150 |
| prctl migrate | 152 |
| prctl mount, umount | 152 |
| prctl move..... | 153 |
| prctl pause, suspend, resume..... | 153 |
| prctl problem-report..... | 154 |
| prctl register, unregister | 154 |
| prctl reset-uptime | 155 |
| prctl set | 155 |

| | |
|--|------------|
| prctl snapshot, snapshot-list, snapshot-switch, snapshot-delete..... | 169 |
| prctl server..... | 170 |
| prctl start, stop, restart, reset, status | 170 |
| pmigrate | 171 |
| Managing Virtual Machine Backups | 174 |
| pbackup..... | 174 |
| prestore..... | 175 |
| prctl backup, backup-list, backup-delete, restore..... | 175 |
| prl_convert..... | 177 |
| prl_disk_tool..... | 178 |
| prl_disk_tool compact | 178 |
| prl_disk_tool configure..... | 178 |
| prl_disk_tool convert | 179 |
| prl_disk_tool encrypt, decrypt..... | 179 |
| prl_disk_tool merge | 180 |
| prl_disk_tool resize | 180 |
| prl_disk_tool validate | 180 |
| pnetstat | 181 |
| Glossary | 182 |
| Index | 184 |

CHAPTER 1

Introduction

Parallels Cloud Server 6.0 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 Parallels Cloud Server and this guide.

- [Parallels Cloud Server basics](#) (p. 8)
- [goals and target audience of the guide](#) (p. 9)
- [guide organization](#) (p. 9)
- [documentation conventions used in the guide](#) (p. 10)
- [resources to consult to get more information on Parallels Cloud Server](#) (p. 11)
- [ways to submit feedback to the Parallels documentation team](#) (p. 11)

In This Chapter

| | |
|---------------------------------------|----|
| About Parallels Cloud Server 6.0..... | 8 |
| About This Guide | 9 |
| Getting Help..... | 11 |
| Feedback..... | 11 |


About Parallels Cloud Server 6.0

Parallels Cloud Server 6.0 is a virtualization solution that allows you to simultaneously run multiple Parallels virtual machines and Containers on a single physical server. With Parallels Cloud Server, you can efficiently share your server's hardware resources among virtual machines and Containers.

Parallels Cloud Server is installed directly on the server hardware and does not need any operating system to function. Once it is installed, Parallels Cloud Server allows you to create virtual machines and Containers and manage them using the following tools:

- **Parallels command-line interface (CLI).** The command-line interface comprises a set of Parallels command-line utilities that you can use to manage virtual machines and Containers, both locally and remotely.
- **Parallels Virtual Automation (PVA).** Parallels Virtual Automation is a remote management tool that allows you to manage physical servers and their virtual machines and Containers with the help of a standard Web browser on any platform.

Graphically, a server with the Parallels Cloud Server software installed can be represented as follows:



About This Guide

This guide is a complete reference on all Parallels Cloud Server configuration files and command-line utilities. It familiarizes you with the way to configure Parallels Cloud Server 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.

Organization of This Guide

This guide is organized in the following way:

Chapter 1, Introduction (p. 8), gives an overview of Parallels Cloud Server and this guide.

Chapter 2, Managing Parallels Cloud Server 6.0 (p. 12), provides instructions on Parallels Cloud Server configuration files, scripts, and command-line utilities.

Chapter 3, Managing Containers (p. 73), describes Parallels Cloud Server command-line utilities that can be used for managing your Containers.

Chapter 4, Managing Virtual Machines (p. 145), focuses on Parallels Cloud Server utilities that can be used for managing your virtual machines.

Documentation Conventions

Before you start using this guide, it is important to understand the documentation conventions used in it.

The table below presents the existing formatting conventions.

| Formatting convention | Type of Information | Example |
|-----------------------|---|---|
| Special Bold | Items you must select, such as menu options, command buttons, or items in a list. | Go to the Resources tab. |
| | Titles of chapters, sections, and subsections. | Read the Basic Administration chapter. |
| <i>Italics</i> | Used to emphasize the importance of a point, to introduce a term or to designate a command-line placeholder, which is to be replaced with a real name or value. | These are the so-called <i>OS templates</i> . To remove a Container, type <code>prlctl delete <i>ctid</i></code> . |
| Monospace | The names of commands, files, and directories. | Use <code>prlctl start</code> to start a Container. |
| Preformatted | On-screen computer output in your command-line sessions; source code in XML, C++, or other programming languages. | <code>Saved parameters for Container 101</code> |
| Monospace Bold | What you type, as contrasted with on-screen computer output. | <code>C:\vzlist -a</code> |
| Key+Key | Key combinations for which you must press and hold down one key and then press another. | Ctrl+P, Alt+F4 |

Besides the formatting conventions, you should also know about the document organization convention applied to Parallels documents: chapters in all guides are divided into sections, which, in their turn, are subdivided into subsections. For example, **About This Guide** is a section, and **Documentation Conventions** is a subsection.

Getting Help

In addition to this guide, there are a number of other resources available for Parallels Cloud Server which can help you use the product more effectively. These resources include:

- *Parallels Cloud Server 6.0 Installation Guide*. This guide provides detailed information on installing Parallels Cloud Server on your server, including the pre-requisites and the stages you shall pass.
- *Installing via PXE*. This guide provides information on installing Parallels Cloud Server 6.0 over a network using a PXE (Preboot Execution Environment) server.
- *Getting Started With Parallels Cloud Server 6.0*. This guide provides basic information on how to install Parallels Cloud Server on your server, create new Containers and virtual machines, and perform main operations on them. As distinct from the *Parallels Cloud Server 6.0 Installation Guide*, it does not contain detailed description of all the operations needed to install and set Parallels Cloud Server to work (e.g., installing Parallels Cloud Server in the text mode).
- *Parallels Cloud Server 6.0 User's Guide*. This guide provides comprehensive information on Parallels Cloud Server covering the necessary theoretical conceptions as well as all practical aspects of working with the product. However, it does not deal with the process of installing and configuring your system.
- *Parallels Cloud Storage Administrator's Guide*. This guide is intended for system administrators interested in deploying Parallels Cloud Storage in their networks.
- *Parallels Cloud Server 6.0 Templates Management Guide*. This guide is meant to provide complete information on Parallels templates - an exclusive Parallels technology allowing you to efficiently deploy standard Linux applications inside your Containers and to greatly save the physical server resources (physical memory, disk space, etc.).

Feedback

If you spot a typo in this guide, or if you have an opinion about how to make this guide more helpful, you can share your comments and suggestions with us by completing the Documentation Feedback form on our [website](http://www.parallels.com/en/support/usersdoc/) (<http://www.parallels.com/en/support/usersdoc/>).

CHAPTER 2

Managing Parallels Cloud Server 6.0

This chapter provides instructions on configuration files, scripts, and command-line utilities that can be used to configure the settings related to the Parallels Cloud Server software and the Parallels server.

In This Chapter

| | |
|---|----|
| Parallels Cloud Server Configuration Files..... | 12 |
| Parallels Cloud Server Scripts | 44 |
| Parallels Cloud Server Utilities | 47 |
| High Availability Utilities..... | 69 |
| Parallels Cloud Server Updates..... | 72 |

Parallels Cloud Server Configuration Files

The table below lists the configuration files available in Parallels Cloud Server 6.0. Most files are located in the `/etc` directory on a Parallels server. Some configuration files are stored on other servers (e.g., on a Backup Node). If a configuration file is stored in a place other than the Parallels server, 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_ID>.conf</code> | Private configuration file of a Container with the ID <code><CT_ID></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 Parallels Cloud Server (e.g., <code>basic</code> and <code>confixx</code>), but you can also create your own samples to meet your demands. |
| <code>/etc/vz/conf/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 Parallels Cloud Server 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. |
| /etc/vz/oom-groups.conf | OOM killer configuration file with task badness adjustments. |
| /etc/vz/conf/networks_classes | Configuration file defining the network classes for traffic shaping and bandwidth management. |
| /etc/cron.d/verebboot | Configuration file for the <code>cron</code> daemon. Using this file, Parallels Cloud Server emulates the “reboot” command running in a Container. |
| /etc/vzvpn/vzvpn.conf | Configuration file defining the parameters for establishing a private secure channel to the Parallels support server. |
| /etc/sysctl.conf | Kernel parameters. Parallels Cloud Server adjusts a number of kernel <code>sysctl</code> parameters and modifies the default <code>/etc/sysctl.conf</code> file. |
| /etc/vzredirect.d/*.conf | Configuration files defining offline management modes administrators use to access their Containers. |
| /etc/vzlmond.conf | Configuration file the <code>vzlmond</code> daemon uses to collect information on how much resources the server consumes at the moment. |
| /etc/vzstat.conf | Configuration file listing the warning and error levels for resource control parameters. If a parameter hits the warning or error value set for it, the <code>pstat</code> utility will display this parameter in yellow or red. |
| /etc/vzbackup.conf | Configuration file located on the Backup Node and determining the global backup settings for Containers. Note that the settings in this file do not apply to virtual machines. |
| /etc/vzpkgproxy/vzpkgproxy.conf | Configuration file used by the <code>vzpkgproxy</code> utility to create special caching proxy servers for OS and application EZ templates. The file can be located on any computer where the <code>vzpkgproxy</code> package is installed. |
| /etc/vztt/vztt.conf | Configuration file used by the <code>vzpkg</code> utility to manage OS and application EZ templates. |
| /etc/vzstatrep.conf | Configuration file stored on the Monitor Node and used by the <code>vzstatrep</code> utility when generating statistic reports and graphics on the Parallels server resource consumption and sending these reports to the server administrator. |
| /etc/shaman/shaman.conf | Configuration file the <code>shaman</code> utility uses to manage and monitor High Availability clusters. |

Global Parallels Cloud Server Configuration File

Parallels Cloud Server keeps its system wide configuration parameters in the `/etc/vz/vz.conf` configuration file. This file is in shell format. Keep in mind that Parallels Cloud Server 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 Parallels Cloud Server.

| Name | Description | Default value |
|----------------------------|---|-----------------|
| VIRTUOZZO | This can be either “yes” or “no”. Parallels Cloud Server System V startup script checks this parameter. If set to “no”, then Parallels Cloud Server 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 Parallels Cloud Server utilities communicating with the outer world through HTTP will use this server for managing all your HTTP messages. | - |
| ACTIONLOGDIR | This is the directory where <code>prlctl</code> keeps a log of its actions in the format suitable for Parallels Cloud Server 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. Parallels Cloud Server keeps lock files in this directory in order to serialize access to one Container. | /vz/lock |
| REMOVEMIGRATED | Specifies whether the private area and the configuration file of the Container moved to a new server with the <code>pmigrate</code> command should be destroyed on the Source Server (the value of the parameter is set to <code>yes</code>) or renamed to have the <code>.migrated</code> suffix (the value of the parameter is set to <code>no</code>). You may wish to leave the Container private area and the configuration file to make migration faster. This configuration value can be overridden by the <code>pmigrate</code> command-line options. | no |
| VE0CPUUNITS | CPU weight designated for the server itself. | 1000 |
| OFFLINE_MANAGEMENT | Specifies whether Containers can be managed by the Container administrator by means of the services indicated in the <code>OFFLINE_SERVICE</code> parameter. | yes |
| OFFLINE_SERVICE | These services correspond to the names of the files in the <code>/etc/vzredirect.d</code> directory, each file defining at what port the service will be accessible and to what Container the requests coming to this port will be redirected. These services will be accessible to those Containers which have the <code>OFFLINE_MANAGEMENT</code> parameter set to “yes”. | vzpp-plesk vzpp |
| BURST_CPU_AVG_USAGE | The CPU usage limit, in percent, set for the Container. This limit is calculated as the ratio of the current Container CPU usage to the CPU limit (i.e to the value of the <code>CPULIMIT</code> parameter) set for the Container in its configuration file. If the limit is not specified, the full CPU power of the server is considered as the CPU limit. Upon exceeding the <code>BURST_CPU_AVG_USAGE</code> limit, the | disabled |

| | | |
|---------------------------|--|-------------------------------|
| | <p>BURST_CPULIMIT limit is applied to the given Container.</p> <p>This parameter can be redefined by the BURST_CPU_AVG_USAGE parameter set in the Container configuration file.</p> | |
| BURST_CPULIMIT | <p>The CPU power limit, in percent, the Container cannot exceed. The limitations set in this parameter are applied to any Container exceeding the limit specified in the BURST_CPU_AVG_USAGE parameter.</p> <p>This parameter can be redefined by the BURST_CPULIMIT parameter set in the Container configuration file.</p> | |
| VEFSTYPE | <p>File system to use when caching OS templates:</p> <ul style="list-style-type: none"> ▪ ext4: Use to create Containers with the new layout. ▪ vz4: Use to create legacy Containers. | ext4 |
| VZMOUNTS | <p>Defines the partitions which will be automatically mounted by the <code>/etc/init.d/vz</code> script after the server boot. This script will check (by calling the <code>fsck</code> utility) and mount all the partitions specified as the value of this parameter, listed in <code>/etc/fstab</code> file on the server, and having the <code>noauto</code> flag set for them in this file.</p> | /vz |
| IPV6 | <p>Defines whether the IPv6 support is enabled on the Parallels server.</p> | yes |
| GOLDEN_IMAGE | <p>Enables (yes) or disables (no) embedding application templates into OS EZ template cache prior to creating Containers based on this cache.</p> | yes |
| PFCACHE | <p>Path to the memory and IOPS deduplication cache with common Container files.</p> | /vz/pfcache |
| PFCACHE_IMAGE | <p>Path to the private area of the memory and IOPS deduplication cache.</p> | /vz/pfcache.hdd |
| PFCACHE_IMAGE_SIZE | <p>Image size (in 1-KB blocks) of the memory and IOPS deduplication cache.</p> | 10485760 |
| PFCACHE_INCLUDES | <p>Directories for which memory and IOPS deduplication is enabled by default.</p> | bin lib lib64 opt sbin usr |
| VZ_TOOLS_MEMLIMIT | <p>Limits the total memory consumption of all vz utilities, e.g., vzbackup, vzmigrate, etc. Set in bytes.</p> | |

Logging parameters

| Name | Description | Default value |
|------------------|---|--------------------|
| LOGGING | This parameter defines whether <code>prlctl</code> should log its actions. | yes |
| LOGFILE | File where <code>prlctl</code> logs its actions. | /var/log/vzctl.log |
| LOG_LEVEL | There are three levels of logging defined in the current version of Parallels Cloud Server. | 0 |

Possible values of the LOG_LEVEL parameter

| Log level | Information to be logged |
|-----------|--|
| 0 | Actions of <code>prlctl</code> on Containers like <code>start</code> , <code>stop</code> , <code>create</code> , <code>destroy</code> , <code>mount</code> , <code>umount</code> . |
| 1 | Level 1 logs events, calls to <code>prlctl</code> helper scripts located in <code>/etc/vz/conf</code> (such as <code>vz-start</code> and <code>vz-stop</code>) and situations when the init process of the Container is killed on Container stop after timeout. |
| 2 | Level 0 and level 1 logging events, plus template version used for Container creation and calls to mount and quota operations with parameters. |

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 |

SLM parameters

Note: In Parallels Containers for Windows 6.0, the SLM system was superseded by the new VSwap memory management scheme, and the SLM parameters are left for compatibility reasons only.

| Name | Description | Default value |
|------------|--|---------------|
| SLM | <p>If set to "yes", the SLM modules are loaded to the server. It means that the <code>slmmemorylimit</code> parameter is supported and can be used to manage the amount of memory consumed by every Container on the server.</p> <p>Note: After changing this parameter, restart the Parallels Cloud Server service for the changes to take effect.</p> | yes |
| SLMPATTERN | Defines the SLM pattern rules for grouping the processes running inside Containers on the server. The default rules are set in the <code>/etc/vzslm.d/default.conf</code> file on the server. | default |

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 "dev1:rate1 dev2:rate2". For example, for two 100 Mbps Ethernet cards, set it to "eth0:102400 eth1:102400". | eth0:102400 |
| TOTALRATE | This parameter sets the size of the bandwidth pool for all | eth0:1:4096 |

| | | |
|-------------------------|---|----------|
| | 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. | |
| 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 Container cannot consume less than this value and more than the sum of this value and <code>TOTALRATE</code> . | eth0:1:8 |
| VZ_TOOLS_BCID | Enables limits for <code>pbackup</code> , <code>prestore</code> , and <code>pmigrate</code> operations. Make sure that the value of this parameter does not coincide with the ID assigned to some of your Containers. For example, if you leave the default value 2, ensure that no Container with ID 2 exists on the Node. | |
| VZ_TOOLS_IOLIMIT | Sets the disk I/O limit for the backup, restore, and migration operations, in bytes per second. Not set by default. | |

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 the templates built by Parallels use the default directory. | /vz/template |

Container default parameters

| Name | Description | Default value |
|-----------------------|--|-------------------|
| VE_ROOT | Path to the Container root directory where the private area is mounted to. | /vz/root/CT_ID |
| VE_PRIVATE | Path to the Container private area. Must reside within a single physical partition. | /vz/private/CT_ID |
| CONFIGFILE | The default configuration file sample to be used for the Container creation; it may be overridden with the <code>--config</code> option of the <code>prlctl create</code> command. | basic |
| DEF_OSTEMPLATE | The default OS template to be used for the Container creation; it may be overridden with the <code>--pkgset</code> command-line option for <code>prlctl create</code> . | centos-6-x86_64 |
| VE_ENVIRONMENT | Additional environment variables to be passed to the Container <code>init</code> process. Should be provided as any number of <code>name=value</code> pairs separated by spaces. | |

Container validation and overcommitment parameters

| Name | Description | Default value |
|--|--|---------------|
| VE_VALIDATE_ACTION | Defines whether the Container configuration should be validated when a Container is started. If this parameter is set to "warning", a warning is displayed in case of misconfiguration. If set to "error", the Container is not started in case of misconfiguration. If set to "fix", the configuration is automatically corrected. | none |
| OVERCOMMITMENT_ACTION | Defines whether the server should be checked for the overcommitment of resources when a Container is started. If this parameter is set to "warning", a warning is displayed in case of overcommitment. If set to "error", the Container that would cause overcommitment is not started. When checking for overcommitment, the following five parameters are checked. | none |
| OVERCOMMITMENT_LEVEL_LOWMEM | The percentage of committed memory residing at lower addresses and directly accessed by the kernel. | 120 |
| OVERCOMMITMENT_LEVEL_MEMSWAP | The percentage of committed memory available for applications including both RAM and swap space. | 90 |
| OVERCOMMITMENT_LEVEL_ALLOCMEM | The allocation memory commitment level is the ratio of the memory size guaranteed to be available for allocation to the capacity of the system. | 100 |
| OVERCOMMITMENT_LEVEL_ALLOCMEM_TOT | The number shows how much memory the applications are allowed to allocate in comparison with the capacity of the system. | 1000 |
| OVERCOMMITMENT_LEVEL_ALLOCMEM_MAX | This allocation memory commitment level is the ratio of the maximal (among all running Containers) amount of allocated memory to the capacity of the system. | 60 |

Supplementary parameters

| Name | Description | Default value |
|--------------------|--|---------------|
| VZWDOG | Defines whether the vzwdog module is loaded on Parallels Cloud Server startup. This module is responsible for catching messages from the kernel. It is needed if you configure the serial Monitor Server for Parallels Cloud Server. | no |
| VZPRIVRANGE | Defines the ID range for the Containers that are allowed to access the <server> ID stored in the /proc/vz/hwid file. | 1 100 |

| | | |
|----------------|--|-------------------------------------|
| DUMPDIR | The directory where the Container dump file created by means of the <code>prlctl suspend</code> command is to be stored. | <code>/vz/private/CT_ID/dump</code> |
|----------------|--|-------------------------------------|

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_ID.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](#) (p. 19)
- [resource management parameters](#) (p. 21)
- [networking](#) (p. 24)
- [backup](#) (p. 25)

Note: In Parallels Cloud Server, you can also configure a number of settings for the server itself by editing the `/etc/vz/conf/0.conf` file. Currently, these settings include the `VERSION` and `ONBOOT` parameters, as well as all parameters listed in the table under the *System parameters* group.

Miscellaneous Parameters

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

| Name | Description |
|---------------------------|--|
| VERSION | Specifies the Parallels Cloud Server version the configuration file applies to. "2" relates to Parallels Cloud Server version 4 and later. |
| ONBOOT | Specifies whether the Container should be started automatically on system startup. Parallels Cloud Server automatically starts all Containers that have this parameter set to "yes" upon startup. Note: If "yes" is specified as the value of this parameter in the <code>0.conf</code> file, all server system management parameters are set on the server boot to the values indicated in this file. |
| OFFLINE_MANAGEMENT | Overrides the <code>OFFLINE_MANAGEMENT</code> parameter from the global configuration file. |
| OFFLINE_SERVICE | Overrides the <code>OFFLINE_SERVICE</code> parameter from the global configuration file. |
| ALLOWREBOOT | Specifies whether the Container may be restarted with the "reboot" command inside. If omitted or set to "yes", reboot is allowed. Note: To make reboot working, you should uncomment the corresponding line in the <code>/etc/cron.d/vereboot</code> file. |
| CAPABILITY | Specifies capabilities inside the Container. You can set the following capabilities: <code>AUDIT_WRITE</code> , <code>CHOWN</code> , <code>DAC_OVERRIDE</code> , <code>DAC_READ_SEARCH</code> , <code>FOwner</code> , <code>FS_MASK</code> , <code>FSETID</code> , <code>IPC_LOCK</code> , <code>IPC_OWNER</code> , <code>KILL</code> , <code>LEASE</code> , <code>LINUX_IMMUTABLE</code> , <code>MKNOD</code> , |

| | |
|--------------------------|--|
| | NET_ADMIN, NET_BIND_SERVICE, NET_BROADCAST, NET_RAW, SETGID, SETFCAP, SETPCAP, SETUID, SYS_ADMIN, SYS_BOOT, SYS_CHROOT, SYS_MODULE, SYS_NICE, SYS_PACCT, SYS_PTRACE, SYS_RAWIO, SYS_RESOURCE, SYS_TIME, SYS_TTY_CONFIG, VE_ADMIN. |
| 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> , <code>vzpkgadd</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 | <p>Determines a set of technologies which should be provided by the Parallels Cloud Server kernel for Container operation. Currently, this parameter can contain the information about the following technologies:</p> <ul style="list-style-type: none"> ▪ The system architecture of the Container (<code>x86</code>, <code>x86_64</code>, or <code>i64</code>). ▪ 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 <code>yes</code> , disables the Container making it impossible to start the Container once it was stopped. You can start the disabled Container by setting the value of this parameter to <code>no</code> or using the <code>--force</code> option with the <code>prlctl set</code> command. |
| DESCRIPTION | <p>Sets the description for the Container.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p>Note: You are allowed to use only symbols in the 'A-z' and '0-9' ranges in your descriptions.</p> </div> |
| NAME | <p>The name assigned to the Container. You can use this name, along with the Container ID, to refer to the Container while performing this or that Container-related operation on the server. Follow the following rules while setting the Container name:</p> <ul style="list-style-type: none"> ▪ The name should contain the <code>A-Z</code>, <code>a-z</code>, <code>0-9</code>, <code>\</code>, <code>-</code>, and <code>_</code> symbols only. ▪ If the name consists of two or more words, it should be quoted (e.g., "My Container 101"). |
| 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 <code>no</code> . |
| UUID | The Container unique identifier. This identifier is used by certain Parallels Cloud Server utilities during their execution. |

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 `prctl set` command and the corresponding option name (in the lower case, e.g., `--kmemsize` for `KMEMSIZE`, etc.). See **Managing Containers** (p. 73) 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.

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 (<code>CPULIMIT</code>) or megahertz (<code>CPULIMIT_MHZ</code>), 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. | |
| 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. | |

| | | |
|----------------------------|--|----------|
| BURST_CPU_AVG_USAGE | CPU usage limit, in percent, set for the Container. This limit is calculated as the ratio of the current Container CPU usage to the CPU limit (i.e to the value of the <code>CPULIMIT</code> parameter) set for the Container in its configuration file. If the limit is not specified, the full CPU power of the server is considered as the CPU limit. Upon exceeding the <code>BURST_CPU_AVG_USAGE</code> limit, the <code>BURST_CPULIMIT</code> limit is applied to the Container. This parameter redefines the <code>BURST_CPU_AVG_USAGE</code> parameter set in the Parallels Cloud Server configuration file. | disabled |
| BURST_CPULIMIT | CPU power limit, in per cent, the Container cannot exceed. The limitations set in this parameter are applied to the Container when it exceeds the limit specified in the <code>BURST_CPU_AVG_USAGE</code> parameter. This parameter redefines the <code>BURST_CPULIMIT</code> parameter specified in the Parallels Cloud Server configuration file. | |

Disk Parameters

| Parameter | Description | Typical value |
|-----------------------|--|---|
| DISKSPACE | Total size of disk space that can be consumed by the Container, in 1 KB blocks. | 204800...10485760– 204800...11534340 |
| 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 <code>quota</code> 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: <ul style="list-style-type: none"> ▪ 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 Parallels Cloud Server, 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 newly created Containers is set to 0, which means that no limits are applied to any Containers. | |

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 by utilities like <code>vzcfgvalidate</code> in order to ensure configuration correctness. | 0...NUMPROC |
| NUMTCPSOCK | Number of TCP sockets (PF_INET family, SOCK_STREAM type). This parameter limits the number of TCP connections and, thus, the number of clients the server application can handle in parallel. | 40...500 |
| NUMOTHERSOCK | Number of sockets other than TCP. Local (UNIX-domain) sockets are used for communications inside the system. UDP sockets are used for Domain Name Service (DNS) queries, as example. UDP and other sockets may also be used in some very special applications (SNMP agents and others). | 40...500 |
| VMGUARPAGES | Memory allocation guarantee, in pages. Applications are guaranteed to be able to allocate memory while the amount of memory accounted as <code>privmmpages</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 |
| KMEMSIZE | Size of unswappable kernel memory, allocated for internal kernel structures for the processes of a particular Container. Typical amounts of kernel memory is 16...50 KB per process. | 798720...13148160-851968...14024704 |
| TCPSNDBUF | The total size of send buffers for TCP sockets, i.e. the amount of kernel memory allocated for data sent from applications to TCP sockets, but not acknowledged by the remote side yet. | 159744...5365760-262144...10458760 |
| TCPRCVBUF | Total size of receive buffers for TCP sockets. Amount of kernel memory, received from remote side but not read by local application yet. | 159744...5365760-262144...10458760 |
| OTHERSOCKBUF | Total size of UNIX-domain socket buffers, UDP and other datagram protocols send buffers. | 61440...1503232-163840...4063232 |
| DGRAMRCVBUF | Total size of receive buffers of UDP and other datagram protocols. | 32768...262144 |
| OOMGUARPAGES | Out-of-memory guarantee, in pages. Any Container process will not be killed even in case of heavy memory shortage if current memory consumption (including both physical memory and swap) until the <code>oomguarpages</code> barrier is not reached. | 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. | 3072...151200-3450...1612800 |

| | | |
|-------------------|---|-------------------------------------|
| 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 |
| DCACHESIZE | Total size of <code>dentry</code> and <code>inode</code> structures locked in memory. As example, application, first opening the <code>/etc/passwd</code> file, locks entries corresponding to <code>etc</code> and <code>passwd</code> inodes. If a second application opens the <code>/etc/shadow</code> file – only entry corresponding to <code>shadow</code> is charged, because <code>etc</code> is charged already. | 184320...3932160 – 196608...4194304 |
| 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 a fraction of a page is charged to each Container. The sum of the <code>physpages</code> for all Containers corresponds to the total number of pages used in the system by all accounted users. | Not limited |
| NUMIPTENT | The number of IP packet filtering entries. | 12...128 |

VSwap Parameters

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

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>pr1ctl</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 | <p>Indicates which <code>iptables</code> modules are allowed for the Container.</p> <p>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.</p> <p>The following modes are available:</p> <ul style="list-style-type: none"> ▪ disabled: none. ▪ stateless: (default) all modules except <code>conntrack</code> and NAT-related. ▪ stateful: all modules except NAT-related. ▪ full: all modules. |
| NETIF | <p>Specifies a number of parameters for the virtual network adapters existing inside the Container. These parameters include:</p> <ul style="list-style-type: none"> ▪ ifname: the name of the <code>veth</code> virtual Ethernet interface inside the Container. ▪ mac: the MAC address assigned to the <code>veth</code> virtual Ethernet interface inside the Container. ▪ host_mac: the MAC address assigned to the <code>veth</code> virtual Ethernet interface on the server. ▪ network: the name of the virtual network where the <code>veth</code> virtual network adapter is included. ▪ ip: the IP address(es) assigned to the <code>veth</code> 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 <code>yes</code> , the bandwidth guarantee is also the limit for the Container, and the Container cannot borrow the bandwidth from the <code>TOTALRATE</code> bandwidth pool. |

Backup Parameters

Backup-related parameters, if present, allow you to specify the number of backups to store. If absent, these parameters are taken from the global backup configuration file or the backup configuration file for a particular server.

| Name | Description | Default Value |
|--------------------------|--|---------------|
| BACKUP_CHAIN_LEN | An incremental backup parameter. After this number of incremental backups, a full backup is performed. | 7 |
| BACKUP_CHAIN_DAYS | An incremental backup parameter. After this number of days a full backup is performed. | 7 |
| BACKUP_KEEP_MAX | The number of backups to store. Only full and plain full backups | 3 |

| | | |
|--|--|--|
| | are accounted. If a regular backup is being performed that exceeds this number, the oldest backup is automatically deleted. This parameter is effective only if the <code>-p</code> option is specified with the <code>vzbackup</code> utility. If there is no <code>-p</code> option, the number of backups to store is not limited whatever the value of this parameter. | |
|--|--|--|

Linux Distribution Configuration Files

Some Parallels Cloud Server utilities (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, Parallels Cloud Server uses special distribution configuration files located in the `/etc/vz/conf/dists` directory on the server.

There are a number of distribution configuration files shipped with Parallels Cloud Server by default (`centos.conf`, `fedora-core.conf`, `gentoo.conf`, etc.). To view all configuration files available on your Parallels Cloud Server, you can go to the `/etc/vz/conf/dists` directory and issue the `ls` command. The distribution configuration files will be displayed in the form of `Linux_Distribution_Name-version.conf` where `Linux_Distribution_Name` and `version` denote the name of the Linux distribution and its version, respectively (e.g., `centos-5.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 Parallels Cloud Server, 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_ID>
prlctl set <CT_ID> --ipadd <IP_address>
prlctl set <CT_ID> --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_ID> --ipdel <IP_address>
prlctl set <CT_ID> --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_ID> --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_ID> --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_ID> --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_ID> --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_ID>
```

- **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:

```
vzp2v [options] --ctid <CT_ID>
```

The scripts specified in distribution configuration files are located in the `/etc/vz/conf/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 `/etc/vz/conf/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 `/etc/vz/conf/dists` directory, the default file from this directory is used.

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.

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. <ul style="list-style-type: none">▪ 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 |
| PURGEHEAD | 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% |

Network Classes Definition File

In Parallels Cloud Server, 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 Parallels Cloud Server 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 - 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
```

vzvpn Configuration File

The `/etc/vzvpn/vzvpn.conf` file is used by the Parallels Support Tool to establish a secure connection (a virtual private network) between your server and the Parallels support server.

Format

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

Parameters

| Name | Description |
|-------------|---|
| REMOTE_HOST | Mandatory. The hostname or the IP address of the Parallels support server. |
| REMOTE_PORT | Mandatory. The port number of the Parallels support server to be used for establishing a virtual private network (VPN). |
| STARTTMO | Mandatory. The time, in seconds, during which there will be attempts to start the Parallels Support Tool if it could not be started immediately after its launching. |
| INACTIVE | Mandatory. The time of inactivity, in seconds, after which the connection between your server and the Parallels support server will be closed. |
| PING | Mandatory. The time, in seconds, at the end of which the port of the Parallels support server will be pinged in case no packets have been received from the support server during the time specified. |
| PING_EXIT | Mandatory. The time, in seconds, after a lapse of which the connection between your server and the Parallels support server will be closed in case no ping signals or other packets have been received from the |

| | |
|---|---|
| | support server during this time. |
| HTTP_PROXY=<hostname>[:port] | Optional. The hostname or the IP address and the port number of the HTTP proxy server through which a VPN between your server and the Parallels support server is to be established. This parameter overrides the HTTP_PROXY parameter set in the <code>/etc/vz/vz.conf</code> file on the server. If the HTTP_PROXY parameter is not specified in either of the files, the Parallels Support Tool looks for the <code>http_proxy</code> environment variable on the server and takes its value for establishing a VPN. |
| HTTP_PROXY_USER | Optional. The user name used by the HTTP proxy server for your authentication. |
| HTTP_PROXY_PASSWORD | Optional. The password of the user specified in the HTTP_PROXY_USER parameter and used for your authentication by the HTTP proxy server. |

Note: You are not recommended to change any of the aforementioned parameters. Modify them only if you are certain of your actions (for example, you have received the corresponding information from Parallels).

Kernel Parameters

There is a number of kernel limits that should be set for the Parallels Cloud Server software to work correctly. Parallels Cloud Server 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 Parallels Cloud Server:

```
# 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 corresponding “Scatter/gather IO” feature is not supported by the `venet` device, used in Parallels Cloud Server 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.

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 Parallels Cloud Server version, two files are accessible: `vzpp.conf` and `vzpp-plesk.conf`. The first file defines the Container offline management by means of Parallels 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>
```

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 Parallels Power Panel or the Plesk control panel. | PORT=8443 |
| DST_VEID | The ID of the Container where the requests coming to the | DST_VEID=1 |

| | | |
|--|------------------------------------|--|
| | specified port will be redirected. | |
|--|------------------------------------|--|

vzlmond Configuration File

The `/etc/vzlmond.conf` file defines the configuration parameters for the `vzlmond` daemon used to periodically check and log the state of your server. The gathered logs can then be used by the `vzstatrep` utility to generate statistic reports and graphics on their basis and to send these reports and graphics to the server administrator's e-mail address(es). For more details on the `vzstatrep` utility, see **vzstatrep** (p. 62).

The parameters in this file are presented on separate lines in the following format:

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

Parameters

| Name | Description | Default Value |
|------------------------------|--|------------------------------|
| STATS_VMSTAT_PERIOD | The periodicity, in seconds, with which the <code>vmstat</code> utility is run on the server and its output is saved to log files in the directory specified as the value of the <code>LOGS_DIR</code> parameter. The <code>vmstat</code> output contains information on the server kernel threads, virtual memory, disks, traps, and CPU activity. For more information on <code>vmstat</code> , see its man pages. | 480 |
| STATS_FULLDUMP_PERIOD | The period, in seconds, at the end of which the complete statistics on the server resources consumption is gathered and logged to the directory specified as the value of the <code>LOGS_DIR</code> parameter. As distinct from the <code>vmstat</code> output, this statistics represents a snapshot of the files contents from the <code>/proc</code> directory on the server and contains information on virtually every server resource: the environment of a certain process, the state and configuration of the CPU(s), the number of I/O ports on the server and their configuration, etc. Keep in mind that the amount of disk space needed to store this information may be considerable (about 0.5 KB per Container). However, you are recommended to set the period to no more than 10 minutes to regularly check and log the current server state and resources consumption. | 480 |
| STATS_NET_PERIOD | The period, in seconds, after which the server network statistics is collected and logged to the directory specified as the value of the <code>LOGS_DIR</code> parameter. The network statistics is gathered separately for each network interface on the server (e.g., <code>eth0</code> , <code>eth1</code>). | 480 |
| LOGS_DIR | The name of the directory on the server where the gathered statistics is to be stored. | <code>/var/log/vzstat</code> |

All the aforementioned parameters are set to their default values during the Parallels Cloud Server installation; so, you do not have to additionally edit any parameter in the `/etc/vz1mond.conf` file to start gathering your server statistics.

pstat Configuration File

This file (`/etc/vzstat.conf`) lists a number of CPU, memory, and disk-related parameters used by the `pstat` utility. The values assigned to these parameters denote either the warning or the error level for the `pstat` utility to start displaying these parameters either in the yellow color (the warning level has been hit) or in the red color (the error level has been hit). Moreover, if a parameter has hit the error level, the **CRIT** warning is displayed instead of **OK** after the name of the corresponding subsystem (CPU, Memory, Swap, Net, or Disks).

The table below provides information on the name and the description of all these parameters, on whether they denote the warning or the error level, whether the real parameter value has to be higher or lower than this level in order to invoke an alert, and on the parameters default values.

Parameters

| Name | Description | Default Value | Alert When | Alert Type |
|------------------|---|---------------|------------|------------|
| LOAD_AVG | Load average. | 30 | Higher | Warning |
| PROC_RUN | Number of running processes. | 20 | Higher | Warning |
| PROC_UNINT | Number of uninterruptable processes (in "D" state). | 20 | Higher | Warning |
| CPU_IDLE | CPU idle time, in percent. | 10 | Lower | Warning |
| CPU_SYS | CPU system time, in percent. | 50 | Higher | Warning |
| CPU_LAT_MAX_WARN | Scheduling latency, in milliseconds (maximum over 5 sec period). | 750 | Higher | Warning |
| CPU_LAT_MAX_ERR | Scheduling latency, in milliseconds (maximum over 5 sec period). | 1000 | Higher | Error |
| CPU_LAT_AVG_WARN | Scheduling latency, in milliseconds (5 sec average). | 500 | Higher | Warning |
| CPU_LAT_AVG_ERR | Scheduling latency, in milliseconds (5 sec average). | 750 | Higher | Error |
| MEM_LAT_MAX_WARN | Memory allocation latency, in milliseconds (maximum over 5 sec period). | 300 | Higher | Warning |
| MEM_LAT_MAX_ERR | Memory allocation latency, in milliseconds (maximum over 5 sec period). | 500 | Higher | Error |
| MEM_LAT_AVG_WARN | Memory allocation latency, in milliseconds (5 sec average). | 250 | Higher | Warning |

| | | | | |
|---|--|------|--------|---------|
| MEM_LAT_AVG_ERR | Memory allocation latency, in milliseconds (5 sec average). | 400 | Higher | Error |
| MEM_ZONE_ACT_INACT_FREE_WARN | Size of available memory (free + active + inactive pages), in percent. | 8 | Lower | Warning |
| MEM_ZONE_ACT_INACT_FREE_ERR | Size of available memory (free + active + inactive pages), in percent. | 4 | Lower | Error |
| MEM_ZONE_ACT_INACT_FREE_ABS_WARN | Size of available memory (free + active + inactive pages), in MB. | 4 | Lower | Warning |
| MEM_ZONE_ACT_INACT_FREE_ABS_ERR | Size of available memory (free + active + inactive pages), in MB. | 2 | Lower | Error |
| MEM_ZONE_ORDER_GT_0 | Number of pages which are gathered in blocks with order > 0. For example, if current memory distribution looks like: 3*1 1*2 3*4 5*8 Then number of pages with order>0 is 1*2 + 3*4 + 5*8 + ... | 100 | Lower | Warning |
| SWAP_FREE_WARN | Free swap space, in percent. | 75 | Lower | Warning |
| SWAP_FREE_ERR | Free swap space, in percent. | 50 | Lower | Error |
| SWAP_IN_WARN | Swap-in activity, in Mb/sec. | 0.5 | Higher | Warning |
| SWAP_IN_ERR | Swap-in activity, in Mb/sec. | 1 | Higher | Error |
| SWAP_OUT_WARN | Swap-out activity, in Mb/sec. | 0.5 | Higher | Warning |
| SWAP_OUT_ERR | Swap-out activity, in Mb/sec. | 1 | Higher | Error |
| SWAP_LAT_MAX_WARN | Swap-in latency, in milliseconds (maximum over 5 sec period). | 750 | Higher | Warning |
| SWAP_LAT_MAX_ERR | Swap-in latency, in milliseconds (maximum over 5 sec period). | 1000 | Higher | Error |
| SWAP_LAT_AVG_WARN | Swap-in latency, in milliseconds (5 sec average). | 500 | Higher | Warning |
| SWAP_LAT_AVG_ERR | Swap-in latency, in milliseconds (5 sec average). | 750 | Higher | Error |
| DISK_FREE_INODES_WARN | Free inodes on the disk, in percent. | 20 | Lower | Warning |
| DISK_FREE_INODES_ERR | Free inodes on the disk, in percent. | 5 | Lower | Error |
| DISK_FREE_SPACE_WARN | Free disk space, in percent. | 20 | Lower | Warning |
| DISK_FREE_SPACE_ERR | Free disk space, in percent. | 5 | Lower | Error |
| CT_FAILCNT_DELTA | Number of failed UBC resource allocations for a particular Container between <code>psstat</code> screen updates (any resource type counts). | 1 | Higher | Error |

vzrmond Configuration File

The configuration for the `vzrmond` daemon is stored in the `/etc/vzrmond.conf` file. The daemon runs on the Monitor Server, provides remote monitoring capabilities for servers registered with it, and sends alerts to the specified e-mail addresses. Alerts can also be sent from external applications (e.g., via ICQ or by SMS). The configuration file lists a number of parameters, some of which should have user-provided values (from `HOSTS` to `CUSTOM_LIST` in the table below). These values must be enclosed in double quotes and separated by white spaces. The remaining parameters are set to default values which may also be changed. These values are not enclosed in quotes.

Note: To install the daemon, install the `vzrmon` package from the official repository.

Parameters

| Name | Description | Default value |
|----------------------------|--|--|
| HOSTS | The list of hosts to monitor delimited by white spaces. Both hostnames and IP addresses are allowed. | |
| EMAIL_ADDRESSES | E-mail address(es) to send alerts to. Must be separated by white spaces. | |
| EMAIL_NOTIFICATIONS | Notification types which can be sent to the specified e-mail address(es). | SYSTEM_UP SYSTEM_DOWN DISK_OK DISK_BAD INODES_NORM INODES_HIGH HDBBUSY_NORM HDBBUSY_HIGH SSH_UP SSH_DOWN VZSTAT_OK VZSTAT_BAD LOADAVG_NORM LOADAVG_HIGH UNINT_NORM UNINT_HIGH MEMLATM_NORM MEMLATM_HIGH MEMLATA_NORM MEMLATA_HIGH CPULATM_NORM CPULATM_HIGH CPULATA_NORM CPULATA_HIGH |

| | | |
|-----------------------------|--|--|
| | | SWAPIN_NORM SWAPIN_HIGH SWAPOUT_NORM SWAPOUT_HIGH |
| CUSTOM_ACTION | A custom program to send alerts from (e.g., via ICQ or by SMS). | |
| CUSTOM_LIST | Command-line options for the program specified in CUSTOM_ACTION . Must be separated by white spaces. | |
| POLL_PERIOD | Registered servers polling interval, in seconds. | 15 |
| CHK_MAX_FAILS | The maximum number of failed attempts to reach a server, after which the "Server is dead" alert will be sent. | 4 |
| LOAD_AVG | The average number of processes on the server. When this value is exceeded, an alert is sent. | 30 |
| PROC_UNINT | The number of uninterruptable sleeping processes (in the "D" state). When this value is exceeded, an alert is sent. | 20 |
| CPU_LAT_MAX_ERR | The maximum process scheduling latency, in milliseconds. When this value is exceeded, an alert is sent. | 1000 |
| CPU_LAT_AVG_ERR | The average process scheduling latency, in milliseconds. When this value is exceeded, an alert is sent. | 750 |
| MEM_LAT_MAX_ERR | The maximum memory allocation latency, in milliseconds. When this value is exceeded, an alert is sent. | 500 |
| MEM_LAT_AVG_ERR | The average memory allocation latency, in milliseconds. When this value is exceeded, an alert is sent. | 400 |
| SWAP_IN_ERR | The swap in activity, in Mb/s. When this value is exceeded, an alert is sent. | 1.0 |
| SWAP_OUT_ERR | The swap out activity, in Mb/s. When this value is exceeded, an alert is sent. | 1.0 |
| DISK_FREE_INODES_ERR | The percentage of free disk inodes. When the actual value becomes less than this value, an alert is sent. | 5 |
| DISK_FREE_SPACE_ERR | The percentage of free disk space. When the actual value becomes less than this value, an alert is sent. | 5 |

To start monitoring a server, you should provide valid values for the `HOSTS` and `EMAIL` parameters. An external program to send alerts from must be installed on the Monitor Server and its name and command-line options must be set in the `CUSTOM_ACTION` and `CUSTOM_LIST` parameters. Alerts will be sent as messages standard for the specified program.

You should increase the value of the `POLL_PERIOD` parameter together with the increase in the number of monitored servers not to create an overload on the Monitor Server. The parameters related to the scheduling latency, memory allocation latency, and swap in/out activity serve to have an alert generated if the system's performance plummets due to the abnormal values of these parameters.

Restart the `vzrmond` daemon after editing its configuration file.

vzstatrep Configuration File

The `vzstatrep.conf` configuration file located in the `/etc` directory on the Monitor Server is used by the `vzstatrep` utility while trying to generate statistic reports and graphics on the server resource consumption and to send them to your e-mail address.

Format

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

Parameters

| Name | Description |
|--------------------|--|
| NODES | The IP address or hostname of the server whose logs are to be analyzed. You can set several servers for being processed with the help of the <code>vzstatrep</code> utility and separate them by spaces. If no server is specified, the logs of the local server (i.e. of the Monitor Server itself) are analyzed. |
| STATS_EMAIL | The e-mail address to send the generated statistic reports and graphics to. You can specify several e-mail addresses and separate them by commas or spaces. |
| GNUPLOT | The path to the <code>gnuplot</code> utility on the Monitor Server. By default, the utility is located in the <code>/usr/bin</code> directory; however, you may specify another directory for its location (e.g., <code>/etc/mydir/gnuplot</code>). <code>gnuplot</code> is used by the <code>vzstatrep</code> utility to present the server resources consumption in the graphical form. The resources whose graphics are to be generated should be set as the values of the <code>STATS_PLOT</code> parameter. For detailed information on the <code>gnuplot</code> utility, see its man pages. |
| MUTT | The path to the <code>mutt</code> utility on the Monitor Server. By default, the utility is located in the <code>/usr/bin</code> directory; however, you may specify another directory for its location (e.g., <code>/etc/mydir/mutt</code>). <code>mutt</code> is used by the <code>vzstatrep</code> utility to send the generated statistic reports and graphics in the form of attached files via e-mail. For detailed information on the <code>mutt</code> utility, see its man pages. |
| LOGS_DIR | The path to the directory on the server where <code>vzstatrep</code> will search for the logs generated by the <code>vzlmnd</code> daemon and containing the information on the server resources consumption. By default, the <code>/var/log/vzstat</code> directory is used. If you have changed the directory where <code>vzlmnd</code> stores the gathered information, you should specify the full path to |

| | |
|-------------------|---|
| | this directory as the value of this parameter (e.g., LOGS_DIR=/my_logs/vzstat). |
| STATS_PLOT | <p>Specify the resources parameters whose graphics are to be generated by means of the <code>gnuplot</code> utility. You can specify several resources and separate them by spaces. Currently, you can create graphics for the following parameters:</p> <ul style="list-style-type: none"> ▪ <code>ve_sum</code>: the information on the CPU usage for all Containers on the server. ▪ <code>ve_top</code>: the information on the CPU usage for 5 Containers with the highest CPU consumption. ▪ <code>loadavg</code>: the average number of active processes for the past 1, 5, and 15 minutes. Active processes can be running, i.e. currently executed by the CPU, or runnable, i.e. waiting in the run queue for the CPU. ▪ <code>io</code>: the amount of data read from and written to all devices on the server, in kilobytes per second. ▪ <code>mem</code>: the total memory consumption on the server. ▪ <code>ints</code>: the number of interrupts and context switches on the server per second. ▪ <code>cpu</code>: the information on the CPU load on the server. ▪ <code>net</code>: the network information for each network interface on the server. ▪ <code>forks</code>: the number of copies of all processes made on the server during one second. <p>By default, all the aforementioned resources except for <code>ve_sum</code> are plotted.</p> |

To start analyzing the logs, creating the server statistic reports and graphics, and receiving e-mail messages with these reports and graphics, you should specify the `NODES` and `STATS_EMAIL` parameters in the `/etc/vzstatrep.conf` file. All the other parameters are automatically set during the `vzrmon` package installation on the Monitor Server.

Backup Configuration File

This file (`/etc/vzbackup.conf`) is in the same format as the global Parallels Cloud Server configuration file and per-Container configuration files. All the parameters define the global backup settings, but some of them may be overridden by the per-server configuration file, if the latter exists. Still, other parameters may be further overridden in the configuration file of a particular Container.

Note: The settings in the `/etc/vzbackup.conf` file are valid only for Containers backed up with the `pbackup` utility. These settings are ignored when you back up individual Containers using the `prlctl backup` command or virtual machines using both the `pbackup` utility and the `prlctl backup` command.

Global Parameters

| Name | Description | Default value |
|-------------------|---|-------------------------------------|
| BACKUP_DIR | The directory where backups are stored. | <code>/var/parallels/backups</code> |

| | | |
|----------------------------------|--|----------|
| | Note: To set the backup directory in the current version of Parallels Cloud Server, use the <code>prlsrvctl set</code> command (p. 52). | |
| BACKUP_TYPE | The backup type. Among the supported types are "plain full (F)", "full (I)", and "incremental (i)". The default is incremental. If it is impossible to do an "incremental" backup, a "full" backup will be made. | i |
| BACKUP_NODES | The hostname of the server whose Containers are to be backed up. You can specify several hostnames of your servers and separate them by spaces. If you wish to back up Containers residing on the Backup server itself, you should specify its hostname as the value of this parameter. | |
| BACKUP_MAX_CHLD | The maximal number of servers to back up in parallel for non-periodic backups. | 1 |
| BACKUP_MAX_CHLD_CRON | The maximal number of servers to back up in parallel for periodic backups. | 3 |
| BACKUP_NOTIFY_EMAIL | The e-mail addresses where to send notifications about the backing up. | |
| BACKUP_COMPRESS | <p>Specifies whether the Containers are to be compressed when being backed up, and with what compression algorithm. When backing up Containers, you can set this option to one of the following values:</p> <p>C0: in this case the Container backup is created without any compression. Using this level of compression, you may greatly reduce the backup creation time; however, the size of the resulting backup file may significantly increase as compared to other compression levels.</p> <p>C1: in this case the Container backup is created with a normal level of compression.</p> <p>C2: in this case the Container backup is created with the high level of compression. The size of the resulting backup file is smaller than that of the backup file compressed in the 'normal' and 'none' modes; however, it takes longer to create the backup file.</p> <p>C3: in this case the Container backup is created with the maximal level of compression. The size of the resulting backup file is the smallest and the time of the backup creation - the longest.</p> <p>This parameter can be overridden by the <code>-Cg</code>, <code>-Cb</code>, <code>-Cn</code> command-line options of the <code>pbackup</code> utility.</p> | none |
| BACKUP_COMPATIBILITY_MODE | If set to "yes" and not commented out, the <code>pbackup/prestore</code> utilities use | disabled |

| | | |
|--------------------|--|----|
| | vzbackup/vzrestore for operations on Containers. Otherwise, <code>prlctl backup</code> and <code>prlctl restore</code> commands are used. This parameter is commented out by default. | |
| CRON_BACKUP | Specifies whether the backing up is performed as a cron job. If set to "yes", the values of the <code>BACKUP_KEEP_MAX</code> and <code>BACKUP_LOADAVG_MAX</code> parameters in the given file are taken into consideration. This parameter can be overridden by the <code>-p</code> or <code>-j</code> command-line switch of the <code>vzbackup</code> utility. | no |

Per-server Parameters

| Name | Description | Default value |
|---------------------------|--|--------------------------|
| BACKUP_SSH_OPTS | Options which are passed to <code>ssh</code> when it is used. | <code>-c blowfish</code> |
| BACKUP_VESTOP | Defines whether the Containers are to be stopped before their backing up. If set to <code>-s</code> , the Containers are stopped by default, otherwise, they are not stopped. | |
| BACKUP_EXCL_VES | Defines those Containers that are to be excluded from the backup list. Container IDs must be given here. | |
| BACKUP_LOADAVG_MAX | The maximal <code>loadavg</code> with which backing up is allowed. This parameter is effective only if the <code>-p</code> option is specified with the <code>vzbackup</code> utility. | 10 |
| BACKUP_FINISH_TIME | The time when the backing up should be stopped and delayed until the next execution (e.g., when running backup scripts at 4am) one can require the backup to be finished before 7am. The backup will continue from the last Container at the next execution. The format is: "HH:MM". This parameter is effective only if the <code>-L</code> option is specified with the <code>vzbackup</code> utility. | none |
| BACKUP_LIMIT_TIME | The number of hours after which the backing up should be stopped and delayed until the next execution. The format is: "HH". This parameter is effective only if the <code>-L</code> option is specified with the <code>vzbackup</code> utility. | none |

Per-Container Parameters

| Parameter | Description | Default Value |
|--------------------------|---|---------------|
| BACKUP_CHAIN_LEN | An incremental backup parameter. After this number of incremental backups, a full backup is performed. | 7 |
| BACKUP_CHAIN_DAYS | An incremental backup parameter. After this number of days a full backup is performed. | 7 |
| BACKUP_KEEP_MAX | The number of backups to store. Only full and plain full backups are accounted. If a regular backup is being performed that exceeds this number, the oldest backup is automatically deleted. This parameter is effective only if the <code>-p</code> option is specified with the <code>vzbackup</code> utility. If there is no <code>-p</code> option, the number of backups to store is not limited | 3 |

| | | |
|--|---------------------------------------|--|
| | whatever the value of this parameter. | |
|--|---------------------------------------|--|

If you want to rewrite the per-server parameters for a particular Parallels server, you should create a new configuration file named `server.conf` and put it to the backup directory (defined by the `BACKUP_DIR` parameter in the global backup configuration file).

vzpkgproxy Configuration File

This file (`/etc/vzpkgproxy/vzpkgproxy.conf`) is the configuration file for `vzpkgproxy` - a special utility which can be used to create special caching proxy servers allowing you to efficiently manage your OS and application EZ templates.

Format

`<parameter_name>=<parameter_value>`

Parameters

| Name | Description |
|----------------------|---|
| REPO_DIR | The path to the directory on the proxy server where the local repository created on the basis of the cached packages is to be stored. By default, this directory has the path of <code>/var/www/html/download</code> . |
| CACHE_DISABLE | The IP addresses of the hosts to be excluded from the caching process. It means that the packages requested by servers and received from these hosts will not be cached on the proxy server. By default, the proxy server is configured to cache all packages from all hosts on external networks. |

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>`

Parameters

| Name | Description |
|-------------------|--|
| VZTT_PROXY | The IP address or hostname of the caching proxy server to be used by the <code>vzpkg</code> 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 <code>HTTP_PROXY_USER</code> 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 <code>vzpkg</code> cache are regarded as 'not obsolete'. During this time, the <code>vzpkg</code> 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 <code>vzpkg</code> execution. The package names should correspond to the name of real packages in the repository and can contain file globs (e.g., * and ?). |

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>
```

Options

| Name | Description |
|------------------|---|
| THRESHOLD | Start compacting a virtual disk if the freed space exceeds the specified value, in percent to the virtual disk size. |
| DELTA | After compacting, leave the amount of free space on a virtual disk equal to "THRESHOLD" - "DELTA", in percent to the virtual disk size. |

shaman.conf

The `/etc/shaman/shaman.conf` file is used by the `shaman` utility to manage and monitor High Availability clusters and cluster resources.

Format

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

Options

| Name | Description |
|---------------------------------------|--|
| CLUSTER_NAME | Set the name of the cluster for shaman to operate on. |
| LOG_LEVEL | <p>Set the verbosity level for printed messages:</p> <ul style="list-style-type: none"> ▪ 0. Print error messages only. ▪ 1. Print warning messages in addition to error messages. ▪ 2. Print information messages in addition to error and warning messages. This is the default verbosity level. ▪ 4. Print debug diagnostic messages in addition to messages from level 0-2. |
| LOCK_TIMEOUT | <p>Set the timeout for shaman-monitor operations (e.g., electing a new master or deciding that a slave node is down). This parameter helps avoid situations when shaman-monitor performs a cluster-related operation if someone pulls out the network cable for just a couple of seconds (not longer than the half of the LOCK_TIMEOUT value).</p> <p>The default value is 60 seconds.</p> <p>The specified value is added to the value of other timeout-related parameters (see below).</p> |
| LEASE_CHECK_TIMEOUT_FOR_MASTER | <p>Set the timeout for electing a new master node when the original master node or the shaman-monitor daemon fails, or high availability support gets disabled.</p> <p>The default value is 10 seconds.</p> |
| LEASE_CHECK_TIMEOUT_FOR_SLAVE | <p>Set the timeout after which the master node will consider a slave node as broken if this node and the shaman-monitor daemon get down, or high availability support gets disabled.</p> <p>The default value is 10 seconds.</p> |
| LEASE_LOST_ACTION | <p>Defines the action to perform if shaman-monitor loses connection to cluster.</p> <p>This may happen when a node goes online after having been disconnected from network for more than LOCK_TIMEOUT seconds. In this case, the watchdog timer has not expired yet, but the cluster is already unavailable, because the master node has prohibited access to the cluster until the node is rebooted.</p> <p>Available values are <code>crash</code>, <code>halt</code>, <code>reboot</code>, and <code>none</code> (do nothing).</p> |
| POOL_CHECK_TIMEOUT | <p>Set the interval for shaman-monitor to check for the resources scheduled for relocation.</p> <p>The default value is 30 seconds.</p> |
| RELOCATION_SKIP_THRESHOLD | <p>Sets the threshold for the number of simultaneously crashed nodes.</p> <p>If the number of simultaneously crashed nodes becomes greater than or equal to the threshold, the master stops relocating resources from the crashed nodes. When the number of simultaneously crashed nodes drops below the threshold, the master automatically resumes relocating resources from the crashed nodes. The threshold can be useful when</p> |

| | |
|-------------------------|---|
| | <p>multiple nodes are being rebooted at the same time. Without it, the master would start relocating resources from all the rebooting nodes.</p> <p>The threshold is set to 3 by default and must be 2 or greater. For clusters with only 3 nodes, the threshold is automatically set to 2.</p> |
| WATCHDOG_TIMEOUT | <p>Set the interval for the watchdog timer.</p> <p>The watchdog timer is responsible for performing the action defined in WATCHDOG_ACTION (see below) if <code>shaman-monitor</code> crashes or hangs up. The <code>shaman-monitor</code> daemon activates the watchdog timer on its start-up and periodically resets it to the specified value. If something goes wrong with <code>shaman-monitor</code> so that it fails to reset the timer, the watchdog timer counts down until it reaches zero and performs the defined action. Setting the interval to zero disables the watchdog timer. Minimal watchdog timer interval that could be set is 10 seconds.</p> <p>The default value is 120 seconds.</p> |
| WATCHDOG_ACTION | <p>Defines a sequence of actions to perform after the watchdog timer expires (happens when <code>shaman-monitor</code> crashes or hangs up).</p> <p>When the watchdog timer expires, the first specified action is attempted. If unsuccessful, the next specified action is attempted and so on. If the last specified action is unsuccessful, then the action specified in the <code>/sys/kernel/watchdog_action</code> file is performed.</p> <p>At least one action must be specified. Multiple actions must be separated with commas. Available actions are listed in the <code>/sys/kernel/watchdog_available_actions</code> file. The default sequence is <code>netfilter, reboot</code>.</p> |

Parallels Cloud Server Scripts

This section provides information on scripts you can use in Parallels Cloud Server 6.0.

Overview

The table below lists the scripts you can use to customize the behavior of Containers and virtual machines and define High Availability actions.

Scripts

| Name | Description |
|--|---|
| <code>/vz/private/<CT_ID>/scripts/<script_name></code> | Container action scripts. These scripts allow you to run user-defined actions on particular events. The currently defined actions are <code>start</code> , <code>stop</code> , <code>mount</code> , <code>unmount</code> . |
| <code>/etc/vz/conf/dists/scripts/<script_name></code> | Scripts to be executed on performing certain Container-related operations (e.g., when adding a new IP address to the Container). These operations should be specified in the corresponding Linux distribution configuration file. |

| | |
|---|--|
| <code>/etc/rc.d/init.d/vz</code> | Parallels Cloud Server start/stop script. This script is responsible for proper Parallels Cloud Server startup and shutdown procedures, including modules loading and Container start/stop procedures. |
| <code>/etc/vmprivate/<vm_name>/scripts/<script_name></code> | Virtual machine actions scripts. These scripts allow you to run user-defined actions when starting and stopping virtual machines. |
| <code>/usr/share/shaman</code> | Scripts used by the shaman utility to manage and monitor High Availability clusters and cluster resources. |

Container Action Scripts

There might be situations when you need to do additional actions when a particular Container is started or stopped. For example, if you want to be able to access the server file system (or part of it) from Container 101, then you can bind mount it inside the Container manually from the server. However, after you restart the Container, your mount disappears, and you should manually type the `mount` command again. Parallels Cloud Server allows you to automate procedures like the above by using *action scripts* below.

Scripts

| Name | Description |
|----------------------|--|
| global mount | This script runs immediately after <code>prlctl</code> mounts the Container private area. The Container itself is not yet running and the script is running in the server context. |
| mount | This script runs immediately after the global mount script. The Container is still not running, and the script is called in the server context. |
| start | After <code>prlctl</code> has started a Container, it runs the Container <code>start</code> script. The script is running already in the Container context. |
| stop | This script runs before the Container is stopped, in the Container context. |
| umount | After the Container has been already stopped, the <code>umount</code> script is executed, and the script runs in the server context. |
| global umount | This script runs when <code>prlctl</code> is about to dismount the Container private area. It also runs in the server context. |

It is important to understand how `prlctl` handles exit codes of action scripts. If exit code is non-zero, then `prlctl` will try to undo the action for the `mount` and `start` scripts. In other words, if the `start` script returns an error, then `prlctl` will stop Container, and if one of the `mount` scripts fails, then `prlctl` will dismount the Container private area. Please note that in this case `prlctl` will not execute the `stop` and `umount` scripts at all.

Caution: When executing `prlctl start`, both `mount` and `start` scripts run. However, if the `start` script fails then neither `stop` nor `umount` scripts will run. As a result, `prlctl` might be unable to dismount the Container private area, if you set up additional mounts in the `mount` scripts and dismount them in the `umount` scripts.

The situation with the `umount` and `stop` scripts is similar. If a script returns an error, then the action will not be taken. Be careful since this allows to create Containers that are not stoppable by `prlctl`.

The global scripts are named `vps.mount` and `vps.umount` and located in the `/etc/vz/conf` directory on the server. These scripts are called when any Container on the server is started or stopped. So, you should include in these scripts those commands that are common for all Containers and leave Container-specific commands for the scripts belonging to a particular Container. Container-specific action scripts are located in the `/vz/private/CT_ID/scripts` directory and have the `mount`, `start`, `stop`, and `umount` names. For example, the scripts specific for Container 101 will have the following names:

- `/vz/private/101/scripts/mount`
- `/vz/private/101/scripts/start`
- `/vz/private/101/scripts/stop`
- `/vz/private/101/scripts/umount`

For the `mount` and `umount` scripts, the environment passed is the standard environment of the parent (i.e. `prlctl`) with two additional variables: `$VEID` and `$VE_CONFFILE`. The first one holds the ID of the Container being mounted (started, stopped, dismounted), and the second one holds the full path to the Container configuration file. It is probably a bit redundant. Parallels introduced both variables for convenience. You can use the following fragment of the code in bash scripts to get access to additional Container information like `$VE_PRIVATE` or `$VE_ROOT` locations:

```
#!/bin/bash
#
# This script sources Container configuration files in the same
# order as prlctl does
#
# if one of these files does not exist then something is
# really broken
[ -f /etc/sysconfig/vz ] || exit 1
[ -f $VE_CONFFILE ] || exit 1
#
# source both files. Note the order, it is important
. /etc/vz/vz.conf
. $VE_CONFFILE
```

The `start` and `stop` scripts are performed in the Container context. If these scripts call any external commands, these commands are taken from the Container itself. Also note that the `start` script runs before any Container tasks (including `init`), thus the `/proc` file system is not mounted inside the Container at this moment – therefore, applications using information from `/proc` may be not functional.

Virtual Machine Action Scripts

In Parallels Cloud Server 6.0, you can create custom scripts and configure them to run when you start or stop virtual machines. The following custom scripts are currently supported:

- *prestart*: this script is executed before a virtual machine is started.

- *poststart*: this script is executed after a virtual machine is started.
- *prestop*: this script is executed before a virtual machine is stopped.
- *poststop*: this script is executed after a virtual machine is stopped.

Action scripts specific for virtual machines must be placed to the `/scripts` subdirectory in the virtual machine's home directory. For example, assuming that you created the `MyVM` virtual machine in the default directory, the paths to the scripts must be as follows:

- `/var/parallels/MyVM/scripts/prestart`
- `/var/parallels/MyVM/scripts/poststart`
- `/var/parallels/MyVM/scripts/prestop`
- `/var/parallels/MyVM/scripts/poststop`

The information to scripts is passed using the `UUID` and `VM_HOME` variables, where `UUID` is the ID of the virtual machine you start or stop and `VM_HOME` is the full path to the virtual machine's directory.

Note: Action scripts you create for a virtual machine cannot contain `prlctl` commands intended for performing operations on this virtual machine (e.g., assigning an IP address to the virtual machine).

High Availability Scripts

Parallels Cloud Server comes with a number of scripts used by the `shaman` utility to manage and monitor cluster resources. There are two types of scripts:

- **Common scripts.** The common scripts are located in the `/usr/share/shaman` directory and used by the `shaman` utility to call resource-specific scripts.
- **Resource-specific scripts.** For each common script, there are one or more resource-specific scripts. Resource-specific scripts are peculiar to each cluster resource and located in separate subdirectories. For virtual machines and Containers, these directories are `/usr/share/shaman/vm-` and `/usr/share/shaman/ct-`, respectively. Resource-specific scripts are used to perform various actions on cluster resources.

For the full list of scripts and their descriptions, see the `shaman-scripts` man page.

Parallels Cloud Server Utilities

This section provides information on utilities that can be used to manage Parallels Cloud Server parameters.

prlsrvctl

The `prlsrvctl` command-line utility is used to perform management tasks on the Parallels server and Parallels Cloud Server. The tasks include getting the Parallels Cloud Server 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]]]
```

Options

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

prlsrvctl info

Displays the Parallels server and Parallels Cloud Server configuration information.

Syntax

```
prlsrvctl info
```

Remarks

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

- Server ID and hostname.
- Parallels Cloud Server version number.
- Default directory for storing virtual machine files.
- Parallels Cloud Server memory limits.
- Parallels Cloud Server minimum allowable security level.
- Default directory for storing virtual machine backups.

- Parallels Cloud Server license information.
- Server hardware configuration information.
- Other miscellaneous info.
- CPU feature mask.

prlsrvctl install-license

Installs the Parallels Cloud Server license on the Parallels server.

Syntax

```
prlsrvctl install-license -k,--key <key> [-n,--name <name>] [-c,--company <name>]
```

Options

| Name | Description |
|----------------------------------|-----------------------|
| -k, --key <key> | License key. |
| -n, --name <name> | License user name. |
| -c,--company <name> | License company name. |

prlsrvctl net

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

Subcommands

| Name | Description |
|-------------------------|---|
| net add (p. 49) | Creates a new virtual network |
| net set (p. 50) | Configures the parameters of an existing virtual network. |
| net del (p. 51) | Removes an existing virtual network. |
| net list (p. 51) | 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>]
```

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 Parallels server to which this virtual network should be bound. |
| <code>-m,--mac <mac_address></code> | The MAC address of a virtual network adapter on the Parallels server 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: <ul style="list-style-type: none"> <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 Parallels server 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-scope-start <IP_address></code> <code>--ip-scope-end <IP_address></code> | Sets the start and end IP addresses for the DHCP pool. The virtual machines and Containers connected to the network you are creating will automatically receive their IP addresses from this DHCP 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>]
                               [-n, --name <new_name>]
```

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 Parallels server to which this virtual network should be bound. |
| <code>-m,--mac <mac_address></code> | The MAC address of a virtual network adapter on the Parallels server 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: <ul style="list-style-type: none"> <code>bridged</code>. A virtual machine and Container connected to |

| | |
|---------------------------------|--|
| | <p>this type of virtual network appears as an independent computer on the network.</p> <ul style="list-style-type: none"> • <code>host_only</code>. A virtual machine and Container connected to this type of virtual network can access only the Parallels server and the virtual machines and Containers connected to the same virtual network. |
| -d, --description <desc> | <p>A user-defined description of the virtual network.</p> <p>Descriptions with white spaces must be enclosed in quotation marks.</p> |
| -n, --name <new_name> | <p>A new name for the virtual network. Use this parameter if you would like to rename the virtual network.</p> |

net del

Deletes an existing virtual network.

Syntax

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

Options

| Name | Description |
|---------------|--|
| <vnetwork_ID> | The name of the virtual network to delete. |

net list

Lists the existing virtual networks.

Syntax

```
prlsrvctl net list
```

Options

None.

prlsrvctl problem-report

Generates and displays problem reports.

Syntax

```
prlsrvctl problem-report
```

Options

None.

Remarks

The command collects technical data about Parallels Cloud Server and the Parallels server and displays the report on screen (the output can also be piped to a file). The report can then be directed to the Parallels technical support for analysis.

prlsrvctl set

Configures Parallels Cloud Server 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>] --backup-timeout <timeout>
                [{--default-encryption-plugin <plugin_ID> |
                --reset-default-encryption-plugin}]
                [--verbose-log <on|off>] [--cluster-mode <on|off>]
```

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: <ul style="list-style-type: none">• 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 Parallels server. The following options are available: <ul style="list-style-type: none">• 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.• 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: <ul style="list-style-type: none"> <code>on</code> — enables CEP. <code>off</code> — disables CEP. |
| --mng-settings <allow deny> | Allows to grant or deny permission to new users to modify Parallels Cloud Server preferences. By default, only administrators of the host OS can modify Parallels Cloud Server preferences. When a new user profile is created (this happens when a user logs in to the Parallels server 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: <ul style="list-style-type: none"> <code>host</code> — assign the device to the Parallels server. <code>vm</code> — assign the device to virtual machines. |
| --backup-storage [<i>user[:passwd]@server[:port]</i>] | 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. |
| --default-encryption-plugin <plugin_ID> | Sets the default encryption plug-in that will be used to encrypt virtual machines. |
| --reset-default-encryption-plugin | Resets the encryption plug-in settings. Virtual machines will be encrypted using the built-in encryption module. |
| --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. |
| --backup-timeout <timeout> | Sets the timeout for backup and restore operations, in seconds. |
| --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. |

prlsrvctl shutdown

Shuts down the Parallels Cloud Server component responsible for managing virtual machines and Containers. No operations on virtual machines and Containers are possible.

Syntax

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

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 Parallels Cloud Server component. |

prlsrvctl statistics

Obtains Parallels Cloud Server statistics.

Syntax

```
prlsrvctl statistics [-a, --all] [--loop] [--filter <name>]
```

Options

| Name | Description |
|------------------------------|---|
| --loop | Subscribes to receive statistics on the periodic basis. Once you execute the command with this option, the statistics will be displayed in your console window every time a new set of values is collected. To unsubscribe, press the Enter key or Ctrl-C in your console window. |
| -a, --all | <i>This parameter is not currently used.</i> |
| --filter <name> | <i>This parameter is not currently used.</i> |

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).

Subcommands

| Name | Description |
|-------------------------|---|
| usb list (p. 55) | Lists USB devices connected to the server together with the information about their virtual machine assignments for the current user. |
| usb set (p. 55) | Permanently assigns a USB device to the specified virtual machine. |
| usb del (p. 55) | Removes a previously created USB device assignment. |

usb list

Lists the USB devices connected to the physical server.

Syntax

```
prlsrvctl usb list
```

Options

None.

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_ID/vm_name>
```

Options

| Name | Description |
|-----------------|---|
| <usb_dev_ID> | The USB device ID. To obtain the list of USB devices connected to the server use the usb list (p. 55) command. |
| <vm_ID/vm_name> | The universally unique ID or 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** (p. 55) 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>
```

Options

| Name | Description |
|--------------|--|
| <usb_dev_ID> | The USB device ID. To see the current USB device assignments for the current user use the usb list (p. 55) command. |

prlsrvctl user list

Displays the list of Parallels Cloud Server 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>]
```

Options

| Name | Description |
|---|--|
| -o,--output <name mng_settings def_vm_home> | <p>Fields to include in the output. The following fields are available:</p> <ul style="list-style-type: none">▪ <code>name</code> — User name.▪ <code>mng_settings</code> — Indicates whether the user is allowed to modify Parallels Cloud Server preferences.▪ <code>def_vm_home</code> — The user default virtual machine folder. <p>The fields must be specified in lowercase.</p> |

prlsrvctl user set

Configures the profile of a Parallels Cloud Server user.

Syntax

```
prlsrvctl user set <UUID|name> [--def-vm-home <path>]  
[--mng-settings <allow|deny>]
```


Options

| Name | Description |
|--|--|
| <code><UUID name></code> | The user UUID (universally unique ID) or name. |
| <code>--def-vm-home <path></code> | The default virtual machine and Container directory name and path. |
| <code>--mng-settings <allow deny></code> | Specifies whether the user should be allowed to modify Parallels Cloud Server preferences. |

privnet

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

Subcommands

| Name | Description |
|---------------------------|---|
| <code>add</code> (p. 57) | Creates a new private network. |
| <code>set</code> (p. 57) | Configures the parameters of an existing private network. |
| <code>del</code> (p. 58) | Removes an existing private network. |
| <code>list</code> (p. 58) | 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>]
```

Options

| Name | Description |
|---|---|
| <code><private_network_ID></code> | The private network ID. |
| <code>-a,--ipadd <addr>[/mask]</code> | Add a subnet to the private network. The network can have multiple subnets. |
| <code>--global <yes no></code> | 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>]
```

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. |
| -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>
```

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
```

Options

None.

vzlicload

Installs and removes Parallels Cloud Server licenses.

Syntax

```
vzlicload [options]
```

Options

| Name | Description |
|--|--|
| -p, --product-key < <i>key</i> > | Installs the specified Parallels Cloud Server license on the server. |
| -f, --license-file < <i>file_path</i> > | The full path to the license file containing the license to be installed on the server. |
| -r, --remove | Removes the license with the specified serial number from the server. You can find out the license serial number using the vzlicview utility (for details, see vzlicview (p. 60)). |
| -i, --stdin | Makes vzlicload use standard input as a license. |
| -h, --help | Prints the usage help and exits. |

vzlicupdate

This utility can be used to perform the following license-related operations:

- Activate your Parallels Cloud Server installation using a special activation code.
- Update the currently installed license on the server.
- Transfer the license installed on the Source Server with the help of an activation code to the Destination Server.

Syntax

```
vzlicupdate [options]
```

Options

| Name | Description |
|---|---|
| -a, --activate < <i>activation_code</i> > | Activates the Parallels Cloud Server installation using the specified activation code. To successfully complete this task, your server must be connected to the Internet. |
| -t, --transfer | Transfers the license activated with the activation code from the Source Server to the Destination Server. Should be run along with the -a option on the Destination Server, i.e. on the server where you are planning to transfer the license. |
| -s, --server < <i>hostname</i> >[: <i>port</i>] | The hostname of the Parallels Key Authentication (KA) server responsible for updating Parallels Cloud Server licenses, activating Parallels Cloud Server installations, and transferring licenses from the Source Server to the Destination Server. If not specified, the ka.parallels.com hostname is used. |
| -n, --no-check | Updates the license currently installed on the server even if it is still |

| | |
|----------------------|---|
| | valid. |
| -v, --verbose | Sets the log level to its maximum possible value. |
| -h, --help | Prints the utility usage and exits. |

When executed without any options, `vzlicupdate` updates the license currently installed on the server. However, you can use the options listed in the table above to complete other license-related tasks.

vzlicview

This utility displays the license contents along with the license status information.

Syntax

```
vzlicview [options]
```

Options

| Name | Description |
|--|---|
| -p, --product-key <key> | Displays the license information contained in the specified Parallels Cloud Server product key. |
| -f, --license-file <file> | Displays the license information from the specified Parallels Cloud Server license file. |
| -i, --stdin | Makes <code>vzlicview</code> use standard input as a license and display its information. |
| -h, --help | Displays the utility usage and exits. |

When executed without any options, the utility returns the contents and status of the license currently installed on the server. The utility can report the following statuses for Parallels Cloud Server licenses:

| Name | Description |
|----------------|--|
| ACTIVE | The license installed on the server is valid and active. |
| VALID | The license the utility parses is valid and can be installed on the server. |
| EXPIRED | The license has expired and, therefore, could not be installed on the server. |
| GRACED | The license has been successfully installed on the server, but it has expired and is currently on the grace period (i.e. it is active till the end of the grace period). |
| INVALID | The license is invalid (for example, because of the server architecture mismatch) or corrupted. |

vznetcfg

The `vznetcfg` utility is used to manage the following network devices on the server:

- physical and VLAN (Virtual Local Area Network) adapters,

- Virtual Networks (VNs).

Syntax

vznetcfg *<command>*

Options

| Name | Description |
|--|---|
| net new <i><VN_name></i> | Creates a new Virtual Network with the name of <i><VN_name></i> on the server. |
| net addif <i><VN_name interface_name></i> | Connects a network device with the name of <i><interface_name></i> to the Virtual Network having the name of <i><VN_name></i> . You can join the following network devices to the Virtual Network: <ul style="list-style-type: none"> physical network interface cards (NICs) installed on the server VLAN adapters bound to NICs on the server |
| net delif <i><interface_name></i> | Disconnects a network device (either a NIC or a VLAN adapter) with the name of <i><interface_name></i> from the corresponding Virtual Network. |
| net change <i><old_VN_name> <new_VN_name></i> | Changes the Virtual Network name from <i><old_VN_name></i> to <i><new_VN_name></i> . |
| net del <i><VN_name></i> | Removes the Virtual Network with the name of <i><VN_name></i> from the server. |
| vlan add <i><parent_interface> <index></i> | Creates a new VLAN adapter, associates it with the VLAN ID of <i><index></i> (where <i><index></i> can be an arbitrary integer number to be used to uniquely identify the VLAN among other VLANs on the server), and ties it to the <i><parent_interface></i> physical network adapter on the server. |
| vlan del <i><vlan_adapter_name></i> | Removes the VLAN adapter with the name of <i><vlan_adapter_name></i> from the server. <div> <p>Note: A VLAN adapter name is automatically generated by Parallels Cloud Server on the basis of the VLAN ID and the name of the physical adapter you specified during the VLAN adapter creation (e.g., eth0.1). You can find out the VLAN name using the <code>vznetcfg if list</code> command.</p> </div> |
| if list | Lists detailed information on all network devices (NICs, VLAN adapters, etc.) available on the server. |
| net list | Displays detailed information on the Virtual Networks currently existing on the server. |
| init all | Initializes all interfaces (e.g., VLANs and bridges) listed in the <code>/etc/vz/vznet.conf</code> file on the server. You may wish to make use of this command when creating startup scripts. |

| | |
|-----------------|---|
| down all | Disables all interfaces (e.g., bridges and VLANs) listed in the <code>/etc/vz/vznet.conf</code> file on the server. |
|-----------------|---|

vzstatrep

`vzstatrep` is run on the Monitor Server and used to analyze the logs collected by the `vz1mond` daemon on one or more servers to generate statistic reports and graphics on the basis of the gathered logs, and to send these reports and graphics to the server administrator's e-mail address(es).

Syntax

```
vzstatrep [options]
```

Options

| Name | Description |
|-------------------------------------|---|
| --plot | Generate graphics for the resources parameters specified as the values of the <code>STATS_PLOT</code> parameter in the <code>/etc/vzstatrep.conf</code> file on the Monitor Server. |
| --sendmail | Send the statistic report and graphics to the e-mail address(es) specified as the value(s) of the <code>STATS_EMAIL</code> parameter in the <code>/etc/vzstatrep.conf</code> file on the Monitor Server. If the <code>--sendmailto</code> option is omitted, you should obligatorily use this option. |
| --sendmailto < <i>mail</i> > | Send the statistic report and graphics to the e-mail address specified as the value of this option. You can set several e-mail addresses and separate them by spaces. If the <code>--sendmail</code> option is omitted, you must use this one. |
| --weekly | Generate statistic reports and graphics on a weekly basis. By default, <code>vzstatrep</code> analyzes the logs and produces the server resources statistics once a day. |
| --nodes < <i>hostname</i> > | Analyze the logs from the server whose IP address or hostname is specified as the value of this option. You can set several servers by separating them by spaces and enclosing them in quotes (e.g., "my_hardware_node1 my_hardware_node2"). If the option is omitted or its value is not specified, the logs from the servers set as the values of the <code>NODES</code> parameter in the <code>/etc/vzstatrep.conf</code> file on the Monitor Server are analyzed. |

The `vzstatrep` utility generates statistic reports and graphics on the basis of the logs gathered by `vz1mond` (by default, the logs are stored in the `/var/log/vzstat` directory on the server) and containing information on the memory and CPU consumption of the server, network resources on the server, etc. You do not need to perform any additional operations to start using `vzstatrep`. All the necessary parameters can be set during the `vzstatrep` execution by using the aforementioned options. However, if you wish to run the `vzstatrep` utility as a `cron` job and/or free yourself from the necessity to manually specify the needed options each time you wish to run `vzstatrep`, you should edit the `/etc/vzstatrep.conf` configuration file on the Monitor Server and set the parameters values contained in this file. For details on the `/etc/vzstatrep.conf` file, see **vzstatrep Configuration File** (p. 37).

pstat

Parallels Cloud Server real-time monitoring utility. Displays the status and load of the system, including disk, network, CPU, and memory parameters (swap included). The status is updated in a preset time interval. `pstat` also provides a sortable list of running virtual machines and Containers together with their resources consumption statistics. Utility's display output can be customized.

Syntax

```
pstat [-l] [-d <X>] [-p <CT_ID>] [-b|-v] [-t] [-a]
```

Options

| Name | Description |
|-------------------------|---|
| -l | Print information once and exit immediately. |
| -d <X> | Specify the delay between screen updates. Can be changed on the fly by the <code>t</code> interactive command. Default is 1 sec. |
| -p <CT_ID> | Monitor only Containers with the specified <i>CT_IDs</i> . This flag can be given up to twenty times as follows: <code>-p CT_ID1 -p CT_ID2 [. . .]</code> . This option is not available interactively. |
| -b | "Brief" mode. Minimal details level. Shows only one summary line about each monitoring subsystem. By default, "standard" details level is in use. Valid levels are "brief", "standard" and "verbose". Can be set on the fly by the <code>b</code> interactive command. See also the <code>-v</code> command-line option and <code>s</code> and <code>v</code> interactive commands. |
| -v | "Verbose" mode. Provides maximum details about all monitored subsystems. Can be set on the fly by the <code>v</code> interactive command. See also the <code>-b</code> command-line option and <code>b</code> and <code>s</code> interactive commands. |
| -t | Text mode, provides information once. It is printed in terse form, suitable for parsing by other programs. All output data are not aligned and numbers are not in a human readable format. In the text mode, there are no colors in the output and only the top 10 Containers sorted by their CPU usage are shown. |
| -a | Display the current disk input and output (I/O) statistics for virtual machines and Containers. |
| -i | Display the IO accounting information for virtual machines and Containers. |
| -m | Display disc statistics for all file system types. By default, the statistics is shown for ext2, ext3, ext4, and reiserfs. |
| -o | Filter the output by the specified parameters. You can specify multiple parameters and separate them by commas. The list of available parameters is given below. |
| -O | Filter the output by the specified parameters. Unlike the <code>-o</code> option, the produced output already contains some default columns (for example, Container ID and IP address). The list of available parameters is given below. |
| -n | Display network statistics. |
| -s | Filter the output by the specified keys. The list of available parameters is given below. |
| -c | Display I/O statistics in the specified units of measurement: B (bytes), K (kilobytes), M (megabytes), or G (gigabytes). |

You can use the following parameters with the `-o`, `-O`, and `-s` options to filter the information related to virtual machines and Containers:

| Name | Column | Description |
|-----------------------|------------|---|
| id | CTID | virtual machine or Container ID |
| st | ST | virtual machine or Container status. |
| vm | %VM | Virtual memory usage, in per cent of the total memory on the Node. It's displayed in the form of "actual usage/barrier". This parameter corresponds to the <code>PHYS_PAGES</code> VSwap parameter. |
| km | %KM | Kernel memory usage, in per cent of the normal zone size. It's displayed in the form of "actual usage/barrier". This parameter corresponds to the <code>kmemsize</code> UBC parameter. |
| sw | %SW | Swap space usage, in per cent of the total swap space on the Node. It's displayed in the form of "actual usage/barrier". This parameter corresponds to the <code>SWAP_PAGES</code> VSwap parameter. |
| swpin | SWAPIN | I/O operations with the swap space, in operations per second. |
| proc | PROC | Process information. It's displayed in the form of "running processes/total processes/barrier". |
| cpu | CPU | CPU usage. It's displayed in the form of "actual CPU usage/guaranteed CPU usage". If more than one processor is installed on the server, all processors are considered as 100%. |
| cpu_iowait | CPU IOWAIT | Percentage of time that Container processes lost waiting for I/O. |
| cpu_iowait_acc | ACC IOWAIT | Percentage of time that Container processes lost waiting for I/O, accumulated since utility start. |
| sock | SOCK | Sockets usage calculated as the sum of UBC <code>numtcpsock</code> and <code>numothersock</code> parameters. It's displayed in the form of "open sockets/barrier". |
| fcnt | FCNT | Fail counters—that is, the number of UBC fail counters for all resources. |
| mlat | MLAT | Maximum process scheduling latency, in milliseconds. In this case, the latency means the maximum time a process in a virtual machine or Container is waiting for the CPU. |
| io | IO | I/O rate and limit, in bytes, kilobytes, megabytes, or gigabytes per second. |
| iops | IOPS | I/O operations rate and limit, in operations per second. |
| iow | IOW | Transfer rate at which data is written to the virtual machine or Container. |
| ior | IOR | Transfer rate at which data is read from the virtual machine or Container. |
| iowt | IOWT | Total amount of data that was written to the virtual machine or |

| | | |
|---------------|---------|---|
| | | Container. |
| ior | IORT | Total amount of data that was read from the virtual machine or Container. |
| ios | IOS | Synchronization rate. |
| iod | IOD | Flushing rate of dirty pages (pages that have been changed but are not yet written to the virtual machine or Container disk). |
| ioused | IOUSED% | Percentage of time the disks are used by the virtual machine or Container. |
| iowait | IOWAIT% | Percentage of time when at least one I/O transaction in the virtual machine or Container is waiting for being served. |
| rx | RX | Incoming traffic rate, in megabytes per second. |
| tx | TX | Outgoing traffic, in megabytes per second. |
| rx | RXT | Total amount of incoming traffic, in megabytes. |
| tx | TXT | Total amount of outgoing traffic, in megabytes. |
| rx | RXP | Incoming traffic, in packets per second. |
| tx | TXPT | Outgoing traffic, in packets per second. |
| rx | RXPT | Total amount of incoming packets. |
| tx | TXPT | Total amount of outgoing packets. |

Displayed Information

`pstat` can display the following information:

| Type | Description | Example | Toggled by |
|--------------------------|--|--|-------------|
| Uptime | This line displays the time for which the system has been up, and three "load averages" for the system. The load averages are the average number of processes ready to run during the last 1, 5, and 15 minutes. This line is just like the output of <code>uptime(1)</code> . | 1:22am, up 1:31, 2 users, load average: 0.00, 0.06, 0.33 | l |
| Containers and processes | Total number of virtual machines and Containers and processes running at the time of the last update. The output is also broken down into the number of tasks which are running, sleeping, uninterruptable, zombie, or stopped. | CTNum 102, procs 467: running 12, sleeping 455, unint 0, zombie 0, stopped 0 | p |
| CPU states | Shows the percentage of CPU time used by all virtual machines and Containers and by the server (shown as CT0), the CPU time spent in the user mode, in the system mode, and being idle, and the maximal/average scheduling latency in ms. The scheduling latency is the time spent by the processes in the system awaiting for scheduling. | CPU [OK]: CTs 43%, CT0 12%, user 41%, sys 13%, idle 45%, lat(ms) 3/2 | c |
| Mem | Statistics on the memory usage, including the total available memory, free memory, and maximal/average memory allocation latency. The | Mem [OK]: total 755MB, free 671MB/0MB (low/high), lat(ms) | m, M |

| | | | |
|-------------------------------|---|--|------|
| | free memory is displayed both for the low and high memory. The low memory is the sum of the DMA and Normal zones memory and the high memory is the High zone memory. Memory allocation latency is the time required to allocate memory inside the kernel in ms. An excessive allocation latency can be a sign of server's overload. | 10/7. | |
| Memory zones information | Information on the memory zones state. This information includes: the total size of the memory zone in MB, the size of active and inactive lists, of the free memory and zone limits. | ZONE1 (Normal): size 752MB, act 29MB, inact 31MB, free 658MB (0/1/2) | m, M |
| Memory zones fragmentation | Information on the memory zones fragmentation. This information describes how much system memory is fragmented and which is the biggest block size possible to allocate atomically. The first number before * is a number of blocks and the second is a block size in pages. | fragm 2*1 3*2 15*4 22*8 25*16 12*32 4*64 0*128 1*256 326*512". | m, M |
| Memory allocation latency | Memory allocation latency is an average time spent in the kernel memory allocator for different memory type requests. Any memory type is coded as XY, where X is A for GFP_ATOMIC, K for GFP_KERNEL and U for GFP_USER, and Y denotes the allocation request order, i.e. Y=0 for order=0 and 1 for order=1. | Mem lat (ms): A0 0, K0 0, U0 1, K1 3, U1 2 | m, M |
| Slab cache information | Slab cache information includes: the total slab cache size/real cache size divided into the inode cache size, dentry cache size, buffer heads cache size and page beancounters cache size. The real cache size is the size to which the cache can be shrunk, i.e. it is always less than the total cache size. | Slab pages: 13MB/13MB (ino 8MB, de 1MB, bh 1MB, pb 0MB) | m, M |
| Swap | Statistics on the used swap space, including the total swap space, the available swap space and the speed of swap-in/swap-out activity in MB/s. | Swap [OK]: tot 1004MB, free 1004MB, in 0.000MB/s, out 0.000MB/s | w, W |
| Swap latency | Swap operations latency. This includes the swap-in count, the swap-in maximal/average latency in ms, the swap-out count, the swap-out maximal/average latency in ms, and the maximal/average CPU time spent for the swap-out. | Swap lat: si 0, 0/0 ms, so 0, 0/0 ms, 0/0 cpu ms | w, W |
| Swap cache | Swap cache information includes the number of addition, deletion, and find operations respecting the swap cache. | Swap cache: add 0, del 0, find 0/0 | w, W |
| Network information | Network statistics summary includes the total incoming traffic speed in MB/s and incoming packets/s, and outgoing traffic speed in MB/s and outgoing packets/s. | Net [OK]: tot in 1.020MB/s 267pkt/s, out 0.001MB/s 1pkt/s | n, N |
| Network interface information | Provides the network statistics summary for a particular Ethernet interface, including its total incoming traffic speed in MB/s and incoming packets/s, and outgoing traffic speed in MB/s and outgoing packets/s. | eth0: in 0.000MB/s 3pkt/s, out 0.001MB/s 1pkt/s | n, N |

| | | | |
|--------------------------|---|--|------------------|
| Disks statistics | Disks statistics summary including the writing and reading activity in MB/s. | Disks [OK]: in 0.000MB/s, out 0.000MB/s | d, D |
| Mounted disks statistics | Information on the mounted disks such as their mount point, free space, and free inodes left on the device. | root(/) free: 3489MB(46%), 511077ino(52%) | d, D |
| Disk I/O statistics | Shows disk input and output statistics for virtual machines and Containers. The following statistics is displayed: <ul style="list-style-type: none"> IOUSED%: the percentage of time the disks are used by the virtual machine or Container. IOWAIT%: the percentage of time when at least one I/O transaction in the virtual machine or Container is waiting for being served. IO: the current speed of disk I/O in the virtual machine or Container and the I/O limit set for this virtual machine or Container, if any. The value can be displayed in bytes, kilobytes, megabytes, or gigabytes per second, depending on the units you used to set the I/O limit. IOPS: the current speed of disk I/O in the virtual machine or Container and the I/O limit set for this virtual machine or Container, if any. The value is displayed in operations per second. CPU_IOWAIT: percentage of time that Container processes lost waiting for I/O. CPU_IOWAIT_ACC: percentage of time that Container processes lost waiting for I/O, accumulated since utility start. SWAPIN: I/O operations with the swap space. The value is measured in operations per second. | IOUSED% 5.00 IOWAIT% 0.00 IO 2/100MB/s IOPS 315/--/s CPU_IOWAIT 25.0% ACC_IOWAIT 25.0% SWAPIN 0.0/s | A A A A |

Interactive Controls

Quite a number of single-key interactive commands can be used while `pstat` is running to change the way the utility displays information. The commands are not available if `pstat` runs with the `-t` or `-l` command-line option.

| Key | Action |
|------|----------------------|
| h, ? | Print a help screen. |

| | |
|--------------|---|
| space | Update display immediately. |
| q | Quit. |
| t | Change the delay between screen updates. You will be prompted to enter new delay time, in seconds. Entering 0 causes continuous updates. See also the <code>-d</code> command-line parameter. |
| b | Set the "brief" details level. See also the <code>-b</code> command-line parameter. |
| s | Set the "normal" details level. |
| v | Set the "verbose" details level. See also the <code>-v</code> command-line parameter. |
| a | "Averaged" mode. Monitoring parameters will be averaged through a minute. This includes: 1. Number of uninterruptible processes; 2. Scheduling max latency; 3. Memory allocation max latency; 4. Size of free/active/inactive memory; 5. Swap-in latency; 6. UBC fail-counters absolute values. |
| e | Toggle display of virtual machine and Container IP addresses/hostnames. |
| i | Toggle display of idle virtual machines and Containers. |
| l | Toggle display of load average. |
| p | Toggle display of processes statistics. |
| c | Toggle display of CPU usage statistics. |
| w | Toggle display of swap information. |
| m, M | Toggle/expand display of memory information. Each subsystem, including memory, network and disk has a number of verbosity levels. In the minimal level no information is displayed. Corresponding interactive lowercase key decreases verbosity level, the same key in uppercase increases it. |
| n, N | Toggle/expand display of network statistics. |
| A | Toggle display of disk input and output statistics for virtual machines and Containers. |
| I | Toggle display of input and output accounting for virtual machines and Containers. |
| X | Toggle display of network statistics for virtual machines and Containers. |
| d, D | Toggle/expand display of disk usage and activity information. |
| 0 | Toggle display of statistics for the server itself. |
| o | Sort key. You can use one of the following sort option keys: |
| | n Sort by Container ID |
| | c Sort by CPU usage |
| | f Sort by UBC failure counters |
| | r Sort by the number of running processes |
| | p Sort by the total number of processes |
| | t Sort by virtual machine or Container status. virtual machines and Containers which probably are unusable or unstable (increasing UBC failure counters or very high scheduling latency) will be shown first. |
| | s Sort by the number of open sockets |
| | m Sort by memory latency |
| | v Sort by virtual memory usage |
| | k Sort by kernel memory usage |

vzreboot

`vzreboot` is used to update Parallels Cloud Server without rebooting the physical server.

Syntax

```
vzreboot [kernel_version]
```

When run without options, the command checks the GRUB configuration file to determine the default kernel, loads this kernel into memory, and then boots the new kernel. You can also specify the kernel version manually (e.g., `vzreboot 2.6.32-042stab054.2`) to load and boot to the indicated kernel instead of the one set as default in the GRUB configuration file.

High Availability Utilities

This section describes the utilities you can use to manage High Availability (HA) support for physical servers and virtual machines and Containers hosted on them.

shaman General Syntax

The `shaman` utility is used to manage and monitor clustering resources (virtual machines and Containers) on servers with Parallels Cloud Server. You can use this utility to enable or disable High Availability (HA) support for servers, monitor global cluster events, and view the overall cluster statistics.

Syntax

```
shaman -c <cluster_name> [options...] [command [command_options]]
```

General Options (can be used with all commands)

| Name | Description |
|--------------------------------------|--|
| <code>-c <cluster_name></code> | Name of the cluster for which to execute the command. |
| <code>-l, --log <path></code> | Full path to the log file. |
| <code>-v, --verbose</code> | Increase the log verbosity level. |
| <code>-q, --quiet</code> | Disable printing of messages to the console. |
| <code>-V, --version</code> | Display the <code>shaman</code> version. |
| <code>-h, --help</code> | Display the list of allowed options for the <code>shaman</code> utility. |

Commands

| Name | Description |
|-------------------|-----------------------------------|
| <code>join</code> | Enable HA support for the server. |

| | |
|--------------|---|
| leave | Disable HA support for the server. |
| stat | Displays information about the current cluster status and clustering resources. |
| top | Displays a dynamic real-time view of the cluster and clustering resources. |

shaman join

The `shaman join` command is used to enable High Availability (HA) support for physical servers.

Syntax

```
shaman -c <cluster_name> join [-r, --ignore-res]
```

Options

| Name | Description |
|--------------------------------|---|
| -c <cluster_name> | Name of the cluster that contains the server for which you want to enable HA support. If you omit the option, the command is executed for the cluster specified in the <code>/etc/shaman/shaman.conf</code> configuration file. |
| join | Tells the <code>shaman</code> utility to enable HA support for the server. Enabling HA support means that <code>shaman</code> will control all virtual machines and Containers that are already hosted on the server and that you will create on it in the future. |
| -r, --ignore-res | Enable HA support for the server but leave it disabled for existing virtual machines and Containers. In this case, <code>shaman</code> will control only the virtual machines and Containers that you will create on the server after enabling HA support. |

shaman leave

The `shaman leave` command is used to disable High Availability (HA) support for physical servers.

Syntax

```
shaman -c <cluster_name> leave [-a, --addr <IP_address>]
```

Options

| Name | Description |
|--------------------------------|--|
| -c <cluster_name> | Name of the cluster that contains the server for which you want to disable HA support. |
| leave | Tells the <code>shaman</code> utility to disable HA support for the server. |

| | |
|--------------------------------|--|
| -a, --addr <IP_address> | Disables HA support for the server with the IP address of <i>ip_address</i> . If you omit this option, the <i>shaman</i> utility will turn off HA support for the local node. |
|--------------------------------|--|

shaman stat

The *shaman stat* command is used to display information about the current cluster status and clustering resources. To display a dynamic real-time view of the cluster and clustering resources, use the *shaman top* command.

Syntax

```
shaman -c <cluster_name> stat [-X, --xml] [-g, --group] [-n, --node-ip <IP_address>] [-N, --node-id <ID>]
```

Options

| Name | Description |
|-----------------------------------|---|
| -c <cluster_name> | Name of the cluster for which to execute the command. |
| -X, --xml | Display cluster statistics in XML format. |
| -g, --group | Group clustering resources according to their state into three groups: <ul style="list-style-type: none"> ▪ Active. Healthy virtual machines and Containers hosted in the cluster. ▪ Broken. Virtual machines and Containers that could not be relocated from a failed server to a healthy one. ▪ Pool. Virtual machines and Containers waiting for relocation from a failed server to a healthy one. |
| -n, --node-ip <IP_address> | Display statistics only for the server with the specified IP address. |
| -N, --node-id <ID> | Display statistics only for the server with the specified ID. |

shaman top

The *shaman top* command is used to display a dynamic real-time view of the cluster and clustering resources.

Syntax

```
shaman -c <cluster_name> top [-g, --group] [-n, --node-ip <IP_address>] [-N, --node-id <ID>]
```

Options

| Name | Description |
|-----------------------------------|---|
| -c <cluster_name> | Name of the cluster for which to execute the command. |
| -g, --group | Group clustering resources according to their state into three groups: <ul style="list-style-type: none">▪ Active. Healthy virtual machines and Containers hosted in the cluster.▪ Broken. Virtual machines and Containers that could not be relocated from a failed server to a healthy one.▪ Pool. Virtual machines and Containers waiting for relocation from a failed server to a healthy one. |
| -n, --node-ip <IP_address> | Display statistics only for the server with the specified IP address. |
| -N, --node-id <ID> | Display statistics only for the server with the specified ID. |

Interactive commands

| Name | Description |
|---------------------|--------------------------------------|
| g | Group or ungroup cluster resources. |
| v | Show or hide additional information. |
| ENTER, SPACE | Update the screen with statistics. |
| q | Quit the command. |
| h | Display the help screen. |

Parallels Cloud Server Updates

Parallels Cloud Server allows quick and easy updates with the `yum` utility standard for RPM-compatible Linux operating systems. For more information on `yum`, see **Updating Parallels Cloud Server** in the *Parallels Cloud Server 6.0 User's Guide* and the `yum` manual page.

CHAPTER 3

Managing Containers

Parallels Containers can be managed using the `prctl` command-line utility. The utility is installed on the Parallels server during the product installation.

In This Chapter

| | |
|---|-----|
| Matrix of Parallels Cloud Server Command-Line Utilities | 73 |
| <code>prctl</code> | 75 |
| Migration Utilities | 99 |
| Backup and Restoration Utilities | 103 |
| EZ Template Management Utilities | 108 |
| Supplementary Tools..... | 132 |

Matrix of Parallels Cloud Server Command-Line Utilities

The table below contains the full list of Parallels Cloud Server command-line utilities and command you can use for managing Containers.

Note: Starting from Parallels Cloud Server 6.0, numeric Container IDs and alphanumeric, case-sensitive Container names are interchangeable. Any command that requires a Container ID can be supplied that Container's name instead.

General Utilities

| Name | Description |
|-------------------------|--|
| <code>prctl</code> | Utility to control Containers. |
| <code>prctl list</code> | Utility to view a list of Containers existing on the server with additional information. |

Container Migration Utilities

| Name | Description |
|--------------------------|---|
| <code>prctl clone</code> | Command for the local cloning of the Containers. |
| <code>pmigrate</code> | Utility for migrating physical servers to Containers and for moving Containers between Parallels servers. |

Container Backup Utilities

| Name | Description |
|----------------------|---|
| pbackup | Utility to back up Containers. |
| prestore | Utility to restore backed up Containers. |
| prctl backup | Command to back up individual Containers. |
| prctl restore | Command to restore individual Containers. |

Template Management Utilities

| Name | Description |
|---------------------|--|
| vzpkg | Utility to manage OS and application EZ templates either inside your Containers or on the server itself. |
| vzmktpl | Utility to create OS and application EZ templates. |
| vzpkgproxy | Utility to create caching proxy servers for handling OS and application EZ templates. |
| vzcreaterepo | Utility to set up repositories for RHEL-like distributions, including Red Hat Enterprise Linux 5. |

Supplementary Utilities

| Name | Description |
|----------------------|---|
| vzps, vztop | Utilities working as the standard <code>ps</code> and <code>top</code> utilities, with Container-related functionality added. |
| pnetstat | Utility that prints network traffic usage statistic by Containers. |
| vzcpucheck | Utility for checking CPU utilization by Containers. |
| vzmemcheck | Utility for checking the server and Container current memory parameters. |
| vzcalc | Utility to calculate resource usage by a Container. |
| vzcheckovr | Utility to check the current system overcommitment and safety of the total resource control settings. |
| pstat | Utility to monitor the server and Container resources consumption in real time. |
| vzpid | Utility that prints Container id the process belongs to. |
| vzsplit | Utility to generate Container configuration file sample, “splitting” the server into equal parts. |
| vzcfgscale | Utility to scale the Container configuration. |
| vzcfgvalidate | Utility to validate Container configuration file correctness. |
| vzhwcalc | Utility to scan the main resources on any Linux server and to save the obtained information to a special file. |
| vzmtemplate | Utility to migrate the installed OS and application templates from the one server to another. |
| pfcache | Memory and IOPS deduplication management utility. |
| vzctl convert | Utility to convert Containers based on VZFS to the Container-in-an-image-file layout. |

| | |
|-----------------|---|
| pcompact | Utility to compact Containers by removing unused blocks from their virtual disks. |
|-----------------|---|

prctl

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

Syntax

```
prctl <subcommand> <CT_ID>
prctl --version
prctl --help
```

Subcommands

| Name | Description |
|--|--|
| create | Creates a new Container. |
| delete, destroy | 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. |

Options

| Name | Description |
|------------------|--|
| --version | Displays the <code>prctl</code> package version currently installed on the server. |
| --help | Displays the usage information about <code>prctl</code> . |

prctl 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 from the console, press Esc and then ".".

Syntax

```
prctl console <CT_ID>
```

Options

| Name | Description |
|---------|---------------|
| <CT_ID> | Container ID. |

prctl create

This command is used to create new Containers.

Syntax

```
prctl create <CT_ID> --vmttype ct [options]
```

With this command, you can create regular Containers. A unique Container ID is required for this command.

Note:

1. Container IDs from 1 to 100 are reserved for internal Parallels Cloud Server needs. Do not use IDs from 1 to 100 when creating Containers.
2. If a Container is created with an alphanumeric name, the next available numeric ID will be used as its private area directory name. If a Container is created with a numeric ID, the same number will be used as that Container's private area directory name.

Options

| Name | Description |
|---------------------|--|
| <CT_ID> | ID to assign to the new Container. |
| --vmttype ct | Tells the <code>prctl create</code> 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 <code>DEF_OSTEMPLATE</code> parameter in the global Parallels Cloud Server configuration file. |
| <code>--config <name></code> | Container sample configuration file to use for creating the Container. Sample configuration files are located in <code>/etc/vz/conf</code> and have names in the format <code>ve-<name>.conf-sample</code> . 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 Parallels Cloud Server configuration file, you will have to set resource control parameters for the Container using the <code>prlctl set</code> command. |
| <code>--uuid <uuid></code> | A custom UUID to assign to the Container. |
| <code>--private <path></code> | <p>Path to the Container private area. This option is used to override default path to private area from the <code>/etc/vz/vz.conf</code> configuration file (<code>VE_PRIVATE</code> variable). The argument can contain <code>\$VEID</code> string which will be replaced by numeric Container ID value.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>--root <path></code> | <p>Path to the mount point of the Container root directory. This option is used to override default path to Container root directory from the <code>/etc/vz/vz.conf</code> configuration file (<code>VE_ROOT</code> variable). The argument can contain <code>\$VEID</code> string which will be replaced by numeric Container ID value.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>--ipadd <addr>[/mask]</code> | <p>IP address and subnet mask to assign to the Container. If you omit this option, you can set an IP address for the Container later using the <code>prlctl set</code> command.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>--hostname <name></code> | <p>Hostname to assign to the Container. If you omit this option, you can assign a hostname to the Container later using the <code>prlctl set</code> command.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>--description <desc></code> | <p>Custom Container description.</p> <p>Descriptions with white spaces must be enclosed in quotation marks.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>--skip_app_templates</code> | <p>Do not install the application templates specified in the Container sample configuration file.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |

| | |
|---------------------------------|--|
| | utility. Use <code>vzctl</code> instead. |
| --velayout <velayout_ID> | <p>Defines the Container layout. <i>velayout_ID</i> can be one of the following:</p> <ul style="list-style-type: none"> 5: create a Container with the new Container-in-a-file layout. The new layout has been redesigned to store all files of a Container in a single image, similar to a virtual machine's hard disk. This is the default and recommended layout in Parallels Cloud Server 6.0. 4: create a Container based on VZFS, the obsolete layout used in earlier versions of Parallels Server Bare Metal. <p>If you omit this option, the layout will be defined by the <code>VEFSTYPE</code> parameter in the global configuration file <code>/etc/vz/vz.conf</code>. For the Container-in-an-image layout, this parameter should be set to <code>ext4</code>.</p> |

prctl delete, destroy

These commands are used to delete Containers from the server.

Syntax

```
prctl delete <CT_ID>
prctl destroy <CT_ID>
```

Options

| Name | Description |
|---------|---------------|
| <CT_ID> | Container ID. |

When executed, `prctl delete/prctl destroy` 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_ID>.conf` to `<CT_ID>.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.

prctl exec, enter

These commands are used to run arbitrary commands in a Container being authenticated as root on the server.

Syntax

```
prctl exec <CT_ID|CT_name> <command>
prctl enter <CT_ID|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.

Options

| Name | Description |
|------------------------------------|-----------------------|
| <code><CT_ID CT_name></code> | Container ID or name. |

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

`prctl enter` 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.

prctl mount, umount

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

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

Syntax

```
prctl mount <CT_ID>
prctl umount <CT_ID>
```

Options

| Name | Description |
|----------------------------|---------------|
| <code><CT_ID></code> | Container ID. |

prctl move

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

Syntax

```
prctl move <CT_ID|CT_name> --dst <path>
```

Options

| Name | Description |
|---------------------|---------------------------------------|
| <CT_ID CT_name> | Container ID or name. |
| --dst <path> | Path to Container's new private area. |

prctl problem-report

Generates a problem report for the specified Container and either sends it to the Parallels technical support or displays it on the screen.

Syntax

```
prctl problem-report <CT_ID> -d,--dump  
prctl problem-report <CT_ID> -s,--send [--proxy [user[:passwd]@]<proxyhost>[:port]] [--no-proxy]
```

Options

| Name | Description |
|---|--|
| <CT_ID> | ID 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 Parallels technical support to analyze your problem. |
| -s, --send | Send the generated problem report to the Parallels technical support. |
| --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. |

prctl register, unregister

The `register` command is used to register a Container with Parallels Cloud Server.

The `unregister` command removes a Container from the Parallels Cloud Server registry.

Syntax

```
prctl register <path> [--preserve-uuid <yes|no>]
prctl unregister <CT_ID>
```

Options

| Name | Description |
|--------------------------|---|
| <path> | Full path to the Container directory. |
| <CT_ID> | ID of the Container to remove from the Parallels Cloud Server 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. |

Remarks

- 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 Parallels Cloud Server. 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 Parallels Cloud Server registry, but does not delete the Container files from the server. You can re-register the Container later using the `register` command.

prctl reinstall

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

Syntax

```
prctl reinstall <CT_ID> [--skipbackup] [--resetpwdb] [--scripts <script> [...]]
prctl reinstall <CT_ID> [--listscripts] [--desc]
```

Options

| Name | Description |
|--------------------------|--|
| <CT_ID> | Container ID. |
| --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 <code>/old</code> 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 <code>--listscripts</code> option. |
|---------------|---|

Notes:

1. If any of the Container application templates cannot be added to the Container in a normal way, the reinstallation process will fail. This may happen, for example, if an application template was added to the Container using the `--force` option of the **vzpkgadd** command.
2. Currently, the `reinstall` command may not be supported by the `prlctl` utility. Use `vzctl` instead.

prlctl set

This command is used for setting Container parameters.

Syntax

```
prlctl set <CT_ID> <option> <value>
```

where <CT_ID> is Container ID.

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

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

Note: The current version of Parallels Cloud Server does not support using some options with the `prlctl set` command (e.g., `--ext_ipadd` or `--disabled`). If you run `prlctl set` with such an option, the command will notify you that the option is not supported. In this case, you can execute the `vzctl set` command instead.

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 Parallels Cloud Server automatically start this Container on next system startup. Note: If "yes" is specified as the value of this parameter in the <code>0.conf</code> file, all server system management parameters are set on the server boot to the values indicated in this file. |
| --offline_management <yes no> | Enabling/disabling the direct managing of the Container through |

| | |
|--|--|
| | a common Internet browser by means of Parallels Power Panels and the Plesk control panel (as defined by the <code>OFFLINE_SERVICE</code> parameter in the global or Container configuration file). |
| <code>--offline_service <service_name></code> | Defines whether the Container can be managed by means of Parallels Power Panel or Plesk or both. Valid only if the <code>OFFLINE_MANAGEMENT</code> parameter is set to "yes". The names of the available services can be taken from the file names (excluding the <code>.conf</code> extension) in the <code>/etc/vzredirect.d</code> directory on the server. |
| <code>--userpasswd <user:password></code> | 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 <code>/etc/passwd</code> and <code>/etc/shadow</code> files inside the Container. In case the Container root is not mounted, it is automatically mounted to apply the changes and then unmounted. |
| <code>--crypted</code> | Used with <code>--userpasswd</code> . Indicates that the specified password is already a hash. |
| <code>--noatime <yes no></code> | Sets the <code>noatime</code> flag (do not update inode access times) on the Container file system. The default is <code>yes</code> for a Class 1 Container, and <code>no</code> otherwise. Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead. |
| <code>--devnodes <device:r w rw none></code> | Lets the Container access the specified devices in the specified mode - read-only, write-only, or read-write - or denies any access. For example: <code>--devnodes hda1:rw</code> The device must be present in the Container <code>/dev</code> directory, otherwise, a new device is automatically created. Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead. |
| <code>--netdev_add <name></code> | Moves the specified network device from the server to the Container. For example: <code>--netdev_add eth0</code> . Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead. |
| <code>--netdev_del <name></code> | Moves the specified network device from the given Container to the server. Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead. |
| <code>--capability <name:on off></code> | Specifies capabilities inside the Container. Setting the following capabilities is allowed: <code>AUDIT_WRITE</code> , <code>CHOWN</code> , |

| | |
|---|--|
| | <p>DAC_OVERRIDE, DAC_READ_SEARCH, FOWNER, FS_MASK, FSETID, IPC_LOCK, IPC_OWNER, KILL, LEASE, LINUX_IMMUTABLE, MKNOD, NET_ADMIN, NET_BIND_SERVICE, NET_BROADCAST, NET_RAW, SETGID, SETFCAP, SETPCAP, SETUID, SYS_ADMIN, SYS_BOOT, SYS_CHROOT, SYS_MODULE, SYS_NICE, SYS_PACCT, SYS_PTRACE, SYS_RAWIO, SYS_RESOURCE, SYS_TIME, SYS_TTY_CONFIG, VE_ADMIN.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>--features <name:on off></code> | <p>Enables/disables the support for the following functionality inside the Container:</p> <ul style="list-style-type: none"> ▪ <code>nfs</code>: mounting NFS shares ▪ <code>ipip</code>: creating IPIP tunnels ▪ <code>sit</code>: using the Simple Internet Transition (SIT) mechanisms ▪ <code>ppp</code>: using the PPP protocol ▪ <code>ipgre</code>: creating IP-GRE tunnels ▪ <code>bridge</code>: using bridges to connect virtual Ethernet devices ▪ <code>nfsd</code>: running an NFS-kernel-space server <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>--root <path></code> | <p>This setting does NOT move the root mount point of your Container to a new path. It simply overrides the <code>VE_ROOT</code> parameter in the Container configuration file.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>--private <path></code> | <p>This setting does NOT move the private area of your Container to a new path. It simply overrides the <code>VE_PRIVATE</code> parameter in the Container configuration file. You should use this option only if you have manually moved the Container private area to a new place and want to update the Container configuration file.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>--setmode <restart ignore></code> | <p>This option tells the utility either to restart or not restart the Container after applying any parameters requiring that the Container be rebooted for them to take effect.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |

| | |
|---|---|
| --disabled <yes no> | <p>If set to <i>yes</i>, disables the Container making it impossible to start the Container once it was stopped. The disabled Container can be started by passing the <i>--force</i> option to <i>prlctl</i> set.</p> <p>Note: Currently, this option may not be supported by the <i>prlctl</i> utility. Use <i>vzctl</i> instead.</p> |
| --name <name> | <p>An arbitrary name assigned to the Container. This name can be used, along with the Container ID, to refer to the Container while performing certain Container-related operations on the server.</p> <p>Names must contain the A-Z, a-z, 0-9, \, -, _ symbols only. Names with white spaces must be enclosed in quotation marks.</p> |
| --description <desc> | <p>Custom Container description.</p> <p>Descriptions must contain the A-z, 0-9 symbols only. Descriptions with white spaces must be enclosed in quotation marks.</p> |
| --bindmount_add [<i>src</i> :] <dst> [, <i>nosuid</i> , <i>noexec</i> , <i>nodev</i>] | <p>Mounts a source directory (<i>src</i>) located on the server to a destination directory (<i>dst</i>) inside the Container. If the source directory is not specified, mounts the directory to the <i>/vz/root/CT_ID</i> directory.</p> <p>Additional options that can be used with <i>--bindmount_add</i> are the following:</p> <ul style="list-style-type: none"> ▪ <i>noexec</i>. Do not allow execution of any binaries on the mounted directory. ▪ <i>nodev</i>. Do not interpret character or block special devices on the mounted directory. ▪ <i>nosuid</i>. Do not allow set-user-identifier or set-group-identifier bits to take effect. <p>Note: Currently, this option may not be supported by the <i>prlctl</i> utility. Use <i>vzctl</i> instead.</p> |
| --bindmount_del <dst all> | <p>Removes the mount point created by using the <i>--bindmount_add</i> option from the Container.</p> <p>Note: Currently, this option may not be supported by the <i>prlctl</i> utility. Use <i>vzctl</i> instead.</p> |
| --vnc-mode <auto manual off> | <p>Enables or disables access to the Container via the VNC protocol.</p> |
| --vnc-port <port> | <p>Sets the VNC port number for the Container. Used with <i>--vnc-mode manual</i>.</p> |
| --vnc-passwd <passwd> / --vnc-nopasswd | <p>Sets the VNC password for the Container or specifies that no password is needed for VNC connections.</p> <p>Either of these options is mandatory for any VNC setup.</p> |
| --autocompact <on off> | <p>Enables or disables compaction for all disks in the Container.</p> |

Note: For details on how to enable or disable compaction for a specific disk in the Container, see **Hard Disk Drive Management Options** (p. 94).

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>-p, --numproc <bar:lim></code> | Number of processes and threads allowed. Upon hitting this limit, the Container will not be able to start new process or thread. In this version of Parallels Cloud Server, the limit shall be set to the same value as the barrier. Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead. |
| <code>--numtcpsock <bar:lim></code> | Number of TCP sockets (<code>PF_INET</code> family, <code>SOCK_STREAM</code> type). This parameter limits the number of TCP connections and, thus, the number of clients the server application can handle in parallel. In this version of Parallels Cloud Server, the limit shall be set to the same value as the barrier. Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead. |
| <code>--numothersock <bar:lim></code> | Number of socket other than TCP. Local (UNIX-domain) sockets are used for communications inside the system. UDP sockets are used for Domain Name Service (DNS) queries, for example. In this version of Parallels Cloud Server, the limit shall be set to the same value as the barrier. Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead. |

| | |
|---|---|
| <code>-e, --numiptent <bar:lim></code> | <p>Number of IP packet filtering entries.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>--vmguarpages <bar:lim></code> | <p>Memory allocation guarantee, in pages (one page is 4 KB). 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 may fail in case of overall memory shortage. In this version of Parallels Cloud Server, the limit shall be set to the same value as the barrier.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>-k, --kmemsize <bar:lim></code> | <p>Size of unswappable kernel memory (in bytes), allocated for internal kernel structures of the processes of a particular Container. Typical amounts of kernel memory are 16...50 KB per process.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>--tcpsndbuf <bar:lim></code> | <p>Total size (in bytes) of send buffers for TCP sockets – amount of kernel memory allocated for data sent from an application to a TCP socket, but not acknowledged by the remote side yet.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>-b, --tcprecvbuf <bar:lim></code> | <p>Total size (in bytes) of receive buffers for TCP sockets. Amount of kernel memory received from the remote side but not read by the local application yet.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>--othersockbuf <bar:lim></code> | <p>Total size in bytes of UNIX-domain socket buffers, UDP and other datagram protocol send buffers.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>--dgramrcvbuf <bar:lim></code> | <p>Total size in bytes of receive buffers of UDP and other datagram protocols.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>--oomguarpages <bar:lim></code> | <p>Out-of-memory guarantee, in 4-KB pages. Any Container process will not be killed even in case of heavy memory shortage if the current memory consumption (including both physical memory and swap) does not reach the <code>oomguarpages</code> barrier. In this version of Parallels Cloud Server, the limit shall be set to the same value as</p> |

| | |
|--|---|
| | <p>the barrier.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>-l, --lockedpages <bar:lim></code> | <p>Memory not allowed to be swapped out (locked with the <code>mlock()</code> system call), in 4-KB pages.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>--shmpages <bar:lim></code> | <p>Total size of shared memory (including IPC, shared anonymous mappings and <code>tmpfs</code> objects), allocated by processes of a particular Container, in 4-KB pages.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>--physpages <bar:lim></code> | <p>The total size of RAM used by processes, in 4-KB pages. This is accounting-only parameter currently. It shows the usage of RAM by the Container. For memory pages used by several different Containers (mappings of shared libraries, for example), only a fraction of a page is charged to each Container. The sum of the <code>physpages</code> usage for all Containers corresponds to the total number of pages used in the system by all accounted users.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>--privvmpages <bar:lim></code> | <p>Size in 4-KB pages of private (or potentially private) memory, allocated by Container applications. Memory that is always shared among different applications is not included in this resource parameter.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>-n, --numfile <bar:lim></code> | <p>Number of files opened by all Container processes. In this version of Parallels Cloud Server, the limit shall be set to the same value as the barrier.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>-f, --numflock <bar:lim></code> | <p>Number of file locks created by all Container processes.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>-t, --numpty <bar:lim></code> | <p>Number of pseudo-terminals. For example, <code>ssh</code> session, <code>screen</code>, <code>xterm</code> application consumes pseudo-terminal resource. In this version of Parallels Cloud Server, the limit shall be set to the same value as the barrier.</p> |

| | |
|---|--|
| | <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>-i, --numsignifo <bar:lim></code> | <p>Number of <code>signifo</code> structures (essentially, this parameter limits size of signal delivery queue). In this version of Parallels Cloud Server, the limit shall be set to the same value as the barrier.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>-x, --dcachesize <bar:lim></code> | <p>Total size, in bytes, of <code>dentry</code> and <code>inode</code> structures locked in memory. Exists as a separate parameter to impose a limit causing file operations to sense memory shortage and return an error to applications, protecting from excessive consumption of memory due to intensive file system operations.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>--cpuunits <units></code> | <p>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.</p> <p>If this parameter is not set, the default value of 1000 is used.</p> |
| <code>--cpulimit <percent megahertz></code> | <p>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.</p> <p>By default, the limit is set in percent. To set the limit in MHz, specify "m" after the value.</p> <p>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%.</p> |
| <code>--cpus <num></code> | <p>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.</p> <p>This option also defines the number of CPUs shown to Container users.</p> <p>This parameter is not set for newly created Containers; so they can consume all free CPU power of the server.</p> |
| <code>--cpumask <num></code> | <p>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).</p> |
| <code>--nodemask <num all></code> | <p>The NUMA node mask defining a NUMA node to bind the Container to.</p> <p>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.</p> |

| | |
|--|---|
| | <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>--diskspace <bar:lim></code> | Total disk space consumed by the Container, in 1-KB blocks. In the current version of Parallels Cloud Server, both values must be equal. |
| <code>--quotauidlimit <0 N></code> | <p>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 <code>quota</code> utility.</p> <p>Keep in mind the following:</p> <p>Enabling per-user and per-group quotas for a Container requires restarting the Container.</p> <p>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.</p> <p>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.</p> |
| <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> | <p>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:</p> <ul style="list-style-type: none"> ▪ G: sets the limit in gigabytes per second. ▪ K: sets the limit in kilobytes per second. ▪ B: sets the limit in bytes per second. <p>In the current version of Parallels Cloud Server, the maximum I/O bandwidth limit you can set for a Container is 2 GB per second.</p> <p>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.</p> |
| <code>--iopslimit <num></code> | <p>The maximum number of disk input and output operations per second a Container is allowed to perform.</p> <p>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.</p> |
| <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>--reset_ub</code> | <p>Resets the current values of all system parameters of the server to the ones set in the <code>0.conf</code> file.</p> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>--memsize</code> | <p>The amount of RAM that can be used by the processes of a Container, in megabytes. You can use the following suffixes to set the RAM in other measurement units:</p> <ul style="list-style-type: none"> ▪ G in gigabytes ▪ K in kilobytes ▪ B in bytes |
| <code>--swappages</code> | <p>The amount of swap space that can be used by the Container for swapping out memory once the RAM is exceeded, in 4-KB pages. You can use the following suffixes to set swap in other measurement units:</p> <ul style="list-style-type: none"> ▪ G in gigabytes ▪ M megabytes ▪ K in kilobytes ▪ B in bytes |
| <code>--swap</code> | <p>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 set swap in other measurement units:</p> <ul style="list-style-type: none"> ▪ G in gigabytes ▪ M in megabytes ▪ K in kilobytes |
| <code>--vm_overcommit</code> | <p>Memory overcommit factor that defines the memory allocation limit for a Container. The limit is calculated as</p> $(\text{PHYSPAGES} + \text{SWAP}) * \text{factor}$ <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> |

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 that Container can use, and other parameters.

| Name | Description |
|---|---|
| <code>--hostname <name></code> | Sets the hostname to the specified name. |
| <code>--ipadd <addr></code> | <p>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.</p> <p>If used with the <code>--ifname</code> option, adds an IP address to the specified Container virtual network adapter.</p> |
| <code>--ipadd <addr/net_mask></code> | <p>Assigns the IP address and network mask to the Container.</p> <div> <p>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 Parallels Containers configuration file is set to <code>yes</code>.</p> </div> |
| <code>--ipdel <addr all></code> | <p>Allows you to revoke IP address from the Container. If “all” is used instead of IP address than all IP addresses will be revoked.</p> <p>If used with the <code>--ifname</code> option, deletes an IP address from the specified Container virtual network adapter.</p> |
| <code>--ext_ipadd <addr></code> | <p>Assigns the external IP address to the Container. External IP addresses are considered valid IP addresses by the <code>venet0</code> adapter, though they are not set as alias addresses inside Containers and are not announced via Address Resolution Protocol (ARP). You can assign the same external IP address to several Containers, irrespective of whether they reside on the same or different Hardware Nodes.</p> <div> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> </div> |
| <code>--ext_ipdel <addr all></code> | <p>Removes the external IP address from the Container. To delete all external IP addresses assigned to the Container, specify <code>--ext_ipdel all</code>.</p> <div> <p>Note: Currently, this option may not be supported by the <code>prlctl</code> utility. Use <code>vzctl</code> instead.</p> </div> |
| <code>--nameserver <addr></code> | <p>The DNS server IP address for the Container.</p> <p>If used with the <code>--ifname</code> option, sets the DNS server for the specified Container virtual network adapter.</p> |
| <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> | <p>Indicates which <code>iptables</code> modules are allowed for the Container.</p> <p>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.</p> <p>The following modes are available:</p> <ul style="list-style-type: none"> ▪ disabled: none. |

| | |
|--|---|
| | <ul style="list-style-type: none"> ▪ stateless: (default) all modules except conntrack and NAT-related. ▪ stateful: all modules except NAT-related. ▪ full: all modules. <p>Note: Currently, this option may not be supported by the <code>prctl</code> utility. Use <code>vzctl</code> instead.</p> |
| <code>--netif_add <name>[,mac,host_mac]</code> | <p>Creates a new <code>veth</code> virtual network adapter and assigns the name of <i>name</i> 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 <code>veth</code> adapter:</p> <ul style="list-style-type: none"> ▪ <i>mac</i>: the MAC address to be assigned to the <code>veth</code> Ethernet interface inside the Container. ▪ <i>host_mac</i>: the MAC address to be assigned to the <code>veth</code> Ethernet interface on the server. <p>Only the Ethernet interface name (<i>name</i>) is mandatory; all the other parameters, if not specified, are automatically generated by Parallels Cloud Server during the <code>veth</code> adapter creation.</p> |
| <code>--netif_del <name></code> | Removes the <code>veth</code> virtual network adapter with the specified name from the Container. |
| <code>--ifname <name></code> | Specifies the name of the <code>veth</code> 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 <code>veth</code> 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 <code>veth</code> 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 <code>veth</code> virtual Ethernet interface on the server. Should be used along with the <code>--ifname</code> option. |
| <code>--network <network_ID></code> | <p>Connects the <code>veth</code> virtual network adapter to the bridge associated with the specified network ID. Should be used along with the <code>--ifname</code> option.</p> <p>You can also use this option to disconnect the <code>veth</code> virtual network adapter from the bridge. To this effect, you should specify " " after the option.</p> |
| <code>--dhcp <yes no></code> | <p>Defines the IP assignment type for the <code>veth</code> virtual network adapter:</p> <ul style="list-style-type: none"> ▪ <code>yes</code> enables the dynamic IP address allocation for the Container. ▪ <code>no</code> turns off the dynamic IP address allocation for the |

| | |
|--|---|
| | Container. Should be used along with the <code>--ifname</code> option. |
| <code>--gw <addr></code> | Set the default gateway for the <code>veth</code> virtual network adapter. Should be used along with the <code>--ifname</code> option. |
| <code>--apply-iponly <yes no></code> | If set to <code>yes</code> , the hostname, nameserver, and search domain settings from the Container configuration file are ignored. |

Hard Disk Drive Management Options

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

Syntax

```
prctl set <CT_ID> --device-add hdd [--image <file>] [--size <size>] [--mnt <path>]
prctl set <CT_ID> --device-set hdd<N> [--image <file>] [--size <size>] [--mnt <path>]
      [--autocompact <on|off>]
prctl set <CT_ID> --device-del hdd<N> [--detach-only | --destroy-image]
```

Options

| Name | Description |
|--|---|
| <code><CT_ID></code> | Container ID. |
| <code>--device-add hdd</code> | <p>Adds a virtual hard disk to the Container.</p> <p>If no other options are specified, the command creates a new unmounted disk with the following parameters:</p> <ul style="list-style-type: none"> name: <code>hdd<N></code> where <code><N></code> is the next available disk index. size: 65536 MB image location: <code>/vz/private/<CT_ID>/disk-<ID>.hdd</code>. <p>Note: In the current version of Parallels Cloud Server, only expanding SATA disks can be added to Containers.</p> |
| <code>--device-set hdd<N></code> | <p>Modifies the parameters of the virtual hard disk <code>hdd<N></code>.</p> <p>Note: For the list of disks, use the <code>prctl list -i</code> command.</p> |
| <code>--image <file></code> | <p>Specifies an image file that will be used to emulate the virtual disk.</p> <ul style="list-style-type: none"> 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. |
| <code>--size <size></code> | Specifies the size of the virtual hard disk, in megabytes. |
| <code>--mnt <path></code> | <p>Specifies the mount point of the virtual hard disk inside the Container.</p> <p>A corresponding entry is also added to Container's <code>/etc/fstab</code> file, so the disk is</p> |

| | |
|---------------------------------|--|
| | 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 a specific disk in the Container, see General Options (p. 82). |
| --device-del hdd<N> | Deletes a virtual hard disk from the stopped Container. |
| --detach-only | Removes the virtual disk from the Container configuration but leaves its image file intact. |
| --destroy-image | Removes the virtual disk from the Container configuration and deletes its image file. |

prctl snapshot, snapshot-list, snapshot-switch, snapshot-delete

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

Syntax

```
prctl snapshot <CT_ID> [-n,--name <name>] [-d,--description <desc>]
prctl snapshot-list <CT_ID> [-t,--tree] [-i,--id <snapshot_ID>]
prctl snapshot-switch <CT_ID> -i,--id <snapshot_ID> --skip-resume
prctl snapshot-delete <CT_ID> -i,--id <snapshot_ID>
```

Options

| Name | Description |
|--------------------------|--|
| <CT_ID> | Container ID. |
| -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> | <ul style="list-style-type: none"> Use with <code>prctl 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>prctl snapshot-switch</code> to specify the ID of the snapshot to revert to. Use with <code>prctl snapshot-delete</code> to specify the ID of the snapshot to delete. |
| --skip-resume | Skips automatic Container resume when switching to snapshots of running Containers. |

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

prctl start, stop, restart, status

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

Syntax

```
prctl start <CT_ID|CT_name> [--wait]
prctl stop <CT_ID|CT_name> [--fast]
prctl restart <CT_ID|CT_name>
prctl status <CT_ID|CT_name>
```

Options

| Name | Description |
|-----------------|-----------------------|
| <CT_ID CT_name> | Container ID or 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 `prctl 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.

When starting a Container, `prctl` executes a number of helper scripts located in the `/vz/private/<CT_ID>/scripts` (the first and last scripts in the table) and `/etc/vz/conf` (all the other scripts in the table) directories

Scripts (in the order of execution):

| Name | Description |
|---------------------|---|
| mount | Optional Container mount script. If it exists, then it is executed immediately after mounting the Container private area. If it exits with a non-zero status, then <code>prctl</code> dismounts the Container private area and returns the error. |
| vz-start | This script sets up IP traffic accounting for the Container. |
| vz-net_add | This script creates the necessary ARP entries and sets up the necessary routing entries for Container IP addresses. |
| ve-alias_add | This script configures the network interfaces inside the Container. |
| ve-veconfig | This script is called by <code>prctl</code> to set a hostname and DNS search domains inside the Container. |
| ve-quota | If the second-level (per-user/group) quota is turned on, then <code>prctl</code> calls this script to form the correct <code>/etc/mtab</code> file inside the Container. |
| start | Optional Container start script. If it exists, then it is executed in the context of a just started Container. |

`prctl 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, `prctl` 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. The normal shutdown sequence of `prlctl stop` is described below in the order of execution:

- 1 stop.** Optional Container stop script. If it exists, then it is executed in the context of the Container prior to any other actions. If it exits with non-zero status, then `prlctl` does not stop the Container.
- 2 umount.** Optional Container unmount script. If it exists, then it is executed after stopping the Container but before dismounting its private area.
- 3 vz-stop.** This script deletes routing and IP traffic accounting for the Container.

You should use action scripts (`mount/umount` and `start/stop`) if you would like to carry out some actions upon the Container startup/shutdown. However, there might be situations when you have to modify other scripts documented above. In this case it is strongly suggested that you create a separate script containing all your modifications and add an invocation of this script to shipped scripts. This will facilitate upgrades to future Parallels Cloud Server versions.

The `prlctl restart <CT_ID>` 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.

prlctl suspend, resume

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

Syntax

```
prlctl suspend <CT_ID>
```

Options

| Name | Description |
|---------|---------------|
| <CT_ID> | Container ID. |

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_ID/dump` directory on the server (or in the directory specified as the value of the `DUMPDIR` parameter in the Parallels Cloud Server global file).

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

Syntax

```
prlctl resume <CT_ID>
```

When executed, `prctl resume` searches for the Dump file in the `/vz/private/CT_ID/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.

prctl list

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

Syntax

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

Options

| Name | Description |
|---|--|
| <code>-a, --all</code> | 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. |
| <code>-t, --template</code> | List available Container templates instead of actual Containers. |
| <code>-o, --output <field>[,...]</code> | Display only the specified fields. For the list of fields, see prctl list Output Parameters (p. 98). Type field names in lower case. Separate multiple fields with commas. |
| <code>-s, --sort <field -field></code> | Sort Containers by the specified field in either ascending or descending order. |
| <code>-i, --info</code> | Display detailed information about the specified Container. |
| <code>-f, --full</code> | Display detailed information about network cards in Containers. Used with the <code>--info</code> option. |
| <code><CT_ID CT_name></code> | The ID or name of the Container for which to display the detailed information. If not specified, the information will be displayed for all registered Containers. |
| <code>-j, --json</code> | Produce machine-readable output in the JSON format. |

prctl list Output Parameters

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

| Name | Output Column | Description |
|-------------------|-------------------|---------------|
| <code>uuid</code> | <code>UUID</code> | Container ID. |

| | | |
|--------------------|-------------|--|
| 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., <code>redhat-el5-x86</code>). |
| 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. |

Migration Utilities

This section describes the utilities you can use to migrate Containers between physical servers or move/clone Containers within the same server.

pmigrate

The `pmigrate` utility is used to perform different kinds of Container migration.

Syntax

```
pmigrate <source_server> <destination_server> [options]
```

<source_server> can be either the physical server hosting the virtual machine or Container to move or the physical computer to migrate. <destination_server> is the server where the resulting Container is to be hosted. If the source and/or destination server is not specified, the operation is performed on the local server.

<source_server> and <destination_server> consist of two parts:

- <type> denotes the type of computer to migrate and can be one of the following:
 - `h` must be specified when migrating a physical server.

- `c` must be specified when migrating a Container.
- `v` must be specified when migrating a virtual machine.
- `<address>` denotes the location of computer to migrate and can be one of the following:
 - The computer location if you are migrating a physical server.
 - The computer location and the virtual machine name or Container ID if you are migrating a virtual machine or Container, respectively. The location must be separated from the virtual machine name/Container ID by the slash (/).

The location format is as follows:

```
[user[:password]@]<dst>[:port]
```

where `<dst>` is the destination server IP address or hostname.

The `[options]` you can use with `pmigrate` depend on whether you are migrating to a virtual machine or a Container. This section describes the parameters for migrating Containers between Parallels servers and for moving virtual machines and physical servers to Containers. For information on performing these operations for virtual machines, see **pmigrate** (p. 171).

Container-related Options

The following options can be used with `pmigrate` when migrating Containers between Parallels servers:

| Name | Description |
|--|---|
| <code>-s, --nostart</code> | Do not attempt to start the Container on the destination server after its successful migration if the Container was running on the source server prior to the migration. This option does not have any effect if the Container was not running on the source server. |
| <code>-r, --remove-area <yes no></code> | This option takes precedence of the <code>REMOVEDMIGRATED</code> setting from the global configuration file. If "yes" is specified, then the Container private area and configuration file will be deleted after successful migration. If "no" is specified, the private area and configuration file will be left on the source server and have the <code>.migrated</code> suffix appended to them. |
| <code>-f, --nodeps</code> <code>[,all][,cpu_check]</code> <code>[,disk_space]</code> <code>[,technologies]</code> <code>[,license][,rate]</code> | <p>During its execution, <code>pmigrate</code> performs a number of checks on the destination server (e.g., it verifies that all OS and application templates required for the Container are present on the destination server) and if some checks fail, exits with an error. This option allows you to bypass all checks and migrate the Container. If you specify this option for a running Container, the Container will not be automatically started on the destination server. You should manually start it after adding the missing templates.</p> <p>You can additionally use one or several of the following parameters with this option:</p> <ul style="list-style-type: none"> ▪ <code>all</code>: do not perform any checks on the destination server. ▪ <code>cpu_check</code>: do not check the CPU capabilities of the Destination Server. ▪ <code>disk_space</code>: do not check the amount of disk space on the destination server. |

| | |
|----------------------------------|--|
| | <ul style="list-style-type: none"> ▪ <code>technologies</code>: do not check a set of technologies provided by the Parallels Cloud Server kernel on the destination server (for the description of the <code>TECHNOLOGIES</code> parameter, see Container Configuration File). ▪ <code>license</code>: do not check the license installed on the destination server. ▪ <code>rate</code>: do not check the value of the <code>RATE</code> parameter in the Parallels Cloud Server global file. |
| -b, --batch | Normally, you do not have to specify this option. It is used by Parallels Cloud Server scripts and changes the screen output to a computer-parsable form. |
| --ssh=<ssh_options> | Additional options to be passed to <code>ssh</code> while connecting to the destination server. Note: Do not specify the destination server hostname as an option of <code>--ssh</code> . |
| --keep-dst | Do not remove the 'synched' Container private area on the destination server if some error occurred during the migration. This option allows you to prevent <code>pmigrate</code> from the repeated 'synching' the Container private area if the first migration attempt failed for some reason or other. |
| --online | Migrates the running Container with zero downtime. By default, the 'iterative online migration' type is used. During the migration: <ol style="list-style-type: none"> 1. The main amount of Container memory is transferred to the destination server. 2. The Container is 'dumped' and saved to an image file. 3. The image file is transferred to the destination server where it is 'undumped'. Using this type of online migration allows you to attain the smallest service delay. To not use the 'iterative online migration' type, supply the <code>--noiter</code> option. |
| --noiter | Can be used only together with the <code>--online</code> option. Sets the migration type to 'simple'. This option cannot be used together with the <code>--require-realtime</code> option. |
| --require-realtime | Can be used only together with the <code>--online</code> option. Forces <code>pmigrate</code> to move the Container by using the 'iterative online migration' type. If this migration type cannot be carried out for some reason, the command will fail and exit. This option cannot be used together with the <code>--noiter</code> option. If the default 'iterative online migration' type cannot be carried out, and this option is omitted, <code>pmigrate</code> will try to move your Container by using the 'simple online migration' type. |
| --readonly | Just copy the specified Container to the destination server without making any changes to the Container on the Source Server. |
| --convert-vzfs | Convert legacy Containers to the Container-in-an-image-file layout. |
| --dry-run | Simulate the same operations as <code>pmigrate</code> completes without specifying this option (connects to the destination server, verifies that all OS and application templates required for the Container are present on the server, etc.); however, the Container itself is not moved to the destination server. |
| --non-sharedfs | Use to migrate Containers from a shared file system to a non-shared one. |

Virtual Machine-related Options

The following options can be used with `pmigrate` when migrating a physical server or a virtual machine to a Container:

| Name | Description |
|----------------------|---|
| -c | Mandatory. The full path to the configuration file on the Parallels server that was created on the physical server by means of the <code>vzhwcalc</code> utility. You can specify only the name of the configuration file if you run the <code>vzp2v</code> utility from the directory where this file is located. |
| -q, --quota | Optional. The partition on your physical server which has any per-user/group quotas imposed on it. This partition will be migrated to the Container together with all quotas imposed on it. Moreover, these quotas will be applied to the entire Container after the server migration. |
| -z, --eztmpl | Optional. The EZ OS template to be used to create the Container. You may list all OS templates installed on the Parallels server together with their updates by executing the <code>vzpkg list</code> command. If an OS template is not specified, the <code>mkvzfs</code> command is executed during the Container creation which makes an empty private area with the name of <code>/vz/private/CT_ID</code> on the Parallels server. This private area is then used to copy all the physical server files to it. |
| -d, --dist | Optional. The Linux version your physical server is running. The name of the version specified should coincide with the name of the corresponding distribution configuration file located in the <code>/etc/vz/conf/dist</code> directory on the Parallels server. For example, if you specify <code>rhel-5</code> as the value of this option, the <code>rhel-5.conf</code> file should be present in the <code>/etc/vz/conf/dist</code> directory on the Parallels server. You must obligatorily set this option, if there is no <code>DISTRIBUTION</code> variable specified in the server configuration file. In case the <code>DISTRIBUTION</code> variable is set in the configuration file and you have specified the <code>-d</code> option, the latter takes precedence. |
| --exclude | Optional. The path to the directories and files which will be excluded from copying to the Container. This option allows you to avoid migrating the data you do not need. To gain more understanding on this option, please consult the man pages for the <code>rsync</code> utility from where it was borrowed. Note: We strongly recommend that you exclude the directories you were informed of while running the <code>vzhwcalc</code> utility on the physical server. |
| -S, --srvstop | Optional. The services to be stopped for the time of the physical server migration. We recommend that you stop all the services on the physical server except for the critical ones (e.g., the <code>sshd</code> service that is needed to provide communication between the physical server and the Parallels server) before the migration. This will prevent the running services from modifying any files being moved. |

prctl clone

Creates an exact copy of the specified Container.

Syntax

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

Options

| Name | Description |
|--------------------------|--|
| <CT_ID> | ID of the Container to clone. |
| --name <new_name> | Name or ID to be assigned to the new Container. |
| --template | Create a Container template instead of a clone. You cannot start a template. Note: Container templates do not have any practical use in Parallels Cloud Server 6.0 and are mainly used in Parallels Automation for Cloud Infrastructure. |
| --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. |

Backup and Restoration Utilities

Any Container is defined by its private area, configuration files, action scripts, and quota information. Backing up these components allows you to restore all the content of a Container on any Parallels Cloud Server-based system at any time if the Container gets broken.

Note: For the `vzabackup/vzarestore` functionality to work, forward and reverse DNS lookups must be correctly configured for both the source and backup servers.

pbackup

The `pbackup` utility is run on the so-called backup server. It connects via SSH to the servers where some or all Containers are to be backed up and puts the tarballs into the directory defined in the `/etc/vzbackup.conf` global backup configuration file (the default directory is `/var/parallels/backups`). Later on, the Container backups may be restored from this directory.

Syntax

```
pbackup [backup_options] server1 ... [CT_options]
```

You may specify any number of servers in the command-line by names or IP addresses. You may also enter these names as the value of the `BACKUP_NODES` parameter in the global backup configuration file to avoid the necessity to specify them in the command-line. In this case, you shall specify the `-a` option instead.

Notes:

1. This section describes backup options for Containers. For backup options that can be used with virtual machines, see **pbackup** (p. 174).
2. While the following two subsections provide the complete reference on the **pbackup** and **prestore** utilities, many of their options can be specified in the `/etc/vzbackup.conf` configuration file to be used as the default ones.

Backup Options

| Name | Description |
|-----------------------------------|---|
| -F, -l | Create a full backup. A full backup contains all virtual machine data. |
| -i | Make an incremental backup or, if no full backups are available, a full backup. An incremental backup contains only the file that were changed since the previous full or incremental backup. |
| --ssh-opts <options> | Options to be passed to <code>ssh</code> . See examples in the global backup configuration file. |

These options define the list of Containers to back up:

| Name | Description |
|---|--|
| -e <CT_ID CT_name> [...] | The Containers to back up on the server. Containers can be specified using both their IDs (e.g., 101 or 102) and their names (e.g., comp1 or comp2). If the <code>--vzcache</code> option is specified, not the Containers themselves, but their caches will be backed up. |
| -x <CT_ID CT_name> [...] | The Containers that need not be backed up (Containers to exclude). Containers can be specified using both their IDs (e.g., 101 or 102) and their names (e.g., comp1 or comp2). If the <code>--vzcache</code> option is specified, not the Containers themselves, but their caches will be excluded from backing up. |

It is sufficient to specify only one Container option: either `-e`, or `-x`; or to do without any Container options if all the Containers from the specified server(s) are to be backed up.

Container-specific Backup Options

These options must be used with the `vzctl` command instead of `prlctl`.

| Name | Description |
|--|--|
| -Cg -Cb -Cn | Compress the resulting Container backups with <code>gzip</code> (<code>-Cg</code>), <code>bzip2</code> (<code>-Cb</code>) or do not use compression (<code>-Cn</code>). This option overrides the <code>BACKUP_COMPRESS</code> parameter in the backup configuration file. |
| -a | Back up all servers specified in the global backup configuration file. |
| -c <config> | Use an alternative backup configuration file. |
| -s | The Containers are to be stopped before their backing up. In this case, if a client |

| | |
|-----------------------------|--|
| | tries to access the Containers during their downtime, a temporary "busy" page is shown. This option takes precedence of the <code>BACKUP_VESTOP</code> parameter in the backup configuration file. |
| -n | The Containers are NOT to be stopped before their backing up. This option takes precedence of the <code>BACKUP_VESTOP</code> parameter in the backup configuration file. |
| -p | Apply time and load restriction rules for periodic backups. These rules are defined by the <code>BACKUP_KEEP_MAX</code> and <code>BACKUP_LOADAVG_MAX</code> parameters in the backup configuration file (<code>/etc/vzbackup.conf</code>). Without this option, these parameters do not take effect. This option is useful if you invoke <code>pbackup</code> in unattended mode as a cron job. It overrides the <code>CRON_BACKUP</code> parameter in the global backup configuration file. |
| --desc <desc> | The description of the backup archive. Descriptions with white spaces must be enclosed in quotation marks. |
| -j | Is opposite to the <code>-p</code> switch. Turns off the periodic backup mode and disregards the <code>BACKUP_KEEP_MAX</code> and <code>BACKUP_LOADAVG_MAX</code> parameters in the backup configuration file (<code>/etc/vzbackup.conf</code>). It overrides the <code>CRON_BACKUP</code> parameter in the global backup configuration file. |
| -L | Make use of the <code>BACKUP_FINISH_TIME</code> and <code>BACKUP_LIMIT_TIME</code> in the backup configuration file (<code>/etc/vzbackup.conf</code>). |
| --rm-tag <tag> | Create a backup and then remove the backup with the specified tag. You can learn the tags of the existing backups on the server by using, for example, the <code>prestore -l</code> command. |
| --rm-old | Create a backup and then remove the oldest backup of the specified server/Container(s). |
| --vzcache | Back up not the Containers themselves, but the cache area of the server (<code>/vz/template/vzcaches</code>). |

prestore

The `prestore` utility is also run on the backup server. It uses the Container backups stored on the backup server to restore them to their original servers (or to any other location if the `-d` option is specified).

Syntax

```
prestore [restoration_options] server [...] [CT_options]
```

You can specify multiple servers (by names or IP addresses) whose Containers were at one time backed up and now need to be restored.

Restoration Options

| Name | Description |
|--------------------------|--|
| -c <config> | Use an alternative backup configuration file. |
| -d <node> | The destination server to restore the Containers to. If no destination server is |

| | |
|-----------------------------|--|
| | specified, the Containers are restored to the server from which they were originally backed up. |
| -t <tag> | Specify the tag of any intermediary incremental backup to be restored. |
| -r <tag> | Remove the backup tagged by the <tag> value. This option is valid only if a single Container is specified and if the <tag> value specifies the last incremental backup. |
| --rm-prev <tag> | Remove the backup tagged by the <tag> value together with all the previous backups. This option is valid only if a single Container is specified and if the <tag> value specifies the last incremental backup. |
| -l | Do not restore any Containers. Show the information on the Containers available to be restored. |
| -f | Show the full information on the backed up Containers (only if the -l option is specified). |
| --chain | Show a tag chain (only if the -l option is specified). |
| --desc <config> | Display the configuration file inside the archive with the specified tag (only if the -l and -t options are specified). |
| --single-tar | Indicates that a single tar is coming (stdin). |
| --skip-check-vzcache | Do not restore the backups of the Container cached files stored in the <code>/var/parallels/backups/CT_ID/vzcache</code> directory on the backup server. For example, you should use this option while restoring a Container with cached files to a destination server other than the backup server. |
| --vzcache | Restore the <code>vzcache</code> template area keeping the Container cached files and located in the <code>/vz/template/vzcache</code> directory on the server. |

The Container options define a list of Containers to be restored:

| Name | Description |
|---------------------------------|--|
| -e <CT_ID/CT_name> [...] | The Containers to be restored on the server. Any Container can be specified using both its IDs (e.g., 101 or 102) and its names (e.g., comp1 or comp2). |
| -x <CT_ID/CT_name> [...] | The Containers that need not be restored (Containers to exclude). Any Container can be specified using both its IDs (e.g., 101 or 102) and its names (e.g., comp1 or comp2). |

It is sufficient to specify only one Container option: either **-e**, or **-x**; or to do without any Container options if all the Containers from the specified servers are to be restored.

prctl backup, backup-list, backup-delete, restore

Creates, lists, deletes or restores Container backups.

Syntax

```
prctl backup <CT_ID> [-f,--full] [-i,--incremental] [--description <desc>]
                    [-s,--storage [user[:passwd]@]server[:port]]
prctl backup-list [CT_ID] [-f,--full] [--localvms]
                  [-s,--storage [user[:passwd]@]server[:port]]
prctl backup-delete {<CT_ID> | -t,--tag <backup_ID>}
```

```

[-s,--storage [user[:passwd]@]server[:port]]
prlctl restore {<CT_ID> | -t,--tag <backup_ID>} [-n, --name <new_name>] [--dst <path>]
[-s,--storage [user[:passwd]@]server[:port]]

```

Options

| Name | Description |
|--|---|
| <CT_ID> | <p>Container ID.</p> <ul style="list-style-type: none"> Use with <code>prlctl backup</code> to create a backup of the specified Container. Use with <code>prlctl backup-list</code> to list backups of the specified Container. Use with <code>prlctl backup-delete</code> to delete all backups of the specified Container. Use with <code>prlctl restore</code> to restore the most recent backup of the specified Container. |
| -s, --storage [user[:passwd]@]server[:port]] | <p>Specifies a remote backup server address, port, and credentials.</p> <p>If this option is omitted, the backup will be saved on the default backup server that can be configured using the prlsrvctl set command.</p> |
| --description <desc> | <p>Backup description.</p> <p>Descriptions with white spaces must be enclosed in quotation marks.</p> |
| -f, --full | <ul style="list-style-type: none"> Use with <code>prlctl backup</code> to create a full backup of the Container. A full backup contains all Container data. Use with <code>prlctl backup-list</code> to display full backup information. |
| -i, --incremental | <p>Create an incremental backup of the Container. An incremental backup contains only the files changed since the previous full or incremental backup.</p> <p>This is the default backup type.</p> |
| --localvms | List local backups only. |
| -t, --tag <backup_ID> | The ID of the backup to restore or delete. |
| -n, --name <new_name> | <p>A new name to assign to the restored Container.</p> <p>If this option is omitted, the Container will be restored with the original name.</p> |
| --dst <path> | <p>Restore the Container to the specified directory on the Parallels server.</p> <p>If this option is omitted, the Container will be restored to <code>/var/parallels/<vm_name></code>.</p> |

EZ Template Management Utilities

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

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_ID|object>  
vzpkg --help
```

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. |
| 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. |

vzpkg install template

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

Syntax

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

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

Options

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

vzpkg update template

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

Syntax

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

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

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.

vzpkg remove template

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

Syntax

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

Options

| Name | Description |
|--|--|
| -F, --for-os < <i>OS_template</i> > | 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.

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_ID|CT_name [...]]
```

If you indicate a Container ID or 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.

Options

| Name | Description |
|----------------------|--|
| -p, --package | Lists the software packages installed in the Container or included in the OS EZ template. |
| -g, --groups | 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. |
| -O, --os | Displays the OS EZ template the Container is based on. |

| | |
|--------------------------------|--|
| -A, --app | Displays the application EZ templates installed in the Container or included in the OS EZ template. |
| -C, --cache | 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. |
| -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 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. |
| -u, --custom-pkg | 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. |
| -i, --pkgid | Displays the ID assigned to the EZ template instead of its name; these IDs are unique within the given system. If the <code><CT_ID></code> or <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. |
| -S, --with-summary | In addition to listing the EZ templates available either in the Container (if the <code><CT_ID></code> or <code><CT_name></code> argument is given) or installed on the server (if the <code><CT_ID>/<CT_name></code> argument is omitted), this option makes <code>vzpkg list</code> display the summary information on the corresponding EZ templates/packages. |
| -c, --cached | This option has no effect if the <code><CT_ID></code> or <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. |
| appcache | Outputs a list of OS EZ template caches with preinstalled applications. |
| -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. |

vzpkg info

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

Syntax

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

Options

| Name | Description |
|---|---|
| <i>CT_ID</i> <i>CT_name</i> | Container ID or name. |
| <i>OS_template</i> | OS EZ template. |
| <i><app_template></i> | Application EZ template. |
| <i><package_name></i> | Software package name. |
| <i><yum_package_group></i> | YUM software group name. |
| -F, --for-os <i><OS_template CT_ID CT_name></i> | Displays information on the application EZ template or the software package (if the -p option is specified) included in the specified OS EZ template or applied to the indicated Container. |
| -p, --package | Displays information about the specified software package. Must be used with the -F option. |
| -g, --groups | Displays information about the packages included in the specified YUM software group. |
| -C, --cache | 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. |
| -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 info</code> get the information on the specified package from the remote repositories set for handling the EZ template where this package is included. |
| -d, --debug <i><num></i> | 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 name or the Container ID. 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.

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 <code>summary</code> , it can contain additional data on the EZ template/package. |
| technologies | <p>Displays the following information:</p> <ul style="list-style-type: none"> ▪ The microprocessor architecture where the EZ template is to be used (<code>x86</code>, <code>x86</code>); ▪ Specifies whether the EZ template can be used only on the servers with the Native POSIX Thread Library (NPTL) support. In this case the <code>nptl</code> entry is displayed after the <code>vzpkg info</code> execution. <p>For EZ templates only.</p> |
| version | The version of the software package. |
| release | The release of the software package. |
| arch | <p>The system architecture where the EZ template/package is to be used. It can be one of the following:</p> <ul style="list-style-type: none"> ▪ <code>x86</code> if the EZ template/package is to be used on 32-bit platforms. ▪ <code>x86_64</code> if the EZ template is to be used on x86-64-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 | <p>The packaging system used to handle the packages included in the specified EZ template. It can be one of the following:</p> <ul style="list-style-type: none"> ▪ <code>rpm</code> for RPM-based Linux distributions (Fedora Core, Red Hat Enterprise Linux, etc.); ▪ <code>dpkg</code> for Debian-based Linux distributions (e.g., Debian and Ubuntu). <p>For EZ templates only.</p> |
| package_manager | <p>The package manager type for managing the packages included in the specified EZ template. It can be one of the following:</p> <p>x86 Linux distributions</p> <ul style="list-style-type: none"> • <code>rpm49db5x86</code>: Fedora 17 • <code>rpm49x86</code>: Fedora 15 and 16 • <code>rpm47x86</code>: Red Hat Enterprise Linux 6 and CentOS 6 • <code>rpm44x86</code>: Red Hat Enterprise Linux 5 and CentOS 5 • <code>rpm43x86</code>: Red Hat Enterprise Linux 3 and 4, CentOS 3 and 4 |

| | |
|--|--|
| | <ul style="list-style-type: none"> • <code>rpmzypp44x86</code>: SUSE Linux Enterprise Server 11 with Service Pack 2 • <code>rpm41x86</code>: SUSE Linux Enterprise Server 10 and SUSE Linux 10.x • <code>rpm41s9x86</code>: SUSE Linux Enterprise Server 9 • <code>rpmzypp49x86</code>: openSUSE 12.1 • <code>dpkg</code>: Debian and Ubuntu <p>x86-64 Linux distributions</p> <ul style="list-style-type: none"> • <code>rpm49db5x64</code>: Fedora 17 • <code>rpm49x64</code>: Fedora 15 and 16 • <code>rpm47x64</code>: Red Hat Enterprise Linux 6 and CentOS 6 • <code>rpm44x64</code>: Red Hat Enterprise Linux 5 and CentOS 5 • <code>rpm43x64</code>: Red Hat Enterprise Linux 3 and 4, CentOS 3 and 4 • <code>rpmzypp44x64</code>: SUSE Linux Enterprise Server 11 with Service Pack 2 • <code>rpm41x64</code>: SUSE Linux Enterprise Server 10 and SUSE Linux 10.x • <code>rpm41s9x64</code>: SUSE Linux Enterprise Server 9 • <code>rpmzypp49x64</code>: openSUSE 12.1 • <code>dpkgx64</code>: Debian and Ubuntu |
|--|--|

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_ID/CT_name/OS_template>
```

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.

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_ID|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.

Options

| Name | Description |
|----------------------|--|
| -p, --package | Installs a software package instead of an EZ template. |
| -g, --groups | Installs a YUM software group instead of an EZ template. The <code>-g</code> option works only for Containers running RPM-based Linux distributions. |
| -f, --force | Forces the EZ template/package installation. |
| -C, --cache | 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 |

| | |
|--------------------------------|---|
| | /etc/vztt/vztt.conf file. |
| -r, --remote | If the elapsed time from the last <code>vzpkg install</code> 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. |
| -n, --check-only | 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. |
| -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 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.

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_ID|CT_name> [object [...]]
```

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.

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_ID|CT_name> <object> [...]
```

This command will remove *<object>* from the Container with the ID of *<CT_ID>* or with the name of *<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.

Options

| Name | Description |
|--------------------------------|--|
| -p, --package | Removes the specified package(s) from the Container. |
| -g, --groups | Removes the specified YUM software group from the Container. The <code>-g</code> 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.

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> [...]]
```

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.

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>] [...]
```

Options

| Name | Description |
|---------------------|---|
| -C, --cache | 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. |
| -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 cache</code> check for the packages updates in the remote repositories set for handling the given EZ OS template. |

`vzpkg update cache` 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., `redhat-el5-x86.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.

vzpkg remove cache

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

Syntax

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

Options

| Name | Description |
|---------------------------------------|--|
| -d, --debug <i><num></i> | 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.

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]
```

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. |

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]
```


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. |

vzpkg remove appcache

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

Syntax

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

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. |

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_ID|CT_name> <rpm_file_path> [...]
```

Options

| Name | Description |
|--------------------------------|--|
| -C, --cache | When handling the package interdependencies, makes the <code>vzpkg localinstall</code> command look for the needed packages 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. |
| -r, --remote | If the elapsed time from the last <code>vzpkg</code> local 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 localinstall</code> look for the packages in the remote repository. |
| -n, --check-only | Simulates the same operations as <code>vzpkg localinstall</code> 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 ID of `<CT_ID>` or with the name of `<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.

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_ID|CT_name> <rpm_file_path> [...]
```

Options

| Name | Description |
|--------------------------------|---|
| -C, --cache | When handling the package interdependencies, makes the <code>vzpkg localupdate</code> command look for the needed packages 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. |
| -r, --remote | If the elapsed time from the last <code>vzpkg</code> local 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 localupdate</code> look for the packages in the remote repository. |
| -n, --check-only | Simulates the same operations as <code>vzpkg localupdate</code> 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 ID of `<CT_ID>` or the name of `<CT_name>` and updates it, if necessary. You may specify a number of packages at once to be updated inside your Container.

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_ID|CT_name>
```

Options

| Name | Description |
|---------------------|--|
| -C, --cache | Makes the <code>vzpkg upgrade</code> command check for the packages included in the OS EZ template in the local <code>vzpkg</code> 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 <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 upgrade</code> will also check the local <code>vzpkg</code> cache only. |
| -r, --remote | If the elapsed time from the last local <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 upgrade</code> check for the packages in the remote repositories set for handling the given EZ OS template. |
| -n, --check-only | Simulates the same operations as <code>vzpkg upgrade</code> 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. |
| -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. |

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>
```

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 <code>vzpkg fetch</code> command look for the metadata in the <code>vzpkg</code> local 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 fetch</code> 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.

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> [...]]
```

Options

| Name | Description |
|--------------------------------|---|
| -k, --clean-packages | Removes the packages, headers, and metadata of the specified EZ OS template from the local <code>vzpkg</code> cache. This is also the default behavior of <code>vzpkg clean</code> . |
| -t, --clean-template | Checks the template area for the specified EZ OS template (the template area has the default path of <code>/vz/template</code>) 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: <ul style="list-style-type: none"> the packages, headers, and metadata of the specified EZ OS template from the <code>vzpkg</code> 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 <code>vzpkg clean</code> execution. |
| -n, --check-only | Simulates the same operations as <code>vzpkg clean</code> 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. |

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 ...]
```

Options

| Name | Description |
|--------------------|--|
| -C, --cache | Makes the <code>vzpkg update metadata</code> command look for available metadata |

| | |
|--------------------------------|---|
| | 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 metadata</code> 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.

vzmktmpl

This utility is used to create new EZ templates.

Syntax

```
vzmktmpl [options] <metafile>
```

Options

| Name | Description |
|------------------------------------|--|
| --pre-cache <file> | The path to the script which will be executed by the <code>vzpkg cache</code> command before installing the packages included in the EZ template on the server. This script is executed in the server context and relevant for OS EZ templates only. |
| --post-cache <file> | The path to the script which will be executed by the <code>vzpkg cache</code> command after installing the packages included in the EZ template on the server. This script is executed in the server context and relevant for OS EZ templates only. |
| --pre-install <file> | The path to the script which will be executed by the <code>vzpkg install</code> command before adding the application EZ template to the Container. This script is executed in the Container context and relevant for application EZ templates only. |
| --post-install <file> | The path to the script which will be executed by the <code>vzpkg install</code> command after adding the application EZ template to the Container. This script is executed in the Container context and relevant for application EZ templates only. |
| --pre-upgrade <file> | The path to the script which will be executed by the <code>vzpkg upgrade</code> command before upgrading the OS EZ template inside the Container. This script is executed in the Container context. |
| --post-upgrade <file> | The path to the script which will be execute by the <code>vzpkg upgrade</code> command after upgrading the OS EZ template inside the Container. This script is executed in the Container context. |
| --pre-update <file> | The path to the script which will be executed by the <code>vzpkg update</code> command before updating the packages included in the application EZ template inside the |

| | |
|-----------------------------|---|
| | Container. This script is executed in the Container context. |
| --post-update <file> | The path to the script which will be executed by the <code>vzpkg update</code> command after updating the packages included in the application EZ template inside the Container. This script is executed in the Container context. |
| --pre-remove <file> | The path to the script which will be executed by the <code>vzpkg remove</code> command before removing the application EZ template from the Container. This script is executed in the Container context and relevant for application EZ templates only. |
| --post-remove <file> | The path to the script which will be executed by the <code>vzpkg remove</code> command after removing the application EZ template from the Container. This script is executed in the Container context and relevant for application EZ templates only. |
| --environment <file> | The path to the file storing a list of environment variables. The variables should be set in the form of <code>key=value</code> . The variables specified in this file are used when running the <code>vzpkg create cache</code> and <code>vzpkg update cache</code> commands and exported to the Container environment during the EZ template scripts execution. |
| -d, --doc <file> | The path to the file containing the information on the EZ template. You can specify several files and separate them by commas. |
| -s, --spec-only | Creates the package specification file only. |
| -r, --srpm | Creates the package source file only. |
| -h, --help | Displays the utility usage and exits. |

The only required parameter for `vzmktmpl` is the path to the metafile containing the template configuration. In most cases, however, you also need to include a number of scripts in the template. These scripts will be executed at different stages of the template life cycle (e.g., when you add the template to a Container) and ensure the correct template operation. To learn what scripts are required for your template:

- 1 Install an official template provided by Parallels for the same Linux distribution for which you are creating the template. For example, if you are making a template for CentOS 5, install the CentOS 5 OS template (use the `vzupdate -z` command, if necessary).
- 2 Go to the configuration directory (`/vz/template/TemplateName/TemplateArch/config/os/default`) of the installed template.
- 3 Look for the files starting with "pre-" and "post-" (e.g, pre-cache or post-install).

Note: The `vzmktmpl` utility is a part of the `vztt-build` package, which is not installed by default. If installed, the package can be found in the `/virtuozzo/RPMS/optional` directory of your Parallels Cloud Server distribution. So, before you can start using the `vzmktmpl` utility, you first need to install the package on your server with the `rpm -i` command.

vzpkg.metafile

This file is used by the `vzmktmpl` utility as the basis for the EZ template creation. The parameters in this file are presented on separate lines in the following format:

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

Parameters

| Name | Description |
|----------------------------|---|
| %osname | Mandatory. The name of the Linux distribution for which you are creating the OS EZ template or under which the application EZ template being created is to be run. |
| %osver | Mandatory. The version of the Linux distribution specified as the value of the %osname parameter. |
| %osarch | Mandatory. The microprocessor architecture where the EZ template is to be run. You can set the value of this parameter to one of the following: <ul style="list-style-type: none"> ▪ x86: this value should be specified if your EZ template is to be used on 32-bit platforms. ▪ x86-64: this value should be specified if your EZ template is to be used on x86-64-bit platforms (e.g., on servers with the AMD Opteron and Intel Pentium D processors installed). |
| %appname | Mandatory, for application EZ templates only. The name of the application EZ template. |
| %setname | Optional. The name of the non-base OS EZ template, if any. This parameter should be specified only while creating non-base OS EZ templates. |
| %upgradable_version | Optional. A list of Linux distribution versions which can be upgraded to the version of the Linux distribution for which you are creating the EZ template. For OS EZ templates only. Should be omitted for application templates. |
| %packages | Mandatory. A list of software packages to be included in the EZ template. The names of the packages listed as the value of this parameter should correspond to the names of real packages that are stored in the repository used for managing your EZ templates and can be specified in one of the following ways: For RPM-based Linux distributions: <ul style="list-style-type: none"> ▪ as a package name only (e.g., <code>wget</code>) ▪ as a package name with the indication of the system architecture on which the package is to be run (e.g., <code>wget.i386</code>, <code>wget.noarch</code>) ▪ as a package name with its versions (e.g., <code>wget-1.9.1</code>) ▪ as a package name with its versions and release number (e.g., <code>wget-1.9.1-12</code>) ▪ as a package name with its version, release number, and system architecture (e.g., <code>wget-1.9.1-12.i386</code>) ▪ as a package name with its version, release number, system architecture, and epoch number (e.g., <code>10:wget-1.9.1-12.i386</code>) For Debian-based Linux distributions: <ul style="list-style-type: none"> ▪ as a package name only (e.g., <code>wget</code>) ▪ as a package name with its version (e.g., <code>wget-1.9.1-12</code>) |
| %packages_0 | Mandatory, for Debian-based OS EZ templates only. A list of packages to be used for creating a minimal Debian/Ubuntu <code>chroot</code> environment. These packages should |

| | |
|-------------------------|--|
| | correspond to those installed on a standalone server on the first stage of the Ubuntu distribution installation. The packages will be installed on the server one by one in the specified order during the OS EZ template caching. If you wish several packages to be simultaneously installed on the server, you should specify the package names on a single line and separate them by spaces. |
| %packages_1 | Mandatory, for Debian-based OS EZ templates only. A list of 'base' packages for the Debian/Ubuntu distribution. These packages are needed to install the packages listed as the value of the %packages parameter. |
| %package_manager | <p>Mandatory. The short name of the package manager to be used for handling the EZ template. Depending on the Linux distribution for which you are creating the template or under which the template will be used, you should set the following values for the PKGMAN parameter:</p> <p>x86 Linux distributions</p> <ul style="list-style-type: none"> ▪ 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 <p>x86-64 Linux distributions</p> <ul style="list-style-type: none"> ▪ 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 <p>This parameter must be specified for all base OS EZ templates and should be omitted for application EZ templates and non-base OS EZ templates.</p> |
| %repositories | Mandatory, for RPM-based Linux distributions only. A list of repositories where the packages comprising the EZ template are stored. |
| %mirrorlist | Mandatory. One or several URLs to the file containing a list of repositories from where |

| | |
|----------------------|---|
| | the packages comprising the EZ template are to be downloaded. This parameter can be omitted if you are creating a metafile for an application EZ template or a non-base OS EZ template. |
| %distribution | Optional. The type of the Linux distribution. Examples of Linux distribution types are <code>centos</code> , <code>debian</code> , <code>fedora-core</code> , <code>gentoo</code> , <code>mandrake</code> , <code>redhat</code> , <code>rhel-3</code> , <code>rhel-4</code> , <code>rhel-5</code> , <code>fedora-core-4</code> , <code>fedora-core-5</code> , <code>slackware</code> , <code>slackware-10.0</code> , <code>suse</code> , <code>suse-9.3</code> , etc. Should be omitted for application templates. |
| %description | Optional. Detailed information on the EZ template package file. |
| %version | Optional. The version of the EZ template package file. |
| %release | Optional. The release of the EZ template package file. |
| %license | Optional. The information about the owner of the EZ template package file. |
| %changelog | Optional. The information about the changes made to the EZ template package file. |

vzpkgproxy

The `vzpkgproxy` utility is used to set up a caching proxy server meant for handling OS and application EZ templates. The `vzpkgproxy` package where the `vzpkgproxy` utility is included can be installed by using the `rpm -i` command on any computer meeting the following requirements:

- The Apache `httpd` server, version 2.0.52 and higher, should be installed on the workstation.
- The `createrepo` package, version 0.4.2 and higher, should be installed on the workstation.

During its installation, the utility performs all the tasks necessary to install, configure, and put into operation your caching proxy server. Detailed information on how to set up caching proxy servers is given in the *Parallels Cloud Server 6.0 Templates Management Guide*.

vzcreaterepo

The `vzcreaterepo` utility is used to set up repositories for RHEL-like distributions, including Red Hat Enterprise Linux 5.

Syntax

```
vzcreaterepo [options] [mount_points]
```

Options

| Name | Description |
|--|---|
| <code>mount_points</code> | Path to the directory with Linux distribution files. If omitted, <code>vzcreaterepo</code> searches for files in the <code>/media/cdrom</code> directory. |
| <code>-d, --dir <repo_path></code> | Path to the directory for setting up the repository. If omitted, the repository is created in the <code>/vz/template/repos</code> directory. |

| | |
|---|--|
| -t, --template <i><template></i> | Name of the template for which to create the repository (e.g., <code>redhat-e15-x86_64</code>). |
| -m, --mirror | Creates a Container and configures it to act as an HTTP mirror for servers in your network. Requires the <code>-v</code> , <code>-n</code> , <code>--nameserver</code> , and <code>--ip</code> parameters to be specified. |
| -q, --quiet | Disables logging to the screen. |
| -u, --update | Creates a repository with updates. At the moment, this parameter is valid for Red Hat Enterprise Linux 5 only. |
| -i, --iso <i><path1>[,path2...]</i> | Path to the an ISO image containing Linux distribution files. You can specify several images and separate them by commas. |
| -a, --cache | Caches the template after setting up the repository. |
| -p, --proxy <code>http://<IP_or_hostname>:<port></code> | Hostname or the IP address and the port number of the HTTP proxy server, if you use any to connect to the Internet. |
| --proxy_user <i><proxy_user></i> | User name used by the HTTP proxy server for your authentication. |
| --proxy_pass <i><proxy_passwd></i> | Password of the user specified in the <code>--proxy_user</code> parameter and used for your authentication by the HTTP proxy server. |
| -e, --email <i><email></i> | Email address for registering in the Red Hat Network. This is the email address you specified when creating an account in the Red Hat Network. |
| --systemname <i><system_name></i> | Name of your system to use for retrieving packages in the Red Hat Network. If omitted, the name <code>updates</code> is used. |
| -r, --repotemplate <i><repo_template></i> | OS template to use for creating the Container that will act as an HTTP mirror. If omitted, the <code>centos-5-x86_64</code> template is used. |
| -v, --ctid <i><CT_ID></i> | ID to assign to the Container. |
| -n, --hostname <i><hostname></i> | Hostname to assign to the Container. |
| --nameserver <i><nameserver></i> | Nameserver to set for the Container. |
| --ip <i><IP_address></i> | IP address to assign to the Container. |

Note: The current version of `vzcreaterepo` does not support creating repositories for Red Hat Enterprise Linux 6 and all versions of openSUSE and SUSE Linux Enterprise Server.

vzmtemplate

The `vzmtemplate` utility is used to migrate the installed OS and application standard templates and OS EZ templates from the Source Server to the Destination Server.

Syntax

```
vzmtemplate [-bhz] [--ssh=<ssh_options>]
               [user_name@]<destination_server_IP_address> <template_name> [...]
```

Options

| Name | Description |
|----------------------------------|--|
| -z, --eztempl | Migrates the specified OS EZ template(s) installed on the Source Server. Without this option specified, <code>vzmttemplate</code> moves standard OS and application templates. |
| --ssh=<ssh_options> | Additional <code>ssh</code> options to be used while connecting to the Destination Server. |
| -b, --batch | This option should be passed to <code>vzmttemplate</code> in scripts if you are going to use these scripts for running the <code>vzmttemplate</code> utility and do not wish them to analyze the <code>vzmttemplate</code> command output. |
| -h, --help | Displays the utility usage and exits. |

To migrate a template, you should execute the `vzmttemplate` command on the Source Server and pass the corresponding options to it. During its execution, the utility will try to connect to the Destination Server with the IP address of *Destination_server_IP_Address* and move the specified template(s) to this server. By default, `vzmttemplate` logs in to the Destination Server as `root` and asks you for the password of this user. However, you can make the utility use other credentials to log in to the Destination Server by appending the corresponding user name with the `@` symbol to the Server IP address (e.g., `user1@192.168.0.123`). Keep in mind that the specified user should have the `root` privileges; otherwise, the command will fail.

Supplementary Tools

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]]
```

Options

| Name | Description |
|----------------------------------|--|
| -v | Increase the command output verbosity. Multiple <code>-v</code> 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. |

pfcache

Memory and IOPS deduplication management utility that enables/disables caching for Container directories and files, verifies cache integrity, checks Containers for cache errors, and purges the cache if needed.

Note: The utility does not support additional disks attached to Container.

Syntax

```
pfcache <command> <file|dir> [options]
```

Commands

| Name | Description |
|------------------------|---|
| mark (p. 134) | Enables caching of the specified files, directories or subdirectories in Containers. |
| unmark | Disables caching of the specified files, directories or subdirectories in Containers. |
| purge (p. 135) | Frees up space in the memory and IOPS deduplication cache image. |
| verify (p. 136) | Verifies the integrity of the specified mounted memory and IOPS deduplication cache and deletes corrupted files. |
| check (p. 133) | Checks for and fixes memory and IOPS deduplication cache errors in the specified Container root directory. |
| stat (p. 135) | Displays inodes summary for a running Container. |
| dump (p. 134) | In addition to the summary provided by stat, shows detailed information on PFCache inodes of a running Container. If no options are specified, outputs full information on checksummed peer inodes. |

pfcache check

Checks for and fixes memory and IOPS deduplication cache errors in the specified Container root directory.

Syntax

```
pfcache check <dir> [--dry-run]
```

Options

| Name | Description |
|------------------|---|
| <dir> | Container root directory. |
| --dry-run | Report errors but do not make changes to the file system. |

pfcache dump

In addition to the summary provided by `stat`, shows detailed information on PFCache inodes of a running Container. If no options are specified, outputs full information on checksummed peer inodes.

Syntax

```
pfcache dump <dir> [--csummed | {--all|-a} | <csum>] [--column <col>...]
```

Options

| Name | Description |
|-------------------|---|
| <dir> | Container root directory. |
| --csummed | Outputs information on checksummed inodes. |
| --all -a | Outputs information on all inodes. |
| <csum> | Outputs information on inodes with the specified checksum. |
| --column <col>... | Displays only the chosen column(s). <col> can be: <ul style="list-style-type: none">▪ a — All▪ h — File handle▪ p — Page cache size, pages▪ c — Checksum▪ f — Filter state▪ s — File size, bytes |

pfcache mark, unmark

Enables or disables caching of the specified files, directories or subdirectories in Containers.

Syntax

```
pfcache mark <file>
pfcache mark <dir> [subdir ...] [--recursive]
pfcache unmark <file>
pfcache unmark <dir> [subdir ...] [--recursive]
```

Options

| Name | Description |
|--------|---|
| <file> | File to enable/disable caching of. |
| <dir> | Directory to enable/disable caching of. |

| | |
|--------------------|---|
| <i>subdir</i> | Subdirectory to enable/disable caching of. |
| --recursive | Process directory's or subdirectory's current contents. |

pfcache purge

Frees up space in the memory and IOPS deduplication cache image. If no options are specified, purges entire cache.

Syntax

```
pfcache purge <cache_dir> [--unused | --size <size> | --expire <date>]
```

Options

| Name | Description |
|------------------------|---|
| <cache_dir> | Memory and IOPS deduplication cache image location. |
| --unused | Remove only files unused at the moment. |
| --size <size> | Attempt to free <i>size</i> bytes in the memory and IOPS deduplication cache image. |
| --expire <date> | Remove files accessed before the specified date. A date can be specified in the ISO 8601 format or in the format defined in your system locale, with optional hours, minutes, and seconds. Examples: <ul style="list-style-type: none"> 05/21/12 [09:50[:33]] 2012-05-21 [09:50[:33]] |

pfcache stat

Displays inodes summary for a running Container.

Syntax

```
pfcache stat <dir>
```

Options

| Name | Description |
|-------|---------------------------|
| <dir> | Container root directory. |

pfcache verify

Verifies the integrity of the specified mounted memory and IOPS deduplication cache and deletes corrupt files.

Syntax

```
pfcache verify <cache_dir>
```

Options

| Name | Description |
|-------------|---|
| <cache_dir> | Memory and IOPS deduplication cache image location. |

pnetstat

This utility outputs traffic usage statistics for Containers.

Syntax

```
pnetstat [-v <CT_ID>] [-c <class>] [-a] [-r] [-h] [-t <ct|vm|all>]
```

The utility displays input and output traffic for Containers for each defined network class, in bytes. The network classes are described in the `/etc/vz/conf/networks_classes` file. `pnetstat` displays statistics only on Containers that were started at least once.

Options

| Name | Description |
|----------------|--|
| -v <ID> | Display statistics for the specified Container. Multiple -v options can be given to a single <code>pnetstat</code> invocation. |
| -c <class> | Display statistics for the <class> class only. |
| -a | Display statistics for all classes. |
| -r | Rounds down the statistics results, which is shown in bytes by default. In this case, the statistics is displayed in the following units of measurement, depending on the amount of traffic: <ul style="list-style-type: none">κ (b) — kilobytesM (b) — megabytesG (b) — gigabytes |
| -t <ct vm all> | Display statistics for <ul style="list-style-type: none">Containers only: ctvirtual machines only: vm |

| | |
|------------------------|--|
| | ▪ both Containers and virtual machines: <code>all</code> |
| <code>-h, -help</code> | Display the utility usage information. |

prl_disk_tool

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

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, FAT16, FAT32, ext2, or ext3. You can also try to compact virtual disks with other filesystems using the `--buildmap` option.

Syntax

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

Options

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

prl_disk_tool merge

Merges all snapshots of the virtual hard disk.

Syntax

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

Options

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

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, FAT16/FAT32, and ext2/ext3.

Syntax

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

Options

| Name | Description |
|--------------------------------------|--|
| <code>--size</code> | New size of the virtual disk. It can be set either in megabytes (M, default) or gigabytes (G). |
| <code>--resize_partition</code> | Resizes the last partition of the specified virtual disk. |
| <code>--hdd <disk_path></code> | Full path to the virtual disk. |
| <code>--force</code> | Forces the resizing operation for suspended virtual disks. |
| <code>--split</code> | Splits the virtual hard disk into 2 GB files. |
| <code>-i, --info</code> | Do not resize the virtual disk; just show the size the disk will have after resizing. |

vzcalc

This utility is used to calculate Container resource usage.

Syntax

```
vzcalc [-v] <CT_ID>
```

This utility displays what part of server resources Container `<CT_ID>` is using. An optional `-v` switch produces verbose output including number of processes, low memory, allocated memory and memory and swap statistics.

For stopped Containers the utility displays promised and maximum values the Container can consume. For running Containers, it also outputs the current values.

The high values of resource usage means that either the server is overcommitted or Container configuration is invalid.

vzcfgscale

This utility is used to “scale” Container configuration. It multiplies Container resource control parameters by the number passed as an argument.

Syntax

```
vzcfgscale [options] <CT_config_file>
```

Container configuration file shall be always the last parameter and the utility uses it to produce scaled Container configuration.

Options

| Name | Description |
|--------------------|---|
| -o <file> | Output scaled configuration into the <file>. By default, utility prints its output to screen. Note that the file specified cannot be the same as <CT_config_file>, otherwise you will loose the configuration file content. |
| -a <factor> | Multiply all Container parameters by <factor>. |
| -c <factor> | Multiply CPU parameters by <factor>. |
| -d <factor> | Multiply disk related parameters by <factor>. |
| -u <factor> | Multiply system resource control parameters by <factor>. |
| -n <factor> | Multiply bandwidth parameters by <factor>. |
| -r | Do not output Container-specific parameters like VE_ROOT, VE_PRIVATE, and so on. This option is useful for producing configuration samples to be used as an argument for the --config option of the prlctl create command. |

At least one multiplying argument is required. It is possible to specify more than one multiplying argument to use different factors for different group of parameters. If both **-a** and a specific group option is used, then the specific option factor takes precedence of the value specified by the **-a** option.

vzcfgvalidate

This utility is used to check the resource management parameters consistency in a Container configuration file.

Syntax

```
vzcfgvalidate <CT_config_file>
```

The utility has a number of constraints according to which it tests the configuration file. If a constraint is not satisfied, the utility prints a message with its severity status. Three severity statuses are defined in Parallels Cloud Server:

| Name | Description |
|-----------------------|--|
| Recommendation | This is a suggestion, which is not critical for Container or server operations. The configuration is valid in general; however, if the system has enough memory, it is better to increase the settings as advised. |
| Warning | A constraint is not satisfied and the configuration is invalid. Applications in the Container with such invalid configuration may have suboptimal performance or fail in a not graceful way. |

| | |
|--------------|--|
| Error | An important constraint is not satisfied and the configuration is invalid. Applications in a Container with such invalid configuration have increased chances to fail unexpectedly, to be terminated or to hang. |
|--------------|--|

It is suggested to use this utility when applications in a Container behave in an unexpected way and there seems to be no resource shortage for the Container.

vzcheckovr

This utility is used to check the current system overcommitment and safety of the total resource control settings.

Syntax

```
vzcheckovr [-v]
```

where `-v` is the option for verbose output. This utility computes the commitment levels of a number of resource management parameters (Low Memory, Memory + Swap, Allocated Memory) and compares them with the values chosen by the Parallels Cloud Server administrator in the `/etc/vz/vz.conf` global configuration file. The utility will produce a warning if these configured values are exceeded. Similarly to `vzmemcheck`, `vzcheckovr` takes into account only those Containers that are currently running and ignores all the others existing on the server.

vzcpucheck

This utility displays the current server utilization in allocated CPU units as well as total server CPU unit capacity.

Syntax

```
vzcpucheck [-v]
```

Without arguments, the utility prints the sum of CPU units of all running Containers and the total server capacity. If the `-v` option is given, the utility prints per Container CPU units information.

vzctl convert

Command to convert Containers based on VZFS, the layout used in earlier versions of Parallels Server Bare Metal, to the Container-in-an-image-file layout used in Parallels Cloud Server 6.0. The new layout has been redesigned to store all files of a Container in a single image, similar to a virtual machine's hard disk.

Syntax

```
vzctl convert <CT_ID> [--velayout <layout_ID>]
```

Options

| Name | Description |
|---|--|
| <code><CT_ID></code> | ID of the Container to convert. |
| <code>--velayout <layout_ID></code> | <p>ID associated with the Container layout. To convert a Container to the Container-in-an-image-file layout, specify 5 after the option.</p> <p>If this option is omitted, the command converts a Container to the layout defined by the <code>VEFSTYPE</code> parameter in the global configuration file <code>/etc/vz/vz.conf</code>. For the Container-in-an-image layout, this parameter should be set to <code>ext4</code>.</p> |

vzhwcalc

`vzhwcalc` is used to scan information on the resources consumption on a server (this can be a physical server or a server) and create a special file on its basis. When launched without any options, it makes a snapshot of the resources consumption and writes down this information to a special file - a server configuration file. The collected information can then be used to create a Container on its basis where the physical server will be migrated. You may also use `vzhwcalc` to collect data on your server resources in one place and be aware of their current consumption.

Syntax

```
vzhwcalc [options]
```

Options

| Name | Description |
|--------------------------------|--|
| <code>-o, --out</code> | The name of the configuration file that will be created by the utility and contain information on the server main resources. |
| <code>-t, --scan-time</code> | The time during which the utility is to be run on the server. The time should be given in the <code>dhms</code> format (e.g., <code>1d2h30m40s</code>). |
| <code>-p, --scan-period</code> | The interval with which the server will be scanned by the utility. |
| <code>--mem-scale</code> | The enlargement factor by which the calculated memory on the server will be increased in the configuration file. |
| <code>--disk-scale</code> | The enlargement factor by which the calculated disk space on the server will be increased in the configuration file. |
| <code>-d, --dist-detect</code> | The path to the file on the server where the <code>distdetect-common.sh</code> script is located. You can specify several scripts and separate them by commas. |
| <code>-h, --help</code> | Print usage information. |
| <code>-v, --version</code> | Print the version of the utility. |

The configuration file created by the `vzhwcalc` utility is placed to the same directory on the server from where you have run this utility and has the default name of `ve.conf` (i.e. in case the `-o` option was omitted during the utility execution).

vziotop

Just like the Linux standard `iostat` utility, `vziotop` monitors I/O usage data output by the Linux kernel and displays the current I/O usage by processes or threads on the system.

In addition to the options of the standard `iostat` utility, the `vziotop` utility has one extra option: `-E <CT_ID>`. When used with this option, `vziotop` shows the current I/O usage by processes or threads running in the specified Container.

Another difference from `iostat` is that `vziotop` shows actual disk read/write alongside the standard total disk read/write. The latter values represent total read/write bandwidth between processes and kernel threads on one side and kernel block device subsystem on the other. In turn, actual disk read/write values represent corresponding bandwidths for actual disk I/O between kernel block device subsystem and underlying hardware (HDD, SSD, etc.). Thus, total and actual values may differ at any given moment due to data caching and I/O operations reordering in the Linux kernel.

vzmemcheck

This utility shows the server memory parameters: low memory utilization, low memory commitment, RAM utilization, memory+swap utilization, memory+swap commitment, allocated memory utilization, allocated memory commitment, allocated memory limit.

Syntax

```
vzmemcheck [-v] [-A]
```

Options

| Name | Description |
|-----------------|---|
| <code>-v</code> | Display information for each Container. |
| <code>-A</code> | Display absolute values (in megabytes). |

It is possible to use any of the available options, both of them, or do without any options.

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.

vzps, vztop

These two utilities can be run on the server just as the standard Linux `ps` and `top` utilities. For information on the `ps` and `top` 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_ID>` command-line switch can be used to show only the processes running inside the Container with the specified ID.

The `vztop` utility has the following functionality added:

- The `-E <CT_ID>` command-line switch can be used to show only the processes running inside the Container with the ID specified. If `-1` is specified as `<CT_ID>`, the processes of all running Containers are displayed.
- The `e` interactive command (the key pressed while `top` is running) can be used to show/hide the `CTID` column, which displays the Container where a particular process is running (0 stands for the server itself).
- The `E` interactive command can be used to select another Container the processes of which are to be shown. If `-1` is specified, the processes of all running Containers are displayed.

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.

Options

| Name | Description |
|--------------------------------------|--|
| -n <i><num></i> | Desired number of Containers to be simultaneously run on the server. |
| -f <i><sample_name></i> | Name of the sample configuration to create. |
| -s <i><swap_size></i> | 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

This chapter describes the utilities that can be used for managing your virtual machines.

In This Chapter

| | |
|--|-----|
| <code>prlctl</code> | 145 |
| <code>pmigrate</code> | 171 |
| Managing Virtual Machine Backups | 174 |
| <code>prl_convert</code> | 177 |
| <code>prl_disk_tool</code> | 178 |
| <code>pnetstat</code> | 181 |

prlctl

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

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_ID/vm_name> [options] [-v, --verbose <number>] [--timeout <sec>]
[-l, --login [user[:passwd]@]server [-p, --read-passwd <file>]]
```

Options

| Name | Description |
|------------------------|--|
| <command> | The name of the command to execute. |
| <vm_ID/vm_name> | The ID or name of the virtual machine to perform the operation on. To obtain the list of the available virtual machines, use the prlctl list command. |
| options | Command options. See individual commands for available options. |
| -v, --verbose <number> | Enables verbose output. The greater the <number>, the higher is the verbosity. |

| | |
|--|--|
| -l, --login [<i>user[:passwd]@server</i>] | Connect to a remote Parallels server <i>server</i> with the specified credentials. If this flag is omitted, the <code>prlctl</code> command is assumed to be run locally. |
| -p, --read-passwd < <i>file</i> > | Use the password from the file < <i>file</i> > to log in to a remote Parallels server, 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, e.g., <code>prlctl change-passwd</code> . In such a case, the supplied passwords will be used one after another: to establish the remote connection, access the encrypted virtual machine, and set a new password for that encrypted virtual machine. Each password must be supplied in a separate file. |

Remarks

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

prlctl capture

Captures the screen of a virtual machine desktop and saves it to a file on the client machine. The data is saved in the Portable Network Graphics (PNG) format.

Syntax

```
prlctl capture <vm_ID|vm_name> --file <name>
```

Options

| Name | Description |
|----------------------|--|
| <vm_ID vm_name> | The virtual machine ID or name. |
| --file <name> | Name and path of the file to which the image should be saved. You should include the file extension (<code>.png</code>) or the file will be saved without one. |

prlctl change-passwd

Changes the virtual machine password.

Syntax

```
prlctl change-passwd <vm_ID|vm_name>
```

Options

| Name | Description |
|------------------------------------|-----------------------------|
| <code><vm_ID vm_name></code> | Virtual machine ID or name. |

Note: The virtual machine should be encrypted. See **prctl encrypt, decrypt** (p. 149).

prctl clone

Creates an exact copy of the specified virtual machine.

Syntax

```
prctl clone <vm_ID|vm_name> --name <new_name> [--template] [--dst <path>] [--changesid]
[--linked] [--detach-external-hdd <yes|no>]
```

Options

| Name | Description |
|---------------------------------------|--|
| <code><vm_ID vm_name></code> | ID or 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, Parallels Tools must be installed in the virtual machine. |
| --linked | Create a linked virtual machine. |
| --detach-external-hdd <yes no> | If set to <code>no</code> , hard disks located outside the source virtual machine are not removed from the configuration of the resulting clone. Setting the parameter to <code>yes</code> removes outside hard disks from the configuration. <div> Note: Outside hard disks are not copied to the cloned virtual machine. </div> |

prctl convert

This command is used to convert third-party virtual machines and disks to Parallels virtual machines and disks. The following third-party virtual machines and disks are supported:

- Microsoft Hyper-V

- Microsoft Virtual PC
- Virtual Box
- VMware

Syntax

```
prlctl convert <path> [--dst <path>] [--force]
```

Options

| Name | Description |
|--------------|---|
| <path> | Full path to the third-party virtual machine's configuration file on the local server. |
| --dst=<path> | Set the destination directory for the resulting virtual machine and its configuration file. If omitted, the default directory (/var/parallels) is used. |
| --force | Convert the third-party virtual machine even if its guest OS cannot be identified. |

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]
```

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. |
| --location <vm_path> | Name and path of the directory where to store the new virtual machine files. If this parameter is omitted, the files will be created in the default virtual machine directory. |
| --uuid <uuid> | A custom UUID to assign to the virtual machine. |

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

Remarks

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 Parallels and are kept in sync with other Parallels management tools such as Parallels Virtual Automation. The best way to find out the default versions used in your Parallels installation is by creating a sample virtual machine.

prctl delete

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

Syntax

```
prctl delete <vm_ID|vm_name>
```

Options

| Name | Description |
|-----------------|--|
| <vm_ID vm_name> | The ID or name of the virtual machine to delete. |

prctl encrypt, decrypt

Encrypts or decrypts a virtual machine.

Syntax

```
prctl encrypt <vm_ID|vm_name>
```

```
prctl decrypt <vm_ID|vm_name>
```

Options

| Name | Description |
|-----------------|-----------------------------|
| <vm_ID vm_name> | Virtual machine ID or name. |

prctl 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. Parallels Tools must be installed in a virtual machine to use this utility.

Syntax

```
prctl enter exec <vm_ID|vm_name>
```

Options

| Name | Description |
|-----------------|--|
| <vm_ID vm_name> | The ID or the name of the virtual machine. |

prctl exec

Executes a command inside a virtual machine. Parallels Tools must be installed in a virtual machine to use this utility. Commands in Linux guests are invoked with `bash -c`.

Syntax

```
prctl exec <vm_ID|vm_name> <command>
```

Options

| Name | Description |
|-----------------|--|
| <vm_ID vm_name> | The ID or the name of the virtual machine. |
| <command> | A command to execute. |

prctl list

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

Syntax

```
prctl list --vmtype vm [-a,--all] [-o,--output <field>[,...]] [-s,--sort <field|-field>]  
                        [-t,--template] [-j,--json]  
prctl list -i,--info --vmtype vm [vm_id|vm_name] [-f,--full] [-t,--template] [-j,--json]
```

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. For the list of fields, see prctl list Output Parameters (p. 151). Type field names in lower case. Separate multiple fields with commas. |
| -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_ID/vm_name> | The ID or 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. |

prctl list Output Parameters

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

| Name | Output Column | Description |
|--------------------|---------------|--|
| uuid | UUID | Virtual machine ID. |
| 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., <code>redhat-e15-x86</code>). |
| 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_PRIO | Virtual machine priority in the High Availability Cluster (0 is the lowest). Higher-priority virtual environments are restarted first in case of failures. |

prctl migrate

Migrates a virtual machine from one host to another.

Syntax

```
prctl migrate [source_server/]<vm_ID/vm_name> <destination_server>[/<vm_ID/vm_name>]  
[--dst <path>] [--keep-src] [--switch-template] [--changesid]
```

Options

| Name | Description |
|-----------------------------------|---|
| < <i>vm_ID</i> / <i>vm_name</i> > | The source virtual machine ID or name. |
| <i>source_server</i> | The source server information. Use the following format to specify this info: [user[:password]@]<server_IP_address_or_hostname>[:port] |
| <i>destination_server</i> | 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 < <i>path</i> > | Name and path of the directory on the destination server where the virtual machine files should be stored. |
| --keep-src | If this option is included, the original virtual machine will be left intact on the source server. If this option is omitted, the original virtual machine will be removed from the source server. |
| --switch-template | Switches the virtual machine to a template and a template to a virtual machine. For example, if the source virtual machine was a template, it becomes a full featured virtual machine after the migration, and vice versa. |
| --changesid | Changes the resulting virtual machine SID. |
| --ssh | Additional options to pass to <code>ssh</code> to connect to the destination server. All standard <code>ssh</code> options are supported. Note: Do not specify the destination server hostname or IP address as an <code>ssh</code> option. |

prctl mount, umount

Mounts or unmounts the hard disks of a virtual machine to the `/vz/mnt/<vm_ID>` directory on the Parallels server.

Syntax

```
prctl mount <vm_ID/vm_name> [-o <ro|rw> | --info]  
prctl umount <vm_ID/vm_name>
```


Options

| Name | Description |
|--------------------------------------|---|
| <code><vm_ID / vm_name></code> | Virtual machine ID or name. |
| <code>-o <ro rw></code> | Sets access rights: <ul style="list-style-type: none"> ▪ <code>ro</code> — read-only, ▪ <code>rw</code> — read-write. |
| <code>--info</code> | Show information about the mounted virtual disks. |

prctl 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

```
prctl move <vm_ID / vm_name> --dst <path>
```

Options

| Name | Description |
|--------------------------------------|---|
| <code><vm_ID / vm_name></code> | Virtual machine ID or name. |
| <code>--dst <path></code> | Path to the new virtual machine files location. |

prctl pause, suspend, resume

Pause, suspend, and resume a virtual machine.

Syntax

```
prctl pause <vm_ID / vm_name>
prctl suspend <vm_ID / vm_name>
prctl resume <vm_ID / vm_name>
```

Options

| Name | Description |
|--------------------------------------|---|
| <code><vm_ID / vm_name></code> | The ID or name of the virtual machine to pause, suspend, or resume. |

Remarks

The `pause` command pauses a virtual machine. To continue the virtual machine operation, use the [prctl start](#) (p. 170) 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.

prctl problem-report

Obtains a problem report for the specified virtual machine and displays it on the screen.

Syntax

```
prctl problem-report <vm_ID/vm_name> <-d,--dump|-s,--send [--proxy  
[user[:password]@]<proxyhost>[:port]] [--no-proxy]>
```

Options

| Name | Description |
|---|--|
| <vm_ID/vm_name> | The ID or 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 Parallels technical support to analyze your problem. |
| -s, --send | Send the generated problem report to the Parallels technical support. |
| --proxy [user[:password]@]<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. |

prctl register, unregister

The `register` command is used to register a virtual machine with Parallels Cloud Server.

The `unregister` command removes a virtual machine from the Parallels Cloud Server registry.

Syntax

```
prctl register <path> --preserve-uuid <yes|no>  
prctl unregister <vm_ID/vm_name>
```

Options

| Name | Description |
|---|---|
| <code><path></code> | An absolute path to the virtual machine directory. |
| <code><vm_ID/vm_name></code> | The ID or the name of the virtual machine to remove from the Parallels Cloud Server registry. |
| <code>--preserve-uuid <yes no></code> | Specifies what to do with the UUID (universally unique identifier). If you specify <code>yes</code> , the UUID is preserved. If you specify <code>no</code> , the UUID is regenerated. Note: By default, IDs are regenerated. |

Remarks

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 Parallels Cloud Server. 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 Parallels Cloud Server registry, but does not delete the virtual machine files from the server. You can re-register such a machine with Parallels Cloud Server later using the `register` command.

prctl reset-uptime

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

Syntax

```
prctl reset-uptime <vm_ID/vm_name>
```

Options

| Name | Description |
|------------------------------------|--|
| <code><vm_ID/vm_name></code> | Virtual machine ID or name. Names consisting of multiple words must be enclosed in quotes. |

prctl set

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

Modifying Virtual Machine Configuration

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

Syntax

```
prctl set <vm_ID/vm_name> [--cpus <number>] [--memsize <number>] [--videosize <number>]
[--memquota <auto|guarantee:limit[:priority[:maxballoon]]>] [--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]]
[--offline-management <on|off>] [--offline-service <service_name>]
[--userpasswd <user:passwd> [--cryptd]]
[--rate <rate>] [--ratebound <on|off>]
[--apply-iponly <yes|no>]
[--efi-boot <on|off>]
[--tools-autoupdate <on|off>]
```

Options

| Name | Description |
|---|---|
| <vm_ID/vm_name> | Target virtual machine ID or 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. |
| --videosize <number> | The amount of video memory available to the virtual machine graphics card. |
| --memquota <auto guarantee:limit[:priority[:maxballoon]]> | Sets the parameters of the memory consumption by the virtual machine: <ul style="list-style-type: none"> <i>guarantee</i>. The amount of memory a virtual machine is guaranteed to get on demand. By default, the guaranteed memory is calculated on the basis of RAM and video memory assigned to a virtual machine and is about a half of its total memory. <i>limit</i>. The maximum amount of memory a virtual machine is allowed to consume. By default, no limit is set for all newly created |

| | |
|--------------------------------------|--|
| | <p>virtual machines, and any virtual machine may consume all free memory on the server.</p> <ul style="list-style-type: none"> ▪ <i>priority</i>. The priority (from 1 to 100) that defines which virtual machine will get memory first. The higher the priority of a virtual machine, the more chances it has to get memory when the Parallels server has insufficient memory resources. By default, the priority is set to 50. ▪ <i>ballooning</i>. The maximum amount of memory the balloon driver in a virtual machine may allocate for its needs. By default, the balloon driver can allocate up to 60% of RAM set for a virtual machine. ▪ <i>auto</i>. Resets all memory parameters to their default values. |
| --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. |
| --distribution <name> | <p>Optimize the virtual machine for use with the operating system <name>.</p> <p>You can get the list of available distributions using the <code>prlctl set <vm_name> -d list</code> command.</p> |
| --description <desc> | <p>Sets virtual machine description.</p> <p>Descriptions with white spaces must be enclosed in quotation marks.</p> |
| --autostart <on off auto> | <p>Defines the virtual machine start-up options:</p> <ul style="list-style-type: none"> ▪ <i>on</i> — the virtual machine is started automatically when the Parallels server starts or the Parallels Cloud Server component responsible for managing virtual machines is disabled. ▪ <i>off</i> — the autostart is off. This is the default virtual machine start-up mode. ▪ <i>auto</i> — resume the virtual machine state prior to shutting down the Parallels server or disabling the Parallels Cloud Server component responsible for managing virtual machines. <p>If you set this option to <i>on</i> or <i>auto</i>, you must additionally specify the <code>--start-as-user</code> option.</p> |
| --autostart-delay <number> | Sets the time delay used during the virtual machine automatic startup. |
| --autostop <stop suspend> | Sets the automatic shutdown mode for the specified virtual machine: |

| | |
|---|--|
| | <ul style="list-style-type: none"> ▪ stop — the virtual machine is stopped when you shut down the Parallels server or disable the Parallels component responsible for managing virtual machines. ▪ suspend — the virtual machine is suspended when the Parallels server is shut down or the Parallels component responsible for managing virtual machines is disabled. |
| --applyconfig <i><conf></i> | <p>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:</p> <ul style="list-style-type: none"> ▪ all memory-related parameters (both RAM and video) ▪ all CPU-related parameters ▪ IO and IOPS limits ▪ disk size |
| --name <i><new_name></i> | Changes the virtual machine name. |
| --start-as-user <i><administrator owner user:passwd></i> | <p>Specifies the account to use to autostart the virtual machine:</p> <ul style="list-style-type: none"> ▪ administrator — start the virtual machine as the administrator of the host operating system. ▪ owner — start the virtual machine as the virtual machine owner. ▪ user:passwd — start the virtual machine as the specified user. |
| --vnc-mode <i><auto manual off></i> | Enables or disables access to the virtual machine via the VNC protocol. |
| --vnc-port <i><port></i> | Sets the VNC port number. |
| --vnc-passwd <i><passwd></i> --vnc-nopasswd | <p>Sets the VNC password or specifies that no password is needed for VNC connections.</p> <p>Either of these options is mandatory for any VNC connection.</p> |
| --vnc-address <i><address></i> | <p>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 Parallels server.</p> <p>By default, you can use any of the IP addresses of the Parallels server to log in to the virtual machine.</p> |
| --cpu-hotplug <i><on off></i> | 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. |
| --cpuunits <i><units></i> | 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> | <p>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.</p> <p>Note: If the server has 2 processors, the total CPU time equals 200%.</p> |
| <code>--ioprio <priority></code> | Disk I/O priority level from 0 to 7. The default is 4. |
| <code>--iolimit <limit></code> | <p>Disk I/O bandwidth limit. The default is 0 (no limit). By default the limit is set in megabytes per second. You can use the following letters following the number to specify units of measure:</p> <ul style="list-style-type: none"> ▪ G — gigabytes per second (e.g., 1G). ▪ K — kilobytes per second (e.g., 10K). ▪ B — bytes per second (e.g., 100B). <p>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.</p> |
| <code>--iopslimit <limit></code> | <p>Maximum number of disk input and output operations per second a virtual machine is allowed to perform.</p> <p>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.</p> |
| <code>--cpumask <N> [,N,N1-N2]</code> | <p>An affinity mask indicating what CPU(s) the virtual machine processes should be run on.</p> <p>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).</p> <p>To make all CPUs available for the virtual machine processes specify <code>--cpumask all</code>.</p> |
| <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. Parallels Tools must be installed in the virtual machine for the command to work. |
| <code>--crypted</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. |

| | |
|--------------------------------------|--|
| --ratebound <on off> | Turns the network traffic rate limitation set by the --rate parameter (above) on or off. The default value is off . |
| --apply-iponly <yes no> | If set to yes , the hostname, nameserver, and search domain settings from the virtual machine configuration file are ignored. |
| --efi-boot <on off> | 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. |
| --tools-autoupdate <on off> | Enables or disables automatic update of Parallels Tools inside Windows virtual machines. If set to on , Parallels Tools are updated automatically on user log in. A reboot is required to complete the update. If set to off , Parallels Tools are not updated automatically, so that you can do it manually at a convenient time. |

Managing Virtual Devices

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

Syntax

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

Options

| Name | Description |
|--|---|
| <vm_ID vm_name> | The virtual machine ID or name. |
| --device-add <dev_type> <options> | Adds a virtual device of the type <dev_type> to a stopped virtual machine. The <dev_type> parameter can be: hdd, cdrom, net, fdd, serial, parallel, sound, usb, pci. |
| --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](#) (p. 161)
- [Optical disk drives](#) (p. 162)
- [Network cards](#) (p. 163)
- [Floppy disk drives](#) (p. 165)
- [Serial ports](#) (p. 165)
- [Parallel ports](#) (p. 166)
- [USB devices](#) (p. 167)
- [Sound cards](#) (p. 167)

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

```
prctl set <vm_ID/vm_name> {--device-add hdd | --device-set hdd<N>}
    [--image <file>]
    [--type <expanded|plain>] [--size <size>] [--split]
    [--iface <sata|ide|scsi>] [--position <pos>]
    [--enable|--disable]
prctl set <vm_ID/vm_name> --device-add hdd --device <dev_name>
    [--iface <sata|ide|scsi>] [--position <pos>]
```

Options

| Name | Description |
|-------------------------|--|
| <vm_ID/vm_name> | The virtual machine ID or name. |
| --device-add hdd | <p>Adds a virtual hard disk drive to the virtual machine.</p> <p>You can connect up to four SATA/IDE devices and up to seven SCSI devices to a virtual machine. This includes hard disks and optical disk drives.</p> <p>New hard disks are created in the virtual machine directory and are automatically named <code>harddisk<N>.hdd</code>, where <N> is the next available disk index.</p> |

| | |
|--|--|
| --device-set hdd<N> | <p>Modifies the parameters of an existing virtual hard disk.</p> <p>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.</p> |
| --image <file> | <p>Specifies an image file that will be used to emulate the virtual disk.</p> <ul style="list-style-type: none"> ▪ 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> | <p>This option is used to connect a physical hard disk on the Parallels Server to the virtual machine. You can obtain the names of the existing hard disks on the server using the <code>prlsrvctl info</code> command.</p> |
| --type <expanded plain> | <p>Sets the disk type for image file based virtual hard disks:</p> <ul style="list-style-type: none"> ▪ <code>expanded</code> — expanding disk. The image file is small initially and grows in size as you add data to it. This is the default virtual disk type. ▪ <code>plain</code> — plain disk. The image file has a fixed size from the moment it is created (the space is allocated to the drive fully). Plain disks perform faster than expanding disks. |
| --size <size> | <p>The size of the virtual hard disk, in megabytes. The default size is 65536 MB.</p> |
| --split | <p>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).</p> |
| --iface <sata ide scsi> | <p>The disk drive Interface type. If omitted, the SATA interface will be used.</p> |
| --position <pos> | <p>The SCSI or IDE device identifier to be used for the virtual disk. The allowed ID ranges are the following:</p> <ul style="list-style-type: none"> ▪ for IDE devices: 0:0, 0:1, 1:0, 1:1; ▪ for SCSI device: 0:0, 1:0, 2:0, 3:0, 4:0, 5:0, 6:0. <p>You can use one of the following formats for specifying IDs: <code>ID:bus</code>, <code>ID-bus</code>, <code>ID</code>. For example, if you specify 3:0 (or 3-0 or 3) as <i>number</i> for a SCSI drive, the guest OS will see the drive as having ID 3 on SCSI bus 0.</p> |
| --enable | <p>Enables the specified virtual disk drive. All newly added disk drives are enabled by default (provided the <code>--disable</code> option is omitted).</p> |
| --disable | <p>Disables the specified virtual disk drive. The disk drive itself is not removed from the virtual machine configuration.</p> |

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_ID/vm_name> {--device-add cdrom | --device-set cdrom<N>}
    {--device <dev_name> | --image <file>}
    [--iface <ide|scsi>] [--position <pos>]
```

[**--enable** | **--disable**] [**--connect** | **--disconnect**]

Options

| Name | Description |
|------------------------------------|--|
| <code><vm_ID/vm_name></code> | The virtual machine ID or name. |
| --device-add cdrom | Adds a DVD/CD drive to the virtual machine. You can connect up to four IDE devices and up to seven SCSI devices to a virtual machine. This includes virtual hard disks and DVD/CD drives. |
| --device-set cdrom<N> | Modifies the parameters of an existing virtual optical disk. The <i>N</i> postfix indicates the drive index number. To obtain the list of the available drives, use the <code>prlctl list</code> command with the <code>--info</code> 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: <code>.iso</code> , <code>.cue</code> , <code>.ccd</code> , and <code>.dmg</code> . The image must not be compressed and/or encrypted. |
| --iface <ide scsi> | Interface type: <ul style="list-style-type: none"> <code>ide</code> — IDE disk. <code>scsi</code> — SCSI disk (default). |
| --position <pos> | The SCSI or IDE device identifier to be used for the DVD/CD drive. The allowed ID ranges are the following: <ul style="list-style-type: none"> for IDE devices: <code>0:0</code>, <code>0:1</code>, <code>1:0</code>, <code>1:1</code>; for SCSI device: <code>0:0</code>, <code>1:0</code>, <code>2:0</code>, <code>3:0</code>, <code>4:0</code>, <code>5:0</code>, <code>6:0</code>. You can use one of the following formats for specifying IDs: <code>ID:bus</code> , <code>ID-bus</code> , <code>ID</code> . For example, if you specify <code>3:0</code> (or <code>3-0</code> or <code>3</code>) as <i>number</i> 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 <code>--disable</code> 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_ID/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]

```

Options

| Name | Description |
|----------------------------------|--|
| <vm_ID/vm_name> | The virtual machine ID or 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 <code>prlctl list</code> command with the <code>--info</code> 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 <code>addr</code> parameter value or use the <code>auto</code> 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. |
| --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. |

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

Options

| Name | Description |
|----------------------------|---|
| <vm_ID vm_name> | The virtual machine ID or name. |
| --device-add fdd | Adds a new floppy disk drive to the virtual machine. You can connect only one floppy disk drive to a virtual machine. |
| --device-set fdd0 | Modifies the parameters of an 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 <code>floppy.fdd</code>) to mount in the virtual machine. |
| --enable | Enables the specified floppy disk drive. All newly added floppy drives are enabled by default (provided the <code>--disable</code> 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_ID|vm_name> {--device-add serial | --device-add serial<N>}
    {--device <dev_name> | --output <file> | --socket <name>}
    [--enable|--disable] [--connect|--disconnect]
```

Options

| Name | Description |
|---|---|
| <code><vm_ID vm_name></code> | The virtual machine ID or name. |
| <code>--device-add serial</code> | Adds a new serial port to the virtual machine. You can connect up to four serial ports to a virtual machine. |
| <code>--device-set serial<N></code> | Modifies the parameters of an existing serial port. |
| <code>--device <dev_name></code> | The name of the physical serial port to which to connect the virtual machine. |
| <code>--output <file></code> | The name and path of the output file to which to connect the virtual serial port. |
| <code>--socket <name></code> | The name of the physical socket to which to connect the virtual serial port. |
| <code>--enable --disable</code> | Enables or disables the virtual serial port. All newly added serial ports are enabled by default (provided the <code>--disable</code> option is omitted). |
| <code>--connect</code> | Automatically connect the virtual serial port during the virtual machine startup process. |
| <code>--disconnect</code> | Do not automatically connect the virtual serial port during the virtual machine startup process. |

Parallel Port Management Options

This group of options is used to manage parallel port in a virtual machine.

Syntax

```
prctl set <vm_ID|vm_name> {--device-add parallel | --device-set parallel<N>}  
    {--device <dev_name> | --output <file>}  
    [--enable | --disable] [--connect | --disconnect]
```

Options

| Name | Description |
|---|---|
| <code><vm_ID vm_name></code> | The virtual machine ID. |
| <code>--device-add parallel</code> | Adds a new parallel port to the virtual machine. You can connect up to three parallel ports to a virtual machine. |
| <code>--device-set parallel<N></code> | Modifies the parameters of an existing virtual parallel port. To obtain the list of ports, use the <code>prctl list</code> command with the <code>--info</code> option. |
| <code>--device <dev_name></code> | The name of the physical parallel port to which to connect the virtual parallel port. |
| <code>--output <file></code> | The name of the output file to which to connect the virtual parallel port. |
| <code>--enable --disable</code> | Enables or disables the specified parallel port. All newly added parallel ports are enabled by default (provided the <code>--disable</code> option is omitted). |
| <code>--connect</code> | Automatically connect the specified virtual parallel port during the virtual machine startup process. |
| <code>--disconnect</code> | Do not automatically connect the specified virtual parallel 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

```
prctl set <vm_ID/vm_name> --device-add usb [--enable | --disable]
```

Options

| Name | Description |
|------------------------------------|---|
| <vm_ID/vm_name> | The virtual machine ID or 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). |

Sound Device Management Options

This group of options is used to manage sound devices in a virtual machine.

Syntax

```
prctl set <vm_ID/vm_name> {--device-add sound | --device-set sound0}  
    [--input <dev_name>] [--output <dev_name>]  
    [--enable | --disable] [--connect | --disconnect]
```

Options

| Name | Description |
|------------------------------------|---|
| <vm_ID/vm_name> | The virtual machine ID or name. |
| --device-add sound | The type of the virtual device to add to the virtual machine (in this instance, a sound device). |
| --device-set sound0 | Modifies the parameters of the existing virtual sound device. |
| --output <dev_name> | The name of a physical output device to connect to the virtual sound device. |
| --input <dev_name> | The name of the physical input device to connect to the virtual sound device. |
| --enable --disable | Enables or disables the specified sound device. All newly added sound devices are enabled by default (provided the --disable option is omitted). |
| --connect | Automatically connect the sound device during the virtual machine startup process. |
| --disconnect | Do not automatically connect the sound device during the virtual machine startup process. |

Removing Devices from Virtual Machines

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

Syntax

```
prctl set <vm_ID|vm_name> --device-del <dev_name> [--detach-only | --destroy-image]
```

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>prctl 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. |

Managing Shared Folders

The `prctl set` command can be used to add shared folders to a virtual machine and to modify and delete existing shared folders.

Syntax

```
prctl set <vm_ID|vm_name> {--shf-host-add <folder> | --shf-host-set <folder>}
--path <path>
[--mode <ro|rw>] [--shf-description <desc>] [--enable | --disable]
prctl set <vm_ID|vm_name> --shf-host <on|off>
prctl set <vm_ID|vm_name> --shf-host-del <folder>
prctl set <vm_ID|vm_name> --shf-guest <on|off>
prctl set <vm_ID|vm_name> --shf-guest-automount <on|off>
```

Options

| Name | Description |
|---------------------------------------|---|
| <vm_ID vm_name> | The virtual machine ID or name. |
| --shf-host-add <folder> | Shares the specified folder on the Parallels server with the virtual machine. |
| --shf-host-set <folder> | Modifies the settings of an existing shared folder. |
| --shf-host-del <folder> | Removes the specified shared folder from the shared folder list. |
| --shf-host <on off> | Turns the host folder sharing on or off. |
| --shf-guest <on off> | Turns the guest folder sharing on or off. |
| --shf-guest-automount <on off> | Mounts or unmounts virtual disks on the Parallels server. |
| --path <path> | Name and path of a folder on the Parallels server to share with |

| | |
|------------------------------------|--|
| | the specified virtual machine. |
| --mode <ro rw> | Sharing mode: <ul style="list-style-type: none"> ro — read-only. rw — read and write. |
| --shf-description <desc> | User-defined shared folder description. Descriptions with white spaces must be enclosed in quotation marks. |
| --enable --disable | Enable or disable the shared folder. |

prlctl snapshot, snapshot-list, snapshot-switch, snapshot-delete

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

Syntax

```
prlctl snapshot <vm_ID/vm_name> [-n,--name <name>] [-d,--description <desc>]
prlctl snapshot-list <vm_ID/vm_name> [-t,--tree] [-i,--id <snapshot_ID>]
prlctl snapshot-switch <vm_ID/vm_name> -i,--id <snapshot_ID> --skip-resume
prlctl snapshot-delete <vm_ID/vm_name> -i,--id <snapshot_ID>
```

Options

| Name | Description |
|--------------------------|---|
| <vm_ID/vm_name> | The virtual machine ID or 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> | <ul style="list-style-type: none"> 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. |
| --skip-resume | Skips automatic virtual machine resume when switching to snapshots of running virtual machines. |

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

prctl server

Obtains the information about the Parallels server and Parallels Cloud Server. It also allows you to disable the Parallels Cloud Server component responsible for managing virtual machines.

Syntax

```
prctl server <shutdown|info>
```

Options

| Name | Description |
|----------|---|
| info | Displays the Parallels Cloud Server information. |
| shutdown | Disables the Parallels Cloud Server component responsible for managing virtual machines. If one or more virtual machines are running, clients are connected, or some tasks are currently in progress, then the operation will be aborted. |

See Also

prlsrvctl info (p. 48)

prlsrvctl shutdown (p. 53)

prctl start, stop, restart, reset, status

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

Syntax

```
prctl start <vm_ID|vm_name>  
prctl stop <vm_ID|vm_name> [--kill]  
prctl restart <vm_ID|vm_name>  
prctl reset <vm_ID|vm_name>  
prctl status <vm_ID|vm_name>
```

Options

| Name | Description |
|-----------------|---|
| <vm_ID vm_name> | The ID or 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. |

Remarks

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 the Parallels Tools package is installed in a virtual machine, the graceful shutdown will be performed using its facilities.
- If the Parallels Tools package is 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 first performs a 'hard' virtual machine shutdown and then starts it again.

The `start` command can be used to start a stopped virtual machine or to resume a [paused virtual machine](#) (p. 153).

pmigrate

The `pmigrate` utility is used to perform different kinds of migration.

Syntax

```
pmigrate <source_server> <destination_server> [options]
```

The `<source_server>` is the source server which can be either the server where the virtual machine and Container to be migrated is residing (if you are migrating a virtual machine and Container) or the physical computer to be migrated (if you are migrating a physical computer). The `<destination_server>` is the destination server—that is, the Parallels server where the virtual machine and Container or the physical server is to be migrated. If the source and/or destination server are not specified, the operation is performed on the local server.

Both `<source_server>` and `<destination_server>` consist of `<type>` and `<address>`:

- `<type>` denotes the type of computer to migrate and can be one of the following:
 - `h` for migrating physical computers
 - `c` for migrating migrating Containers
 - `v` for migrating Parallels virtual machines
 - `x` for migrating Xen virtual machines
- `<address>` denotes the location of computer to migrate and can be one of the following:
 - The computer location if you are migrating a physical computer.

- The computer location and the virtual machine name or Container ID if you are migrating a virtual machine or Container, respectively. The location must be separated from the virtual machine name/Container ID by the slash (/).

The location format is as follows:

```
[user[:password]@]<destination_server_IP_address_or_hostname>[:port]
```

The `[options]` you can use with `pmigrate` depend on whether you are migrating a virtual machine or a Container. This section describes the options you can use to migrate Parallels and Xen virtual machines, Containers, and physical computers.

General Options

| Name | Description |
|---------------------------------|--|
| <code>--dst=<path></code> | Specifies the name and path of the directory on the destination server to use for storing the virtual machine files. If this option is omitted, the default directory is used. |
| <code>-h, --help</code> | Displays the information on the utility usage. |

Virtual Machine Migration Options

| Name | Description |
|--------------------------------|--|
| <code>--keep-src</code> | If this option is included, the original virtual machine is not removed from the source server after migration. By default, the original virtual machine is removed from the source server. |
| <code>--switch-template</code> | Switches the virtual machine to a template and a template to a virtual machine. For example, if the source virtual machine was a template, it becomes a full featured virtual machine after the migration, and vice versa. |
| <code>--changesid</code> | Changes the resulting virtual machine SID. |

Container-to-Virtual Machine Migration Options

| Name | Description |
|---|--|
| <code>-s, --size=<size></code> | Sets the limit of the resulting virtual machine disk capacity. If you omit this option or specify 0, no limit is set. In this case disk quotas set for virtual machines on the Parallels server will be used. The following size modifiers can be used: G — gigabytes, M — megabytes (default), K — kilobytes. |
| <code>-r, --reg[=<y n>]</code> | Specifies whether to register the resulting virtual machine on the Parallels server. By default, the virtual machine is registered. |
| <code>-d, --remove[=<y n>]</code> | Removes the source Container after migration. By default, the Container is left intact. |
| <code>--key=<auth_key></code> | Set the authentication key for passwordless access to the Parallels server. To set this key, use the <code>parallels-c2v-agent --regkey <auth_key></code> command on the server where the Container is residing (<code><auth_key></code> can be any alphanumeric value). |

Physical Computer-to-Virtual Machine Migration Options

| Name | Description |
|--------------------------------------|---|
| <code>-r, --reg[=<y n>]</code> | Specifies whether to register the resulting virtual machine on the Parallels server. By default, the virtual machine is registered. |
| <code>--osdata=<path></code> | Sets the path to the directory with data files required for the OS reconfiguration and not found on the source computer. |
| <code>-a, --all</code> | Migrate all computer's disks and partitions with all available data. If this option is omitted: <ul style="list-style-type: none"> For a Windows source computer, only boot and system volumes are migrated. For a Linux source computer, only the partitions specified in the <code>/etc/fstab</code> file on the source are migrated. |
| <code>--key=<auth_key></code> | Set the authentication key for passwordless access to the physical computer you want to migrate. To set this key on the computer, use the <code>parallels-transporter-agent --regkey <value></code> command (<code><auth_key></code> can be any alphanumeric value). |

Xen Virtual Machine-to-Virtual Machine Migration Options

| Name | Description |
|--------------------------------------|--|
| <code>-r, --reg[=<y n>]</code> | Specifies whether to register the resulting virtual machine on the Parallels server. By default, the virtual machine is registered. |
| <code>--osdata=<path></code> | Sets the path to the directory with data files required for the OS reconfiguration and not found on the source virtual machine. |
| <code>--key=<auth_key></code> | Set the authentication key for passwordless access to the Xen server. To set this key on the Xen server, use the <code>parallels-transporter-agent --regkey <value></code> command (<code><auth_key></code> can be any alphanumeric value). |
| <code>--ip=<IP>[/mask]</code> | Sets the IP address for the resulting virtual machine. If this option is omitted, the DHCP support is enabled for the virtual machine. |
| <code>--gw=<IP></code> | Sets the default gateway for the resulting virtual machine. If this options is omitted, the DHCP support is enabled for the virtual machine. |
| <code>--nameserver=<IP></code> | Sets the DNS server for the resulting virtual machine. |

| | |
|--|---|
| <code>--searchdomain=<domain></code> | Sets the DNS search domain for the resulting virtual machine. |
|--|---|

Managing Virtual Machine Backups

This section describes the utilities you can use for creating and managing virtual machine backups.

pbackup

The `pbackup` utility is run on the so-called Backup Server. It connects via SSH to the servers where some or all virtual machines are to be backed up and puts the tarballs into the `/vz/vmprivate/backups` directory. Later on, the virtual machine backups can be restored from this directory.

Syntax

```
pbackup [backup_options] server [...] [vm_options]
```

You may specify any number of server names or IP addresses in the command-line.

Notes:

1. This section describes only backup options for virtual machines. For backup options that can be used with Containers, see **pbackup** (p. 103).
2. The backup settings in the `/etc/vzbackup.conf` file do not apply to virtual machines. To change some of the backup-related settings for virtual machines, you can use the `prlsrvctl` utility.

Backup Options

| Name | Description |
|---|--|
| <code>-n <credentials></code> | Set the Backup Server, i.e. the server that will be used for storing the resulting virtual machine backups. If this parameter is omitted, the specified virtual machines will be backed up to the server where they are hosted. The Backup Server can be specified in this format: <code>user[:password]@server_IP_address_or_hostname[:port]</code> |
| <code>--ssh-opts <options></code> | Options to be passed to <code>ssh</code> . See examples in the global backup configuration file. |
| <code>-F, -l</code> | Create a full backup. A full backup contains all virtual machine data. |
| <code>-i</code> | Make an incremental backup or, if no full backups are available, a full backup. An incremental backup contains only the file that were changed since the previous full or incremental backup. |

Virtual Machine Options

| Name | Description |
|-------------------------------------|--|
| -e <vm_ID/vm_name> [. . .] | The virtual machines to back up on the server. Virtual machines can be specified using both their IDs and their names. |
| -x <vm_ID/vm_name> [. . .] | The virtual machines that need not be backed up (virtual machines to exclude). Virtual machines can be specified using both their IDs and their names. |

prestore

The `prestore` utility is run on the Backup Server. It uses the virtual machine backups stored on the Backup Server to restore them to their original servers (or to any other location if the `-d` option is specified).

Syntax

```
prestore [restoration_options] server [...] [vm_options]
```

You can specify any number of servers (their names or IP addresses) whose virtual machines were at one time backed up and now need to be restored.

Restoration Options

| Name | Description |
|---------------------------|---|
| -l | Do not restore any virtual machines. Show the information on the virtual machines available for restoring. |
| -n <backup_server> | The Backup Server where to look for existing backups. If this option is omitted, <code>prestore</code> searches for the backups on the server from which they were originally backed up. The format for the Backup Server is as follows: [user[:password]]@<server_IP_address_or_hostname>[:port] |

Virtual Machine Options

| Name | Description |
|-------------------------------------|---|
| -e <vm_ID/vm_name> [. . .] | The virtual machines to restore on the server. Virtual machines can be specified using both their IDs and their names. |
| -x <vm_ID/vm_name> [. . .] | The virtual machines that need not be restored (virtual machines to exclude). Virtual machines can be specified using both their IDs and their names. |

prlctl backup, backup-list, backup-delete, restore

Creates, lists, deletes or restores virtual machine backups.

Syntax

```
prctl backup <vm_ID/vm_name> [-f,--full] [-i,--incremental] [--description <desc>]
                               [-s,--storage [user[:passwd]@]server[:port]]
prctl backup-list [vm_ID/vm_name] [-f,--full] [--localvms]
                  [-s,--storage [user[:passwd]@]server[:port]]
prctl backup-delete {<vm_ID/vm_name> | -t,--tag <backup_ID>}
                    [-s,--storage [user[:passwd]@]server[:port]]
prctl restore {vm_ID/vm_name | -t,--tag <backup_ID>} [-n, --name <new_name>] [--dst <path>]
               [-s,--storage [user[:passwd]@]server[:port]]
```

Options

| Name | Description |
|--|--|
| <vm_ID/vm_name> | Virtual machine ID or name. <ul style="list-style-type: none"> Use with <code>prctl backup</code> to create a backup of the specified virtual machine. Use with <code>prctl backup-list</code> to list backups of the specified virtual machine. Use with <code>prctl backup-delete</code> to delete all backups of the specified virtual machine. Use with <code>prctl restore</code> to restore the most recent backup of the specified virtual machine. |
| -s, --storage [user[:passwd]@]server[:port]] | Specifies a remote backup server address, port, and credentials. If this option is omitted, the backup will be saved on the default backup server that can be configured using the prlsrvctl set command. |
| --description <desc> | Backup description. Descriptions with white spaces must be enclosed in quotation marks. |
| -f, --full | <ul style="list-style-type: none"> Use with <code>prctl backup</code> to create a full backup of the virtual machine. A full backup contains all virtual machine data. Use with <code>prctl backup-list</code> to display full backup information. |
| -i, --incremental | Create an incremental backup of the virtual machine. An incremental backup contains only the files changed since the previous full or incremental backup. This is the default backup type. |
| --localvms | List local backups only. |
| -t, --tag <backup_ID> | The ID of the backup to restore or delete. |
| -n, --name <new_name> | A new name to assign to the restored virtual machine. If omitted, the virtual machine will be restored with the original name. |
| --dst <path> | Restore the virtual machine to the specified directory on the Parallels server. If this option is omitted, the virtual machine will be restored to |

| | |
|--|---------------------------|
| | /var/parallels/<vm_name>. |
|--|---------------------------|

prl_convert

This utility is used to convert third-party virtual disks to Parallels virtual machines and disks. The following third-party disks are supported:

- Microsoft Hyper-V
- Microsoft Virtual PC
- Virtual Box
- VMware

Syntax

```
prl_convert <src> [options]
prl_convert <src> --estimate
```

If the disk is a data disk, `prl_convert` converts it to a Parallels virtual disk. If a disk is a system disk, `prl_convert` converts it to a Parallels virtual machine. If the utility cannot create a virtual machine for the disk (for example, it fails to detect the operating system on the disk), the disk is converted to a Parallels virtual disk.

Options

| Name | Description |
|---------------------------------------|--|
| <code>--dst=<path></code> | Set the destination directory. If omitted, the default directory (/var/parallels) is used. |
| <code>-r, --reg=<y n></code> | Register the resulting virtual machine on the Parallels server. By default, the virtual machine is registered. |
| <code>--os-files=<path></code> | Set the path to the operating system installation files. These files may be needed when reconfiguring the resulting virtual machine or disk. |
| <code>--vbox-home=<path></code> | Specify the path to the directory storing virtual machine configuration files. This option should be used only when converting VirtualBox virtual machines. |
| <code>--err-code</code> | Show error codes instead of text messages. |
| <code>--allow-no-os</code> | Convert data disks or disks whose operating system cannot be detected. |
| <code>--no-reconfig</code> | Skip the reconfiguration step when converting a virtual disk. |
| <code>--no-src-check</code> | Do not check the source virtual disk state. Use this option with caution because disks of virtual machines that were running before the conversion may be in an inconsistent state after conversion. |
| <code>--estimate</code> | Estimate the disk size required for conversion, but do not perform the conversion. |
| <code>-h, --help</code> | Display the utility usage information. |

prl_disk_tool

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

Syntax

```
prl_disk_tool <command> [options] --hdd <disk_path> [options]  
prl_disk_tool --help
```

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, FAT16, FAT32, ext2, or ext3. You can also try to compact virtual disks with other filesystems using the `--buildmap` option.

Syntax

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

Options

| Name | Description |
|--------------------------|---|
| --hdd <disk_path> | Full path to the virtual disk. |
| --buildmap | Compacts virtual disks with unsupported filesystems. |
| --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. |

prl_disk_tool configure

Prepares a disk or a real Boot Camp partition for booting into a virtual machine.

Syntax

```
prl_disk_tool configure --hdd <disk_path> [--para <paravirt_driver>] [--boot <boot_driver>]
```

Options

| Name | Description |
|--------------------------|---|
| configure | Prepares a disk or a real Boot Camp partition for booting into a virtual machine. |
| --hdd <disk_path> | Full path to the virtual disk. |

| | |
|---------------------------------|---|
| --para <paravirt_driver> | Specifies the full path to the Parallels virtualization driver. |
| --boot <boot_driver> | Specifies the full path to the Parallels boot driver. |

prl_disk_tool convert

Converts the virtual disk between plain and expanding formats; splits the virtual disk into 2GB files or merges them back.

Syntax

```
prl_disk_tool convert --hdd <disk_path> {--plain|--expanding}
prl_disk_tool convert --hdd <disk_path> {--split|--merge}
prl_disk_tool convert --hdd <disk_path> --extend
prl_disk_tool convert -i,--info --hdd <disk_path>
```

Options

| Name | Description |
|--------------------------|--|
| --hdd <disk_path> | Full path to the virtual disk to convert. |
| -i, --info | Do not convert the virtual disk; just show the size the disk will have after resizing. |
| --plain | Converts the disk to the plain format. A plain virtual hard disk has a fixed size. |
| --expanding | Converts the disk to the expanding format. An expanding virtual hard disk is small initially. Its size grows as you add applications and data to it. |
| --split | Splits the virtual hard disk into 2 GB files. |
| --merge | Merges all parts of the split disk into one file. |
| --extend | Enables sizes over 2TB for the specified expandable disk. |

prl_disk_tool encrypt, decrypt

Encrypts or decrypts data on the specified virtual hard disk.

Syntax

```
prl_disk_tool encrypt --hdd <disk_path> [--force]
prl_disk_tool decrypt --hdd <disk_path>
```

Options

| Name | Description |
|--------------------------|---|
| --hdd <disk_path> | Full path to the virtual disk to convert. |
| --force | Forcibly encrypt the virtual hard disk even if it is already encrypted. |

prl_disk_tool merge

Merges all snapshots of the virtual hard disk.

Syntax

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

Options

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

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, FAT16/FAT32, and ext2/ext3.

Syntax

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

Options

| Name | Description |
|--------------------------------|--|
| --size | New size of the virtual disk. It can be set either in megabytes (specify M after the value) or in gigabytes (specify G after the value). By default, the size is set in megabytes. |
| --resize_partition | Resizes the last partition of the specified virtual disk. |
| --hdd <disk_path> | Full path to the virtual disk. |
| --force | Forces the resizing operation for suspended virtual disks. |
| --split | Splits the virtual hard disk into 2 GB files. |
| -i, --info | Do not resize the virtual disk; just show the size the disk will have after resizing. |
| --units | Displays the disk size in kilobytes (K), megabytes (M, default), or gigabytes (G). |

prl_disk_tool validate

Validates the operating system on the virtual hard disk.

Syntax

```
prl_disk_tool validate --hdd <disk_path>
```

Options

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

pnetstat

This utility outputs traffic usage statistics for virtual machines.

Syntax

```
pnetstat [-v <ID>] [-c <class>] [-a] [-r] [-h] [-t <ct|vm|all>]
```

The utility displays input and output traffic for virtual machines for each defined network class, in bytes. The network classes are described in the `/etc/vz/conf/networks_classes` file. `pnetstat` displays statistics only on virtual machines that were started at least once.

Options

| Name | Description |
|-----------------------|---|
| -v <ID> | Display statistics for the specified virtual machine. Multiple <code>-v</code> options can be given to a single <code>pnetstat</code> invocation. |
| -c <class> | Display statistics for the <class> class only. |
| -a | Display statistics for all classes. |
| -r | Rounds down the statistics results, which is shown in bytes by default. In this case, the statistics is displayed in the following units of measurement, depending on the amount of traffic: <ul style="list-style-type: none"> ▪ <code>K(b)</code> — kilobytes, ▪ <code>M(b)</code> — megabytes, ▪ <code>G(b)</code> — gigabytes. |
| -t <ct vm all> | Display statistics for: <ul style="list-style-type: none"> ▪ <code>ct</code> — Containers only, ▪ <code>vm</code> — virtual machines only, ▪ <code>all</code> — both Container and virtual machines. |
| -h, -help | Display the utility usage information. |

CHAPTER 5

Glossary

This glossary defines terms and spells out abbreviations used in Parallels Cloud Server documentation. References to terms defined elsewhere in the glossary appear in *italics*.

Application template. A template used to install a set of applications in *Containers*. See also *Template*.

Container (or regular Container). A virtual private server, which is functionally identical to an isolated standalone server, with its own IP addresses, processes, files, users database, configuration files, applications, system libraries, and so on. Containers share one *Parallels server* and one OS kernel. However, they are isolated from each other. A Container is a kind of ‘sandbox’ for processes and users.

Guest operating system (Guest OS). An operating system installed inside a virtual machine and Container. It can be any of the supported Windows or Linux operating systems.

Hardware virtualization. A technology allowing you to virtualize physical servers at the hardware level. Hardware virtualization provides the necessary environment for creating and managing Parallels virtual machines.

Operating system virtualization (OS virtualization). A technology allowing you to virtualize physical servers at the operating system (kernel) level. OS virtualization provides the necessary environment for creating and managing Parallels Containers.

OS template (Operating System template). A template used to create new *Containers* with a pre-installed operating system. See also *Template*.

Package set. See *Template*.

Parallels server (physical server or server). A server where the Parallels Cloud Server software is installed for hosting Parallels virtual machines and Containers. Sometimes, it is marked as Container 0.

Parallels Cloud Server license key. A license key that you should install on the physical server to use Parallels Cloud Server. Every physical server must have its own key installed.

Parallels Cloud Storage license key. A license key for Parallels Cloud Storage to use its functionality.

Memory and IOPS deduplication. A feature introduced in Parallels Cloud Server 6.0. By caching identical files in multiple Containers it helps eliminate memory and I/O bottlenecks and reduce the starting times of multiple similar Containers.

Parallels Virtuozzo Containers for Linux. An operating system virtualization solution allowing you to create multiple isolated Containers on a single physical server to share hardware, licenses, and management effort with maximum efficiency.

Private area. A part of the file system storing *Container* files that are not shared with other *Containers*.

Template (package set). A set of original application files (packages) repackaged for mounting. There are two types of templates. OS Templates are used to create new *Containers* with a pre-installed operating system. Application templates are used to install an application or a set of applications in *Containers*.

UBC. An abbreviation of *User Beancounter*.

User Beancounter. The subsystem of the Parallels Cloud Server software for managing *Container* memory and some system-related resources.

Virtual Environment (VE). An obsolete designation of a *Container*.

Virtuozzo File System (VZFS). A virtual file system for legacy Containers.

Virtual machine (VM). A computer emulated by Parallels Cloud Server. Like a Container, a virtual machine is functionally identical to an isolated standalone computer, with its own IP addresses, processes, files, users database, configuration files, applications, system libraries, and so on. However, unlike Containers, virtual machines run their own operating systems rather than sharing one operating system kernel.

Index

A

About Parallels Cloud Server 6.0 - 8
About This Guide - 9
add - 57

B

Backup and Restoration Utilities - 103
Backup Configuration File - 38
Backup Parameters - 25

C

Container Action Scripts - 45
Container Configuration File - 19

D

del - 58
Displayed Information - 65
Documentation Conventions - 10

E

EZ Template Management Utilities - 108

F

Feedback - 11
Floppy Disk Drive Management Options - 165

G

General Options - 82
General Syntax - 145
Getting Help - 11
Global Parallels Cloud Server Configuration
File - 13
Glossary - 182

H

Hard Disk Drive Management Options - 94,
161
High Availability Scripts - 47
High Availability Utilities - 69

I

Interactive Controls - 67
Introduction - 8

K

Kernel Parameters - 30

L

Linux Distribution Configuration Files - 26
list - 58

M

Managing Containers - 73
Managing Parallels Cloud Server 6.0 - 12
Managing Shared Folders - 168
Managing Virtual Devices - 160
Managing Virtual Machine Backups - 174
Managing Virtual Machines - 145
Matrix of Parallels Cloud Server Command-
Line Utilities - 73
Memory and IOPS Deduplication
Configuration File - 28
Migration Utilities - 99
Miscellaneous Parameters - 19
Modifying Virtual Machine Configuration - 156

N

net add - 49
net del - 51
net list - 51
net set - 50
Network Adapter Management Options - 163
Network Classes Definition File - 28
Network Options - 91
Networking Parameters - 24

O

Offline Management Configuration Files - 31
Optical Disk Drive Management Options -
162

Organization of This Guide - 9
Overview - 44

P

Parallel Port Management Options - 166
Parallels Cloud Server Configuration Files - 12
Parallels Cloud Server Scripts - 44
Parallels Cloud Server Updates - 72
Parallels Cloud Server Utilities - 47
pbackup - 103, 174
pcompact - 132
pcompact.conf - 42
pfcache - 133
pfcache check - 133
pfcache dump - 134
pfcache mark, unmark - 134
pfcache purge - 135
pfcache stat - 135
pfcache verify - 136
pmigrate - 99, 171
pnetstat - 136, 181
prestore - 105, 175
privnet - 57
prl_convert - 177
prl_disk_tool - 137, 178
prl_disk_tool compact - 137, 178
prl_disk_tool configure - 178
prl_disk_tool convert - 179
prl_disk_tool encrypt, decrypt - 179
prl_disk_tool merge - 137, 180
prl_disk_tool resize - 138, 180
prl_disk_tool validate - 180
prlctl - 75, 145
prlctl backup, backup-list, backup-delete, restore - 106, 175
prlctl capture - 146
prlctl change-passwd - 146
prlctl clone - 102, 147
prlctl console - 76
prlctl convert - 147
prlctl create - 76, 148
prlctl delete - 149
prlctl delete, destroy - 78
prlctl encrypt, decrypt - 149
prlctl enter - 150
prlctl exec - 150
prlctl exec, enter - 78
prlctl list - 98, 150

prlctl list Output Parameters - 98, 151
prlctl migrate - 152
prlctl mount, umount - 79, 152
prlctl move - 79, 153
prlctl pause, suspend, resume - 153
prlctl problem-report - 80, 154
prlctl register, unregister - 80, 154
prlctl reinstall - 81
prlctl reset-uptime - 155
prlctl server - 170
prlctl set - 82, 155
prlctl snapshot, snapshot-list, snapshot-switch, snapshot-delete - 95, 169
prlctl start, stop, restart, reset, status - 170
prlctl start, stop, restart, status - 96
prlctl suspend, resume - 97
prlsrvctl - 48
prlsrvctl info - 48
prlsrvctl install-license - 49
prlsrvctl net - 49
prlsrvctl problem-report - 51
prlsrvctl set - 52
prlsrvctl shutdown - 53
prlsrvctl statistics - 54
prlsrvctl usb - 54
prlsrvctl user list - 56
prlsrvctl user set - 56
pstat - 63
pstat Configuration File - 33

R

Removing Devices from Virtual Machines - 168
Resource Management Options - 86
Resource Management Parameters - 21

S

Serial Port Management Options - 165
set - 57
shaman General Syntax - 69
shaman join - 70
shaman leave - 70
shaman stat - 71
shaman top - 71
shaman.conf - 42
Sound Device Management Options - 167
Supplementary Tools - 132

U

USB Controller Management Options - 167

usb del - 55

usb list - 55

usb set - 55

V

Virtual Machine Action Scripts - 46

vzcalc - 138

vzcfgscale - 138

vzcfgvalidate - 139

vzcheckovr - 140

vzcpucheck - 140

vzcreaterepo - 130

vzctl convert - 140

vzhwcalc - 141

vziotop - 142

vzlicload - 58

vzlicupdate - 59

vzlicview - 60

vzlmnd Configuration File - 32

vzmemcheck - 142

vzmktmpl - 126

vzmtemplate - 131

vznetcfg - 60

vzpid - 142

vzpkg - 108

vzpkg clean - 125

vzpkg create appcache - 120

vzpkg create cache - 118

vzpkg fetch - 124

vzpkg info - 111

vzpkg install - 115

vzpkg install template - 109

vzpkg list - 110

vzpkg localinstall - 122

vzpkg localupdate - 122

vzpkg remove - 117

vzpkg remove appcache - 121

vzpkg remove cache - 119

vzpkg remove template - 110

vzpkg status - 114

vzpkg update - 116

vzpkg update appcache - 120

vzpkg update cache - 119

vzpkg update metadata - 125

vzpkg update template - 109

vzpkg upgrade - 123

vzpkg.metafile - 127

vzpkgproxy - 130

vzpkgproxy Configuration File - 41

vzps, vztop - 143

vzreboot - 69

vzrmnd Configuration File - 35

vzsplit - 143

vzstatrep - 62

vzstatrep Configuration File - 37

vztt Configuration File - 41

vzvnp Configuration File - 29