

---

# Parallels Business Automation - Standard

## Advanced Configuration Manual

Release 3.3



(c) 1999-2008

---

# Contents

<b>Preface</b>	<b>5</b>
Documentation Conventions.....	5
Typographical Conventions.....	5
Shell Prompts in Command Examples .....	6
General Conventions .....	6
Feedback.....	7
<b>Parallels Business Automation - Standard Profiling</b>	<b>8</b>
Turning Profiler On and Off.....	8
Viewing Profiler Statistics.....	9
<b>Parallels Business Automation - Standard Cluster</b>	<b>10</b>
Introduction to Clustering.....	11
Database Splitting And Clustering.....	13
Backend Clustering.....	13
<b>Backing Up Management Node</b>	<b>15</b>
Using Management Node Backup Tool.....	16
<b>Index</b>	<b>18</b>

---

# Table of Figures

<b>Figure 1: HSPcomplete Cluster .....</b>	<b>12</b>
--	-----------

*ISBN: N/A*  
*Parallels*  
*13755 Sunrise Valley Drive*  
*Suite 600*  
*Herndon, VA 20171*  
*USA*  
*Tel: +1 (703) 815 5670*  
*Fax: +1 (703) 815 5675*

*© 1999-2008 Parallels. All rights reserved.*  
*Distribution of this work or derivative of this*  
*work in any form is prohibited unless prior*  
*written permission is obtained from the*  
*copyright holder.*

## CHAPTER 1

# Preface

This manual describes the tools that allow improving Parallels Business Automation - Standard performance.

## In This Chapter

Documentation Conventions.....	5
Feedback .....	7

---

## Documentation Conventions

Before you start using this guide, it is important to understand the documentation conventions used in it. For information on specialized terms used in the documentation, see the Glossary at the end of this document.

## Typographical Conventions

The following kinds of formatting in the text identify special information.

Formatting convention	Type of Information	Example
Triangular Bullet(➤)	Step-by-step procedures. You can follow the instructions below to complete a specific task.	<i>To create a Container:</i>
Special Bold	Items you must select, such as menu options, command buttons, or items in a list.	Go to the <b>QoS</b> tab.
<i>Italics</i>	Titles of chapters, sections, and subsections.	Read the <b>Basic Administration</b> chapter.
	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>EZ templates</i> . To destroy a Container, type <code>vzctl destroy <i>CTid</i></code> .
Monospace	The names of commands, files, and directories.	Use <code>vzctl start</code> to start a Container.

<b>Preformatted</b>	On-screen computer output in your command-line sessions; source code in XML, C++, or other programming languages.	Saved parameters for Container 101
<b>Monospace Bold</b>	What you type, contrasted with on-screen computer output.	<b># rpm -V virtuozzo-release</b>
<b>CAPITALS</b>	Names of keys on the keyboard.	SHIFT, CTRL, ALT
<b>KEY+KEY</b>	Key combinations for which the user must press and hold down one key and then press another.	CTRL+P, ALT+F4

## Shell Prompts in Command Examples

Command line examples throughout this guide presume that you are using the Bourne-again shell (bash). Whenever a command can be run as a regular user, we will display it with a dollar sign prompt. When a command is meant to be run as root, we will display it with a hash mark prompt:

Bourne-again shell prompt	\$
Bourne-again shell root prompt	#

## General Conventions

Be aware of the following conventions used in this book.

- Chapters in this guide are divided into sections, which, in turn, are subdivided into subsections. For example, **Documentation Conventions** is a section, and **General Conventions** is a subsection.
- When following steps or using examples, be sure to type double-quotes ("), left single-quotes ('), and right single-quotes (') exactly as shown.
- The key referred to as RETURN is labeled ENTER on some keyboards.

The root path usually includes the `/bin`, `/sbin`, `/usr/bin` and `/usr/sbin` directories, so the steps in this book show the commands in these directories without absolute path names. Steps that use commands in other, less common, directories show the absolute paths in the examples.

---

## Feedback

If you spot a typo in this guide, or if you have thought of a way to make this guide better, we would love to hear from you!

If you have a suggestion for improving the documentation (or any other relevant comments), try to be as specific as possible when formulating it. If you have found an error, please include the chapter/section/subsection name and some of the surrounding text so we can find it easily.

Please submit a report by e-mail to [userdocs@swsoft.com](mailto:userdocs@swsoft.com).

## CHAPTER 2

# Parallels Business Automation - Standard Profiling

Profiling of web requests in Parallels Business Automation - Standard is very useful. This allows understanding, which requests are slow (from some IP address, or to some specific URL, etc.).

You can measure time needed to access the requested web page with the help of HSPC profiling. This allows you to check the performance and spot slow pages needing improvements.

You are not recommended to use profiler all the time as it is likely to put an additional load on your system resources and disk space. When profiling is turned on, performance will reduce, mostly due to operations with file system (for saving log files).

## In This Chapter

Turning Profiler On and Off .....	8
Viewing Profiler Statistics .....	9

---

## Turning Profiler On and Off

Profiler cannot be turned on from Parallels Business Automation - Standard web interface. To activate or deactivate profiler, it is necessary to edit the Parallels Business Automation - Standard configuration file and restart the backend web server.

➤ *To turn profiler on:*

- 1 Edit the Parallels Business Automation - Standard configuration file:

```
/etc/hspc/hspc.conf
```

Set the `CORE_PROFILER` parameter as follows:

```
CORE_PROFILER = 1
```

- 2 Restart the backend web server:

```
/etc/init.d/hspcd restart
```

➤ *To turn profiler off:*

- 1 Edit the Parallels Business Automation - Standard configuration file and set the CORE\_PROFILER to zero:

```
CORE_PROFILER = 0
```

- 2 Restart hspcd.

---

## Viewing Profiler Statistics

When profiling is turned on, it collects all the necessary information and stores it in the Parallels Business Automation - Standard database (logs are saved in files).

Debug/sql logs are saved into files in the `/var/log/hspc/profiler` directory.

In addition, all profiling statistics are shown in the Provider Control Center web interface, on the special page (Action Log -> Profiler Statistics).

The profiler statistics columns include:

**ID** - Process ID;

**Location** - Request URL;

**Wallclock** - real time spent on the request processing from the very start to its finish. Wallclock time may be not necessarily the time occupied only by the requested process. It may record the time during which another interrupting process engaged the system and prevented the requested process from execution.

**User time** - in UNIX terms, the user time of the process;

**System time** - in UNIX terms, the system time of the process;

**Initiator's IP** - the IP of the remote host initiating the process;

**Started** - Start date and time;

**Finished** - Finish date and time.

## CHAPTER 3

# Parallels Business Automation - Standard Cluster

Hosting provider business based on Parallels solutions grows rapidly day by day. Increasing accounts number and resources volume used by them calls for loading of MySQL server. Rising accounts-related activity and extending set of nodes (Parallels Virtuozzo Containers, Plesk, SiteBuilder, etc.) increases Parallels Business Automation - Standard back-end and front-end parts load. Hardware resources consumption exceeds a reasonable limit. The only one solution of the problem is increasing hardware resources volume available both for Parallels Business Automation - Standard management server and MySQL server. However, upgrading hardware resources Parallels Business Automation - Standard is hosted on is not very efficient and scalable solution. The real way out is to use clustering technologies.

Another problem is ensuring an uninterrupted and reliable access to service. In the situation when a client base is massive, any minimal downtime of the management node provokes loss of clients and therefore, loss of returns. Parallels Business Automation - Standard management node and/or MySQL database server maintenance (upgrade, backup/restore operations, hardware problems and so on) produces such undesirable downtimes. The cluster technology will shorten or even allow avoiding downtimes and minimize related risks.

## In This Chapter

Introduction to Clustering .....	11
Database Splitting And Clustering .....	13
Backend Clustering .....	13

---

## Introduction to Clustering

The clustering solution is based on splitting Parallels Business Automation - Standard parts by different hardware nodes. There are three common parts of Parallels Business Automation - Standard software that produce high load:

- front-end (less loading)
- back-end
- database (most loading)

In case of front-end that produces less load, a node will be combined with load balancing system and synchronization server. The MySQL database server can be clustered natively (<http://dev.mysql.com/doc/refman/4.1/en/ndbcluster.html>). Parallels Business Automation - Standard back-end nodes can be replicated many times without any limitations.

Parallels Business Automation - Standard can be installed and successfully used without any clustering. After the node becomes overused the good solution is to move out MySQL server from the management node to a separate server. In case a new MySQL server resources are insufficient a database cluster can be created to prevent this. The time for next step comes when the Parallels Business Automation - Standard management node becomes overloaded. The best solution to solve the problem is to set a cluster for Parallels Business Automation - Standard back-end and load balancing between them.

Here are a few words about load balancing system. The key purposes of the system are to track users' requests and distribute them uniformly between all of available back-end systems. In addition, the load balancing system monitors back-end nodes availability and latency. This allows a balancing mechanism to provide more flexible and powerful loading results.

Another situation is when a hosting business is planned to be huge, so Parallels Business Automation - Standard cluster should be set first to prevent any downtimes. Actually, this situation does not call for any actions different from last stages of clustering described above, so we are not focus on any individual use-cases.

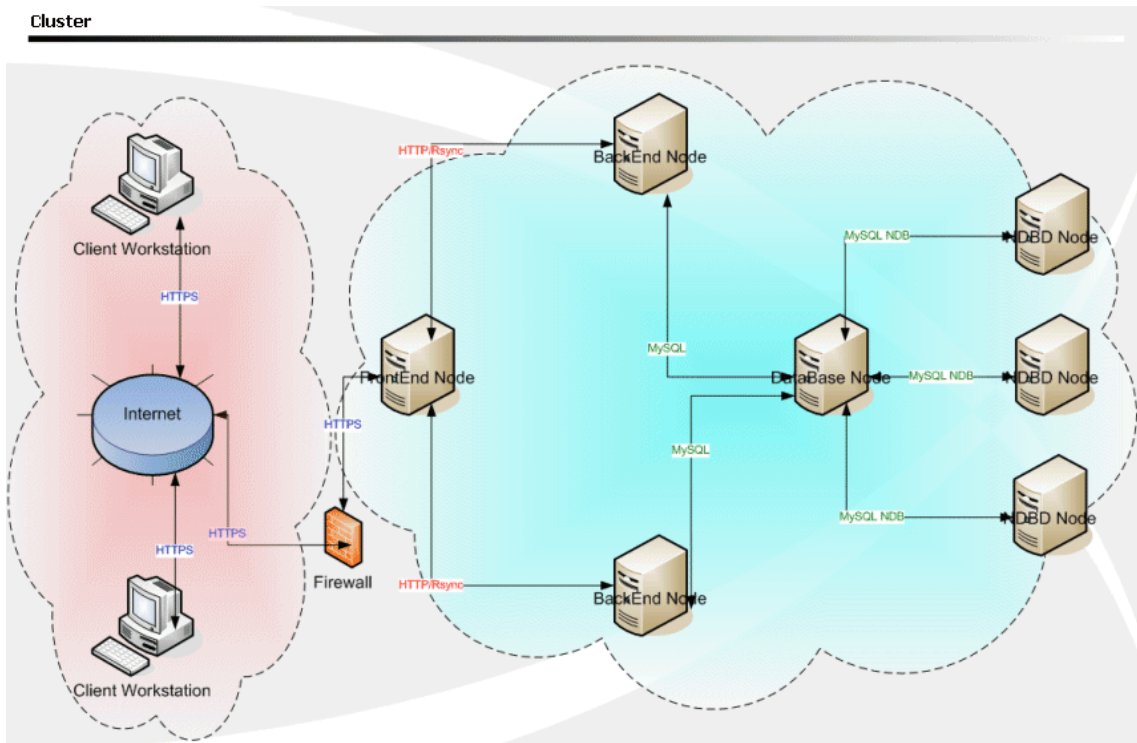


Figure 1: HSPcomplete Cluster

Here, client workstations are located in the “untrusted” Internet area. Parallels Business Automation - Standard front-end node is only accessible through firewall. All other Parallels Business Automation - Standard nodes are located in the “trusted” Ethernet area so there are no any additional resources spending for security.

In a few words, the clustering process may consist of the following steps:

- 1 Move the database from the Parallels Business Automation - Standard Management Node to a separate server.
- 2 Create Parallels Business Automation - Standard backend node(s). If Parallels Business Automation - Standard is installed at a Virtuozzo Container (the most widely used approach), this can be done by cloning the Parallels Business Automation - Standard Container. Allocate IPs to backend nodes and assign hostnames. The original Parallels Business Automation - Standard Container will later also be used as a backend node.
- 3 Create the frontend node. To this effect, create a 'light' Container and install Parallels Business Automation - Standard at it. Then run `hspc-cluster.pl --frontend` at the frontend node and specify all the backend nodes when prompted.

---

## Database Splitting And Clustering

Parallels Business Automation - Standard team has prepared a tool set that will help building the cluster. The tool is shipped in the form of Perl script, its usage is quite simple and we describe it step-by-step below.

First of all, it is necessary to prepare a clean Linux installation with installed MySQL server version not lower than 4.1 with opened SSH access for user `root`.

➤ *To split database from the back-end node:*

- 1 Start the `hspc-cluster.pl` tool in the back-end configuration mode ( with `-b` or `--backend` flag).
- 2 Reply yes to the Do you want to move out database on another server? question.
- 3 Enter a new server's address. After this the script will generate (if needed) and put a DSA SSH public key to a remote server.
- 4 The script will stop the MySQL service and start migrating its data file. For a bulky dataset this can take some time.
- 5 After the database is migrated, the script will be automatically cloned at a remote server and transparently started in the database configuration mode (with `-d` or `--database` flag). The remote script will exit and its local copy will continue to work.

After the steps enlisted above are finished, the database node will be completely split from Parallels Business Automation - Standard. Any unnecessary services will be stopped to prevent parasitic loading of the hardware node. Do not forget to restart related services (will be prompted by the configuration tool as well).

---

## Backend Clustering

Back-end clustering procedure is quite simple. It is possible to add a new back-end nodes on-demand. It is supposed that the database node is already moved out on a remote server. In addition, it is required to have a number of clean Parallels Business Automation - Standard installations which will be the next back-end nodes in a cluster plus another one (and only one!) for the front-end server. It is recommended to use original Parallels Business Automation - Standard installation as a future front-end node.

➤ *This procedure can be repeated many times:*

- 1 Log in to the front-end node and start the `hspc-cluster.pl` tool in the front-end configuration mode (with `-f` or `--frontend` flag)
- 2 Enter a list of back-end nodes available for the cluster. It is recommended but not required to enter the most recently used (production) management node first.
- 3 Select whether you want to set up database access parameters generally or individually for all the back-end nodes. Those settings will be cloned over all the Parallels Business Automation - Standard cluster nodes.

- 4** The script will generate (if needed) a SSH DSA public key and set it for every back-end server specified for future access, so the script will prompt for a root password for each of back-end nodes.
- 5** Rsync synchronization server will be configured by scheduling synchronization every 15 minutes. In this case, the front-end server will be a primary storage node for data synchronization.
- 6** Configuration for every back-end will be done by cloning the script for every node and transparently starting it in an automated or interactive mode (depending on the choice made on step 3).
- 7** Front-end node configuration itself starts. This will install a certain number of configuration files and scripts for a load balancing system.
- 8** All the new services will be registered and configured. The script will prompt also for a necessary services restart. The following services will be run on the front-end node: Task Manager, Virtuozzo Collector and Power DNS services. The rest will be cloned among the back ends.

**CHAPTER 4****Backing Up Management Node**

This chapter describes how to back up the server that runs Parallels Business Automation - Standard (Management Node).

**In This Chapter**

Using Management Node Backup Tool..... 16

## Using Management Node Backup Tool

Management Node can be backed up using the `hspc-mn-backup` utility shipped with Parallels Business Automation - Standard distribution. After Parallels Business Automation - Standard installation `hspc-mn-backup` is located at the Management Node in the `/usr/sbin` directory.

The `hspc-mn-backup` utility allows you to:

- Create Management Node backup.
- Restore Management Node from backup.
- Migrate Management Node from one server to another. In this case, `hspc-mn-backup` creates a Management Node backup, moves it to a server specified, and restores Management Node on a new place.

The following Management Node directories are backed up:

Directory	Description
<code>/var/lib/mysql/</code>	Parallels Business Automation - Standard database
<code>/etc/my.cnf</code>	MySQL configuration file
<code>/var/www/.ssh/</code>	Virtuozzo hardware nodes access
<code>/etc/hspc/</code>	Parallels Business Automation - Standard configuration files and license
<code>/etc/httpd/</code>	Frontend configuration files
<code>/etc/hspcd/</code>	Application server configuration files
<code>/var/opt/hspc-root</code>	Customization files
<code>/var/opt/hspc-event-manager/</code>	Event manager templates
<code>/vz/</code>	Parallels Virtuozzo Containers integration files

Management Node backup is a compressed tarball and by default, it is put into a current directory. A backup file is named `hspc-mn-backup_YYYY.MM.DD` (as design), where `YYYY.MM.DD` is a backup creation date in the Year.Month.Day format (for example, `hspc-mn-backup_2006.09.06.tar.bz2`). A backup creation date is also used as a backup identifier (ID) and you can use this ID in the `hspc-mn-backup` command line to specify a backup file to deal with.

Command line syntax:

```
hspc-mn-backup [options]
```

Options:

Option	Description
<code>-u</code> or <code>--usage</code>	Print a brief help and exit.
<code>-?</code> or <code>-h</code> or <code>--help</code>	Print the manual page and exit.
<code>-b</code> or <code>--backup</code>	Start Management Node backup process.
<code>-r</code> or <code>--restore</code>	Start restore process.
<code>-m</code> or <code>--migrate</code>	Start migration process to the node specified.
<code>-v</code> or <code>--verbose</code>	Increase the verbose level. Single use of the option sets the output to debug level.
<code>-q</code> or <code>--quiet</code>	Decrease the verbose level. Single use of the option sets the output only for serious problems.
<code>-C</code> or <code>--directory</code>	Put backup files into a directory specified instead of a current one.
<code>-i</code> or <code>--id</code>	Use a specified backup (ID) instead of a dated of today.
<code>-p</code> or <code>--no-pre</code>	Do not perform pre-backup actions. Pre-backup actions consist in stopping all the services on a Management Node.
<code>-P</code> or <code>--no-post</code>	Do not perform post-backup actions. Post-backup actions consist in starting all services.
<code>-f</code> or <code>--force</code>	Do not perform the check of the ability of restoring backup on a current node.

# Index

## **B**

Backend Clustering • 13

Backing Up Management Node • 14

## **D**

Database Splitting And Clustering • 12

Documentation Conventions • 5

## **F**

Feedback • 7

## **G**

General Conventions • 6

## **I**

Introduction to Clustering • 11

## **P**

Parallels Business Automation - Standard  
Cluster • 10

Parallels Business Automation - Standard  
Profiling • 8

Preface • 5

## **S**

Shell Prompts in Command Examples • 6

## **T**

Turning Profiler On and Off • 8

Typographical Conventions • 5

## **U**

Using Management Node Backup Tool • 15

## **V**

Viewing Profiler Statistics • 9