

# Comparison of Virtuozzo 4.0 for Linux and Solaris 10 Zones Technology

White Paper

SWsoft, Inc.©

\$Revision: 1.9 \$  
2006

# Table of Contents

<b>1. Preface</b> .....	<b>3</b>
<b>2. Executive Summary</b> .....	<b>3</b>
2.1. High Level Features Comparison .....	4
<b>3. What is Virtuozzo?</b> .....	<b>6</b>
<b>4. What is Solaris 10 Zones?</b> .....	<b>7</b>
<b>5. Complete Feature Comparison Table</b> .....	<b>7</b>
<b>6. Comments on Feature Comparison</b> .....	<b>9</b>
6.1. Management Tools .....	9
6.1.1. Command Line Tools .....	9
6.1.2. GUI Tools .....	10
6.1.3. Operating System and Application Templates .....	12
6.2. Virtualization .....	13
6.3. 'Fair' CPU Scheduler .....	13
6.4. Virtual Networking .....	14
6.5. Filesystem and Data Sharing .....	14
6.6. Disk Usage Limits (quota) .....	15
6.7. Service Level Management .....	15
6.8. Zero-Downtime Migration .....	15
6.9. Memory Management .....	15
6.10. Security .....	16
6.11. Quality-of-Service Management .....	16
<b>7. Tests</b> .....	<b>17</b>
7.1. Hardware .....	17
7.2. Software .....	17
7.3. Basic operations .....	18
7.4. Unixbench benchmark .....	18
7.5. VolanoMark benchmark .....	19
7.6. vzt-static test. Web-server with static pages test .....	19
7.7. vzt-dynamic test. Web-server with dynamic pages test .....	19
7.8. vzt-latency test. Latency of apache web server replies .....	20
<b>8. Future Development</b> .....	<b>20</b>
<b>9. Virtuozzo Unique Features</b> .....	<b>20</b>
<b>10. Summary</b> .....	<b>21</b>

# 1. Preface

This document is a comparison between Virtuozzo for Linux and Solaris 10 Zones Technology as alternative solutions for enterprise server consolidation and for hosting service provider business.

## 2. Executive Summary

Virtuozzo and Solaris 10 Zones are two competing products, allowing to run multiple virtual environments on one physical host. Both have similar underlied technologies, called virtualization on the OS level.

This comparison shows that Virtuozzo product outperforms Solaris Zones by management tools and supported functionality.

Virtuozzo consists of high-level tools (GUI, web-based), low level (CLI, APIs, SDK) and core technology.

Virtuozzo high-level tools are designed specifically for reducing administration costs. These tools have samples management and templates management subsystems, allow to mass-manage VEs, cluster management. They have also a lot of features, requested by Virtuozzo customers during more than 5 years in production, for example, features for higher availability - backups management subsystem, migration, repair/reinstall, snapshots.

Core Virtuozzo technology development is active as well - a number of unique advanced features were developed since first release. SLM - third generation of Virtuozzo resource management, VZFS - copy-on-write filesystem, ZDTM - migration of VEs between hardware nodes without downtime and others. All of them are unique and important, see details in respective sections.

Additionally, Virtuozzo could be easily integrated in existent infrastructure - there are a number of APIs, as well, as already built-in support of other virtualization solutions.

Solaris Zones consists of Solaris Container Manager (web-based add-on to Solaris Management Center), simple command-line tools and core technology.

Solaris Container Manager allows only base operations - VE create/destroy/start/stop/modify. Command-line tools have the same functionality and are complex in use - it is required to issue about 10 commands in order to create zone. Besides that, there is no APIs for their extension and customizing for customer needs. Above makes Solaris Zones technology **non-suitable for enterprise** market.

Solaris core technology is not perfect as well. Solaris resource management for Zones leaves a lot of abilities for possible Denial-of-Service attacks, making its **non-suitable for hosting** market too.

Solaris Zones is a technology and a limited set of basic tools now, while Virtuozzo consists of more robust technology, full set of tools, including command-line tools, web-based and native GUI tools, and API and SDK, allowing easy integration to existent infrastructure.

## 2.1. High Level Features Comparison

**Table 1. General Features**

Feature	Virtuozzo 4.0	Solaris 10 Zones
<b>GUI/Web-Based Tools</b>		
Web-based for HN administrators	VzCC, allows to easily provide required VEs and monitor them, VE samples management, OS/application templates, application-ready VEs, backups/restore, migration, cluster management, VE mass-management, VE package management, etc.	Solaris Container Manager, basic VE operations - create, destroy, start, stop, modify, clone, 2 alerts, usage reports
Web-based for VEs owners	VzCC, allows VE owners to manage their VE through Web-interface, even if VE is offline or broken. VE repair mode, VE reinstallation, restore from backup, package management, monitoring, etc	-
Rich GUI	VZMC, more advanced than VzCC tool for hardware nodes administrators, rich alerting, reports, monitoring systems	-
conversion tools	V2V, P2V tools - allow to convert physical hosts or VEs with other base virtual technology to Virtuozzo VEs.	-
<b>Low Level Tools/API</b>		
Provision	samples, application-ready samples and OS/application templates allow fast(~5 sec) and easy(simple one-line command) VE creation, easy applications add/remove even for stopped VEs	no templates, no samples, only clone, creation tooks significant time - ~15 minutes for VE without inherited dirs, - required additional configuration before first start

Comparison of Virtuozzo 4.0 for Linux and Solaris 10 Zones Technology

Feature	Virtuozzo 4.0	Solaris 10 Zones
Software updates	vzup2date allows to keep Virtuozzo system up-to-date, mass-management combined with package management from outside allows to keep VEs up-to-date	Sun Update Connection
Availability	backuping, online migration, snapshots, high availability(DBD), etc	Sun Cluster HA
Monitoring	standard monitoring tools extensions, own monitoring toolkit, HN health monitoring, VEs resource usage, more than 40 parameters controlled	number of processes, CPU usage, memory usage are monitored
APIs	XML interface, RPC XML, SOAP, SDKs	-
<b>Advanced Core Features</b>		
Copy-On-Write Filesystem	VZFS, allows to save disk space and memory - code and data are shared between VEs	-
Service Level Management	allows VE owners to give advance in resources distribution to important services; simplifies resource management by providing simple way of managing complex resource pools	-
Zero-Downtime Migration	allows to migrate VEs from one hardware node to another without any downtime	-
Name-Based Hosting	removes requirement of real IP address for VEs with Internet-accessible services (www, ftp, smtp, pop)	-
<b>Core Features</b>		
Resource Management	complete, more than 30 parameters with guarantees and upper limits, CPU affinity/share/limits, virtual/physical/kernel memory limits, network bandwidth/accounting, disk space limits, quotas in environments	incomplete, CPU affinity/share, network bandwidth, number of processes

Feature	Virtuozzo 4.0	Solaris 10 Zones
Virtualization (isolation)	complete, virtualized networking, devices, firewall, IPCs, signals, etc.	incomplete, functionality restrictions (SOCK_RAW, SOCK_PACKET), possible DoS attacks on system-wide limits, opened file descriptors, for example.
<b>Platforms</b>		
host OS	Linux, Windows	Solaris
environment OS	Linux, Windows	Solaris
hardware	x86, x86_64, ia64, (pa-risc, ppc*)	x86, sparc, sparc64
other virtualization technologies support by tools	VMware, XEN, Parallels and Microsoft Virtual Server	-

From the above table it is seen, that Virtuozzo has much more features, than Solaris Zones. In terms of the source code, about 0.7% of overall Vz code exists in Solaris Zones.

But beside that

- 0.9% of overall Vz code is basic functionality code,
- 1.4% of overall Vz code is advanced functionality code.

Overall VzLinux core code is now just 6% of all Virtuozzo/Tools related code, the rest:

1. 35% - Backend Tools, Command Line tools and P2v/V2P and such,
2. 40% - GUI/Web UI,
3. 19% - Windows virtualization kernel/drivers,

Most active area of development now is actually backend tools/ui for them, second most active - windows virtualization and 3rd most active is advanced Virtuozzo functionality. None of it is a part of Zones.

### 3. What is Virtuozzo?

**Virtuozzo** (<http://www.virtuozzo.com/>) is a patent-pending technology and a product from the SWsoft (<http://www.swsoft.com/>) company. SWsoft (<http://www.swsoft.com/>) has been designing it from the scratch specifically for the server consolidation and hosting service uses.

Virtuozzo implements partitioning of the commodity hardware by the virtualization of the operating system. Virtuozzo currently works on the Linux and Windows systems, and the same technology can be applied to other operating systems.

The virtualization component of Virtuozzo technology allows to run simultaneously from 1 to 5,000 Virtual Environments (VE) on a single computer, depending on the hardware and the performance requirements for the Virtual Environments.

Virtuozzo Virtual Environments are functionally identical to standalone servers, with each of the environments having its own processes, files and users. **Virtuozzo technology** provides high levels of

- isolation,
- manageability of resources, and
- efficiency

and can be used for server consolidation both in the enterprise and hosting service provider environments.

**Virtuozzo** as a **product** combines this advanced technology with

- multi-tiered tools, including command-line tools, a graphical Management Console and a web-based Power Panels for Virtuozzo administrators and enterprise users, and
- support infrastructure, including documented APIs, SDKs, Administrator and User Guides, technical support, and training programs.

## 4. What is Solaris 10 Zones?

**Solaris 10 Zones** (<http://www.sun.com/bigadmin/content/zones/>) partitioning technology developed by Sun Microsystems Company (<http://www.sun.com>) and is the part of Solaris 10 Containers ([http://www.sun.com/software/whitepapers/solaris10/grid\\_containers.pdf](http://www.sun.com/software/whitepapers/solaris10/grid_containers.pdf)). Solaris Containers consist of zones technology and resource management over all Solaris 10 instances. Zones are used to virtualize operating system services and provide an isolated and secure environment for running applications.

Zones can be used on any machine that is running the Solaris 10 release. The practical limit for the number of zones on a system with 4Gb of RAM is about 100 zones which is much less than about 500 Virtuozzo VEs on the same node.

Sun Microsystems plans to use Zone technology for SUN Grid Computing<sup>1</sup> to run resource-intensive applications and special attention is paid to Oracle Computing<sup>2</sup>. As for now zones are used for development and testing purposes.

## 5. Complete Feature Comparison Table

Table 2.

	Virtuozzo 4.0	Solaris 10 Zones
--	---------------	------------------

Comparison of Virtuozzo 4.0 for Linux and Solaris 10 Zones Technology

	Virtuozzo 4.0	Solaris 10 Zones
<b>GUI tools</b>		
Web-based for HN administrators	+, VzCC	±, Solaris Container Manager <sup>a</sup>
Web-based for VEs owners	+, VzCC	-
Rich GUI	+, VZMC	-
conversion tools	+, V2V, P2V	-
<b>Virtualized infrastructure management tools</b>		
OS Templates	~45	-
Application Templates	~500	-
Patch management	+	+
Online backup/restore	+	-
Online migration	+	-
Provisioning	+	± <sup>b</sup>
Monitoring	+	+
XML interface and SDKs	+	-
High Availability	+	+
Continuous Data Protection	+	-
Services API - command line tools	+	-
<b>core functionality</b>		
OS virtualization	+	± <sup>c</sup>
fully virtualized network	+	± <sup>d</sup>
filesystem with copy-on-write	+	-
disk usage limit	+	-
user/group disk quotas in VE	+	-
'fair' CPU scheduler	+	+
quality-of-service management	+	-
name-based VEs with kernel-level proxy	+	-
service level management	+	-
zero-downtime migration	+	-
<b>performance</b>		
memory management efficiency	high	low
filesystem performance	no overhead	extra copy for each VE
filesystem buffers in memory	single copy	extra copy for each VE
<b>stability</b>		
first release	in 2000	in 2005
percentage of core code	3%	80% <sup>e</sup>

	Virtuozzo 4.0	Solaris 10 Zones
percentage of tools code	97%	20%
available test suite	extensive	absent
<b>platforms</b>		
host OS	Linux, Windows	Solaris
environment OS	Linux, Windows	Solaris
hardware	x86, x86_64, ia64, (pa-risc, ppc*)	x86, sparc, sparc64
other virtualization technologies support by tools	VMware, XEN, Parallels and Microsoft Virtual Server	-
Notes:		
a. Only basic operations, including create/destroy/start/stop/modify/clone.		
b. Takes about 15 minutes, complex - require to issue more than 10 commands.		
c. See Virtualization section for details.		
d. See Virtual Networking section for details.		
e. Estimated.		

## 6. Comments on Feature Comparison

### 6.1. Management Tools

#### 6.1.1. Command Line Tools

The following table summarizes the features of Virtuozzo and Solaris Zones command line tools.

**Table 3.**

	Virtuozzo 4.0	Solaris 10 Zones
VE creation	+	± <sub>a</sub>
VE start	+	+
VE stop	+	+
VE instant restart	+	-
VE reinstallation	+	-
node overcommit calculation tools	+	-
monitoring toolkit	+	-
ps & top utilities expansions with VE IDs	+	-
on-boot startup scripts	+	-
VE online migration	+	-

	Virtuozzo 4.0	Solaris 10 Zones
changing VE resource assignments	+	± <sup>b</sup>
VE resources calculation tools	+	-
changing VE network configuration	+	-
software installation and update in the VE from outside	+	± <sup>c</sup>
execution of commands inside the VE from outside	+	+
tools for creation own template	+	-
tools for "templating" common VE files	+	-
backup tools	+	± <sup>d</sup>
backup compression	+	-
per-VE backup configuration	+	-
providing access to hardware devices for the VE	+	+

Notes:

- a. Takes about 15 minutes, complex - require to issue more than 10 commands.
- b. Not all resources could be changed without zone restart
- c. It's unable to update package set from outside for specific zone, only for all zones.
- d. Backup is performed for whole node, could not be performed for specific zone only.

As it can be seeing from the table above, Solaris Zones has only a simple set of tools designed for basic VE operations. These tools are significantly smaller and simpler than the respective Virtuozzo tools. Furthermore, Virtuozzo has a set of unique tools, which were requested by customers during previous Virtuozzo versions usage and which are important for enterprise and hosting markets.

Solaris Zones has no command line tools for VE downtime reducing, quick VE provision, simplification of administration of multiple environments.

### 6.1.2. GUI Tools

Virtuozzo has two full-fledged GUI tools, VzCC and VzMC. VzCC is a web-based Control Center for hardware nodes administrators and VEs owners, while VzMC is GUI console, concerned more on the requirements of hardware nodes administrators. Both of them use APIs, provided by VzAgent - daemon, which runs in special VE on each managed hardware node.

Solaris Zones can be managed by means of Solaris Container Manager, which is add-on to Solaris Management Center. The last has 3-level infrastructure - Server, Agents and Consoles. Consoles communicate with server, which in turn collect information and manipulate with agents. The neck of

such approach is the server node, which should be powered enough. It is impossible to understand, why Sun also restricts such nodes to be SPARC only, while Solaris 10 may be installed on x86 nodes also. This fact makes it impossible to use Management Center for x86-only nodes cluster management.

The following table shows VzCC, VzMC and Solaris Container Manager features in more details.

**Table 4. GUI tools features**

	VzCC 4.0	VzMC 4.0	Solaris Container Manager 3.6.1
VE creation	+	+	+
VE start	+	+	+
VE stop	+	+	+
VE destroy	+	+	+
VE clone	+	+	+
VE modification	+	+	+
VE instant restart	+	+	-
VE reinstallation	+	+	-
VE snapshots	+	+	-
action scripts	+	+	-
access to hardware devices for VE	+	+	-
VE self-management	+	-	-
VE offline self-management	+	-	-
VE repair self-management	+	-	-
VE mass-management	+	+	-
monitoring toolkit	+ <sup>a</sup>	+	+
alerting tools	+ <sup>b</sup>	+	± <sup>c</sup>
VE online migration	+	+	-
VE packages management	+	+	-
execution of commands inside VE	+	+	-
requests for VE provision	+	+	-
automated VE provision	+	+	-
VE samples management	+	+	-
application-ready VE	+	+	-

	VzCC 4.0	VzMC 4.0	Solaris Container Manager 3.6.1
HN/VE software updates	+	+	-
templates management	+	+	-
templates mass-management	+	+	-
tools for "templating" common VE files	+	+	-
backup tools	+	+	-
backup compression	+	+	-
per-VE backup configuration	+	+	-
clusters management	+	+	-
support of other virtualization solutions	+	+	-
Notes: a. >30 params for each VE and HN. b. >60 alerts types. c. 2 alerts types, CPU and memory usage.			

This table shows, that all operations, available by means of command line tools, are supported by Virtuozzo GUI tools as well. So, almost all operations in Virtuozzo may be performed by multiple ways, allowing to use the most applicable in the each case.

Solaris Container Manager, which is the one GUI tool for Solaris Zones, is significantly smaller and simpler than the respective Virtuozzo tools. It doesn't allow to customize zones, backup, migrate, restore, reinstall them, manage zone when it is stopped, has no interface for zone owner and many other useful features, available in Virtuozzo Control Center and Virtuozzo Power Panels.

### 6.1.3. Operating System and Application Templates

Virtuozzo provides a comprehensive system of operating system and application templating. Virtuozzo templates provide

- high efficiency of hardware resources usage allowing the maximal sharing of common executables and libraries on the disk and in the memory,
- automation of management and reduction of manual operations: the selected applications can be easily installed or updated in multiple or all Virtual Environments via Virtuozzo Management Console or by simple scripts.

Solaris Zones has no templates. During zone creation, full set of packages, installed on the hardware node is copied inside zone root. This is time-consuming and not efficient as many of the hardware node packages are not required for zone.

## 6.2. Virtualization

Both Virtuozzo and Solaris Zones allow to run a number of virtual OS instances on a single computer. The instances are isolated from each other and from the host system and can't see any processes, files, inter-process communication objects, and network connections outside of their scope.

Virtuozzo virtual environment — is functionally identical to standalone servers and can't break other environments or hardware node because of complete objects isolation or per-environment limits on system-wide limited objects. For example, Virtuozzo has per-environment limits on number of opened files, sockets, IPCs objects, number of network filter entries, etc.

Solaris 10 Zone — **is not** functionally **identical** to standalone servers, because:

- Zone disallowed to:
  - administration of IP firewall, masquerading and accounting,
  - modification of routing tables (route),
  - use of RAW sockets (traceroute),
  - use of PACKET sockets (tcpdump),
- Zone can break other zones or hardware node due to limited resources isolation. It can exhaust one of the following system-wide resource:
  - user/kernel memory,
  - number of IPCs objects,
  - number of file descriptors/sockets.

These restrictions make Solaris Zones to be useful only for trusted environments and to be **completely useless** for hosting purposes, where environments are not trusted at all.

## 6.3. 'Fair' CPU Scheduler

From the Quality of Service point of view, each environment should have a guaranteed share of CPU time. Basically, two thresholds are involved here:

- minimally guaranteed CPU share

- maximally allowed CPU share

Both Virtuozzo and Solaris 10 have such schedulers, which guarantee CPU share, cap allocated CPU power and have ability to assign CPU sets to environments.

## 6.4. Virtual Networking

Virtuozzo has the concept of a virtual network device. Virtual network device is accessible inside environments and let packets pass between the host system and environments.

Such approach allows to have full-fledged virtual network including:

- packet filtering in VE and the host system
- bandwidth management
- isolated routing
- name-based hosting

Virtuozzo name-based hosting allows to have VE with http, smtp, ftp and pop3 services without assigning them a real IP address. The benefits of this is un-necessity to purchase IP addresses and diminution management expenses.

Solaris Zones are assigned by logical interfaces — aliases to physical network device with special flags. Such approach is **inflexible** and follows in additional complexity in implementation of per-zone firewall, routing. More than, it is also **insecure** by default, and lead to zone management restrictions (prohibition of RAW sockets, ifconfig, PACKET sockets, etc.) in order to fix this issue.

## 6.5. Filesystem and Data Sharing

One of the important points for efficient server consolidation is an efficient method of storing data on the disk, so that the common files (such as executables, libraries, and identical data files used by multiple virtual servers) are stored only once.

Virtuozzo allows to keep the common data in only one place.

Virtuozzo implements copy-on-write when data is being modified. The copy-on-write mechanism intercepts the attempts to modify the common data and transparently copies the data into the private area of the virtual server.

Virtuozzo file system has been specially designed and optimised for a high number of virtual servers.

Solaris has no copy-on-write filesystem. This leads to **code and data duplicating** in case of zones with private trees and **significant management restrictions** (unable to customize package sets, modify configuration files as they are shared).

## 6.6. Disk Usage Limits (quota)

Virtuozzo disk quota enables the administrator to limit disk resources available to an environment. Disk quota for users and groups inside environments is also supported and allow to configure it with non-modified Linux quota tools.

Solaris Zones disk usage can be limited only by assignment of whole disk partition to one zone, which is **inflexible and increases administration costs**.

## 6.7. Service Level Management

Virtuozzo Service Level Management is a third generation of Virtuozzo resource management. It allows to manage environment's resources on basis of service importance for VE's owner. For example, VE's owner may assign higher priority to web server, sacrificing resources for mail and DNS servers. Another one useful feature of SLM is simplification of resource management by decreasing a number of managed resources. This decrease is achieved by combining of similar resources in one, providing the same level of isolation as earlier.

## 6.8. Zero-Downtime Migration

Virtuozzo Zero-Downtime Migration allows to migrate VE from one hardware node to another without networks communication break. This is unique Virtuozzo feature, based on checkpoint/restoring full environment state, including processes, connections and all related in-kernel objects states.

## 6.9. Memory Management

Solaris Zones memory management has following disadvantages:

- **memory usage is not limited** - each zone may exhaust all physical memory or swap space,
- **data and code duplicating in memory** - this lead to ineffective memory usage.

These disadvantages **significantly decrease a number of possible usage scenarios** due to strong requirement to run only trusted environments and **decrease a maximal number of environments per-node** due to more memory usage per environment.

## 6.10. Security

The security features of a technology allowing to run multiple virtual servers on a single computer need to be analysed from 2 points of view:

- the quality of isolation among different virtual servers and between a virtual server and the host system, and
- the quality of protection of the virtual servers and the host system from Denial-of-Service attacks caused by an accidental or deliberate misbehaviour of programs and users.

Virtuozzo completely isolates Virtual Environments from each other and from the host system by namespace separation. The isolation covers all operating system interfaces, including filesystem, devices, all kinds of inter-process communications, network, and management interfaces (e.g., sysctl). Virtual Environments are functionally identical to standalone servers, each having its own configuration, data and executable files, applications, user and administrative accounts. Virtual Environment administrators can manage their Virtual Environments like standalone servers, install and configure software, add users, configure network and firewall parameters.

On top of complete isolation of Virtual Environments, Virtuozzo implements unique resource control and Denial-of-Service protection mechanisms, covering the isolation of CPU, disk, memory, and internal data, such as network connection buffers.

Solaris Zones have only simplistic resource control covering a number of processes, network bandwidth, memory and CPU usage. This is **insufficient** for effective DoS attacks protection.

## 6.11. Quality-of-Service Management

Virtuozzo provides a wide set of features to differentiate the quality of service provided for different Virtual Environments. The controlled quality-of-service parameters are

- CPU time shares,
- a set of memory-related parameters,
- disk space and inodes,
- network bandwidth,
- the number of processes, network sockets, and other operating system resources.

The total number of Virtuozzo resource control parameters is more than 20.

All Virtuozzo quality-of-service parameters can be changed on the fly. When the environment resource demand grows, the allowed resources may be increased with no service interruption.

Solaris technology provides a limited resource isolation of Zone and allows to control:

- number of processes inside Zone,
- CPU share,
- network bandwidth,
- physical memory usage,

The crucial QoS parameters for Zones lack in Solaris 10 concerns IPC, file descriptors. Solaris allows from inside Zone to allocate all of these resources thereby to block functioning of other Zones and even physical server. Poor resource isolation can be thought as critical security hole because it enables DoS attacks on other Zones and hardware node.

Another disadvantage of Solaris QoS for Zones consists in hard limits. For example, if zone exceeds its process number limit, then it can't be stopped graceful, because during stop scripts execution it may be required to create additional process, which will fail. This may lead to incorrect service shutdown and possible loss of in-memory data.

In contrast to Solaris, Virtuozzo resource management allows to set:

- guaranteed limits - guaranteed amount of a resource, available for use in environment;
  - upper limits - amount which can be consumed during the 'grace' period;
- so the problem with additional resources for the short time-period is avoided.

## 7. Tests

### 7.1. Hardware

- Intel(R) Celeron(R) CPU 2.80GHz, 256MB RAM, SATA 82Gb HDD

## 7.2. Software

- Solaris 10
- Virtuozzo 4.0
- CentOS 4.2

Tests were run for Virtuozzo VEs and Solaris Zones.

## 7.3. Basic operations

**Table 5. Basic operations, time**

Operation	Virtuozzo for Linux 4.0	Solaris 10 Zones
VE create	2 sec	15 min
VE destroy	instant	1.5 min
VE start	3 sec	5 sec
VE stop	2 sec	2 sec
VE reboot	6 sec	8 sec

These tests clearly show an advance of the copy-on-write filesystem, used by Virtuozzo, over its absence. In Virtuozzo, files are not copied during VE creation, but only links are setuped. This follows in significant smaller times of VE create and destroy operations.

## 7.4. Unixbench benchmark

For testing different aspects of the CPU, memory & file systems performance, we used well-known Unixbench test, designed by Byte Magazine. It includes both CPU tests and tests of typical UNIX tasks. All tests have base results and current mark is evaluated as division of current results by these bases. The more value - the better result is.

**Table 6. Unixbench**

Test	Virtuozzo for Linux 4.0, HN	Virtuozzo for Linux 4.0, VE	Solaris 10 Zones, HN	Solaris 10 Zones, Zone
Drystone 2	322.9	322.8	372.6	386.6
Double-Precision Whetstone	162.4	162.4	162.8	162.7
Pipe Throughput	440.7	440.5	332.0	331.9
Pipe-based Context Switching	308.9	309.3	138.3	140.7
Process Creation	359.4	348.7	63.6	63.6

Test	Virtuozzo for Linux 4.0, HN	Virtuozzo for Linux 4.0, VE	Solaris 10 Zones, HN	Solaris 10 Zones, Zone
System Call Overhead	352.2	350.0	186.3	186.7

## 7.5. VolanoMark benchmark

Industry standard benchmark for evaluating the performance of Java applications. Creates a number of chatrooms and simulates intensive messages exchange. The greater number the better.

**Table 7. VolanoMark results**

	Virtuozzo for Linux 4.0, HN	Virtuozzo for Linux 4.0, VE	Solaris 10 Zones, HN	Solaris 10 Zones, Zone
Download rate, messages/s	11707.2	11426	8647	7913.25

## 7.6. vzt-static test. Web-server with static pages test

Apache server with set of static pages is running inside VE. From the remote host http\_load application is running and tries to download static web-pages from VEs.

**Table 8. Web server performance under static workload**

	Virtuozzo for Linux 4.0, HN	Virtuozzo for Linux 4.0, VE	Solaris 10 Zones, HN	Solaris 10 Zones, Zone
Download rate, MB/s	10.3757	9.77207	8.7399	8.80317

## 7.7. vzt-dynamic test. Web-server with dynamic pages test

Apache server and MYSQL database server are running inside VE. From the remote host http\_load application is running and tries to download dynamic web-pages from VEs, which are generated by perl scripts by retrieving some data from proper database server.

**Table 9. Web server performance under dynamic workload**

	Virtuozzo for Linux 4.0, HN	Virtuozzo for Linux 4.0, VE	Solaris 10 Zones, HN	Solaris 10 Zones, Zone
Download rate, MB/s	0.366814	0.351695	0.025187	0.024602

## 7.8. vzt-latency test. Latency of apache web server replies

Apache server with the set of static and dynamic pages (perl scripts) is running inside VE/Zone. From the remote host http\_load application is running and tries to download static and dynamic web-pages from VEs/Zones. The latency of replies is measured.

**Table 10. Web server replies latencies**

	Virtuozzo for Linux 4.0, HN	Virtuozzo for Linux 4.0, VE	Solaris 10 Zones, HN	Solaris 10 Zones, Zone
Static pages, response time, msec	1.2596	1.282135	1.35201	1.3492525
Dynamic pages, response time, msec	3.25364	3.3696875	15.085	16.049

## 8. Future Development

In the future Virtuozzo will provide

- redundant distributed filesystem (TorFS),
- SLM extension,
- easy sandbox,
- transparent package management, simplifying the process of massively updating the common software in multiple Virtual Environments,
- improved roles/permission management,
- smart installation, allowing a fully automated installation of PXE-compatible servers,
- new conversion tools, V2P, P2P
- new security features, including encrypted filesystem and Mandatory Access Control.

## 9. Virtuozzo Unique Features

Virtuozzo has a lot of unique features, which makes it faster, more dense and robust than Solaris Zones.

**Table 11.**

	Virtuozzo	Solaris Zones

	Virtuozzo	Solaris Zones
shared file system with copy-on-write support	+	-
complete resource management	+	-
firewall inside VEs	+	-
cross-VE memory sharing	+	-
VE online migration tools	+	-
VE state checkpoint/restore	+	-
VEs mass-management	+	-
Quality-of-Service support	+	-
performance isolation	+	-
name-based hosting	+	-
a variety of different OS templates	+	-
template mass-management	+	-
ability to create own template	+	-
ability to template common VE files	+	-
per-VE backup configuration	+	-
backup compression	+	-
tools for 'repairing' VEs accidentally made not functional	+	-
localisation-ready GUI interfaces	+	-
open XML-based management API	+	-
complete Plesk integration	+	-
always available Power Panels for VE owners	+	-
auto VE limits calculation	+	-
instant VE reboot	+	-
other virtualization solutions support by tools	+	-

Shared file system with copy-on-write support is important for performance and efficiency, yet keeping VE functionally identical to standalone servers, as discussed in section Filesystem and Data Sharing; cross-VE memory sharing is important for scalability, allowing to have more VEs per host and providing better resource utilisation, as described in Memory Management section; VE migration tools are important for easy and transparent moving VE across servers without downtime; and so on.

## 10. Summary

The major benefits of using Virtuozzo other than Solaris Zones technology are:

- Comprehensive set of management tools significantly decreases management costs;
- VEs mass management and a lot of templates significantly decreases administration costs;
- Multi-level APIs an SDK simplifies integration to existent infrastructure;
- Support of others than Virtuozzo virtualization solutions allows to leveraging existent infrastructure;
- Proven stability and security - more than 400,000 VEs over the world and more than 5 years in production;
- Flexible and complete VE resource control: CPU, processes, memory, network, disk space;

Virtuozzo provides a comprehensive set of management tools, including GUI and web-interface, cover all regular and extensive management operations. These tools let the administrator create and destroy VEs, manage VEs' configurations, install and manage software inside VEs, manage software templates for a fast and trouble-free installation of an OS and applications inside VEs, back up and restore VEs, migrate VEs between computers without downtime, monitor the VEs and system state.

Virtuozzo tools are designed for management 100s to 100,000s VEs installed on 100s of physical servers.

Solaris Zones do not cover all needs of **hosting market** because of:

- **not complete resource isolation** of zones;
- **limited network configurability** from inside zone;
- **rudimentary tools**;

At the same time, Solaris Zones do not cover all needs of **enterprise market** because of:

- **rudimentary tools**;
- **absence of published APIs**;
- **absence of published SDK**;
- **lack of other virtualization technologies support**;

Currently Solaris Zones hardly can be used in production environments mostly due to incomplete virtualization and resource management systems. But even after their completion, it will be years behind Virtuozzo in the feature set. Essentially, Solaris Zones is a technology rather than a product as of now, and a technology is only a small part of the virtualization solution.