

Parallels® Virtuozzo Containers

White Paper

An Introduction to OS Virtualization and Parallels Virtuozzo Containers

www.parallels.com/virtuozzo

Version 2.0



Table of Contents

Introduction	3
IT Overview	4
-Enterprise It Today	4
Parallels Virtuozzo Containers Overview	5
The Complete Server Virtualization Solution	6
How is Parallels Virtuozzo Containers Different?	8
Parallels Virtuozzo Containers Performance	9
Addressing the Challenges of IT and Conclusion	10

Introduction

Today's IT organizations are dealing with the consequences of exploding IT infrastructure growth and complexity. While computing resources continue to increase in power, organizations are unable to fully utilize them in single application deployments and cannot change computing resource assignments easily when application or business requirements change. At the root of the problem is uncontrolled server sprawl, servers provisioned to support a single application. Organizations that implemented hardware virtualization have unwittingly created a new problem: OS sprawl. While hardware remains a considerable cost component, software and management continue to be the largest cost considerations. The daily management and operations functions are daunting, and adding in business continuity requirements, the costs and complexity are overwhelming. Moreover, few tools provide the management and automation to ease the burden on IT departments.

In order to address these critical challenges, IT organizations have to find ways to accomplish the following:


- Improve the flexibility of computing resource assignment
- Decrease complexity to improve manageability of systems
- Automate routine tasks
- Reduce overall management costs through efficiency
- Provide cost-effective data availability and recovery
- Increase the return from their infrastructure investment by better utilizing resources

Operating-system level virtualization and Virtuozzo empower IT organizations to meet these challenges and the increasing demands on IT infrastructure..

IT Overview

ENTERPRISE IT TODAY

During the last two decades, enterprises experienced explosive IT infrastructure growth, both in number of servers and complexity of configurations. The typical organization currently has to manage a complex heterogeneous environment consisting of 100s to 1000s of Windows, Linux, UNIX and other servers.



Historically, most enterprises have deployed a dedicated server for each application or department.

The dedicated server model is extremely costly due to:

- **Infrastructure and Power Costs** Housing, cooling, connecting and providing power to servers can be extremely expensive as servers grow to large numbers, especially as power costs continue to rise. Server electricity requirements alone can be a large cost component; a rack of servers can cost \$17,000 a year in power.¹
- **Hardware Cost** Every year servers increase in capacity and computing power. As servers become more powerful it becomes increasingly difficult to maximize utilization of these powerful resources. IT organizations and application owners have become accustomed to dedicating a server to each application to ensure complete control over the application. In almost every case, dedicating these powerful servers results in over provisioning from 50-500%.
- **Software Cost** Each server requires licenses and support from the OS and application software vendors.
- **Management Cost** Management costs are by far the largest, yet most intangible, portion of server costs; analysts estimate management costs at 50-70% of the total cost of server ownership.² The IT staff has to upgrade, patch, back-up and fix hardware and software, provision new servers and applications, maintain user accounts and perform many other tasks. Often times personnel only have the skill set to manage a single platform, so the overhead increases even further. With all of these considerations, IT departments find it challenging to meet the demands of the associated server management. In addition to the aforementioned costs, IT organizations today face the following issues:
- **Limited Benefit of Infrastructure Assets** There is only so much real estate in the world, and as organizations try to work within limited or costly data centers, the assets must be fully utilized. The obvious assets include rack space and servers, but there are several other associated assets to evaluate including networking equipment and bandwidth, power and even staff resources.
- **Manual, Complex and Time Consuming Management** Most IT departments support numerous OSs. Add applications into IT department support requirements and the complexity is overwhelming just for software. Consider provisioning, patch management, configuration changes, backup and audit tasks and it becomes clear why management is the largest cost component of the data center. Although there are existing tools for many of these, nothing is comprehensive and many are cost prohibitive.

¹ <http://hightech.lbl.gov/DCTraining/opportunity.html>

² IDC The Evolution of the Virtual Environment April 2003

- **Low Flexibility** Any change in the existing network can create a service breakdown or outage. An application may be subjected to an unusual or unexpected processing demand. A hardware issue may necessitate an upgrade. Moving an application server to different or more powerful hardware involves significant planned downtime and administrative effort. Rather than risk interruptions in service or frequent application migrations, servers are consistently heavily over-provisioned, greatly decreasing server utilization and increasing associated costs as more servers are deployed.
- **Unacceptable or Uncertain Data Availability** Data availability guarantees and demands are becoming a reality for many IT organizations. Managers and planners are expected to be able to answer the question “What if..”. While all data has its relative value, in the past there have only been very limited offerings for data availability and recovery, either at the very low end of the cost and recovery time cycle or at the very top. IT organizations need more options for both disaster recovery and data availability.

Each IT department faced with server sprawl is impacted by common symptoms negatively impacting the infrastructure. What are the symptoms of these challenges?


- Increasing numbers of under-utilized servers and other assets, both from processing and capacity perspectives
- Ballooning administrative staff requirements and an inability to stay ahead of workload demands
- Over-provisioned servers, ready for the worst-case workload scenarios
- Slow response time to IT requests including server provisioning and updates
- Incomplete information on hardware and software status and ownership
- No disaster recovery plan for the majority of server and data resources

Most enterprises continue to struggle with the high cost and complexity and low flexibility of IT infrastructure. Virtualization and automation software can help enterprises solve these critical problems.

Parallels Virtuozzo Containers Overview

Parallels Virtuozzo Containers is an operating system level server virtualization solution. Virtuozzo creates isolated partitions or virtual environments (VEs) on a single physical server and OS instance to utilize hardware, software, data center and management efforts with maximum efficiency.

Virtuozzo’s OS virtualization solution has higher efficiency and manageability, making it the best solution for organizations concerned with containing the IT infrastructure and maximizing resource utilization. Virtuozzo’s complete set of management tools and unique architecture makes it the perfect solution for easily maintaining, monitoring and managing virtualized server resources for consolidation and business continuity configurations.



The Virtuozzo low-overhead, efficient architecture maximizes server resources. Parallels created a portable layer to existing operating systems which adds a dynamic partition or Virtual Environment (VE) that resides on a common OS. The single thin Virtuozzo layer introduces only a small percentage of overhead and allows up to 100s of VEs to run on a physical server.

This architecture allows a much more flexible, efficient and cost-effective solution for server management.

The enhanced design of the virtualization technology also enables any VE on any network to be easily and transparently moved to another server with near-zero downtime, enabling IT departments to more fully utilize existing servers and minimize or eliminate planned downtime.

The Complete Server Virtualization Solution

Parallels Virtuozzo Containers was designed to enable IT organizations to maximize server resources through a unique virtualization technology and manage the IT infrastructure easily with an advanced set of administration tools. In order to meet the virtualization and automation requirements of today's IT departments, a virtualization solution must provide the following components:

INTELLIGENT PARTITIONING

Intelligent partitioning divides a server into as many as hundreds of VEs with full server functionality. A VE has its own processes, users, files, root or administrator access, full networking, system libraries and everything it needs to appear to be a standalone server. While Virtuozzo uses only a single instance of an operating system, each VE has links back to the standard operating system. These links can be represented in a capability in the software called templating, which creates a basic blueprint of an OS or application. The links make a VE footprint very small, and associated operations very fast. Because of the low overhead and single operating system required, potentially 100's of VEs may be loaded on a single physical server. The partitions are also flexible, as they may be changed in real-time without affecting the VE or its workload.

COMPLETE ISOLATION

Complete isolation ensures that the VEs are secure and have full functional, fault and performance isolation.

Isolation is achieved through multiple layers of security designed to ensure that each VE is secure, isolated and unaffected from other VEs on the same physical server. Virtuozzo's patent-pending technology, Kernel Service Abstraction Layer (KSAL), mediates activity to the kernel and prevents any single VE from taking the entire server down. Virtuozzo File System (VZFS) is a proprietary file system that, ensures that a user cannot access any other VE or part of the server. Additionally, Virtuozzo has copy-on-write technology, which makes a local copy of anything unique in the VE. Lastly, the software also ensures performance isolation with a complete set of resource management capabilities and controls.

DYNAMIC RESOURCE ALLOCATION

Dynamic resource allocation of key system resources such as CPU, memory, network, disk and I/O enables near real-time changes without affecting a VE or its workload. Virtuozzo's sophisticated resource management controls CPU, memory, network I/O, disk space and other system resources. Resources can be controlled with simple minimums (guarantees), maximums (limits), shares or more sophisticated allocations that allow for specified overages. Virtuozzo assigns, monitors, accounts and controls resource levels in real-time and alerts administrators according to pre-defined criteria. Should adjustments in resources be required, simply change the resources in real-time without affecting a VE or application performance. A unique feature of Virtuozzo resource management is the ability to "burst" or allow VEs to use resources that are currently unused on the server. In this respect, a VE can get at least its guaranteed resources on the server.

LIVE MIGRATION

Live Migration and other business continuity capabilities ensure data is available and recoverable. Virtuozzo virtualization provides separation of the VEs from the physical server hardware; any VE and application can reside on any x86 hardware. Migrating a VE to another piece of hardware is simple as long as Virtuozzo is the underlying infrastructure. Virtuozzo moves the VEs and any corresponding application(s) between any networked servers transparently with almost no downtime. The abstraction from the hardware also makes the underlying hardware irrelevant, no application testing on different hardware is required.

As VEs begin to multiply, the dependence on a single point of failure - the underlying server -increases. With Virtuozzo, the underlying OS is always running; therefore, VEs are always accessible and recoverable through VE owner self-management Virtuozzo Power Panel. Virtualization also allows the migration of VEs immediately upon threat of a hardware problem. While hardware accounts for only 10% of all downtime³, this can be further minimized by justifying and purchasing a better quality of server with the higher utilization rate provided by Virtuozzo.

MASS MANAGEMENT

Virtuozzo has a full suite of tools/ templates for automated, multi-VE and multi-server administration.

Virtuozzo complete roles and permissions-based management toolset has three options for the Virtuozzo administrator: command line, Virtuozzo Management Center (VZMC) and Virtuozzo Control Center (VZCC). The VZMC manages both Linux and Windows servers from a single GUI interface providing wizards and an intuitive interface for easy management and monitoring. The VZCC is a web-based interface for anywhere access without loading software. Capabilities of these tools include conversion, provisioning, monitoring, backing up, recovering, analyzing, troubleshooting, repairing and migrating VEs. The command line interface is used heavily by large Virtuozzo shops to automate tasks. Using the interfaces, the server administrator can efficiently manage, script and monitor 100s of VEs easily.

VE owners/administrators can use the browser-based Virtuozzo Power Panel (VZPP) to start/stop, backup/restore, repair and remotely reinstall the VE without the Virtuozzo administrator's support. IT administrators at any skill level can provision multiple VEs at once, move VEs to other physical servers and start/stop and reboot as many VEs as required. Even better, using the VZPP, VE owners can perform these same functions without ever requesting the help of IT resources.

The VZMC helps administrators manage multiple servers efficiently. Applications can be deployed and updated on many servers simultaneously. OS updates and patches can also be managed centrally. The VZMC provides a centralized look at server information and software versions across all server resources, facilitating versioning and patch management. Virtuozzo replaces the most popular server resource management tool- the spreadsheet. Virtuozzo has a well documented API and accessible to other infrastructure management tools, which see and manage Virtuozzo VEs because VEs operate as standalone servers.

³ Gartner October 2000

How is Parallels Virtuozzo Containers Different?

Hardware virtualization is well-known in the server virtualization market. OS-level virtualization expands the opportunities for server virtualization and is a complementary or replacement technology in many organizations.

How does Parallels Virtuozzo Containers compare to hardware virtualization technologies such as VMWare, Xen, Virtual Server and Parallels?

Hardware virtualization solutions all have the same premise: resource consolidation. The end goal of the consolidation is to be able to place many different and separate operating systems on the same server. Each of these technologies, whether it is a hypervisor or an OS, has an underlying piece of software that is the basis of the relationship between the hardware and the residing software. Above the hypervisor or virtual machine monitor level is the virtualized hardware level. In order to assign resources, this type of technology must virtualize every single piece of hardware, in essence recreating the work of the hardware vendors and the OS software vendors that already support the hardware. Finally, inside the guest virtual machines, there is a complete copy of an additional operating system and the residing workloads or applications.

Hardware virtualization has its benefits. At least on the leading market solution, just about any operating system is supported. It is possible to load Windows next Linux next to Solaris, and it still supports older OSs. Despite the benefits, the flexibility of this type of solution comes with a lot of overhead and inefficiencies.

Many organizations find that after deploying a hardware virtualization solution, they loaded an entire server with virtual machines all running the same operating system. This happens for a variety of reasons, mainly because most organizations have enough servers to consolidate like operating systems and the system administrator supporting the server has the skill set of that OS. Suddenly, that goal of loading any operating system on a single server doesn't warrant the associated overhead of this type of flexibility.

Further limitations inherent to hardware virtualization lead to additional overheads and inefficiencies, as addressed in the following points:

- A single application has two OSs to traverse; the guest-level OS and the hypervisor or host OS. More processing equates to slower responses and more overhead.
- Each OS takes space in memory, and memory is always the most constrained resource on a server. While some vendors have taken steps to maximize memory usage, the inefficiencies still exist.
- The duplicate OSs consume hard drive space and must be licensed and managed separately. Because virtual machines don't need hardware justification, many IT organizations are now starting to see OS sprawl instead of server sprawl.
- Hardware support and interoperability for all of the hardware on the market is difficult to emulate well, so it is often a source of slower response times and higher processing overhead.

As an OS-level virtualization solution, Virtuozzo uses a single standard Operating System on a server. Not only does this model eliminate the inefficiencies of hardware virtualization, but there are advantages inherent to this approach:

- Performance
 - Parallels Virtuozzo Containers uses all of the real hardware and OS software. There are no rewritten drivers or OS technology. There is also very little or no delay in supporting new technologies and it performs at near-native server levels.
 - There is no duplicate OS, so it again performs very efficiently.
 - Because of the near native server performance, Parallels Virtuozzo Containers is used to virtualize high I/O applications such as data bases and email servers.
- Scalability
 - The flexible design allows any VE to seamlessly scale to the resources of the entire server.

- The lean architecture ensures that nearly all of the system resources are available for use by VEs.
- Parallels Virtuozzo Containers fully supports native 32 bit and 64 bit systems, 16 CPUs SMP support and 64 GB RAM on the host as well in as in VEs.
- Density
 - Because of the single OS, memory and other technology efficiencies, Parallels Virtuozzo Containers has the highest density of any virtualization solution. It is possible to deploy up to 100's of VEs on a regular two-way x86 server. This high density equates to the fastest ROI of any virtualization solution.
- Manageability
 - There is only a single OS to manage, maintain and license. Even if OS vendors forget they are managing the hardware below rather than the applications above, and change their licensing models, the management savings will make the extra licensing costs worthwhile.
 - Application management may also be centralized and managed as a single instance.
 - One update can patch multiple VEs on multiple servers.
 - Many operations are automated and scriptable, making management of many servers and VEs simple.
 - OSs and other software can be created in a VE as simple links making the footprint small operations extremely fast. For example, a VE is created in seconds, and clones or migrations take less time because only the active or different data is moved rather than an entire operating system.
 - The comprehensive resource management ensures that all required resources are provided as configured for each VE. It also allows VEs to use available resources and make changes in near real-time.

Parallels Virtuozzo Containers Performance

Parallels Virtuozzo Containers's unique form of virtualization not only eliminates complications that arise with hardware virtualization, but its impact on server performance is so minimal that the server operates at near-native levels. The performance factors include:

- **System Calls** Parallels Virtuozzo Containers has a common kernel, no additional systems calls are created between excess layers.
- **Filesystem** Common data is stored only once and is available to multiple environments, significantly improving the efficiency of both the disk space and system memory. The disk space for caching is also used efficiently, caching payload rather than duplicated data.
- **Filesystem Performance** Parallels Virtuozzo Containers's filesystem and efficient system call design create very small CPU usage overhead.
- **Filesystem buffers/cache in memory** Parallels Virtuozzo Containers places a single instance of data and cache across the server, maximizing memory efficiency.
- **Memory Management** Parallels Virtuozzo Containers handles memory allocation requests dynamically. Virtuozzo environments can use a lot of memory during load peaks, and users will fully benefit from running Virtuozzo on high-end servers. Applications in Virtuozzo environments can show very high peak performance. Also, when environment resource demand grows gradually in time, Virtuozzo's dynamic resource management allows increasing the resource allocation with no downtime.

Addressing the Challenges of IT and Conclusion

IT organizations are tasked with implementing powerful and reliable services at low a cost. Due to exploding infrastructure growth and complexity, this task has been increasingly difficult until now. Virtuozzo addresses each of the following critical challenges faced by IT organizations and empowers them to grow their business while cutting overhead and inefficiencies.

- Improve the flexibility of computing resource assignment- Virtuozzo abstracts software and workloads from the hardware. It makes all resource assignments changeable in real-time and movement between hardware seamless.
- Decrease complexity to improve manageability of systems- Templating, single OSs and applications and other technology advances make management of Virtuozzo VEs and servers much simpler than regular servers or virtual machines.
- Automate routine tasks- The automation tools and scripts possible with Virtuozzo make managing VEs across servers simple and potentially a single action.
- Reduce overall management costs through efficiency- Many manageability advances enable much better and efficient server management.
- Increase the return from their infrastructure investment by better utilizing resources- By consolidating servers to higher utilization, every other IT resource is better used, from networking infrastructure to administrative personnel.

Virtualization and automation technology is a major advance in enterprise IT infrastructure management. By leveraging the capability of abstracting software and data from the hardware and treating the hardware as a pool of resources, virtualization and automation technology is a key enabler for IT organizations.

Parallels Virtuozzo Containers is the only virtualization and automation solution that was developed specifically for virtualizing production servers. Virtuozzo allows enterprises to implement server consolidation, increase manageability and service levels, and dramatically lower total cost of ownership (TCO).