

Parallels® Cloud Server

White Paper

An Introduction to Operating System Virtualization and Parallels Cloud Server

Table of Contents

| | |
|--|---|
| Introduction | 3 |
| Hardware Virtualization | 3 |
| Operating System Virtualization | 4 |
| Parallels Cloud Server: Optimized for the Cloud | 5 |
| Integrated with Parallels Cloud Storage | 5 |
| Key Virtualization and Management Features of Parallels Cloud Server | 6 |
| Intelligent Partitioning | 6 |
| Complete Isolation | 6 |
| Resource Limit Controls | 6 |
| CPU Resource Scheduling | 7 |
| Memory | 7 |
| Disk I/O | 7 |
| Disk Space | 7 |
| Dynamic Resource Allocation | 7 |
| Live Migration | 8 |
| Rebootless Updates | 8 |
| Mass Management | 8 |
| Key Performance-Enhancing Features of Parallels Cloud Server | 8 |
| System Calls | 8 |
| Memory and IOPS Deduplication | 8 |
| File System Performance | 8 |
| Memory Management | 9 |
| How to Get Parallels Cloud Server | 9 |
| Conclusion | 9 |
| Contact Us | 9 |

Introduction

Two types of virtualization are in use today: hardware virtualization (generally known as hypervisors) and operating system virtualization (commonly referred to as containers). Both types of virtualization are broadly deployed, and each has its inherent benefits and drawbacks. This paper begins with an overview of both technologies, describing the benefits and drawbacks of each for cloud service providers. It then provides an overview of Parallels Cloud Server, an industry-leading operating system virtualization technology that has been optimized specifically for cloud service providers.

Hardware Virtualization

A hypervisor solution virtualizes at the hardware level, creating a duplicate of all system resources – including the operating system, CPU, memory, and configuration files – in each virtual machine. This approach results in up to 20% more overhead on the server. As a result, it reduces the number of virtual servers that can be supported by a single physical server, as well as negatively affects overall system performance. Figure 1 shows what a hypervisor-based system looks like.




Figure 1. Components of a hypervisor-based virtualization solution.

The overhead and performance issues associated with hypervisor-based virtual machines become particularly acute in cloud deployments because such deployments usually consist of many very small virtual machines, with each machine generally requiring just 100 MHz of a single CPU, 1 GB of RAM, and 5 GB of disk space. Because a fixed amount of resources is needed for each virtual machine, regardless of its size, the relative overhead of each virtual machine gets bigger as the virtual machine gets smaller.

Hypervisors further reduce efficiency in several ways:

- A single application has two operating systems to traverse: the guest-level operating system (used inside the virtual server) and the host operating system. More processing equates to slower responses and more overhead.
- Each operating system 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, inefficiencies still exist.
- Duplicate operating systems consume hard drive space and must be licensed and managed separately, increasing maintenance costs.

- It's difficult to accurately emulate hardware support and interoperability for all of the hardware on the market, often resulting in slower response times and higher processing overhead.
- Hypervisors may require guest and host operating system drivers or guest tools to enable important cloud functionality, such as CPU hotplugs, memory ballooning, and disk resizing.

Despite the inefficiency inherent in hardware virtualization, this approach does have its benefits for certain types of workloads. For example, because hypervisors can support just about any operating system, they enable different operating systems to be deployed on the same hardware node, making it possible to load Windows next to Linux next to Solaris, as well as to support older operating systems. However, this flexibility generally isn't applicable to providers of hosted cloud services, who typically base large cloud deployments on a single operating system – and whose profitability depends on maximizing efficiency.

Operating System Virtualization

With **operating system virtualization**, a single operating system kernel can support multiple isolated virtualized instances, instead of just one. Such instances, called containers, look and feel like a real server from the point of view of their owners (the service provider's customers). Figure 2 shows what operating system virtualization looks like.

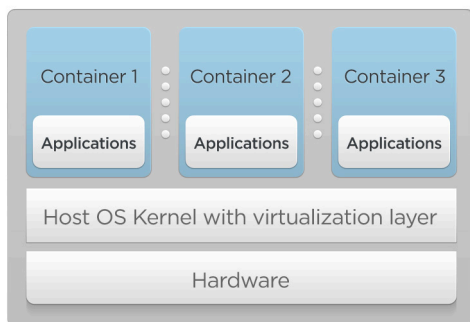


Figure 2. Components of an operating system virtualization solution.

This approach provides a common virtualization layer that allocates a single set of system resources across all containers. The result is a leaner, more efficient virtualization layer, with an overhead of less than 2%. For service providers, this approach translates into higher ratios of virtual servers to physical servers, near-native server performance, and unique advantages for managing the virtual environment. These advantages include dynamic reallocation of resources, which allows service providers to accomplish more with fewer resources, and the ability to create links back to the server host operating system, resulting in lightning-fast management operations.

Using a single standard operating system on a server also offers additional advantages, besides eliminating the inefficiencies of hardware virtualization. For example, it provides:

- **Performance enhancements.** Because there is no need to duplicate the operating system, there is also no need to duplicate drivers, further enhancing performance. And the near-native server performance available with operating system virtualization makes containers particularly efficient at virtualizing high I/O applications, such as databases and email servers.
- **Scalability and elasticity.** The flexible design of solutions based on operating system virtualization allows any container to seamlessly scale to the resources of the entire server; and the lean architecture ensures that nearly all of the system resources will be available for use by the containers.

- **Density.** Because containers use only one operating system and one set of memory, as well as offering other technology efficiencies, this approach has the highest density of any virtualization solution. It is possible to deploy hundreds of containers on a standard dual-CPU x86 server. This high density gives service providers the fastest return on investment (ROI) of any virtualization solution.
- **Manageability.** Operating system virtualization simplifies server management in several ways:
 - With operating system virtualization, there is only one host operating system to manage, maintain, and license – and a single update to this host operating system can patch all the containers the on the physical servers.
 - Application management can be centralized and managed as a single instance.
 - Many operations can be scripted and automated, making it easy to manage many servers and containers.
 - Operating system templates and other software can be created in a container as simple links, making the footprint small and operations extremely fast.
 - A new container can be created in seconds.
 - Clones and migrations take less time, because you’re only moving the container-specific data, not the entire operating system.

Parallels Cloud Server: Optimized for the Cloud

Parallels Cloud Server is an operating system virtualization solution that creates isolated partitions on a single physical server, using a single operating system instance – thereby maximizing your investment in server hardware. Optimized for the cloud, it enables near-instant provisioning and on-the-fly modification of hosting and cloud server plans, while also delivering maximum density, cost efficiency, and application performance. In addition, its complete set of management tools makes it easy to maintain, monitor, and manage virtualized server resources.

INTEGRATED WITH PARALLELS CLOUD STORAGE

Parallels Cloud Server is delivered as a component of Parallels Cloud Server 6, which includes Parallels Cloud Storage. (See section on How to Get Parallels Cloud Server, below, for more information on Parallels Cloud Server 6.) Parallels Cloud Storage is a distributed, shared storage solution that decouples computation from storage, enabling virtual machines and containers to be instantly migrated to an operational physical server if the original physical server becomes unavailable. And with the storage cluster distributed across multiple physical machines, if one data source becomes unavailable, virtual machines can rapidly get their data from another physical machine.

Key Virtualization and Management Features of Parallels Cloud Server

The following sections describe in more detail how Parallels Cloud Server meets the specific virtualization and automation requirements of today's cloud service providers.

INTELLIGENT PARTITIONING

Intelligent partitioning divides a server into as many as hundreds of containers, each with full server functionality. That is, each container has its own processes, users, files, root (administrator) access, full networking, and system libraries – in short, everything it needs to appear to its owner to be a standalone server. Because the partitions between containers are flexible, you can change them in real time (for instance, to increase or decrease memory) without affecting either the virtual environment or its workload.

Although Parallels Cloud Server uses only a single instance of an operating system, each container has links back to that operating system. These links can be represented as templates, which serve as blueprints of the operating system or application. The links make the container's footprint very small, which in turn makes its operations very fast.

COMPLETE ISOLATION

Complete isolation ensures that the containers are secure and have full functional, fault, and performance isolation. This isolation is achieved through multiple layers of security, which ensure that each container is unaffected by other containers on the same physical server. Parallels Cloud Server' patent-pending technology, Kernel Service Abstraction Layer (KSAL), mediates activity to the kernel, preventing the actions of any single container from taking the entire server down.

Parallels Cloud Server also improves isolation with a new storage technique, called ploop (Parallels loopback device). Ploop stores all of a container's data in a single image, with its own private file system journal. These individual file systems enhance isolation by eliminating the problems that can occur when multiple container operations request file journal access at the same time.

In addition to improving isolation, storing all of a container's data in a single image provides several additional benefits:

- It enables you to take consistent snapshots of the container file system, which you can use for incremental backups or provide to the user for other applications.
- It allows you to make backups on top of snapshots using conventional cp or tar, rather than requiring special backup tools like Acronis True Image.
- It significantly reduces the time required for container backups and migration, since sequential reading of an image is an order of magnitude faster than reading separate small files on rotational drives.
- It enables deduplication (caching identical files that exist in multiple containers), which saves memory and increases input/output operations per second (IOPS).

RESOURCE LIMIT CONTROLS

Parallels Cloud Server comes with a variety of controllers – for CPU usage, memory, disk I/O, disk space, network bandwidth, and more. These controllers ensure that services in one container can't affect other containers on same the physical system.

In most cases, services running inside containers will consume resources below their assigned limit. However, should they approach or try to go over the limit, Parallels Cloud Server will prevent this – for instance, by pushing pages into VSwap in the case of memory limits, or by throttling disk transactions in the case of I/O bandwidth limits.

CPU RESOURCE SCHEDULING

CPU-related parameters include `cpuunits` and `cpulimits`. Both of these parameters affect the CPU resources available to a container, but the two are not directly related to each other.

The `cpuunits` parameter defines the weighting of CPU resources between containers. If you don't define this parameter and no other containers are competing for the CPU, an individual container will be able to consume up to 100% of the CPU resources. If you do define this parameter and multiple containers are requesting CPU resources, the container with the higher `cpuunits` value will have a better chance of getting more CPU resources.

The `cpulimit` parameter defines a hard limit that controls the total percentage of CPU resources an individual container can use. Once this parameter is defined for a container, the container will not be able to exceed it, even if no other containers are competing for the CPU resources.

MEMORY

Parallels Cloud Server uses the RHEL6-based OpenVZ kernel, which includes a memory management model called **VSwap**. VSwap uses two primary parameters – **physpages** and **swappages** – to control the amount of allocated memory and swap space. The sum of the `physpages` and `swappages` limits is the maximum amount of allocated memory that a container can use. When the `physpages` limit is reached, memory pages belonging to the container are pushed out to virtual swap space. Once the total amount of allocated memory is consumed, the container's performance will start to degrade.

DISK I/O

Two disk I/O parameters, I/O limits and IOPS limits, are very granular and powerful controls that let you prevent performance degradation in situations when high disk I/O activities in one container could slow down the performance of other containers.

DISK SPACE

To avoid performance degradation in high-density environments, it's important to make sure that disk quotas aren't exceeded. Parallels Cloud Server lets you assign two types of disk quotas: one limiting disk space for the entire container, and the other limiting the space available to individual users inside the container. The disk space quota for containers is managed using the Parallels Container Command Line Interface (CLI) or Parallels Virtual Automation. Disk space quotas for individual users are managed using quota utilities in the same way as in a non-virtualized Linux kernel, enabling you to continue to use your existing Linux management tools.

DYNAMIC RESOURCE ALLOCATION

Dynamic allocation of key system resources – such as CPU, memory, network, disk, and I/O – lets you make near-realtime changes without affecting a container or its workload. You can control these resources either by specifying simple minimums (guarantees), maximums (limits), or shares (percentages), or by using more sophisticated allocations that allow for specified overages.

In either case, Parallels Cloud Server assigns, monitors, and controls the resource levels in real time and alerts administrators according to the criteria you assign. Should adjustments in resources be required, you can simply change the resources in real time, without affecting the performance of a container or application.

A unique feature of Parallels Cloud Server resource management is the ability to “burst” – that is, to allow containers to use resources that are currently unused on the server, even if that means exceeding their preset limits.

LIVE MIGRATION

Because Parallels Cloud Server virtualization separates containers from the physical server, any container or application can reside on any x86 hardware – and with the software’s live migration capabilities, moving a container to another piece of hardware is simple. Both the container and any applications on it can be moved between any networked servers with complete transparency – and with no downtime. In fact, because the abstraction from the hardware makes the underlying hardware irrelevant, you don’t even need to do any application testing on the new hardware. Live migration also enables you to migrate containers immediately if a hardware problem degrades the performance of the host server.

REBOOTLESS UPDATES

Parallels Cloud Server supports rebootless updates, suspending and resuming all operational virtual machines and containers rather than shutting them down and restarting them. This approach virtually eliminates any service outage for end users, since in most cases the suspension is so brief that the end user doesn’t even notice that the application was suspended.

MASS MANAGEMENT

Parallels Cloud Server comes with both a command line interface and a GUI-based management tool called Parallels Virtual Automation allowing a single server administrator to efficiently manage, script, and monitor hundreds of containers.

Parallels Cloud Server also comes with a browser-based Parallels Power Panel, enabling container owners and administrators to start, stop, back up, restore, repair, and remotely reinstall the container without the support of the Parallels Cloud Server administrator (i.e., without support from the hosting provider’s staff). End users at any skill level can provision multiple containers at once; move containers to other physical servers; and start, stop, and reboot containers as required.

Parallels Cloud Server also has a well-documented API and is accessible to other infrastructure management tools, allowing them to monitor and manage Parallels Cloud Server.

Key Performance-Enhancing Features of Parallels Cloud Server

Not only Parallels Cloud Server eliminate the complications that can arise with hardware virtualization, but it also has such a minimal impact on server performance that the server is able to operate at near-native levels. The following sections describe these performance-enhancing features in more detail.

SYSTEM CALLS

With Parallels Cloud Server, because all containers share a common kernel, no additional systems calls are needed between multiple operating system layers.

MEMORY AND IOPS DEDUPLICATION

Parallels Cloud Server locates and efficiently caches identical files that exist in multiple containers, thereby reducing I/O and memory bottlenecks and increasing the number of containers that can be supported on a single server.

FILE SYSTEM PERFORMANCE

The efficient file system and system call design of Parallels Cloud Server results in very small overhead in terms of CPU usage.

MEMORY MANAGEMENT

Parallels Cloud Server handles memory allocation requests dynamically, enabling performance to remain high even during peak loads. And in cases where resource demand grows slowly over time, its dynamic resource management capability allows it to gradually allocate increasing additional resources in real time – with no downtime.

How to Get Parallels Cloud Server

Parallels Cloud Server is delivered as part of Parallels Cloud Server 6, a complete virtualization and cloud storage solution for service providers that includes Parallels Cloud Storage and Parallels Hypervisor, as well as Parallels Cloud Server. Parallels Cloud Server 6 provides you with a wide range of information-as-a-service (IaaS) capabilities – from simple virtual private servers and cloud storage to complete virtual data centers and usage-based cloud computing services. It's designed to help you increase control, availability, and flexibility of your hosted infrastructure services, while also reducing costs.

Conclusion

While hypervisors, based on hardware virtualization, and containers, based on operating system virtualization, both have their strengths, container-based approaches offer many advantages that specifically address the needs of cloud service providers. Parallels Cloud Server, in particular, has been optimized to meet hosters' needs. Its density advantages – allowing a single physical server to support hundreds of containers, each with full server functionality – means that your hardware investments will go much further. Its support for near-instant provisioning and on-the-fly modification of hosting and cloud server plans provides the flexibility that you need to address changing customer requirements. And its complete set of management tools, ranging from templates to automation tools and scripts, take the complexity out of maintaining, monitoring, and managing virtualized server resources.

For more information about Parallels Cloud Server, please visit www.parallels.com/products/pcs

Contact Us

For more information about Parallels hosting and cloud solutions, please contact:

Parallels, Inc.
500 SW 39th St. Suite 200
Renton, WA 98057
+1 425 282 6448
www.parallels.com

Copyright© 2012 Parallels IP Holdings GmbH. All rights reserved. Parallels and the Parallels logo are registered trademarks of Parallels IP Holdings GmbH. Other product and company names are the trademarks or registered trademarks of their respective owners.