

- **Reliable and robust**, since it removes the need to download old templates (which may no longer be available) in order to restore older versions of containers.
- **Configurable**, with the ability to apply parameters to specify factors such as cache size, directories to be used for caching, and expiration times for cached files.
- **Always up to date**, because a daemon that tracks file access ensures that the cache always has the most frequently accessed data.
- **Easy to keep at top performance**, because it comes with a configuration tool that lets you modify various options to enhance performance.

How Parallels Container Templates Improve Memory Efficiency

A template (or package set) in Parallels Containers is a set of original application files repackaged for rapid container creation. Using templates within containers can greatly improve the efficiency of server memory utilization. Parallels Containers provides tools for creating templates, installing them, upgrading them, adding them to a container, and removing them from a container.

Using templates lets you:

- Share RAM among similar applications running in different containers, saving hundreds of megabytes of memory.
- Share the files that make up a template among different containers, saving gigabytes of disk space.
- Reduce IOPS by allowing containers which access the same file from a template to read it only one time from the disk and then allow all other containers to access it from memory

Parallels Containers has two types of templates:

- **OS templates**, which consist of an operating system plus the standard set of installation applications. Parallels Containers uses OS templates to create new containers with a preinstalled operating system.
- **Application templates**, which consist of a set of repackaged software packages – optionally accompanied by configuration scripts. Parallels Containers uses application templates to add extra software to an existing container. For example, you can create a container on the basis of the CentOS 5 OS template and add MySQL to it using the MySQL application template.

For detailed information on PCS templates, see the *“Parallels Cloud Server Templates Management Guide.”*

Summary of Best Practices for Maximizing Density

Wrapping up what we’ve covered so far, here are the best practices you’ll want to follow in order to maximize resource utilization and scale your Parallels Containers servers most effectively:

- **Determine the appropriate overcommit level for each type of VPS offering.** Because larger, more expensive VPS offerings are typically used at or near their resource limits much more frequently than smaller offerings, you should segment your offers and related service plans based on CPU and memory allocations. That way, you can provision larger VPSs on servers with lower overcommit levels, and smaller VPS on servers with higher overcommit levels.
- **Determine what level of overcommit you’re comfortable with.** The higher the degree of overcommitment, the more intensively you’ll need to monitor the performance of your physical hardware and containers.
- **Make sure your implementation plans include all of the following:**
 - Monitoring usage of CPU, RAM, and disk I/O on your host hardware.
 - Ensuring that the CPU cores and limits assigned to each container are consistent with the actual CPU configuration installed in the server hardware.
 - Managing your disk I/O limits as described in Parallels Disk I/O Limits Best Practices white paper.
 - Using VSwap to set the amount of physical memory and virtual swap space allocated per container (see VSwap Best Practices white paper for guidelines).
 - Provisioning VPSs via templates to ensure the most efficient use of memory.
- **Monitor the server as the number of deployed VPSs increases to ensure that you’re maintaining quality of service.** As you approach the server’s resource limits, you have several choices:
 - Stop provisioning additional VPSs on the server.
 - Upgrade the server with additional resources.
 - Move containers with a high level of resource usage (e.g., high disk I/O) to another hardware node with lower overcommit levels.
 - Contact customers whose VPS are approaching their allocated limits and upsell them, adding more RAM or CPU to their subscriptions.

Real-World Examples

While maintaining the optimal number of virtual containers on each physical node is critical to being able to both deliver a high quality of service and leverage your infrastructure investment efficiently, there is no one answer as to what is the optimal number. Our partners deploy a wide range of density levels, depending on the type of service they're offering. Some focus on reselling large, powerful VPSs to web hosters who want to maximize the number of shared accounts on each container, and therefore maintain lower densities per node. Others use smaller VPSs for testing applications, offering hosted desktops, or providing backup services, and therefore are able to support much higher densities.

Typical Server Configurations

We recently conducted a survey of Parallels partners to determine the range of hardware configurations they are currently deploying. Although the results differed depending on the target markets each service provider was addressing, we found many service providers achieving success with the configurations described in Table 2.

Table 2: PCS Best Practice Configurations

CPU	GB RAM	GB HDD	No. of Containers	VPSoffering and Parameters
One dual-quad AMD Operon	16	4x500GB RAID 10	7-30	L: 1GB RAM, 20 GB HDD XL:2GB RAM, 40 GB HDD XXL: 4 GB RAM, 80GB HDD
One dual-quad CPU AMD Operon	6-32	4x500GB RAID 10	10-30	L: 1GB RAM, 20 GB HDD XL: 2GB RAM, 40 GB HDD XXL: 4GB RAM, 80GB HDD
i7 (2x8) Xeon X5550 or i7 (2x8 or 2x4) Xeon Q9550	12-32	4x500GB RAID 10	20-60	HVD Standard:1 GB RAM, 2GB HDD HVD Professional: 2GB RAM, 10GB HDD
i7 (2x8 of 2x6) Xeon X5550	128	4x500GB RAID 10	80-110	L: 1GB RAM, 50GB HDD XL: 2 GB RAM, 100GB HDD XXL: 2 GB RAM, 100GB HDD
i7 (2x8) Xeon X5550	48	4x500GB RAID 10	40-50	2GB RAM, 20GB HDD
i7 (2x8) 9200	72	4x500GB RAID 10	70-80	L: 1GB RAM, 30GB HDD XL: 2GB RAM, 50GB HDD XXL: 4GB RAM, 100GB HDD

Optimizing Your Server Configuration and Management

Providing customers with reliable service – which means an absolute minimum of service disruptions – is one of the most important factors in service providers' success. At the same time, to enhance profitability, you have to minimize your maintenance costs. Achieving both goals can be difficult – but not when you use Parallels Cloud Server.

Best Practices for Maximizing Quality of Service and Profitability

The following best practices will help you ensure a high quality of service while also enhancing your profitability:

- **Optimize your environment.** Key steps you can take to do so include:
 - Using compatible OS environments on your host machine – for example, deploying Parallels Cloud Server – to ensure optimum performance and reliability.
 - Using the same hardware in all your nodes. Including disk controllers and other peripherals, will reduce the risk of driver-related problems.
 - Avoiding the use of third-party drivers on any of your nodes. In our experience, many third-party drivers have not been validated across all the variations of hardware and software that can be deployed on a machine, and therefore can negatively affect system reliability.
 - Keeping all configurations synchronized across all your Parallels Containers servers, to simplify troubleshooting and maintenance.
- **Streamline deployment.** Parallels Cloud Server (PCS) streamlines deployment by enabling you to use a single disk to deploy the complete software stack, including the operating system and containers, on a bare-metal server. This approach not only speeds installation, but also improves reliability, by ensuring that all machines are configured consistently. PCS also lets you provide VMs for those customers that want to maintain their own operating system on the server. You can also streamline deployment by automating network-based installation – a step made possible by the compatibility between PCS and PXE boot.
- **Reduce the impact of disk failures.** Hard drives are typically the most unreliable server component. Using hot-plug-compatible disk bays and battery backup on RAID controllers will reduce the impact of hard drive failures. Also, avoid the use of fans on

controllers, as they decrease reliability – and when they fail, they can generate data failures.

- **Optimize hard drive configuration.** The most commonly deployed hard drive configuration today is a 4x500GB RAID 10 array, which provides you with about 1TB of usable storage on the server. For most of our partners, this is a cost-effective configuration for initial deployment and allows for significant VPS storage growth. You should also configure the swap space so that when the resource load on the machine begins to approach the limit of allowable memory, the swap space will be big enough to handle the extra load. Table 3 shows our recommendations for configuring swap space.

Table 3: Recommended Swap Space Configurations

Memory	Recommended Swap space
4GB of RAM or less	a minimum of 2GB of swap space
4GB to 16 GB of RAM	a minimum of 4GB of swap space
16GB to 64GB of RAM	a minimum of 8GB of swap space
64GB to 256GB of RAM	a minimum of 16GB of swap space
256GB to 512GB of RAM	a minimum of 32GB of swap space

Real-Word Examples of VPS Configurations

Many different VPS configurations are deployed in the market, and no one configuration, price, infrastructure, or service level is appropriate for every application. Table 4 provides a

few real-world examples from our service provider partners, illustrating the wide range of options, hardware sizes, and densities that can be effectively deployed with Parallels Containers. The small VPS configuration achieves very high densities with a high level of service by carefully managing and monitoring resource usage on the server. The large VPS configuration achieves lower VPS densities but still requires careful management of overcommitted resources to ensure a high quality of service. The third configuration is unique in that it allows customers to deploy as many VPS instances as they like without a per-container fee, charging instead based on total RAM usage across all containers.

Table 4: Three Real-World Offerings and Their Underlying VPS Configurations

Component	Small VPS	Large VPS	VPS Based on Innovative Pricing Paradigm
VPS Offering & Features	512 MB RAM 20 GB storage 100 GB network transfer per month	4 GB RAM 650 GB storage 1 TB network transfer per month	50 GB storage 1TB network transfer per month No per-VPS fee Free load balancing and Control panel
VPS Pricing	\$30 per VPS	\$90 per VPS	\$35 for 256 MB RAM \$70 for 512 MB RAM \$105 for 1 GB RAM
Processor	Dual-quad core	16-core Intel Xeon	Dual-quad-core
CPU	2.4 GHz	2.13 GHz	2.4 GHz

Server Memory	192 GB RAM	32 GB RAM	64 GB RAM
Server Disk	4 TB	500 GB	2 TB
Server Costs	\$7,500	\$5,000	\$7,000
No. of virtual containers	200+	30	60
Techniques Employed for Driving Higher Density	Overcommitment of resources, carefully monitored and managed	Overcommitment of resources, carefully monitored and managed	Overcommitment of resources, carefully monitored and managed

Conclusion

Leveraging the best practices detailed in this report – both for achieving maximum density per server and for optimizing server configuration and management – will enable your business to increase the revenue per server that your VPS offerings generate. And by diversifying your VPS offerings and aligning them optimally with your hardware configurations, management processes, and capabilities, you will differentiate yourself from your competition, enabling you to grow your business as well as maximize your profitability.

Appendices: Checklist for Optimizing Your VPS Infrastructure

- Have you set CPU limits, disk I/O, memory parameters, and VSwap settings appropriately?
- Have you reviewed the “VSwap Best Practices” whitepaper?
- Have you limited disk throughput per container and aligned it to your network bandwidth?
- Have you reviewed the “Disk I/O Best Practices” white paper?
- Have you set disk space and network I/O appropriately for your environment?
- Are your VPS offerings mapped appropriately to the underlying container configuration?
- Have you documented your level of overcommitment and mapped it to your specific hardware?
- Have you set up the necessary tools and approaches for monitoring usage of key resources?
- Do you continue to monitor the server as you scale your VPS levels?
- Have you reached out to customers approaching their allocated limits and upsold them on additional VPS resources?
- Have you minimized driver-related problems by using the same hardware in all nodes?
- Have you leveraged PSBM to streamline deployment?
- Do you provide VMs for customers that want to maintain their own operating system on the server?
- Do you use PXE Boot to automate network-based installation?
- Have you avoided fan usage in your controllers and implemented hot-plug-compatible disk bays and battery backup on your RAID controllers?
- Have you optimized your hard drive configuration and swap levels in accordance with best practices?
- Have you set the appropriate swap space for your memory levels?
- Have you implemented clustering to improve reliability and let you deliver high-performance services?
- Have you reviewed “PCS Cluster Guide”?
- Do you have multiple VPS offerings, aligned to your underlying technology capabilities?
- Have you introduced innovative pricing paradigms or other unique approaches?