



Virtual Business Servers technical description

Virtual Business Servers have been developed under an extremely advanced computing platform. The product combines, within a single solution, Parallels Cloud Storage technologies for the storage in SANs, Parallels Cloud Server (Virtuozzo for Linux) for available Virtual Servers using container-type (CT) virtualization, and Parallels Cloud Server (Parallels Hypervisor) for Cloud Servers using Bare-metal-type (VM) virtualization, in addition to automation and orchestration provided by Odin Automation for Cloud Infrastructure (OACI), a module installed in a structure that is already automated, also used for Parallels.

The virtual hardware, network and storage elements (also known simply as “resources”) comprising the Virtual Servers offering are the following:

Virtualization Type (vType)

Computational platform virtualization technology compatible with the Intel / AMD x86-64 architecture technology used by a virtual server, which may be:

- **Bare-metal (VM):** “Bare-metal” virtualization, also known as Type 1, which emulates the x86-64 architecture without requiring an operating system between the physical machine (host) and the virtual machine (guest), since the Virtual Server (guest) has direct access to the intermediary physical hardware controlled by a monitor system (hypervisor) and instructions, circuits and resources to assist with hardware virtualization (hardware-assisted), thus creating an instance of a complete and isolated virtual machine. Further details at <http://goo.gl/WZAJE4>
- **Container (CT):** Virtualization at the operating system level, also known as kernel virtualization, creates separated user spaces and processing that share a single operating kernel, which directly controls the physical server (host), thus creating an instance of a complete and isolated virtual machine. Further details at <http://goo.gl/X4UBS5>. The Cloud Server virtualization platform has Docker native support (<http://www.docker.com>), as the client may opt to use it in the automation and rollout of its applications on its Virtual Servers running in container-type (CT) virtualization.



Virtual Processing Units (vCPU)

These consist of virtual central processing units assigned for the use of a Virtual Server. vCPUs may range from 1 vCPU (minimum) to 20 vCPU (maximum) per server unit, with the minimum indivisible unit being 1 vCPU.

Virtual Processing Speed (vClock)

This consists of the virtual processing speed (or virtual clock - "clock rate") assigned for the use of each Virtual Processing Unit (vCPU). vClock can range from 1.0 Ghz (minimum) to 2.5 Ghz (maximum) per vCPU unit, with the minimum indivisible unit being 1.0 Ghz.

Permitted vClock values vary by type of virtualization chosen. For containers (CT), the vClock of each virtual processing unit (vCPU) may be 1.0 Ghz or 2.0 Ghz, with no intermediate values (the interval is 1 Ghz). For Bare-metal (VM)-type virtualization, the vClock of each vCPU may be 1.5 Ghz or 2.5 Ghz, also with no intermediate values (interval of 1 Ghz).

Virtual RAM Memory (vRAM)

This consists of the volume or quantity of virtual primary memory (RAM) assigned for the use of a Cloud server. When container-type (CT) virtualization is used, it may range from 0.5 GB (minimum) to 256 GB (maximum) per Virtual Server unit. When Bare-metal (VM)-type virtualization is used, it may range from 1 GB (minimum) to 96 GB (maximum) per Virtual Server unit. Regardless of the virtualization type used, the minimum indivisible unit (for increase and decrease) of vRAM is 0.5 GB.

Virtual Disks (vDisk)

These consist of logical storage units ("volumes") in blocks assigned for a virtual server. They are allocated to the storage of any file and data, including those belonging to the basic operating system and any programs (software) whatsoever needed for the functioning of the virtual server. All virtual hard disks are predetermined to be redundant, mounted in a SAN server structure using replication and duplication technologies (only in containers - CT), in which all the blocks (or "chunks") of data are stored on a minimum of three different physical disks, which in turn overlap each other on three different physical servers. Thus, special configurations applicable to physical environments such as RAID for hardware or software (at the operating system level) are unnecessary, and even not recommended.



Disk space (vHDD)

This consists of the volume or quantity of virtual secondary memory (HDD) assigned for the use of a virtual hard disk (vDisk). When container-type (CT) virtualization is used, it may range from 20 GB (minimum) to 5 TB (maximum) per unit of Virtual Hard Disk (vDisk). When using Bare-metal-type (VM) virtualization, it may range from 50 GB (minimum) to 5 TB (maximum) per unit of Virtual Hard Disk (vDisk), with the minimum indivisible unit being 10 GB.

External Transfer (vTransfer)

This consists of the total volume, in GB, of files and data transferred via the public Internet network to outside ("outbound transfer") for all virtual servers belonging to a subscription during a specified period. Keep in mind that consumption is related to a subscription, and not to a specific Virtual Server.

Transfer volume may range from 0 (zero) GB (minimum) to unlimited GB (maximum) in each subscription; regardless of the quantity of Virtual Servers you possess, you will have access to an External Transfer rate (vTransfer) valid during the payment period and renewed automatically for identical periods as long as your contract is in effect.

External Bandwidth (vBandwidth)

This consists of a measurement of transmission capacity from the connection to the external network - public Internet - from a virtual server. This rate determines the speed by which packets and data originating from or sent to a virtual server are routed through the network interface specific to it (FrontNet - technical name of this network in our infrastructure). Access by Virtual Servers to the Internet is guaranteed by a high-speed network and availability, connected to the Claro Internet backbone. We have our own, completely dedicated edge network, including independent routers with BGP support, isolated from other solutions created at the same Data Center. However, the client must be aware that several factors external to the solution may influence the final quality and speed of access to the Virtual Servers. Specifically, attention must be paid to the geographic location and choice of access provider (final mile) of both the direct client and its users.

Virtual servers may be configured with a bandwidth ranging from 1 MB (minimum) to 1 GB (maximum). Through the control panel, the client may administer the bandwidth of each of its virtual servers, and change it at any time.

External Bandwidth (vBandwidth) is closely related to External Transfer (vTransfer) resources, in that the greater the bandwidth, the greater the potential transfer volume (vTransfer), for a given period (one month, for example).

This function must be used as a principal control instrument for the external transfer volume (vTransfer) of a virtual server. The greater the vBandwidth, the greater the vTransfer capacity / power.



Backup Tool

The offering includes a backup tool operated via the control panel or API. The client may create manual (*ad hoc*) backup copies or choose a preexisting routine (or policy). Manual (*ad hoc*) backups are always Complete (FULL), while scheduled backup copies may range from Complete (FULL) to Incremental (INCREMENTAL), depending on the number of cycles. In a backup copy routine, the first task (copy) is always FULL, and all subsequent tasks are INCREMENTAL. Thus, the number of cycles determines the number of times that INCREMENTAL-type backup copies will be executed before a FULL-type task is performed, always at the start of the cycle.

We can take as an example a “Daily” backup routine. This routine always performs a backup of a Virtual Server every day at 12:00 noon; it has a number of cycles equal to 6 and a retention equal to 30 copies. Thus, on the first day of validity of the routine, a FULL-type backup copy will be made at 12:00 noon, on a daily basis (also always at 12:00), for 6 INCREMENTAL-type backup copies. Since the number of cycles is equal to 6, this means that a new FULL-type backup copy will be executed on the 7th day, to be followed by another 6 INCREMENTAL backup copies, and so on. Since the routine provides for the retention of 30 copies, the client will have the last 30 backups made, whether FULL or INCREMENTAL, at its disposal--several FULL and INCREMENTAL backup copies, regardless of the number of cycles--for executing a restore. When the number of copies reaches the predetermined value (30 in the case of the daily routine), the oldest copies are automatically and sequentially deleted (from oldest to newest).

A virtual server can be associated with only one backup copy routine / policy. You cannot associate more than one backup copy routine / policy on a virtual server.

Storage for Backups (vBackup)

This consists of the volume, measured in GB, of external storage space for making backup copies of the Virtual Servers pertaining to a single subscription. This storage volume can only be accessed and used through the integrated backup copy tool.

Image creation tool (ImageTool)

This offers a predetermined image creation tool for Virtual Servers operated via the control panel or API. The image functionality allows clients to configure and customize their Virtual Servers with everything that may be stored on the initial hard drive, such as configurations and changes on the operating system or in any software installed under it. This allows these modifications and/or customizations to be reused for the provisioning of new Virtual Servers, which may have virtual hardware configurations different from the original.



Image storage (vImage)

This consists of the volume, measured in GB, of external storage space for making backup copies of the Virtual Servers pertaining to a single subscription. This storage volume can only be accessed and used through the image creation tool on the subscriber's own platform.

Public IPv4 addresses (IPv4Pub)

This consists of a numeric identification through which any Virtual Server is unequivocally referenced on the public Internet network. Each subscription will receive a pool of public IPv4s, which may be used on one or more Virtual Servers belonging to the same subscription.

IPv4s are supplied randomly, with no guarantee that CIDRs (Classless Inter-Domain Routings) will be contiguous with other existing Public IPv4s belonging to the same subscription. This does not mean that the public IPv4 is dynamic, but rather that its numbering will be maintained as long as the IPv4 is linked to an active or inactive (suspended) virtual server. By unlinking a public IPv4 from a virtual server, or even at the same time deleting a virtual server completely, the IPv4s pertaining to it will be automatically returned to the platform's general IPv4 pool, with no guarantee that the same number may be reused on new Virtual Servers or even on existing Virtual Servers of the same subscription or client.

If the client needs to reserve a public IPv4 independently from the Virtual Servers it possesses, it may use a Load Balancer, which can cover a single virtual server (using a public IPv4 for this Virtual Server in this case; see the description of the Load Balancer functionality) to guarantee that the IPv4 is never returned, since it will remain associated with a Load Balancer, and not with a Virtual Server.

Private Local Virtual Network (PVLAN)

This product allows clients to use a private VLAN (PVLAN), which creates partitions on Layer 2 of the OSI model, separating the network packets contained in a single distribution domain, such that the traffic within is considered secure, even against packet-sniffing-type practices. Data on the vLAN is not subject to tariffs, and the average speed of this network is 1 Gbit/s, sufficient for the vast majority of application arrangements in which more than one server is needed (e.g., web server and database server). Virtual Servers may have access to the Internet and the vLAN simultaneously (one vNIC connected to the PVLAN and another to the Internet). In general, a vLAN works exactly like a private local network (LAN), connected to a private virtual switch. Thus, the vLAN allows for the creation of varying infrastructure arrangements, such as having database servers connected to web servers without having to use Internet bandwidth.

The client will be entitled to one PVLAN regardless of the number of Virtual Servers registered and the number of



interconnected hosts and/or Virtual Servers; this PVLAN corresponds to the number of private IPv4s available by subscription, i.e., 2,048 usable private IPv4s, equivalent to one block /21 for the CIDR (Classless Inter-Domain Routing).

Private IPv4 Addresses (IPv4Prv)

This consists of a numeric identification through which any Virtual Server is unequivocally referenced in the private local network. Each Server, when provisioned, will be given a private IPv4. This IPv4 may be used in the context of a private network (PVLAN). Private IPv4 numbering is contiguous and reserved once only within the subscription.

Virtual Load Balancer (vLB)

This consists of a virtual TCP/HTTP traffic balancing unit (Layers 4 and 7 of the TCP/IP model), allowing clients to distribute the workload uniformly (inbound accesses) between two or more Virtual Servers to optimize the use of resources, maximizing throughput, minimizing response time and avoiding surcharges on virtual servers and applications installed on it. Virtual Servers under a vLB have no need of a Public IPv4, and each vLB already includes the assignment of an exclusive public IPv4 address that corresponds to this vLB (Virtual IP). They are limited to a maximum of 128 Virtual Servers per vLB (instance). They are based on the dedicated open-source solution

HaProxy, and are managed through the simplified web interface available on the client's control panel or via API. The base and standardized balancing algorithm that has been adopted is round robin, with health checking (heart beat) of its members. If a Cloud server does not respond to the beat-check rhythm, the traffic will stop being directed to it. The VIP (Virtual IP) used by users The load balancers will be discounted from the subscription's public IPv4 pool.

Virtual Firewall (Firewall)

Any provisioned Virtual Server can be protected by a firewall. The individual and specific rules for each Virtual Server may be configured through the control panel. By default, Virtual Servers are not supplied with any configured firewall rule, which means that all inbound or outbound connections are permitted. At the time the first rule is created, the pattern is reversed, denying any inbound or outbound connection except those configured as rules, which are thus of the "allow" type. Available protocols for rule creation are TCP and UDP. A unique IP or even a network may be specified (through the network masking specification) for access from origin (remote) or even the entire Internet, in addition to the port(s). Internally, for each Virtual Server, the desired internal port(s) may be specified, including that a Port Address Translation (PAT) be executed.



DNS Hosting

This consists in the supply of primary (NS1) and secondary (NS2) authoritative resolution infrastructure for domain names (DNS - Domain Name System) belonging to the client that are directly linked to one or more virtual servers.

Scaling

It is possible to scale server capacity vertically (or scale-up). Through the control panel, the client can easily change the processing speed (vClock) and/or RAM memory (vRAM) values in a Virtual Server, thereby improving the throughput of the applications under peak access and/or demand conditions. The number of virtual processors (vCPU) cannot be changed, only the processing speed (vClock) of the vCPU(s) already provisioned in the Virtual Server.

Throughput and consumption monitoring (Monitor)

For each Virtual Server, the product offers an interface for monitoring the (a) vCPU, (b) vRAM and (c) External Transfer resources. The interface is provided graphically, in real time or in accordance with the selected period. It is provided as a data generator to assist the client in the processes of optimization (tunning) and monitoring of its Virtual Server.

Remote Web Access Console (WebConsole)

To allow for Virtual Server administration independently of client software (VNC, Remote Office, SSH, etc.), as well as local network restrictions (corporate networks that block specific ports), it offers a web interface, accessible through the control panel, that allows clients to access their Virtual Server dashboard. This dashboard uses the VNC protocol and accesses the virtual server through the platform administration network, which is not subject to interference and/or configuration errors.

To use this feature, the client must have installed Java JRE on its equipment. The browser used must permit the use of Java Applets. The interface will automatically direct the client to install this plug-in if it detects that it has not been installed or configured.

Application Programming Interface (API)

This feature allows clients to manage their resources through an API (Application Programming Interface) based on the REST standard. This API may be securely accessed (HTTPS) by using a secure key (token) for authentication. It allows the client to create and integrate its own offers, products and/or platforms into our cloud infrastructure.



Through a consistent and complete REST-ful architecture, clients can automate a series of actions and interactions with the infrastructure provided in order to, for example, obtain a list of existing Virtual Servers, their details, such as IP addresses, etc. They can create new Virtual Servers, delete existing ones, as well as start or suspend them. They can also reconfigure the virtual hardware, change IP addresses, create clones, make backup copies and many other actions, all on a scheduled basis. Compatible elements, features and functions that can be configured through the API may be found in the Virtual Server's API Restful User Guide, available on the control panel.

Operating Systems Offered

For Bare-metal (VM)-type virtualization:

- Windows Server 2012 R2 (x64)
- Windows Server 2008 R2 with Service Pack 1 (x64)
- CentOS 6.x (x64)
- Ubuntu Server 14.04 LTS (x64)

For container-type (CT) virtualization:

- CentOS 6.x (x64)
- Ubuntu Server 14.04 LTS (x64)

For any further information or to sign up, visit: <https://www.usclarocloud.com/>