CloudSigma

CloudSigma Platform Capabilities

Technical Capabilities, Architecture, and Service Overview

August, 2026

CloudSigma Platform Capabilities

CloudSigma is one of the most customizable cloud providers on the market with a focus on open design and flexibility with regards to computing deployments. The cloud platform is designed to provide an environment with the same degrees of freedom that private in-house environments end-users might have are able to offer. All functionality is available via an API or the end-user WebApp.

Cloud Architecture Overview

The CloudSigma cloud platform is built around a modular three-layer architecture comprising the infrastructure layer, the orchestration and management layer, and the customer-facing interface layer. This separation of responsibilities enables the platform to deliver a highly available, scalable, and fully programmable cloud environment while allowing each layer to evolve independently.

Infrastructure Layer

The infrastructure layer provides the physical compute, storage and networking resources that underpin the cloud platform. Each physical compute host runs a Linux operating system with the KVM hypervisor, providing hardware-assisted virtualisation for customer workloads. Resource isolation and scheduling are managed using Linux kernel resource management mechanisms, while virtual machine lifecycle management is handled through libvirt.

CloudSigma extends the standard KVM/libvirt stack with proprietary platform components, including the internally developed npax management agent. Together with a custom storage agent, these components provide functionality beyond that offered by the standard virtualisation stack, enabling advanced resource management, orchestration and monitoring capabilities across the platform.

Persistent block storage is delivered through a distributed, highly available NVMe/SSD-based storage platform designed to provide high IOPS, low latency and resilience against hardware failures. The storage architecture is engineered to continue operating through component failures, maintenance activities and software upgrades, ensuring continuous availability for customer workloads.

To maximise storage performance, the platform utilises high-performance networking technologies, including RDMA where supported, together with CloudSigma's custom storage components to minimise latency and maximise throughput between compute and storage nodes.

Orchestration and Management Layer

The orchestration and management layer is responsible for resource allocation, provisioning, monitoring, billing, customer account management and overall platform control. This layer has been developed entirely by CloudSigma and provides the flexibility required to expose fine-grained infrastructure resources directly to customers while supporting automated management of the underlying cloud infrastructure.

Communication between the infrastructure and management layers is implemented through an asynchronous messaging architecture based on RabbitMQ and distributed task processing. This messaging layer coordinates provisioning requests, infrastructure state changes and operational tasks across the platform, contributing significantly to overall platform reliability and resilience.

Platform metadata and operational records are maintained within a clustered PostgreSQL database, while individual platform services—including monitoring, billing and operational management—are implemented as independent containerised microservices. This modular architecture enables services to scale independently, simplifies maintenance and supports high availability across the management layer.

CloudSigma's flexible resource allocation model allows CPU, memory, storage and networking resources to be provisioned independently, while the platform's usage-based billing engine enables customers to pay only for the resources consumed.

Customer Interface Layer

The customer interface layer provides the public interfaces through which customers manage their cloud infrastructure. These consist primarily of the CloudSigma WebApp and the CloudSigma REST API, both of which expose comprehensive platform functionality for provisioning, configuration, monitoring and lifecycle management of cloud resources.

The WebApp provides an interactive management environment that delivers near real-time updates of infrastructure state, enabling customers to monitor and manage their resources efficiently through a responsive graphical interface.

The REST API follows standard HTTP methods and exposes virtually all functionality available through the WebApp, allowing complete automation of infrastructure deployment and management. This API-first approach enables integration with a broad ecosystem of cloud management platforms, infrastructure automation frameworks, software development kits (SDKs) and DevOps tooling.

Open Cloud Platform

CloudSigma provides a secure multi-tenant cloud environment that allows customers to deploy and manage workloads without imposing proprietary requirements on guest

operating systems or applications. Customers retain full control over their virtual infrastructure and are free to import, export and migrate workloads as required, minimising vendor lock-in and supporting hybrid and multi-cloud deployment strategies.

The platform combines carefully selected open-source technologies with proprietary CloudSigma software components to deliver a flexible and robust cloud environment. This hybrid approach enables CloudSigma to leverage mature open technologies while developing specialised capabilities where required to support platform performance, scalability and operational flexibility.

The architecture has been designed to support modern infrastructure automation, orchestration and service integration, enabling customers and technology partners to deploy, monitor and manage complex cloud environments through standard APIs and automation frameworks while maintaining high levels of availability, resilience and operational efficiency.

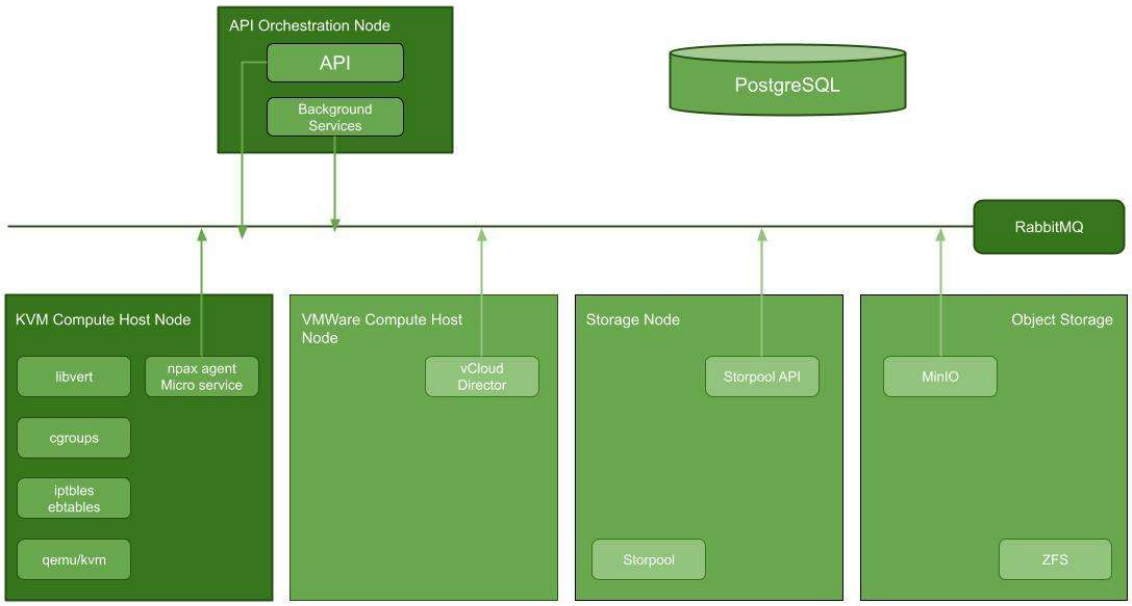

Software Components in the Cloud Architecture

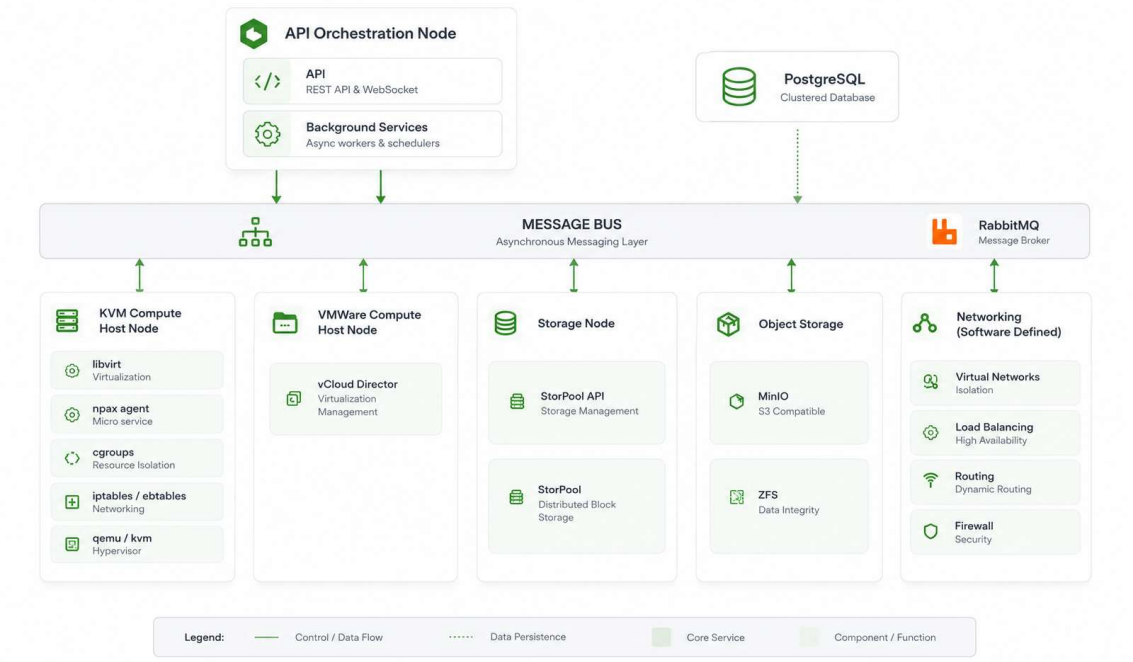

Underlying Hardware Overview

Outlined below is a description of the equipment capabilities used for the CloudSigma Services. Please note that some hardware specifications may vary depending on the cloud location.

Racks

Standard 42U racks are used with dual HPE power strips with external network based monitoring and management capabilities. Each rack is lockable using a keycode lock with codes controlled and subject to the relevant CloudSigma ISO27001 compliant security policies.

Compute hosts

The KVM compute hosts are hosting the end user KVM VMs and some non critical system VMs. This includes only system resources such as CPU and MEMORY, the storage is physically separated and distributed.

1. OS Ubuntu 18.04 (to be updated to Ubuntu 22.04)

2. System reservation via Cgroups
3. Libvirt - daemon and management tool for managing platform virtualization.
4. Npax - hypervisor agent, which acts to execute API server control calls, as well as to provide logging data to the API server
5. QEMU/KVM - virtualization technology/hypervisor
6. Overload protection mechanism via OOM killer
7. VMs running on top:
    - Uptime VMs - VMs which allow us to monitor the cloud from the inside.
    - NS - DNS management
    - Barman - backup server for the database
    - Penetration tests VM
    - Migration orchestrator - DR and Migration tool
    - Customer VMs - all end user VMs
8. Ebtables - firewall which protect us from the VMs inside the cloud
9. Features:
    - CPU resources
    - MEMORY resources
10. Redundant OS drives - Each Compute host has two drives with a mirror (RAID1) built on top
11. Redundant power - Same as the rest devices, the Compute hosts have two PSUs, connected to PDU-A and PDU-B
12. Redundant networking - Same as the rest devices, the Compute hosts are connected via logical aggregated interfaces with two physical ports, one to each switch in the IRF stack. Two of the ports are used for the VMs traffic and two for the storage traffic

Storage Nodes

CloudSigma has a clustered storage system leveraging three full copies of all data (one on SSD and two on magnetic). It is a genuine scale-out storage system that utilizes the entire cluster of servers to give superior IOPS performance compared to other software storage solutions (e.g. ZFS, Ceph, etc.). The new system works much closer to the hardware, so we are able to achieve actual performance levels in line with the underlying hardware potential. It combines the IOPS performance of all drives in the cluster and optimizes drive access patterns to provide low latency and handling of storage traffic bursts. Load is distributed equally between all servers through striping and sharding and avoiding the overhead of RAID based systems. It also achieves unmatched performance through its superior network protocol (bypassing the standard Linux networking stack) and its copy-on-write on-disk format, while also accelerating random writes and mitigating the "I/O blender effect".

The clustered storage system uses a replication mechanism that slices and stores copies of the data on different Servers. For primary, high performance storage this solution has many advantages and provides considerably higher levels of reliability and availability. In case of drive, server or other component failure, the system automatically uses another copy of the data located in another server (or rack), so that no customer data is lost or even temporarily unavailable.

Redundancy is provided by three copies (replicas) of the data written synchronously across the cluster - one whole copy on SSD and two on magnetics. The replication provides copies on different servers (not copies on one server). This technology is superior to RAID in both reliability and performance:

- The replication provides copies in different servers (not copies in one server). So even in the case of a server or component failure the data that sits on this server does not get lost
- Rebuilding a failed drive/node is done from many drives (and not from one particular drive as in RAID). This means rebuild times are significantly shorter and they do not impact the overall system performance

Node failure mode is superior to other systems based on the distribution and redundancy in the system. For example, in a twenty node system, failure of a single node would result in no impact on storage availability. Further, write performance would be totally unaffected with writes being sent directly to the SSD on available nodes, read performance would deprecate only for the 1/20 of the virtual drive whose copy on SSD was on the missing node as this portion is now served from magnetic storage on the other nodes. The result is the only impact of a node failure is a marginal slowing of some of the drive data for read only operations.

The clustered storage system used by CloudSigma offers high availability and an 'always live' environment that can survive hardware and software life cycles uninterrupted. This is combined with a very high performance profile and all offered at a competitive price per GB.

1. OS Ubuntu 18.04
2. System reservation via Cgroups
3. Overload protection mechanism via OOM killer
4. SP agent - storage provider agent
5. Hypervisor support - KVM and VMWARE simultaneously
6. Clustering - block level data clustering
7. Replication - triple
    - SSD - primary copy is on SSD drives and the secondary and tried copies on HDD
    - HDD - all copies are on HDD drives
8. Supports the VMs hosted on the Compute nodes
9. Holding customer's data

10. Redundant drives
11. Hot node replacement
12. Hot drive swap
13. Data live migration
14. Redundant OS drives - Each Storage host has two drives with a mirror (RAID1) built on top
15. Redundant power - Same as the rest devices, the Storage hosts have two PSUs, connected to PDU-A and PDU-B
16. Redundant networking - Same as the rest devices, the Storage hosts are connected via logical aggregated interfaces with two physical ports, one to each switch in the IRF stack. Two of the ports are used for the KVM traffic and two for the VMWARE traffic.

Routers

These servers host our routing software that is used to maintain the BGP sessions for external connectivity to the cloud. An open source FreeBSD based system is used to allow speeds up to 10Gbps for external connections to be achieved. The FreeBSD host OS is installed on RAID (two drives). Each router runs on dual power from two independent power supplies.

Each router consists of an HPE E5-2643v4 processor with 3.4GHz speed cores. As external connection speeds are single thread activities, the router uses this CPU choice in order to achieve higher maximum single connection throughput speeds. Each router also consists of one Intel X710 for 10GbE SFP+ dedicated NIC and 8x8GB HPE 1Rx8 PC4-2400T-R RAM.

Switches

The cloud platform utilises a flat redundant dual 10GigE networking fabric for customer compute and storage traffic. Traffic is separated using VLANs as well as KVM's traffic separation layer at the boundary of each cloud VM. HPE FF 5940 48SFP+ 6QSFP28 48 port 10GigE switches are used in IRF failover mode in pairs at the top of each rack. Both switches are dual powered and the networking is designed to survive a full switch failure without interruption. Each switch runs on dual power from two independent power supplies.

Parallel to this, a management switching fabric is used for monitoring and remote access to all physical components of the cloud. The HPE Aruba-2930M-48G is used for the management networking fabric offering firewalled IPMI and other cloud management access to certified engineers.

Cabling

The cloud utilises 1GigE copper network cabling in pairs to two separate switches. The top of rack switches are connected between racks with dual 40GigE cabling. The

management network uses 1GigE interfaces and 1GigE copper cabling.

Infrastructure-as-a-Service overview

Customers are able to provision processing, storage, networks and other fundamental computing resources in an unbundled way meaning CPU, RAM, storage and bandwidth can be purchased and combined independently to allow the best combination of cloud resources without the limitation of fixed sizes. So any combination of CPU and RAM can be achieved along with multiple drives mounted to a server and multiple networking interfaces. Each resource is billed separately and transparently as either a subscription or as pay-as-you-go in 5-minute billing segments, enabling customers to track exactly how much their cloud servers are costing over time and be billed for it accurately.

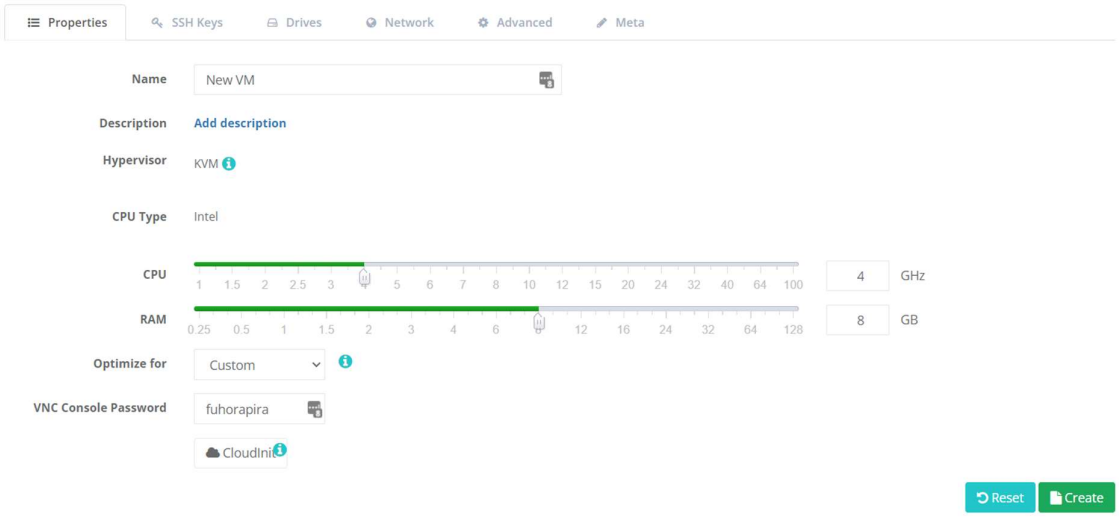

VM Creation

Uniquely, any x86 based operating system and software can be used with complete administrator/root control including all variants of BSD, Linux and Windows. End-users or in this particular case the managed service providers can upload raw ISO image, attach CPU and RAM to it and boot it up. This allows full backwards compatibility from the platform.

All our cloud servers and drives are persistent and modeled around the same methodology as physical dedicated server equivalents - i.e. drives, NICs etc. VLANs and IP addresses are also controlled using standard behavior and support all types of traffic including multicast and broadcast traffic, which is critical for high availability infrastructure in failover.

By creating a tighter link between the application and infrastructure layer through the virtual abstraction of the hypervisor, the cloud platform exposes a number of power tools allowing end-users to achieve greater performance levels from their cloud servers. For performance sensitive workloads this is especially important. These power tools include

the ability to define the virtual core size, to expose NUMA topology and to tweak hypervisor timer settings for maximum performance. Additional power features for example include exposing the full CPU instruction set to the virtual machine, allowing much faster processing of certain calculations. This is ensured by enabling the 'host CPU' setting as can be observed below.

For example, an end-user may have a virtual machine of size 100GHz and 125GB of RAM*. For an application that benefits from parallel processing, the number of virtual cores can be set to 20 thus giving twenty CPU threads of 1GHz each, for example. On the other hand for an application that's bound by core speed, the virtual core number would be set to eight, giving eight threads of 2.3GHz each. In this way, small changes can have a profound impact on performance without costing the end-user anything.

**It is theoretically possible to create a virtual machine with much higher amounts of RAM. This limit can be adjusted under certain circumstances (e.g. private cloud implementation).*

A screenshot from the provisioning portal shows an example of some of these power tweaking tools:

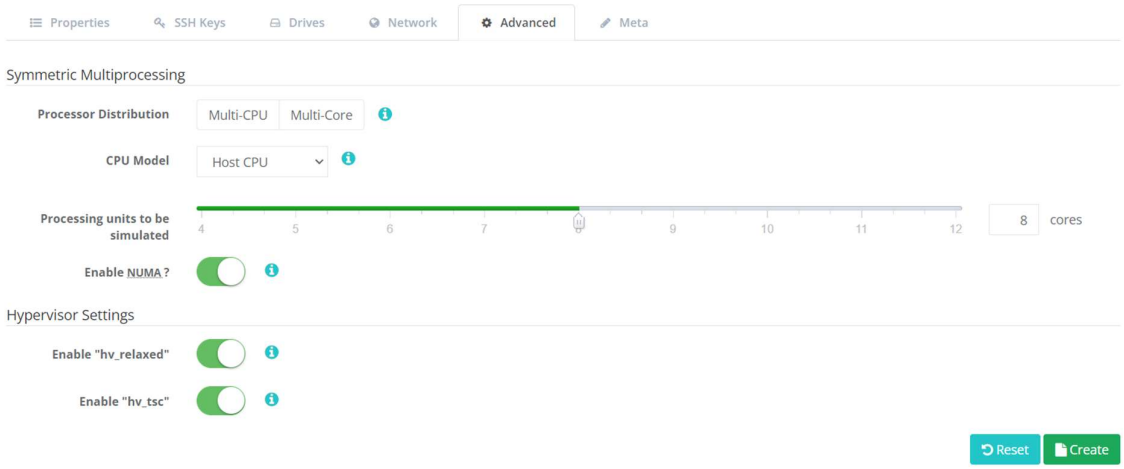

Advanced Tools

At an account level, account administrators are able to assign specific access and control rights over certain account related operations using access control policies. This allows one account to grant certain rights to another account (or user). This is done through a 'labeling' feature. The account administrator may for example label a set of servers 'project x' and through an access control policy grant full access to 'project x' labeled infrastructure to one or more other users. On the other hand, the account administrator might wish to grant read-only access to billing information for the accounting department. The flexibility and powerful possibilities achieved by pairing labeling with ACLs is clearly evident.

Virtual machines (VMs) can be provisioned in a matter of seconds with a high degree of flexibility and control. The average provisioning time for a new cloud server/VM is 15 seconds. This is critical for running end-user facing web services that need to flexibly provision new resources in line with fluctuating end-user demand.

Users have the option of full API access with all account actions available, i.e. 100% API coverage, allowing complete automation and remote infrastructure monitoring, or the option of a feature-rich, yet intuitive web browser based GUI. The CloudSigma WebApp has been designed to allow easy resource management via any web browser. CPU and RAM can be specified to the nearest MHz and MB.

VM provisioning is achieved via a simplified wizard or custom server creation tool. The end-user can choose from a wide selection of ready system images including a number of BSD, Linux and Windows based operating systems as well as being able to quickly and easily upload their own ready ISO image. End-users can customise these marketplace images using cloud initialisation frameworks such as cloudinit (see https://help.ubuntu.com/community/CloudInit), allowing them to take a standard installation and contextualise/customise it to their specific requirements on-the-fly on first boot-up.

Virtual Machine Templates

CloudSigma offers two approaches to VM creation - a quick and simple creation process and a detailed, flexible and extremely granular custom one. Both are outlined below.

- *Simple Server/VM Creation Process*

Using this process, an end-user can quickly create a VM from a predefined list of sizes and available operating systems and applications. VM settings are optimised automatically by the cloud based on the user's choice of operating system.

Next, the user can choose from the operating system they wish to deploy. After this final step, the virtual machine will automatically be created with associated storage and networking and will be booted up. Details such as the username to be used when logging in to the server, and the VM's public IP address is succinctly communicated with the end-user. Where appropriate, a VNC connection in the browser will be opened to allow the user immediate access.

The quick server/VM creation process takes less than 30 seconds with a provisioning time of 15 seconds. A new VM can be configured, provisioned and accessible in under one minute using this method.

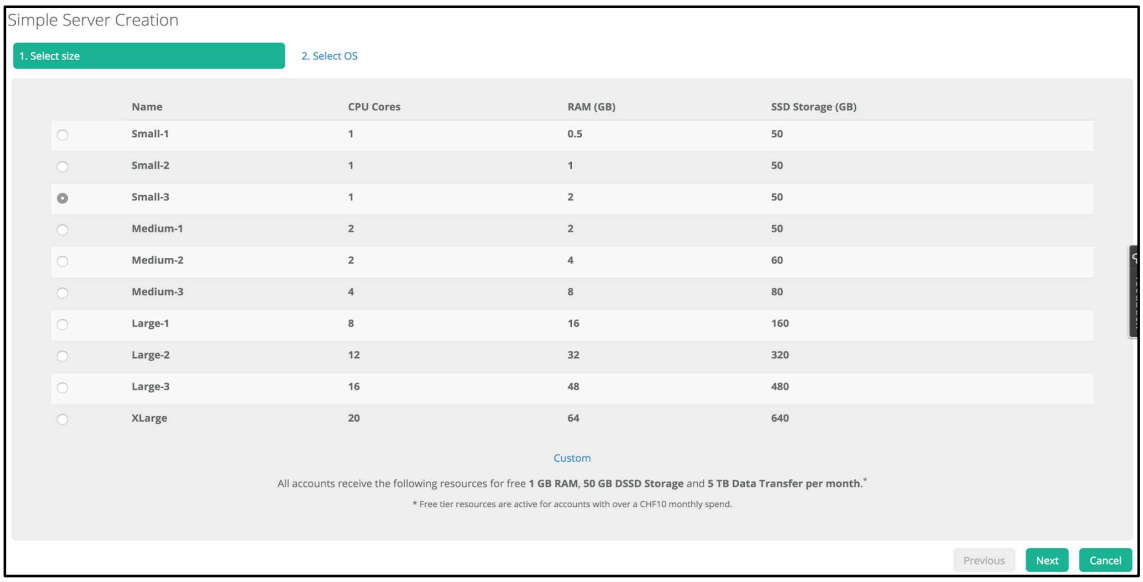

Simple Server Creation Sizing Menu

- *Custom Server/VM Creation*

For users wishing to specify their VMs in more detail, the custom server creation tool is available, allowing users to specify every aspect of their VM exactly how they would like it. Firstly, they are able to choose the exact CPU / RAM sizing and give the VM a name:

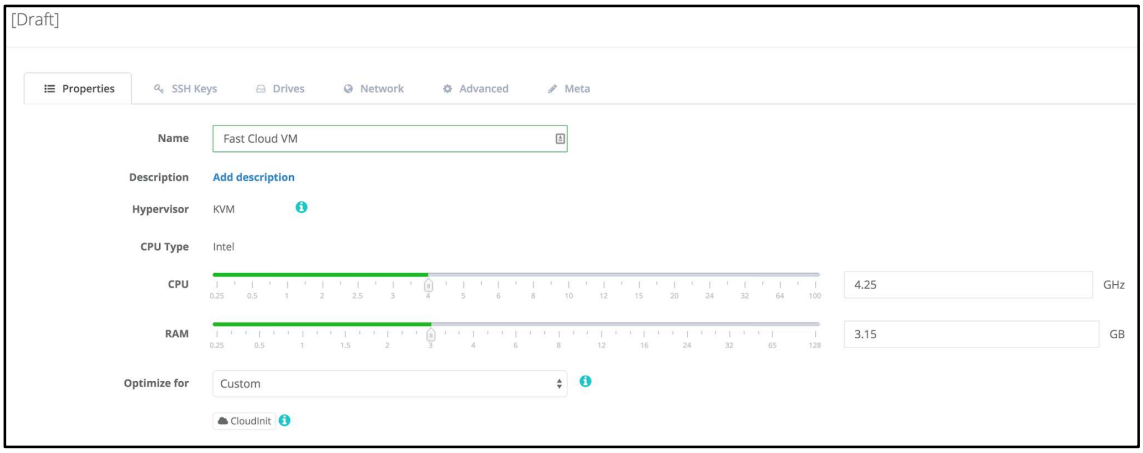

Customer Server Creation Sizing and Options

Next, the user can add one or more SSH keys in order to enable secure access from from the get-go (in the case of Linux and BSD based systems). Once this step is completed, it is possible to create and/or attach and mount one or more drives of the user's choosing to the new VM. After this, the networking set-up can be defined. As outlined in the networking section in more detail, multiple public and private networking interfaces can be assigned to a single VM. Finally, the end-user can tweak the advanced settings of the virtual machine to maximise its performance in relation to the specific application requirement that they have.

Virtual Machines Provisioning Times

Virtual provisioning times for virtual machines do not vary by size. The average provisioning time is 15 seconds. The cloud provisioning layer is able to provision multiple virtual machines in parallel. The bulk aggregate provisioning times are estimated below:

Number of New VMs being Provisioned	Estimated Delivery Time
1	15 seconds
100	2.5 minutes
500	12.5 minutes

Virtual machine creation, spin-up and spin-down times are identical. The above table can be used as a guideline for typical lead times to be expected when provisioning VMs in volume. As you can see, VM provisioning times scale linearly on CloudSigma's high-performance stack. The billing system is based on five minute billing segments, so VMs once shutdown will stop being billed for within a maximum of five minutes.

As outlined above, the cloud is built from modular components. Through capacity management, the capacity of each respective resource is managed over time to enable sufficient resources to always be available to end-users. In general when any resource reaches over 80% utilisation, additional capacity is purchased. Additionally, due to the simple module nature of all systems, additional computing capacity beyond aggregate cloud capacity at any one time can be delivered within 10-15 working days. Therefore, with a relatively modest amount of notice larger scale requirements can be serviced.

VM sizing range and configuration

The following resource ranges are available in relation to virtual machine and drive sizing:

Resource Type	Minimum	Maximum
CPU	1 core / 1GHz	100 cores / 100GHz
RAM	1GB	128GB
Block Storage (one drive)	1GB	10TB

Note: A VM can have up to sixteen separate block storage devices attached to it.

The solution exposes in a utility manner CPU and RAM to users. End-users can freely combine CPU and RAM to the nearest MHz and MB respectively subject to the minimum and maximum ranges outlined in the table above. It is therefore possible to size CPU independently of RAM on a VM by VM basis. The initial VM configuration can be easily scaled up and down as required by the end user, i.e. adding or removing CPU or RAM.

The CloudSigma cloud supports on demand and subscription based consumption of computing resources. In the case of on-demand resource consumption, the cloud billing system monitors usage per compute resource (i.e. CPU, RAM, storage etc.) and compares this against the end user's subscribed resource level (which may be 0). This check is conducted every five minutes. The billing system then charges the user for any resource consumption amount above the subscribed level in five minute segments. In this way a user will get charged for burst usage within five minutes of spinning up the VM and, within five minutes of shut down the system will detect this change and stop charging for those burst resources. Actual burst prices vary depending on the level of utilization of the cloud. The cloud platform uses an advanced pricing algorithm that allows our pay-as-you-go burst pricing to change over time depending on how busy our cloud is and whether the resources are being provisioned in peak or off-peak times.

Subscription and Resource Commitment Flexibility

CloudSigma's resource model allows compute, memory, storage, networking, and other applicable resources to be subscribed and consumed independently rather than being permanently tied to a fixed virtual machine configuration.

Subscribed resources may be allocated or reallocated across workload configurations as requirements change, while additional capacity may be consumed on demand where applicable.

CloudSigma supports multiple subscription commitment periods. Subject to applicable commercial terms and service availability, resource subscriptions may be renewed, extended, or transitioned to another supported commitment period without requiring the associated workloads to be rebuilt solely because of the change in subscription term.

GPU-as-a-Service

CloudSigma offers customers the ability to add GPUs to their virtual machines and take advantage of high-performance, cost-effective computing that can meet the most demanding workloads. At the heart of CloudSigma's GPU offering is the NVIDIA RTX A6000 card, which is optimised for HPC, AI and data analytics. We enable customers to easily build NVIDIA RTX A6000 optimized VMs in passthrough mode so VM instances have direct control over the GPU/s and their built-in memory.

CloudSigma is working on GPU-integration at several cloud locations globally. Several graphics cards have been validated including the NVIDIA RTX A6000, NVIDIA A100 and V100 Tensor Core GPUs. The NVIDIA RTX A6000 GPU is ideal for data analytics workloads and applications such as VDI, HPC, and AI/Deep learning, while the NVIDIA A100 and

V100 Tensor Core GPU are optimised for HPC, AI and data analytics. More cards will be validated according to demand.

Multi-Location Deployment Capability

CloudSigma operates cloud infrastructure through multiple locations globally. Equivalent compute configurations may be provisioned across available CloudSigma locations, subject to capacity and applicable service arrangements, providing customers flexibility for workload relocation, resiliency, and changing deployment requirements.

Supported Software & Tools

CloudSigma's stack was chosen as the cloud stack to facilitate hosted processing as it offers an open API and a very high degree of compatibility with end users who have legacy requirements. There are absolutely no restrictions on the type of operating system or application that the customer can deploy, as long as it is x86 compatible. The following section presents the options available in CloudSigma cloud platform, which customers or their SI/MSPs can deploy and configure instantly.

Operating Systems Support

The cloud platform provides an extensive library of ready-to-use public drive images that include pre-installed systems that can be deployed instantly, alongside install CDs for users preferring to install their own systems from scratch. The following table provides a sample of the available drive images in the marketplace:

Linux distributions	CentOS (desktop, server) 7.x, 8x 9x Debian 10x, 11x, 12x Fedora (desktop, server) 38 Ubuntu LTS 16.x, 18.x, 20.x, 22.x, 23.x Ubuntu 21.x Server CoreOS, OpenVPN, OwnCloud, pfSense, Vyatta
Windows distributions	Microsoft SQL Server Standard Edition* Microsoft SQL Server Enterprise Edition* Microsoft SQL Server 2014 Web Edition* *(2014, 2016, 2017, 2019, 2022) VirtIO Drivers for Windows Windows Server Datacenter (2016, 2019, 2022) Windows Server Standard (2016, 2019, 2022)
Other drive images	FreeBSD 12.x, 13.x NetBSD 8.x, 9.x OpenBSD 7.x

In addition, instead of restricting users to a curated list of predetermined images, we have provided customers with a simple and easy way to import and export data and ready operating system images into and out of the system. End-users can use their own images and upload them (or download them) directly via the end-user provisioning portal or over the API. This method also allows for pausing and resuming uploads and downloads.

End-users always have sole root/administrative access over their VMs. This approach not only eliminates vendor lock-in but allows the chosen cloud system to be fully backwards compatible with end-users able to run any version of an operating system they so wish rather than choosing from a fixed list and be forced to upgrade periodically as that list changes.

Middleware Components / Available Tools

CloudSigma offers a range of ready middleware components for a variety of end user markets. These middleware tools are exposed via a middleware services marketplace that's integrated on top of the cloud platform's infrastructure cloud layer.

- In the case of applications, they deploy into an auto-scaling platform-as-a-service environment already patched for the latest updates making them ideal for web services. All the application templates are instantly deployable into a CloudSigma Infrastructure account, offering auto-scaling and load balancing as the demand on computing resources varies through their use. The applications include Apache, Docker, Hadoop, Joomla, MongoDB, Ruby and others.
- In the case of libraries, drivers and other integrations, they are used in conjunction with other services to build wider deployments. These include: pycloudsigma Library, Go integration, Terraform integration, Ansible integration, OpenStack Heat, CloudInit, CoreOS, Apache jclouds Driver, Apache Libcloud Driver and others.

API Support

The solution exposes 100% of all functionality available via APIs allowing full integration into wider software deployments. The core infrastructure API is provided by the CloudSigma cloud stack, however this is complemented by interoperability with a number of open source and other industry standard APIs for wider user engagement and accessibility.

As outlined earlier in this section, a number of commercial and open source cross cloud drivers are available that allow the management of CloudSigma in conjunction with other leading public and private cloud stacks, both commercial and open source. Overall the vast majority of end users will find a tool, integration or direct API interface that fits within their programming language and legacy computing requirements in order to minimise friction to onboarding new users to the platform.

Compute Performance

Please note that performance results may vary depending on the specific equipment installed. The following example should be considered a baseline. In most cases, even higher performance can be achieved.

The deployment uses the latest generation HPE E5-2697v4 with 3.6 Ghz speed/core CPU which form part of the high performance range. These are coupled with 32GB DDR4 RAM for some of the highest performance available from any public cloud. The table below shows the raw computing performance available from a 16 core (40 GHz), 32GB cloud server of the sort commonly deployed for big data workloads:

16 CPUs in system; running 1 parallel copy of tests

Dhrystone 2 using register variables	24927164 lps	(10.0 s, 7 samples)
Double-Precision Whetstone	3237 MWIPS	(9.7 s, 7 samples)
Execl Throughput	3602 lps	(30.0 s, 2 samples)
Pipe Throughput	1683243 lps	(10.0 s, 7 samples)
Pipe-based Context Switching	341015 lps	(10.0 s, 7 samples)
Process Creation	8962 lps	(30.0 s, 2 samples)
Shell Scripts (1 concurrent)	11326 lpm	(60.0 s, 2 samples)
Shell Scripts (8 concurrent)	5318 lpm	(60.0 s, 2 samples)
System Call Overhead	2967605 lps	(10.0 s, 7 samples)

System Benchmarks Index Values	BASELINE	RESULT	INDEX
Dhrystone 2 using register variables	116700	24927164	2136
Double-Precision Whetstone	55	3237	588
Execl Throughput	43	3602	837
Pipe Throughput	12440	1683243	1353
Pipe-based Context Switching	4000	341015	852
Process Creation	126	8962	711.3
Shell Scripts (1 concurrent)	42.4	11326	2671
Shell Scripts (8 concurrent)	6.0	5318	8863
System Call Overhead	15000	2967605	1978

16 CPUs in system; running 16 parallel copies of tests

Dhrystone 2 using register variables	316260286 lps	(10.0 s, 7 samples)
Double-Precision Whetstone	46442 MWIPS	(10.0 s, 7 samples)
Execl Throughput	30049 lps	(30.0 s, 2 samples)
Pipe Throughput	21011265 lps	(10.0 s, 7 samples)
Pipe-based Context Switching	3368796 lps	(10.0 s, 7 samples)
Process Creation	69130 lps	(30.0 s, 2 samples)
Shell Scripts (1 concurrent)	69615 lpm	(60.0 s, 2 samples)
Shell Scripts (8 concurrent)	9452 lpm	(60.0 s, 2 samples)
System Call Overhead	5179760 lps	(10.0 s, 7 samples)

System Benchmarks Index Values	BASELINE	RESULT	INDEX
Dhrystone 2 using register variables	116700	316260286	27100
Double-Precision Whetstone	55	46442	8444
Execl Throughput	43	30049	6988
Pipe Throughput	12440	21011265	16890
Pipe-based Context Switching	4000	3368796	8422
Process Creation	126	69130	5486
Shell Scripts (1 concurrent	42.4	69614	16418
Shell Scripts (8 concurrent)	6	9452	15753
System Call Overhead	15000	5179760	3453

Block Storage Environment

The cloud platform solution provides the ability to users to create and (re)size multiple persistent virtual drives. It is an NVMe SSD-based storage system perfectly suited to the high-performance requirements associated with big data mining and scientific computing requiring high performance storage. Below is a screenshot of the end-user provisioning portal showing the current drives in the user's account:

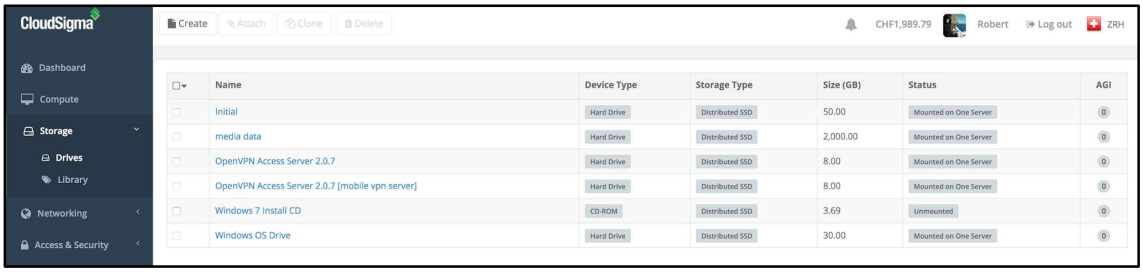

WebApp Provisioning Portal

Drives can be mounted or unmounted. The system supports the creation of regular drives as well as CDs / DVDs. Users can directly upload RAW images of CDs and hard drives, complementing the extensive selection of install CDs and pre-installed systems available from the marketplace.

There is no fixed drive sizing, users can create drives subject to a minimum and maximum size in a granular way. All storage is persistent and the provisioning time in creating a new drive is 10-20 seconds with an average provisioning time of 15 seconds, similar to starting or creating a new virtual machine. Once created, a drive can be resized subject to the maximum drive size as outlined in the figure below. Drive resizing is executed instantly with a provisioning time of 10-20 seconds for the increased drive size.

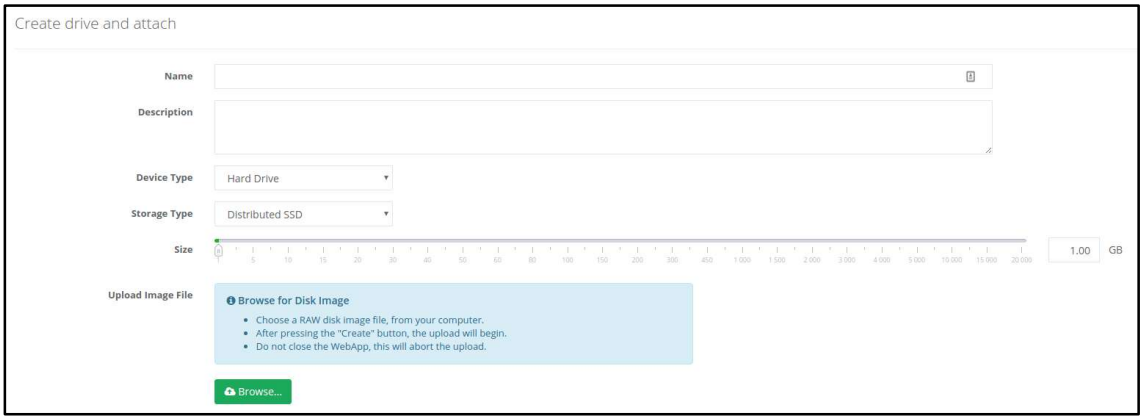

WebApp Storage Creation Section

Unlike less flexible cloud systems that bundle storage and compute resources together in pre-defined sizes, the cloud platform allows multiple drives to be mounted to a single VM. In the example VM below, two drives are attached to the same cloud server. One smaller one (50GB) for the operating system and a second drive for media data (2TB). Users can also track read and write activity as shown in the pop-up from the screenshot.

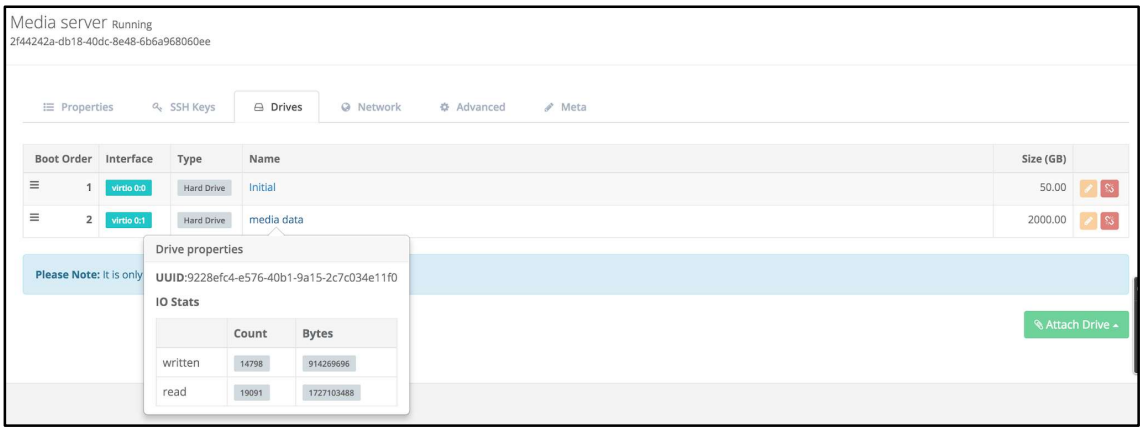

Drives Section

The block storage system is driven by StorPool - an advanced proprietary high-performance distributed storage technology chosen by CloudSigma after extensive research into best-in-breed block storage solutions. It pools the attached storage (hard disks or SSDs) of standard servers in order to create a single pool of shared block storage. The software is installed on each server in the cluster and combines the performance and capacity of all drives attached to the servers into one global namespace.

The system provides standard block devices which can then be attached to virtual machines/cloud servers. In this implementation, the system has been configured with no file system in order to use volumes to store VM images directly.

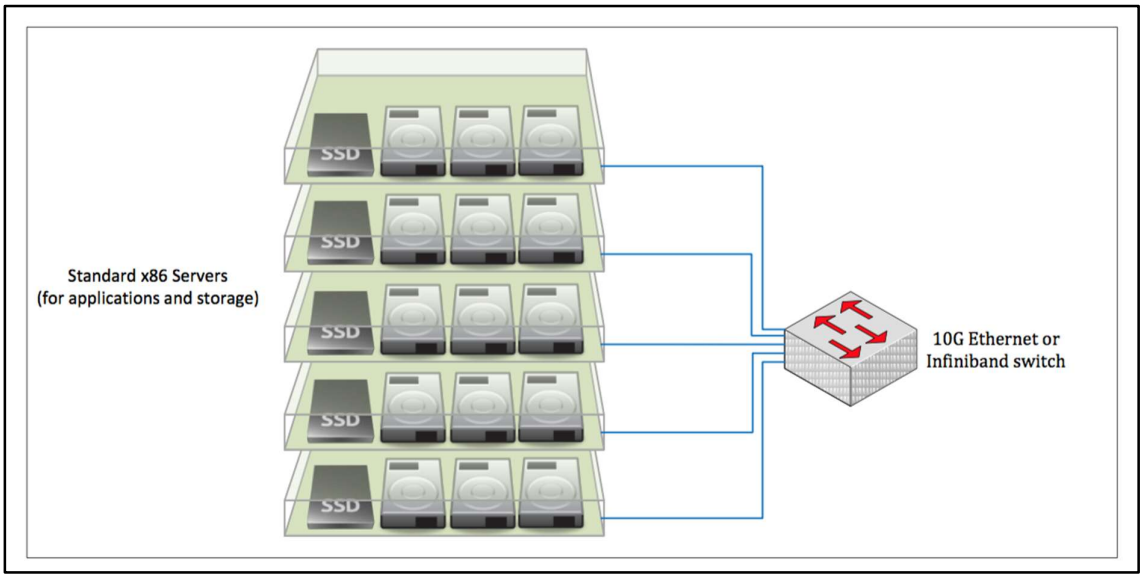

StorPool Shared Block Storage Solution

Redundancy is provided by multiple copies (replicas) of the data written synchronously across the cluster. Users can set the number of replication copies, with the CloudSigma cloud configured to store three copies of all data. This technology is superior to RAID in both reliability and performance.

Unlike RAID, our system replication distributes copies across different servers. As such, in the case of a server or component failure, data that is stored on this affected server is not lost.

Rebuilding a failed drive/node is done from multiple drives (and not from one particular drive as is the case with RAID). What this means is that rebuild times are significantly shorter and they do not impact overall system performance. The system also protects data and guarantees data integrity via a 64-bit checksum and version for each sector maintained by the system. The system provides a very high degree of flexibility in volume management. Unlike other storage technologies, such as RAID, the system does not rely on device mirroring (pairing drives for redundancy). So every disk that is added to the cluster adds capacity to the cluster, not just for new data but also for existing data. Provided that there are sufficient copies of blocks, drives can be added or taken away with no impact to the storage service. Unlike rigid systems like RAID, the system does not impose any strict hierarchical storage structure that links and reflects to the underlying disks. The system simply creates a single pool of storage that utilises the full capacity and performance of a set of commodity drives. This allows customers to easily scale out storage requirements as the cloud data ecosystem grows over time.

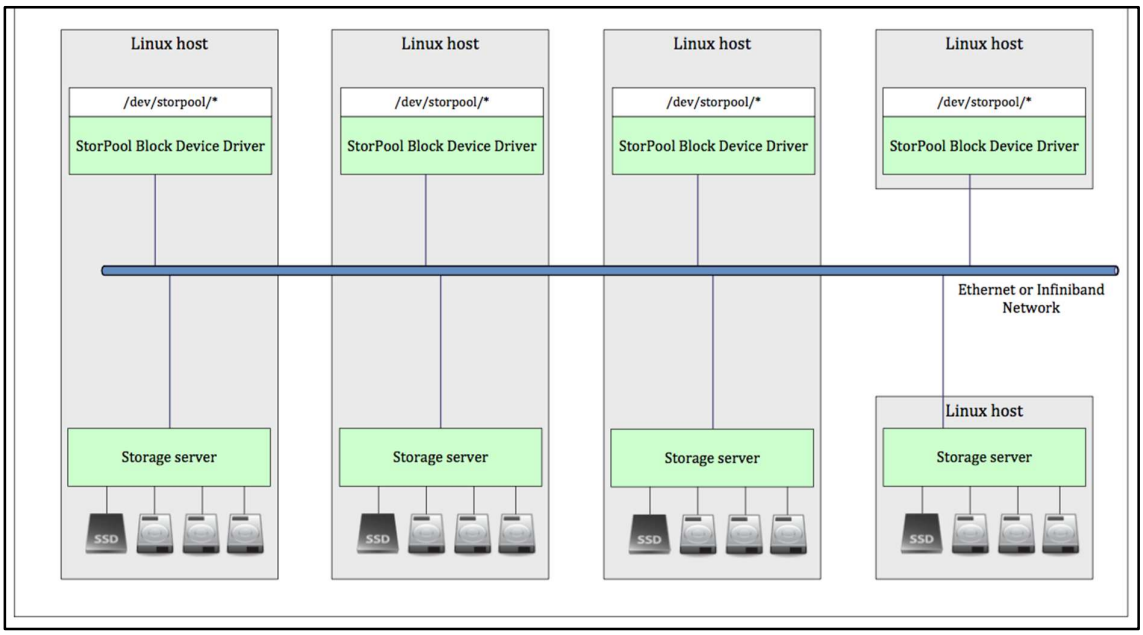

Data Replication and Redundancy

In summary, the proposed block storage system provides a flexible yet simple interface to users whilst combining advanced technology on the back-end to delivery high performance and extreme storage resilience.

Block Storage Performance

One of the primary challenges facing a multi tenant public cloud is delivering sufficient storage performance in the face of near random blended IOPS from many different end-users. Below are the results of a customer test performed on the cloud platform stack to be used for the cloud data ecosystem and using the same SSD storage solution. The comparison is made against Atlassian indicators in nanoseconds. The system delivers, by Atlassian defined standards, excellent performance levels on all operations except deletion, which scores near-excellent:

Operation	Clustered SSD	Excellent	Good	Bad
Open	35,770	< 40,000	< 80,000	> 150,000
Read/Write	26,836	< 40,000	< 60,000	> 100,000
Close	15,392	< 20,000	< 40,000	> 100,000
Delete	87,512	< 50,000	< 200,000	> 300,000

Block Storage Backup Solution

The block storage system used for the proposed cloud solution by CloudSigma is used for all the cloud storage utilised by user VMs. This block storage solution is able to provide implicit non-disruptive backups at the storage block level for all user data. This includes any data contained within virtual drives including application data, databases, all operating system information etc. It provides full drive level backup of customer data. The backup system backs up all end-user computing data each night and retains seven days of rolling snapshots. End-users can roll back one or more of their virtual drives to any of the seven off-site data snapshots at any time. Additionally, they can manually request a backup of a virtual drive at any time via the API or WebApp.

Due to the clustered three copy block storage system used, the backup process does not impact service availability or performance, but instead utilises free capacity to sync data for the backups in the background. End-users can therefore have a continuously running computing environment and at the same time have rolling backups of all their data.

In addition to the automatic backup system, end-users are able to create point-in-time snapshots of their drives, which can later be cloned and upgraded to create stand-alone drives. A snapshot can be created on-demand while the server is running, thus in no way affecting the performance or availability of the systems. By using snapshots, customers can protect themselves from data corruption or use them for auditing purposes.

Furthermore, there is an advanced snapshot management feature, allowing customers to create snapshot management policies and apply them to one or more drives. In this manner, customers are able to automate the snapshot process.

Customers can define and apply backup and snapshot policies to one or more virtual drives through the CloudSigma WebApp and REST API. These policies support automated backup and snapshot scheduling, retention management and recovery-point administration. Customers can create, update or remove policies as their operational, cost-management and compliance requirements change, while retaining the ability to create point-in-time snapshots on demand and restore or clone them into independent drives. Where enabled for the relevant CloudSigma location and service, remote snapshot and disaster-recovery workflows can also be used to support resilient data protection and recovery.

Block Storage Deletion Lifecycle

The platform applies different deletion policies to drives, snapshots, and backups. These policies balance customer recovery options, storage capacity, and operational reliability.

Drive Deferred Deletion

When a customer deletes a storage drive, the drive enters a configurable deferred-deletion period rather than being removed immediately. During this period, its data may remain available for recovery, subject to the applicable platform policy.

The duration can vary according to the customer's account state. Active, suspended, deleted, or otherwise restricted accounts may have different retention periods based on operational and business requirements.

After the applicable period expires, the platform permanently removes the drive. If a storage system approaches its capacity or drive-count limits, deletion may occur earlier to preserve service availability and sufficient free capacity.

Snapshot Deletion

Snapshots do not have a deferred-deletion period. When a customer deletes a snapshot, the platform removes it regardless of the customer's account state.

Once deletion is complete, the snapshot is no longer available for recovery, and its associated storage resources can be reclaimed.

Backup Deferred Deletion

When a customer deletes a backup, the platform applies a configurable deferred-deletion period. Its duration can vary according to the customer's account state and may also be configured for an individual customer.

After the applicable period expires, the platform confirms that the backup is no longer

present in the remote storage system before removing its remaining platform record.

Automated and Resilient Cleanup

Cleanup runs automatically and independently across storage locations. A failure affecting one location does not prevent the platform from continuing the process on others.

Object Storage Environment

The object storage approach allows easy granular growth in capacity in a highly responsive way at the hardware level, enabling us to deliver additional capacity cost effectively on relatively short lead times whilst at the same time providing file systems support and Amazon S3 compatible cloud storage service which can easily be integrated into many different applications. In this way the end-user or application is abstracted from the underlying storage complexity. Files can easily be organised into buckets/directories using whatever logic is desired. Additionally, it is very simple to expose individual or sets of files externally with a full user access control framework. Finally the architecture makes it easy to update or replicate files from the primary to secondary sites or in reverse for file recovery.

CloudSigma supports customer-controlled data lifecycle management for object and file-storage workloads through the CloudSigma WebApp, REST API and S3-compatible interface. Customers can organise data into buckets and directories; control access to individual files or sets of files through account-level access controls; and automate storage-management workflows through standard APIs and S3-compatible tools. These workflows can be used to manage data upload, replication, recovery, retention processes and deletion in accordance with the customer's operational, cost-management and compliance requirements.

Object Storage is accessed through an Amazon S3-compatible interface and the CloudSigma API, allowing customers to integrate standard S3-compatible tools, applications and automation workflows. Data can be organised into buckets and directories, while access to individual files, groups of files and storage resources can be controlled through CloudSigma's account-level access-control framework. This enables customers to apply access policies and organise data according to application, project, retention and compliance requirements. Customers retain control over their data-management processes, including data replication, recovery and deletion, through the available API and WebApp interfaces.

The object storage environment offers a highly scalable, open source environment capable of storing large scale data archives accessed via a simple API over HTTPS or mounted directly as a filesystem. It also offers an S3 compatible API providing even wider compatibility and integration possibilities. The simple API interaction allows easy integration into various applications, as well as an ability to upload and download files

quickly and easily. By using a popular, open standard, the cloud platform offers a solution that is portable in the future to new deployments.

The solution exposes 100% of all functionality available via APIs allowing full integration into wider software deployments. The core infrastructure API is provided by the CloudSigma cloud stack, however this is complemented by interoperability with a number of open source and other industry standard APIs for wider user engagement and accessibility. The object storage environment is exposed using the main CloudSigma's API in a unified way. The object storage environment also provides additional accessibility via the S3 compatible API.

The object storage archive is part of CloudSigma's public cloud architecture, where virtual machines and containers can be provisioned in a matter of seconds with a high degree of flexibility and control.

Networking Capabilities

CloudSigma has extensive experience in managing the challenging networking workloads of the cloud. Having multiple suppliers of a high quality allows us to ensure internet connectivity is always available to each infrastructure location. As a result, all VMs regardless of their size benefit from an excellent networking performance both internally and externally within the services.

VM External Access: Internet & Remote Access

The proposed solution leveraging CloudSigma's cloud platform offers an open networking environment to end users. As such end users can connect to their VMs using any networking protocol they desire including SSH, remote desktop protocol (for Windows), FTP, SFTP, HTTP, HTTPS etc. Users in turn can secure and lock down their networking connections to only the services/ports they desire and against specific whitelisted IP ranges for added security. Additionally for convenience and system recovery, CloudSigma's cloud offers built-in VNC access within the browser allowing secure remote system visibility directly from within the end user provisioning portal. This can be enabled and disabled on running machines as required. One VNC connection per VM is possible at any one time.

In the case of SSH access, the solution additionally includes a framework for storing SSH keys and injecting them into VMs upon boot-up to allow users immediate secure access via SSH, this is particularly useful for Linux based operating environments. More information on SSH key management is included in the security focused section of this document.

There are no restrictions on the number of SSH or other protocol connections per VM or per user account with the exception of the single VNC connection concurrently per VM already outlined above (this is a protocol limitation).

CloudSigma has taken an unrestricted approach to networking meaning that end user VMs can bring up as many concurrent networking sessions per VM as they desire. The only performance limitation on virtual machines is there sizing, this becomes the bottleneck on the number of concurrent connections they can sustain and the total throughput. Based on existing customer deployments with CloudSigma, customers wishing to sustain hundreds of concurrent connections to VMs will be able to do so without stability or performance issues.

Due to the large amount of deployed capacity in excess of 20Gbps, VMs do not have reserved networking capacity as there is no network contention on the external or internal networking anticipated in order to require such a measure.

Included in the drives library and services marketplace are a number of ready appliances allowing customers to immediately deploy VPN and SFTP servers to secure external access and move data into and out of the cloud.

VM Networking Configuration & Options

End-user VMs can be assigned public and private networking interfaces. Multiple interfaces may be assigned to a single VM, allowing for example gateway servers to be created between separated networks, as well as many more more sophisticated networking configurations.

Public VM Networking

End-user VMs may have a public networking interface attached. This is exposed as a regular networking card to the VM operating system. A number of network card emulations are available to ensure maximum compatibility across many operating systems.

All compute nodes housing virtual machines are connected via dual 10Gbps lines to two top of rack switches in failover, which in turn switch traffic via the routers to the external 10Gbps layer 3 lines. We do not rate limit VM public networking traffic. As such, end-users can transfer data in and out of the cloud at multiple Gbps speeds with ease.

A public networking interface can be assigned a dynamic IP or a static IP. In the case of the static IP, the user subscribes for this IP and assigns it to the specific VM. A security-conscious user not wishing to expose a virtual machine to public networking connectivity can choose to have no public networking interface at all on the VM.

Customer-Owned Public IP Address Space

CloudSigma supports the use of customer-owned public IP address ranges within applicable cloud locations, subject to technical validation, routing requirements, and applicable service arrangements.

Approved customer-provided IP prefixes may be routed to the customer's cloud environment and used by hosted virtual machines, network appliances, gateways, and other supported services. The customer retains ownership of the applicable address space, while CloudSigma provides the required cloud-network integration and routing support.

Availability and implementation requirements may vary by cloud location and are validated as part of the network design.

Virtual Firewalling

CloudSigma offers a custom firewall policy creation feature to facilitate the improvement of security and traffic in and outside of the cloud. If you use the API, we have made available a template requiring you to complete the rules. Management is achieved via policies which are applied to single or groups of VMs, allowing easy management and application across both small and large scale infrastructure in a convenient way.

You can also choose the sequence in which you want the firewall rules to be applied. In order for the firewall policy to be active, it needs to be applied to your public network interface.

Order #	Source IP	Destination IP	IP protocol	Source port	Destination port	Direction	Action
0		178.22.66.228/32	tcp		80	both	accept
1		178.22.66.228/32	tcp		443	both	accept
2	46.218.123.34/32	178.22.66.228/32	tcp		22	both	accept
3			tcp	53		in	accept
4			udp	53		in	accept
5			tcp		0:1024	in	drop

Examples of Firewall Policies

Policies range from a single rule that blocks all external public IP traffic, to complex schemes that only allow connections to certain ports from a set of IPs. Network policies are saved and then applied to one or more virtual servers as required. Furthermore, network policies can be reconfigured and re-applied to running servers without service disruption.

Additionally, customers are free to implement their own virtual firewall appliances within their cloud-based service stack (for example SmoothWall) and / or make use of each VM's own firewall policies, for example via IPTables. End-users can create networking policies/firewalls virtually to secure public networking interfaces at the hypervisor level. Please see the networking policy section for more information below.

Private Networking Interfaces

The CloudSigma cloud supports the ability for end-users to attach one or more private networking interfaces to their virtual machines. These networks are exposed as network card interfaces to the operating system of the virtual machine.

By placing multiple virtual machines in the same private network, traffic can be passed securely between them. These private networks offer a 'virtual wire' implementation supporting all kinds of networking traffic including broadcast and multicast traffic which most cloud stacks do not support. Private networking traffic is passed over dual 10Gbps networking within the cloud. VM to VM networking performance figures have been provided above within the cloud performance section.

Internal Networking Performance of connectivity between VMs

The current and following subsections include customer generated results on CloudSigma cloud for internal and external networking performance. The results show that the cloud networking capabilities outreach the requirements for incoming data transfer stated in the SOW.

The table below provides a customer generated report for connectivity between cloud servers/VMs in CloudSigma cloud. Please be advised that the performance results may vary depending on the specific equipment installed. The figures below should be considered as a baseline only. In most cases even performance can be achieved.

Interval	Transfer	Bandwidth	Total Datagrams
0.00-1.00 sec	472 MBytes	3.96 Gbits/sec	60366
1.00-2.00 sec	466 MBytes	3.91 Gbits/sec	59647
2.00-3.00 sec	474 MBytes	3.97 Gbits/sec	60652
3.00-4.00 sec	476 MBytes	3.99 Gbits/sec	60890
4.00-5.00 sec	485 MBytes	4.07 Gbits/sec	62086
5.00-6.00 sec	464 MBytes	3.90 Gbits/sec	59449
6.00-7.00 sec	463 MBytes	3.88 Gbits/sec	59227
7.00-8.00 sec	473 MBytes	3.97 Gbits/sec	60568
8.00-9.00 sec	475 MBytes	3.98 Gbits/sec	60776
9.00-10.00 sec	475 MBytes	3.99 Gbits/sec	60857

Interval	Transfer	Bandwidth	Jitter	Total Datagrams
0.00-10.00 sec	4.61 GBytes	3.96 Gbits/sec	0.056 ms	604241

CloudSigma offers a full 'virtual wire' implementation for private networks allowing low latency high throughput connectivity between cloud servers/VMs as demonstrated above.

External Networking Performance of connectivity between VMs

CloudSigma's cloud will utilise multiple external layer 3 lines to deliver reliable connectivity to end-users. Through its experience in cloud networking from both hardware and software perspectives, and is able to deliver high throughput rates externally to end-users.

Virtual Private Network (VPC)

A virtual private cloud (VPC) is a private cloud solution provided within public cloud infrastructure. CloudSigma provides VPC by isolating a portion of our public cloud infrastructure to be provisioned for private use. Because VPC networks interface only through private connections, the network's traffic is guarded against exposure to the public internet, where that information could be vulnerable to interception or exposure. VPC networks can also be used to isolate execution environments, as well as tenants. Internet gateways can also be set up as a single point of access between the public internet and the resources on your VPC network.

With CloudSigma's VPCs you can provision on-demand a configurable pool of dedicated computing and/or storage resources. Each VPC represents a private pool of resources - physically isolated from all other users in the cloud. Customers can take advantage of various compute node sizes depending on the cloud location. Billing is simple and transparent with no hidden costs.

The advantages of VPC include the following:

Performance - Use all the CPU power on your dedicated compute hosts and/or enjoy maximum IOPS on your fast dedicated storage volumes. Enjoy the same high performance consistently.

Scalability - Create a VPC environment quickly and intuitively. Provision new dedicated nodes and volumes on-demand with just a few clicks.

Compatibility - Run any x86 OS or application in your VPC unmodified. Use the same WebApp and API as on CloudSigma's public cloud.

Enhanced Security - Use all dedicated capacities in the VPC that are physically isolated from all other cloud users. Manage private networks for additional network-level isolation.

Compliance - Meet your legal requirements for a dedicated environment. Connect to any on-premises or data center deployments using VPN.

100% SLA - Enjoy highly available private cloud resources, covered by our 100% SLA for uptime and connectivity.

Security Capabilities

CloudSigma endeavours to deliver a high degree of security and privacy for customers in accordance with the various aspects of their computing. CloudSigma has attained the international auditable standard of ISO 27001 by setting best practices for data privacy, security, and information governance that are applied to processes, IT systems and people, by establishing and maintaining a company wide Information Security Management System (ISMS). We are committed to openness and transparency with respect to our security procedures and policies. Legal documentation for each cloud location is available publicly on the CloudSigma website. In addition to ISO 27001, the cloud platform is also ISO 27017 and ISO 27018 certified. CloudSigma applies the highest standards regarding users security, data protection, business continuity and payment processing. Third party audits are also supported. The latter must be covered by either the partner or the third party.

Network Security & Traffic Separation (Data in Transit)

CloudSigma's cloud leverages the open source KVM hypervisor to provide full separation of all traffic between client accounts below the virtual machine level. No end user can view traffic from any other end user. This is achieved through full packet inspection of all incoming and outgoing packets to VMs by Linux KVM. KVM implements a virtual switch for every networking interface of each VM. Acceptable traffic courses (i.e. other VMs in the user's account) are instantiated on boot and updated as VMs are added and removed from various networks in (i.e. end user private networks in the cloud). In addition, end users can apply virtual firewalls at the hypervisor level that apply additional rules.

Storage Separation (Data at Rest)

Users can easily keep data private and secure by ensuring the operating system/file structure is fully encrypted using technologies such as KVM for Linux distributions or Truecrypt for Windows environment. While this approach doesn't eliminate the potential for data leakage, it does render any leaked data completely unusable to others. However, this approach can be somewhat disruptive if for example an encrypted server crashes, as it will require manual procedures to enable access to encrypted data on reboot. Customers can apply encryption to the drive on creation. This does eliminate the possibility of data leakage and ensures the automatic encryption of any new data as it is written. Encryption can be enabled via the API or WebApp when creating a new drive. It should be noted that this approach may have a small impact on performance. Customers can always configure their servers to have a system drive with no encryption and a data drive that is fully encrypted.

Two-Factor Authentication

CloudSigma customers are able to use Google's Two-step authentication in order to log onto their accounts. Two-step verification increases the security for access to their cloud platform account by providing a six to eight-digit unique password, which users must provide in addition to their username and password in order to log into the cloud platform UI. The feature is available via an API call and the WebApp. The default status of the feature is disabled and can be activated by individual customers if they want to.

Root Access & Operating System Security

Customers retain full sole access to their data at the file system level; the CloudSigma system does not have access inside VMs or drives. All customer data is handled automatically by our system. This includes activities such as drive deletion and scheduled deletion (for deprecated accounts). CloudSigma makes no copies of client drive data and therefore the sole copy resides in our cloud unless the customer chooses to clone the drive to another storage system or location.

Via the drives marketplace preinstalled systems of a large selection of operating systems are provided, these operating systems are correctly patched regularly to ensure security vulnerabilities are patched enabling end users to deploy secure virus and vulnerability free operating systems for their VMs on first boot.

Security Management

CloudSigma is also ISO-27001 certified including all areas of sales, operations and support as well as being PCI-DSS compliant. A copy of the latest ISO 27001 (Information Security), ISO 27017 (Cloud Security), ISO 27018 (Privacy Protection) certificates can be obtained upon request. CloudSigma is also certified by Canonical as a certified Ubuntu Public Cloud meaning Ubuntu server images in CloudSigma's marketplace are optimized and patchedd regularly.

Quality Management

At present CloudSigma applies internal quality management procedures to processes relating to the creation and quality control of the products and services offered by the company. We use a combination of methodologies and management tools to ensure customer requirements and expectations are continuously monitored and met. The heads of each department are responsible for the implementation of all quality management procedures and to ensure the management system is compatible with ISO-27001.

An integrated management interface is the centralized system we use to manage and monitor the cloud, from both an operations and account management perspective. There are different access levels defined by the separate user roles and rights. Team members are trained and kept up to date on the different components and metrics used, and are then granted access level based on their roles.

An agile framework provides us with a group of software development methods in which requirements and solutions evolve through collaboration between self-organizing, cross-functional teams. Retaining short term flexibility through an agile approach reduces the risk of failure and surfaces issues earlier before they threaten the success of the proposal. The iterative sprint process provides the ability over time forecast the work effort required for each for each deliverable allowing the product owner to fine tune their product roadmap. Being agile also moves the trade off between completeness of product and release timing. It is possible to release more frequently and to iterate faster.

The second facet of our engineering approach, are the systems in place to manage software deployment in a secure and reliable manner, complementing the agile methodologies discussed. Deployment is managed across three separate environments: development, acceptance testing, and production. The main source code repository is managed through the Mercurial Source Code Management tool. The updated codebase is verified through the Jenkins Continuous Integration tool, which tests each check-in via an automated build, and running a sequence of integration tests and unit tests on the code.

On the integration servers we run a suite of user level acceptance tests, that primarily monitor performance. If these tests pass successfully, the code is added to the Mercurial Production Repository. At this point the code becomes subject to an internal code review, by a developer who has not been involved with this code base. When this is signed off, the code is sent to a third and final mercurial repository, ready for deployment into the production environment.

Risk Management is applied in tandem with our agile approach and assigned the following four elements: risk description, probability, size of loss measured in days or story points and exposure. The risks are reevaluated at each sprint, with a single consolidated risk value created.

Technical Audit

All customers of the CloudSigma platform are entitled to perform security, operations and processes auditing in relation to the services that we provide. The audit can be performed by the customer or a third party authorized by the customer. Please note the following:

- any audits shall be executed at the cost of the customer, including but not limited to charges that we have incurred during this process;
- the data center can be visited and access can be granted only after an advance notice of two weeks prior to the day of visit;
- in order to conduct the audit, the customer or their third-party auditor shall be accompanied by a CloudSigma staff member.

Data Encryption

The CloudSigma cloud supports the encryption of partial or full (boot level) encryption of virtual drives. We recommend as a best practice that end users perform boot-level

encryption of sensitive data and retain the keys outside our cloud. The cloud platform currently supports a number of customers running fully encrypted data storage in conjunction with their services in the cloud. End users can also connect to their VMs using encrypted protocols to ensure the integrity of login and other data they transmit to and from their servers.

Typical end user use cases where encryption would be used would be when a hosted processing provider is storing sensitive end user information or when a service provider themselves wishes to store proprietary data that they wish to be secured additionally. In these cases an encrypted partition can be created for that specific data or a separate virtual drive with full file system encryption used. In this way the end user providing the service can combine best performance from data not needing encryption with high security for the data that does.

CloudSigma has extensive experience of encrypting drive data using numerous encryption approaches, such as Cryptsetup, dm-crypt, FDE, TrueCrypt (VeraCrypt), as well as lower-level block storage encryption via ZFS and is happy to work with end users to ensure the right encryption is implemented to reflect their requirements.

CloudSigma supports customer-managed encryption of data at rest for virtual drives, including full-disk, boot-volume and partition-level encryption. Using industry-standard encryption technologies such as Cryptsetup and dm-crypt, customers can implement AES encryption with 128-bit keys or stronger, including AES-256, according to their security, compliance and key-management requirements. Encryption keys can be retained and managed by the customer outside the CloudSigma environment.

This approach enables customers to protect sensitive data at rest while retaining control over the encryption configuration and associated key-management processes. CloudSigma can support customers in selecting and implementing an appropriate encryption approach for their workloads, including full-disk encryption, encrypted filesystems and lower-level block-storage encryption where applicable.

Virtual Router

Virtual Router functionality is an integrated part of CloudSigma's cloud platform that offers customers an effective Network-as-a-Service tool accessible via the user interface and API. A high level of granularity provides for a redundant and protected network configuration, deep level of control of access, and set up of preferred connections and routing. The Virtual Router tool grants unlimited virtual domains, firewall policies and registered endpoints, as well as a rich set of additional features. The Virtual Router tool allows customers to build very complex network solutions. Customers are also able to manage virtual network deployments on VMware avoiding the need for VMware's NSX hypervisor.

Customers can use the virtual router functionality to define and manage traffic flows between any two networks they wish. A multi-tiered networking topology can be defined within the cloud by creating multiple virtual routers. The following diagram shows the first Virtual Router controlling the traffic passing between the Internet and Private Network A, while the second Virtual Router controls traffic between Private Network A and B. Each virtual router can be deployed in a redundant way to ensure resiliency by utilizing dedicated, reserved underlying hardware capacity. This ensures stable and reliable performance as well as independence between any two network pairings.

Intel-SGX

Intel Software Guard Extensions (Intel-SGX) helps protect data in use via application isolation technology. By protecting selected code and data from modification, developers can partition their application into hardened enclaves or trusted execution modules to help increase application security. With Intel-SGX application developers have the ability to protect select code and data from disclosure or modification. Enclaves are trusted execution environments (TEE) that utilize a separate portion of memory that is encrypted for TEE use. Customers are able to select Software Guard Extension when provisioning a server and allocate RAM to that server. Intel-SGX is an additional security measure that can be beneficial to companies working with sensitive and confidential data. Intel-SGX ensures the integrity and confidentiality of computations in such systems where privileged processes are deemed unreliable. The data in the enclave remain protected even if the cloud servers are compromised.

Keys Management

Secure access to end-user VMs is facilitated using SSH key pairs. This allows users to run commands on a machine's command prompt without them being physically present near the machine. This enables users to establish a secure channel over an insecure network.

The SSH key creation covers the following three scenarios:

- CloudSigma support team can generate a public and a private SSH key for the customers.
- Customers can generate the SSH keys themselves and upload only the public key in their CloudSigma account. In this scenario customers take the responsibility for the protection and access of the private key. This option is provided for customers that are especially concerned about security in the cloud.
- Customers can generate the SSH keys themselves and upload both SSH keys in the CloudSigma account. Currently, this scenario doesn't provide additional benefits, but in the near future an SSH console (similar to the VNC console today) will be opened automatically in the WebApp. This option will be only available for customers that have uploaded both their public and private SSH keys to their CloudSigma accounts.

Access Control Lists (ACLs)

Access control lists (ACLs) are meant to segment account control rights and access to the different operational aspects. With this feature the account administrators can allow access to different resources or a group of resources across the account. The account administrator delegates permissions to each account and lets each user log in to the web console with their own user credentials. Examples of delegated abilities:

- Provide accounting with access to billing, but not to edit any server/networking resources.
- Give junior sysadmins access to start/stop servers, but not to create or delete anything.
- Provide senior sysadmins access to fully manage the architecture, but not being able to access billing.
- Provide the operations team with access to firewall policies and networking, but not to servers.
- Provide a team with full access to their servers (using server tagging), but not not any of the other resources.

ACLs enable a very granular control over the account's permissions and budget, resulting in higher levels of transparency and security. For each module, it is possible to delegate either read-only or read-write permission. It is also possible to delegate permission on individual resources, for example a server or set of drives.

Event Logging

CloudSigma implements comprehensive logging against all its infrastructure deployments. All infrastructure components contain logging information against all critical system functions (including access or data impacting actions for example) and by user. Logs are retained locally on the infrastructure component and replicated to a central repository using the logging service tool Kibana. Logs include networking activity, as well as key application and operating system events. Logs are retained for a minimum of one year onsite with logs retained for up to two years upon request.

Patching Service

Software upgrades and system patches at both the operating system and application layer are achieved without service disruption due to the redundant and clustered architecture of the solution. System patching including security updates are subject to our security and change management procedures covered by CloudSigma's ISO27001:2013 certified processes.

DDoS Protection Measures

The following measures are used to prevent Distributed Denial of Service (DDoS) attacks:

- Implement additional rules for fraud payment prevention (Number of tries per new account eg. 5. This should only apply if the account age is less than a week)

- Apply an ISP approach for safety - Traffic shaping (put in place a policy limiting the number of packets and throughput). Upon request that policy will be editable for a particular client or set of clients
- blacklisting of IP addresses in the event of an attack
- maintenance of significant spare external IP connectivity to absorb malicious traffic
- additional firewall measures both at our edge and internally
- obfuscation of and removal (in some cases) of public IP connectivity from core cloud infrastructure where possible to avoid targeting of key cloud infrastructure assets
- externally hosted cloud status page allowing status updates even during a potential total outage (see http://status.cloudsigma.com/)
- using IP proxies on core services and other measures that can't be shared publicly
- automatic blocking of DDOS attacks against our clouds.

Monitoring Capabilities

Every component in the cloud delivery chain is monitored by multiple overlapping services. Any incident can therefore be viewed from various angles and it is the combination of the feedback from those various services that allows fast and effective mitigation and root cause analysis of incidents. We outline each monitoring system that we use in turn below with a high level overview of it's primary function.

Monitoring Tools

- **Pingdom** provides network availability information. We use this service to monitor not only external networking availability, but also to measure file system responsiveness and other basic availability metrics of components remotely.
- **New Relic** provides in-depth application and process level information. We use New Relic for key software services such as the API and web console components, as well as for more basic resource utilisation levels across the various physical hardware. New Relic includes software service scoring called Appdex which alerts against degraded responsiveness of key services. After installing the Server Monitoring agent, customers can monitor and get alerts for servers or applications running within them against performance or availability issues. The Application Monitoring tool provides greater insight into application behaviors, which tremendously simplifies the process of identifying bottlenecks and application health issues such as capacity issues, server health issues, server availability, CPU utilization, memory utilization, disk I/O utilization, and many others.
- **Zabbix** provides a very comprehensive monitoring framework which is used across almost all components of the cloud, using our own custom implemented monitoring regime. Zabbix monitors components from the physical layer all the way through to the virtual layer including the performance and availability of test virtual machines sat on every compute node in the cloud. Zabbix is additionally used for monitoring networking traffic flows, which are also visualised using

MRTG. Zabbix alerts are implemented to alert system engineers when critical thresholds are breached.

- **IPMI** - All hardware used within our cloud (except networking) includes IPMI management abilities. We additionally integrate information harvested from IPMI cards to complement other available system information. This includes metrics such as CPU fan status, temperature measurements and similar.
- **Mixpanel** - Mixpanel is a user event driven data analytics platform. All events on the web console are captured via this tool by customer allowing us to identify and respond to errors and issues that may be affecting customers.
- **Selenium** - Selenium provides automated user agent testing for the front end web console, we use this to regularly test the web console and ensure that functionality is correctly functioning at the user interface level.
- **MRTG** - It monitors SNMP network devices and illustrates the data in graphs showing how much traffic has passed through the devices.
- **NetFlow** - This is the Open Source service providing two tools called Nfsen and Nfdump which we are using to closely monitor the network components of the cloud equipment, i.e. the switches, routers and anti-DDoS machine. The instrument allow CloudSigma to keep a very close eye on the network processes running on individual IPs and packets.
- **OpsGenie** - This tool represents a platform where the information from all above mentioned systems is streamlined and provided in the form of graphs, alerts and others. The consolidation of the information in one place is time-saving, enabling the CloudSigma operations team to assess the incident and take appropriate actions rapidly.

The systems are combined to provide ongoing performance measurement and reporting. Consolidated system information is available in read-only form via a dedicated password protected portal to monitor in real time system status and performance.

The graph below depicts a sample CloudSigma infrastructure deployment to demonstrate the monitoring tools implementation:

- Light green components represent software level services.
- Grey components represent pure networking hardware only.
- Blue components represent orchestration and management layer physical equipment.
- Dark green components represent physical systems and equipment being virtualized and sold to end-users.

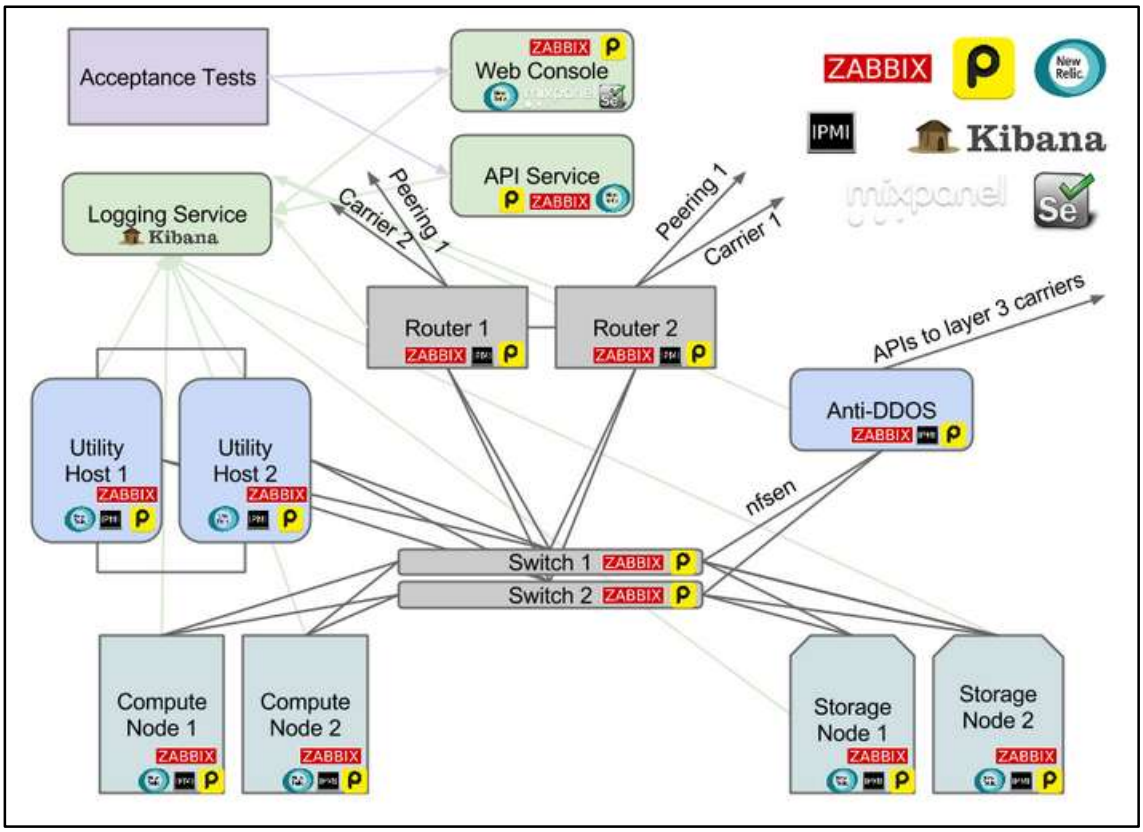

A sample monitoring tools implementation

VMWare on CloudSigma

CloudSigma offers a cloud platform that supports both KVM and VMware deployments in an integrated and efficient way. VMWare is offered alongside a wider suite of services that includes the CloudSigma PaaS platform and the existing KVM environment. VMWare features on CloudSigma include:

- Full support for VMware hypervisor and VMware cloud director
- Advanced features such as NSX
    - vCloud RDS
    - vCloud site recovery manager
    - vCloud availability
    - vCenter monitoring and automation
- Support for certified VMware environments like SAP HANA, 3PAR, NFS

CloudSigma reduces the VMWare Licencing footprint by:

- Only requiring the core operational components of VMware without the need to rely on additional components that severely escalate costs
- Using our own resource allocation engine, avoiding the need for expensive options. Additionally we can deploy a zero RAM guarantee on all VMs. This alone HALVES the licensing footprint

- Advanced networking services are delivered via our virtual routing solution which avoids the need for NSX

Whilst this approach avoids a huge amount of licensing costs, a customer requiring a certain feature can still utilise it but it will attract additional licensing costs

Platform-as-a-Service overview

With the CloudSigma PaaS we offer an automated cloud hosting platform with support for all popular stacks including Java, PHP, Node.js, Python, Goland, Ruby, .NET, Docker, Kubernetes, and many more. Customers can add any other language or stack with the help of Docker containers. Automated installation, scaling, replication and updates simplify the management and support of applications & clusters. Within the PaaS Marketplace customers can choose among over 250 performance-optimized and highly available applications and 15000+ Docker templates ready for one click install. There are a number of features that help to increase interoperability and eliminate vendor lock-in. They include:

- **Language agnostic -** Provision out-of-the-box Java, PHP, Ruby, Python, Node.js, .NET or Go production or dev environments as well as Docker and Kubernetes clusters. Choose between user-friendly GUI, SSH access, CLI and API.
- **Zero code change -** Deploy and run traditional monolithic applications and cloud-native microservices with no code changes and zero application redesign. Deploy via archives, FTPS/SFTP, GIT/SVN or via integrated plugins.
- **Private Git Repository Hosting and DevOps Integration** - CloudSigma's IaaS and PaaS environments support the deployment and operation of customer-selected private Git repository and DevOps platforms, including self-managed Git services deployed on virtual machines, Docker containers or Kubernetes clusters. Customers can use standard Git-based deployment workflows through SSH, CLI, API and integrated plugins, and can deploy applications directly from Git repositories. The number of private repositories is governed by the capacity, configuration and applicable licensing of the customer-selected Git platform rather than by an artificial CloudSigma platform repository limit. This enables customers to implement a scalable private source-code-management environment with access controls, persistent storage, backup and disaster-recovery configurations appropriate to their requirements.
- **Application neutral -** Run both cloud-native microservices and legacy applications. Run both application and system containers. Enjoy easy entry point to the cloud, reduce your go-to-market time and secure no vendor lock-in.
- **Anything-as-a-Service -** Full marketplace with a rich set of preconfigured applications and one click installs for common clustered environments. All available as performance-optimized and highly available prepackaged templates.

The PaaS delivers computing resources in a highly efficient way. Resources are measured in cloudlets - a measurement unit that includes 128 MiB of RAM and 400 MHz of CPU

power simultaneously.

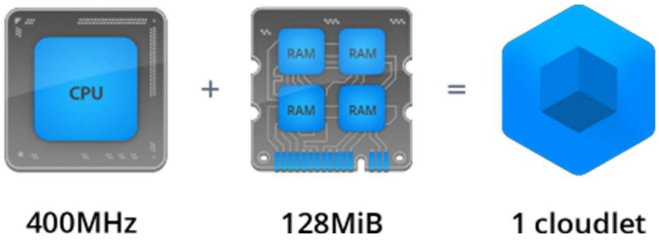

Being highly granular, it allows customers to accurately determine the required capacity for each server in their environment, which ensures a truly fair usage-based pricing to avoid overpaying for the unused resources. Cost efficient resource consumption is further enhanced with the following feature set:

- **Elastic auto-scaling -** Automatic vertical & horizontal scaling for granular load-based resource allocation without container restart. Guaranteed high performance, application availability and uptime.
- **Pay-as-you-use -** New generation pay-per-use pricing based on real usage instead of server size. Eliminated right-sizing problem and guaranteed cost efficiency by paying only for actual consumption.
- **Monitoring -** Monitor your application performance and get real-time reports on resources consumption and application load. Set up custom alert notifications and get notified about traffic spikes.

Load Balancing - Increase application performance and capacity with out-of-the-box load balancers based on NGINX, HAProxy and more. Available for all environments on both infrastructure and application level.

Service Support & SLAs

CloudSigma Support Channels

Live support is provided 24/7 via Chat, Email and our Zendesk online ticketing system. CloudSigma's Technical Support Team can be reached directly by emailing support@cloudsigma.com. Live Chat available from the WebApp (UI) or the company website.

Response time is usually immediate if Live Chat is used. Otherwise, we can guarantee a response time of less than 3 hours for all support channels. CloudSigma infrastructure is continuously monitored using an extensive and intelligent monitoring platform, ensuring that all systems are operational. All monitoring systems comply with our security and confidentiality protocols.

CloudSigma uses a number of other channels to communicate and provide updated information. These channels include the following:

1. Service Status Page: https://status.clousigma.com/

2. Constant notification from our Technical Support Team (support@cloudsigma.com) to your registered email address "xx@yy.zz". We send out an email in the case of events, including scheduled maintenance, explaining the situation and delivering detailed root cause analysis.
3. Twitter account especially for automatic update
4. CloudSigma YouTube channel provides a number of useful how-to-guides.

CloudSigma Infrastructure Support Service Levels

The following table provides an overview of the support service levels for the cloud infrastructure:

Item	Service Level
Uptime guarantee	100%
Chat and Email Customer Service - 24x7x365	Yes
Web based ticket system	Yes (Zendesk)
Phone Customer Service (office hours)	Yes
Documentation, Best Practice Guides & Webinars	Yes
Max. Response Time	30 minutes
Dedicated Resolution Engineer	Yes & continuous efforts
SLA credits for qualifying downtime after 15min	x50

The following table outlines the escalation policies for the cloud infrastructure:

Escalation policy	
First level Customer Relation	3 hours continuous effort then escalation
Second level - IT operations	6 hours continuous effort then escalation
Third level - developers	Final escalation level

CloudSigma Service Level Agreement (SLA)

- 100% virtual server availability and network uptime guarantee - we guarantee 100% availability of virtual servers as defined by their availability on our network and their responsiveness in a non-degraded way. This guarantee covers the hardware and virtualization hypervisor layers only and not the software (including but not limited to operating systems and applications) running within virtual servers. The network will be deemed 'available' if CloudSigma border

routers and switches are available and responding to CloudSigma monitoring tools in a non-degraded manner.

- 1ms or less network latency guarantee - we guarantee a network latency of 1ms or less for data packets between servers within CloudSigma services and network. The network latency refers to network latency times between the boundary layer of one virtual server to the boundary layer of another virtual server and excludes internal latency times resulting from software running within a virtual server at either end of the data transit.

CloudSigma Penalty Scenario / Credit

If we fail to meet the guarantees detailed above, you will be able to request a credit as detailed below up to a maximum of 100% of your fee for capacity used during the previous 30 calendar days:

- Credit of 50 times the fees for any period of lack of availability for a virtual server or network uptime lasting more than 15 minutes as measured from the time at which you validly inform us at support@cloudsigma.com or the time at which our monitoring systems detect the lack of availability, whichever is earlier;
- Credit of 50 times the fees for any period of network latency as defined above, with greater than 1ms lasting for more than 15 minutes as measured from the time at which you validly inform us at support@cloudsigma.com or the time at which our monitoring systems detect the lack of availability, whichever is earlier;
- Credit of your entire fee for the previous 30 calendar days in case of permanent loss of your stored data resulting from hardware or software failure of CloudSigma systems. This provision entirely excludes data loss or corruption resulting from software running within a virtual server. In the event that we fail to meet the guarantee on more than one occasion within a period of 30 calendar days, then the credit that you may claim for any incident will be limited to the maximum of 100% of your fee for capacity used since the previous incident or 100% of your fee for capacity used during the previous 30 calendar days, whichever fee is lower. To receive a credit, you must contact us at support@cloudsigma.com within 30 calendar days of the incident, specifying the start time, date and duration of the qualifying period which forms the basis of your claim and the amount of credit claimed. We will be the sole arbiter regarding the award of credit and our decision will be final and binding.

CloudSigma SLA Limitations and exemptions

The following items or situations are exempt from CloudSigma guarantees:

- Acts or omissions of you or your users;
- Software running within your virtual servers;
- Scheduled maintenance which we have announced at least 24 hours in advance;
- Factors outside our control, including but not limited to any force majeure events, failures, acts or omissions of our upstream providers or failures of the internet;

- Actions of third parties, including but not limited to security compromises, denial of service attacks and viruses provided CloudSigma makes reasonable efforts to keep its software and systems up to date;
- Violations of our Acceptable Use Policy;
- Any product currently in Beta as per our Terms of Service;
- Law enforcement activity.

A user must be up to date with all payments and have sufficient pre-pay balance where appropriate to cover current usage levels to be eligible for the credits outlined in this Service Level Agreement. No credits will be extended if a user is delinquent on any payments or has insufficient balance to continue using CloudSigma services at usage levels during the qualifying claimed credit period for at least 10 calendar days.

The award of credit by CloudSigma to you as described in this Service Level Agreement will be the sole and exclusive remedy for unavailability or performance degradation of CloudSigma services. Credits will only be provided against future service and for the avoidance of doubt may not be exchanged for cash or other forms of payment.

Notwithstanding anything in this Service Level Agreement to the contrary, the maximum total credit for the monthly billing period, including all guaranties, shall not exceed 100% of your fee for the previous 30 calendar days. Credits beyond your fee for the previous 30 calendar days will not be carried forward for use against future fees.