Hyperconverged Cloud Architecture
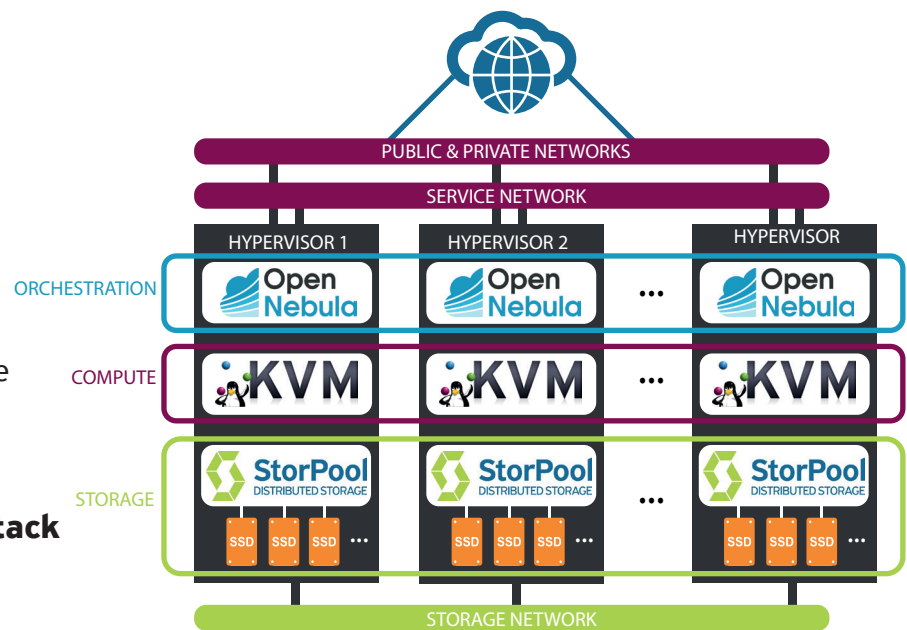
Open Nebula

StorPool
DISTRIBUTED STORAGE

OpenNebula is a turnkey, enterprise-ready solution that includes all the features needed to provide an on-premises (private) cloud or public cloud services. With tens of thousands of deployments, OpenNebula is powering industry leaders like Akamai, CentOS, BBC, BlackBerry, etc. OpenNebula Systems has a global presence with offices in Europe and USA.

StorPool is a software-defined storage for building high-performance public and private clouds. It runs on standard servers, drives and network and turns them into an outstanding storage system. Compared to traditional SANs, all-flash arrays or other storage software StorPool is faster, more reliable and scalable.

Hyperconverged cloud

StorPool and OpenNebula joined efforts and technologies to help companies easily build a simple, efficient and high performance cloud. The Hyperconverged Cloud is using standard servers and network. Every server within the cloud combines computing, storage and virtual networking functions to achieve better efficiency and simplified management.

The final result is turnkey software stack for building public or private cloud.

PUBLIC & PRIVATE NETWORKS
SERVICE NETWORK

| HYPERVISOR 1 | HYPERVISOR 2 | HYPERVISOR |

ORCHESTRATION — Open Nebula ... Open Nebula ... Open Nebula

COMPUTE — KVM ... KVM ... KVM

STORAGE — StorPool DISTRIBUTED STORAGE ... StorPool DISTRIBUTED STORAGE ... StorPool DISTRIBUTED STORAGE

SSD SSD SSD ... SSD SSD SSD ... SSD SSD SSD ...

STORAGE NETWORK

Hyperconverged cloud benefits

Better performance and reliability

Scalability

Optimized infrastructure costs

Simplified hardware inventory

Simplified management

Cloud orchestration

Open Nebula

All resources of the hyperconverged cloud are orchestrated by OpenNebula. It is responsible for creating and managing VMs, networks,as well as the management of the hardware resources and the storage. It provides a seperate, account-based GUIs for the cloud administrator and for the end-users.

Summary of the implementation

Operating system	Supported host OS (Ubuntu, CentOS or RHEL) on all hosts
Hypervisor	KVM
Networking	Redundant 10Gbps or faster storage network; 1Gbps or 10Gbps network shared by service network, public and private virtual networks
Storage	StorPool Storage
Authentication	Native authentication or Active Directory

Hyperconverged nodes

In hyperconverged cloud architecture, each node implements all cloud functions - Virtualized Compute, Storage and Networking.
This provides high level of availability with optimum resource utilization.

Compute

Compute layer is responsible for providing VMs with execution resources (e.g. CPU, memory, network access). Hypervisor in the architecture is KVM.
To keep the cloud simple and efficient, it is recommended to have a smaller number of nodes with a higher number of CPU cores and frequencies.

Storage

StorPool provides a shared storage pool combining all the available storage capacity and performance.
It is built using a distributed, shared-nothing architecture.
All functions are performed by all nodes on an equal peer basis.
All nodes are connected over a dedicated 10G/25G/40G Ethernet for minimum latency and maximum throughput.

Networking

Three networks are defined: instance networks (public, private) - providing connectivity to the VMs across different hosts; service network, used by OpenNebula front end to manage hypervisors and storage, and storage network used by StorPool.
All hardware nodes are using redundant physical connections to avoid single point of failure.

High Availability

High availability design is implemented in each plane of the Hyperconverged Cloud Architecture. In case of a failure of a hypervisor, all VMs are automatically relaunched on the remaining operational hypervisors. Planned maintenance operations can be performed without service downtime, by using live migration of the virtual machines off the hypervisor, before it is shut down or restarted.

Hyperconverged cloud combines the reliability and flexibility of SAN-based cloud with the simplicity and speed of local storage. It is fast to deploy, cost efficient and easy to manage.
Your fastest way to launch a cloud!

Download the full Hyperconverged Cloud Reference Architecture at
https://opennebula.org/ecosystem/storpool/
https://storpool.com/opennebula

Open Nebula
www.opennebula.org
contact@opennebula.org

StorPool DISTRIBUTED STORAGE
www.storpool.com
info@storpool.com