

# Hálózatba kapcsolt erőforrás platformok és alkalmazásaik

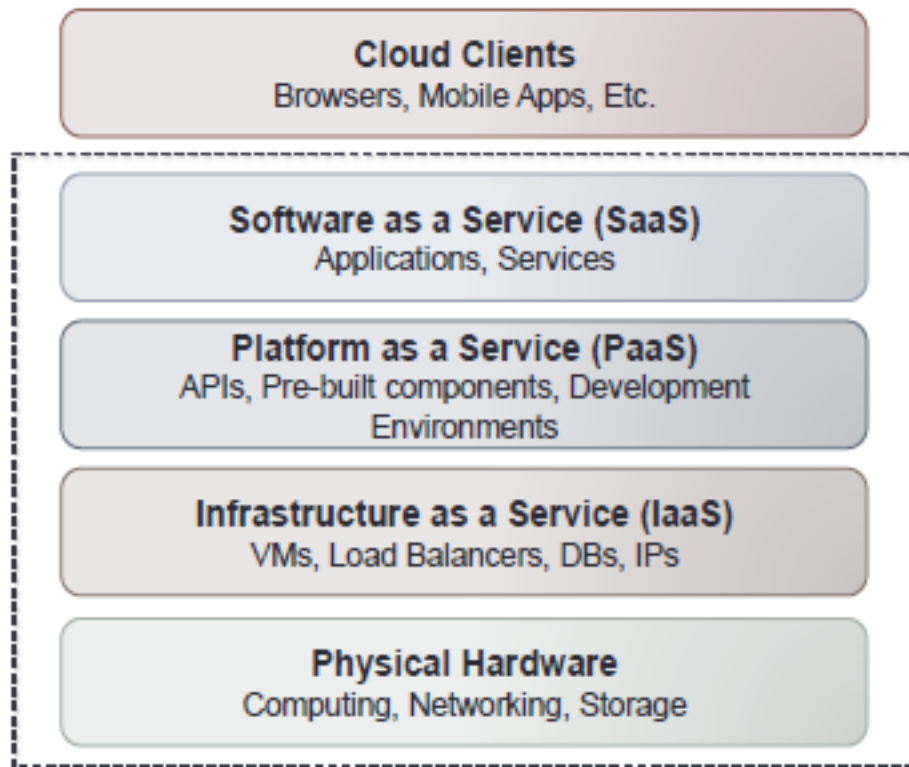
Maliosz Markosz

TMIT

2018

# Cloud Computing

- Layered model
- Offload management and maintenance tasks
- Service rather than a product
- Billing based on time and utility



# OpenStack

# OpenStack

- Collection of open source software for building private and public IaaS clouds.
- IaaS Cloud Services allows users to manage:
  - VMs, Virtual networks, storage resources
- OpenStack is a cloud platform that controls large pools of compute, storage, and networking resources throughout a datacenter.
- Provides on-demand increase or decrease of resource allocation
- Designed to run on commodity hardware
  - runs on common Linux platforms such as RHEL, SUSE, or Ubuntu
- Deliver Self-Service IT Rapidly and At Scale
  - via command-line interface (CLI) or a web-based dashboard

# Automation and Orchestration of Resources

A new management layer that adds automation and control



Connects to apps via APIs

Self-service Portals for users



USERS




ADMINS



CLOUD OPERATING SYSTEM

Creates Pools of Resources



Automates The Network

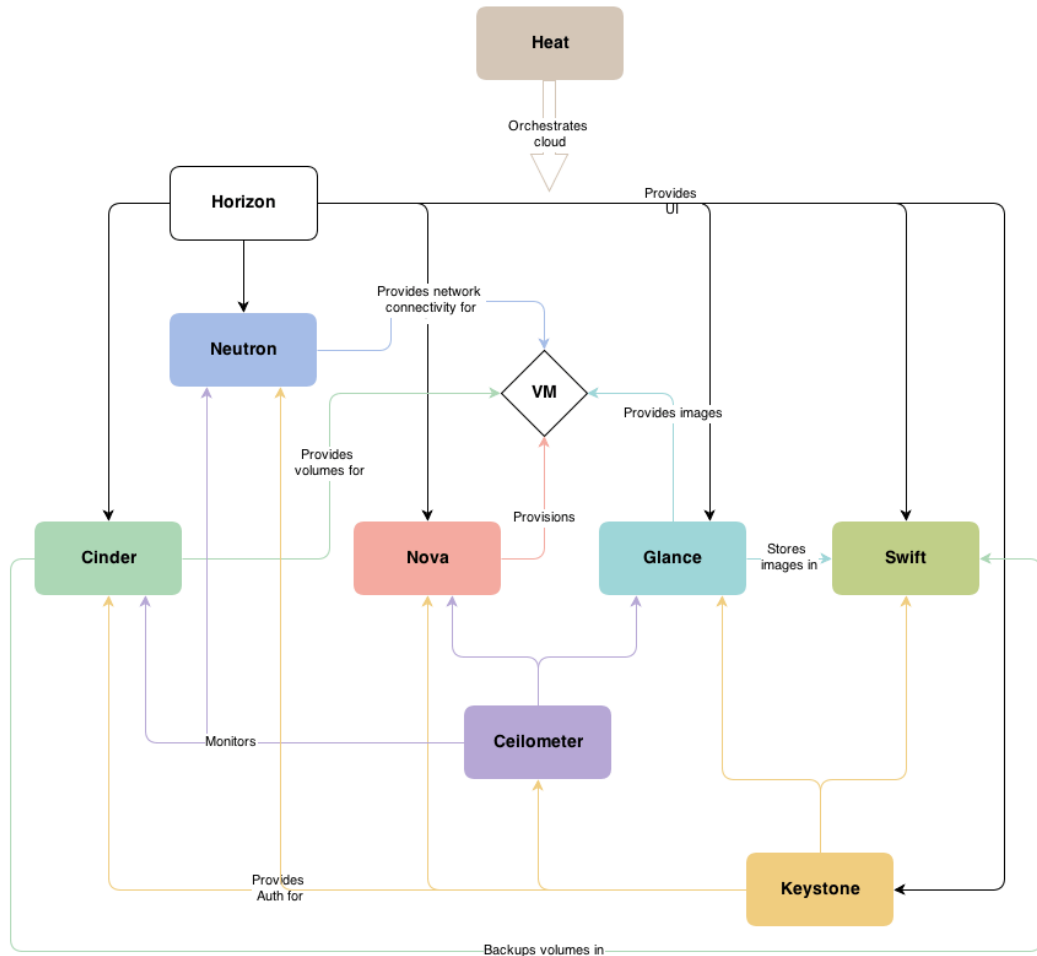


# Components

Service	Project name	Description
Dashboard	Horizon	Provides a web-based self-service portal to interact with underlying OpenStack services, such as launching an instance, assigning IP addresses and configuring access controls.
Compute	Nova	Manages the lifecycle of compute instances in an OpenStack environment. Responsibilities include spawning, scheduling and decommissioning of virtual machines on demand.
Networking	Neutron	Enables network connectivity as a service for other OpenStack services, such as OpenStack Compute. Provides an API for users to define networks and the attachments into them. Has a pluggable architecture that supports many popular networking vendors and technologies.
Storage		
Object Storage	Swift	Stores and retrieves arbitrary unstructured data objects via a RESTful, HTTP based API. It is highly fault tolerant with its data replication and scale out architecture. Its implementation is not like a file server with mountable directories.
Block Storage	Cinder	Provides persistent block storage to running instances. Its pluggable driver architecture facilitates the creation and management of block storage devices.
Shared services		
Identity service	Keystone	Provides an authentication and authorization service for other OpenStack services. Provides a catalog of endpoints for all OpenStack services.
Image Service	Glance	Stores and retrieves virtual machine disk images. OpenStack Compute makes use of this during instance provisioning.
Telemetry	Ceilometer	Monitors and meters the OpenStack cloud for billing, benchmarking, scalability, and statistical purposes.
Higher-level services		
Orchestration	Heat	Orchestrates multiple composite cloud applications by using either the native HOT template format or the AWS CloudFormation template format, through both an OpenStack-native REST API and a CloudFormation-compatible Query API.
Database Service	Trove	Provides scalable and reliable Cloud Database-as-a-Service functionality for both relational and non-relational database engines.

# Architecture

- Main node types (logical components)
  - Control
  - Network
  - Compute
- built as a set of distributed services
- services communicate with each other



# Pluggable Framework

- allows vendors to write plug-ins to implement a solution using their own technology
  - and which allows users to integrate their technology of choice
- Supported Hypervisors
  - <https://docs.openstack.org/developer/nova/support-matrix.html>



# Why use OpenStack and not just hypervisor?

- A hypervisor abstract the resource from the physical hardware
- OpenStack pushes this one step further by providing an elastic, self-service, and measurable infrastructure for managing a pool of compute, storage, and networking resources.
- With traditional virtualization technology, the state of the virtual machine is persistent
  - OpenStack can support both persistent and ephemeral models
    - Ephemeral: When an instance is terminated, the original image remains intact, but the state of the terminated instance is not retained.

# OpenStack Overview

- Tenant (CLI) = Project (dashboard)
- End User Guide: <https://docs.openstack.org/horizon/pike/user/>
  - Log in to the dashboard
  - Upload and manage images
  - Configure access and security for instances
  - Launch and manage instances
  - Create and manage networks
  - Create and manage object containers
  - Create and manage volumes
  - Launch and manage stacks
  - Create and manage databases
  - View and manage load balancers v2

# Launch a VM Instance

- Parameters
  - VM name
  - Image (OS type)
  - Flavor (CPU, memory, disk)
  - Network
  - Optional
    - SSH keypair, persistent volumes, etc.

# Instance-specific information

- Post-boot configuration facility: cloud-init
  - a script that connects to metadata service
- Metadata service
  - via a server on a link-local address (169.254.0.0/16)
    - To test: `curl http://169.254.169.254/latest/meta-data/`
- Configuration information, e.g.
  - IP address
  - Hostname
  - SSH public key
  - cloud-init script
  - user-data
  - Static routing information