


Introduction to AWS – OpenStack

Dr. Assane Wade
2023 – 2024

Introduction to *AWS*

- Amazon Web Services (AWS) is a collection of remote infrastructure services.
- Introduced in 2006/2007, it is considered the first real cloud computing offering.
- In IaaS the main services offered by AWS are
 - **Compute** — for example *Elastic Compute Cloud*
 - **Storage** — for example *Simple Storage Service*
 - **Database** — for example *Relational Database Service*
 - **Networking** — for example *Virtual Private Cloud*
 -

- What was Amazon's motivation in creating AWS?
 - Typical weekly traffic to Amazon's e-commerce web site in 2007



Amazon Web Services - Motivation

- What was Amazon's motivation in creating AWS?
 - Typical weekly traffic to Amazon's e-commerce web site in 2007



Amazon Web Services - Motivation

Traffic in the month of November 2007



Amazon Web Services - Motivation



November Traffic to Amazon.com



Families of services

- Compute
- Storage
- Database
- Migration & transfer
- Networking & Content Delivery
- Developer Tools
- Management & Governance
- Media Service
- Machine Learning
- Analytics
- Security, Identity & Compliance
- AWS Cost Management
- Mobile
- AR & VR
- Application Integration
- Customer Engagement
- Business Applications
- End User Computing
- Internet of Things
- Game Development
- ...

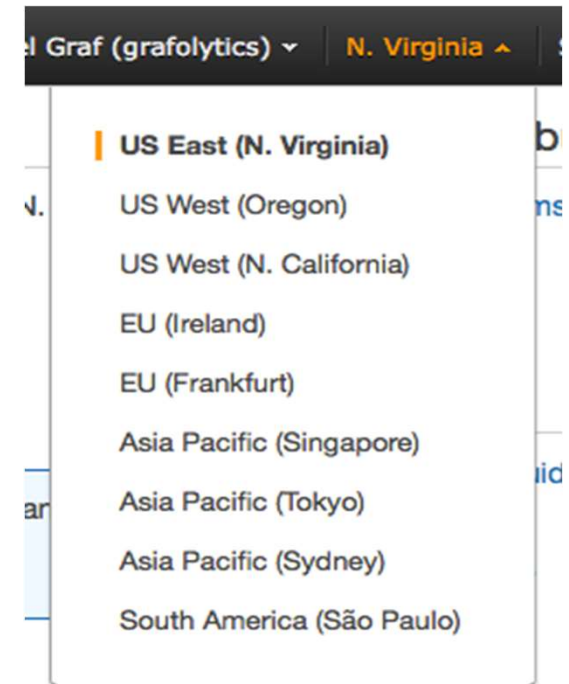
AWS - Datacenters

Amazon Web Services are currently available in 32 regions, 102 availability zones. 12 planned Availability zones and 4 new regions (Canada, Malaysia, New Zealand, and Thailand.)



Regions & Availability zones

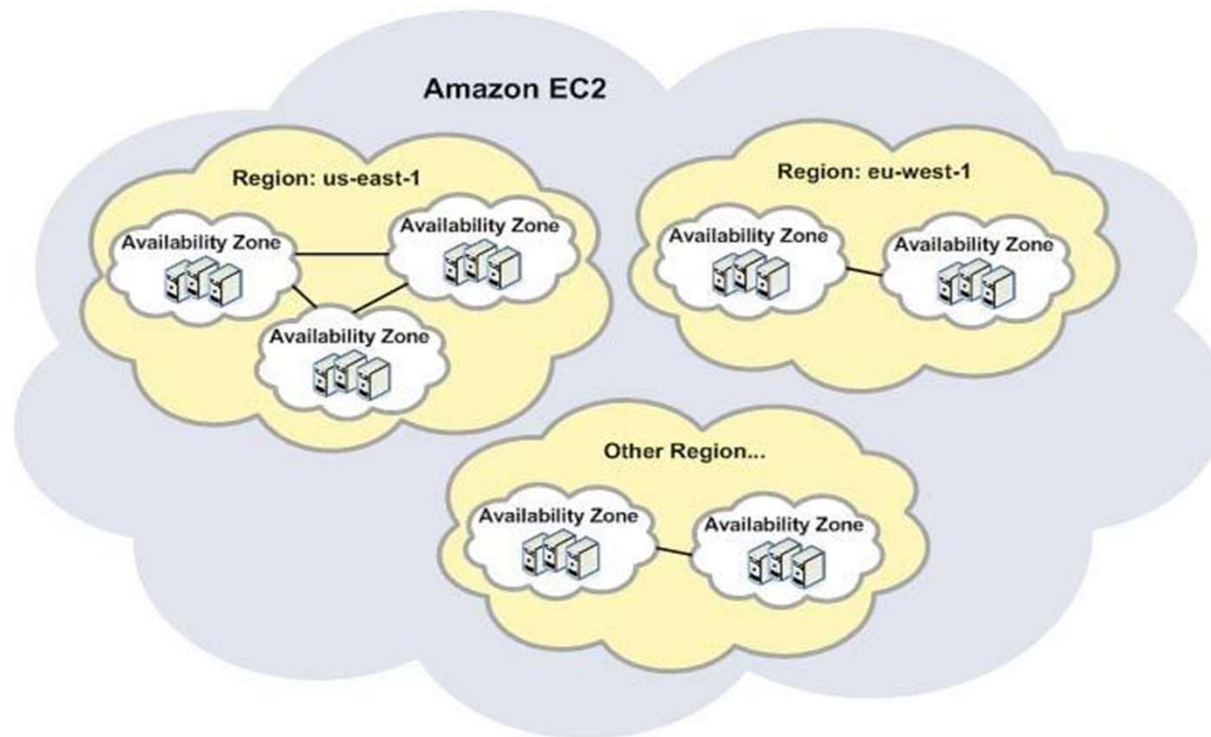
- When allocating a cloud resource (e.g., a virtual machine), the AWS customer can choose the *region*, and within the region, the *availability zone*.
- **Regions** are distributed globally and enable a developer to place his application and/or data in a particular country / region ...
 - to be closer to his customers,
 - so that the data resides in a particular jurisdiction to be compliant with regulations (e.g., data privacy laws).
- **Availability Zones** are separate datacenters inside a Region
 - Each datacenter has its own independent infrastructure for power, cooling.
 - Availability Zones in a region are connected with low-latency network links.



Availability Zone Status:

- ✓ us-east-1a:
Availability zone is operating normally
- ✓ us-east-1b:
Availability zone is operating normally

Regions & Availability zones



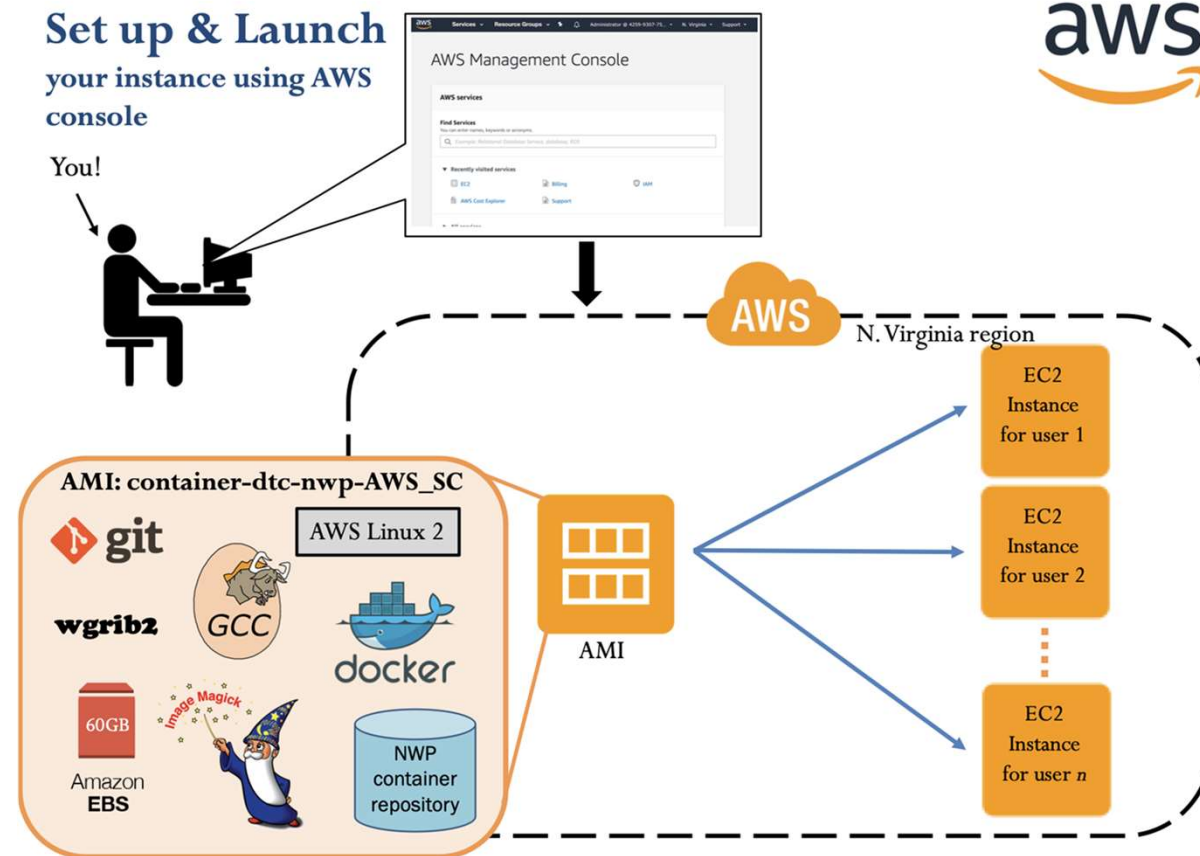
Regions & Availability zones

- Amazon *Elastic Compute Cloud* (EC2) is a web service that provides resizable compute capacity in the cloud.
- Publicly available since 2006
- Is considered the first real cloud computing product.
- Developers can rent virtual machines (called *EC2 instances*).
- Many *instance types* are available.



Elastic Computer (EC2)

- Amazon EC2 presents a virtual computing environment, allowing you to:
 - Use a web interface to launch instances with a variety of operating systems
 - which are bundled into *Amazon Machine Images (AMI)*.
 - Load your instances with your custom application environment .
 - Manage your network's access permissions.



- **On-Demand Instances**

- Pay-by-the hour
- Start and stop as you wish

- **Reserved Instances**

- Pay a yearly upfront fee and receive a discount on the hourly charge
- Amazon EC2 Reserved Instances (RI) provide a significant discount (up to 72%) compared to On-Demand pricing and provide a capacity reservation when used in a specific Availability Zone

- **Savings Plans**

- Savings Plans is a flexible pricing model that can help you reduce your bill by up to 72% compared to On-Demand prices, in exchange for a one- or three-year hourly spend commitment

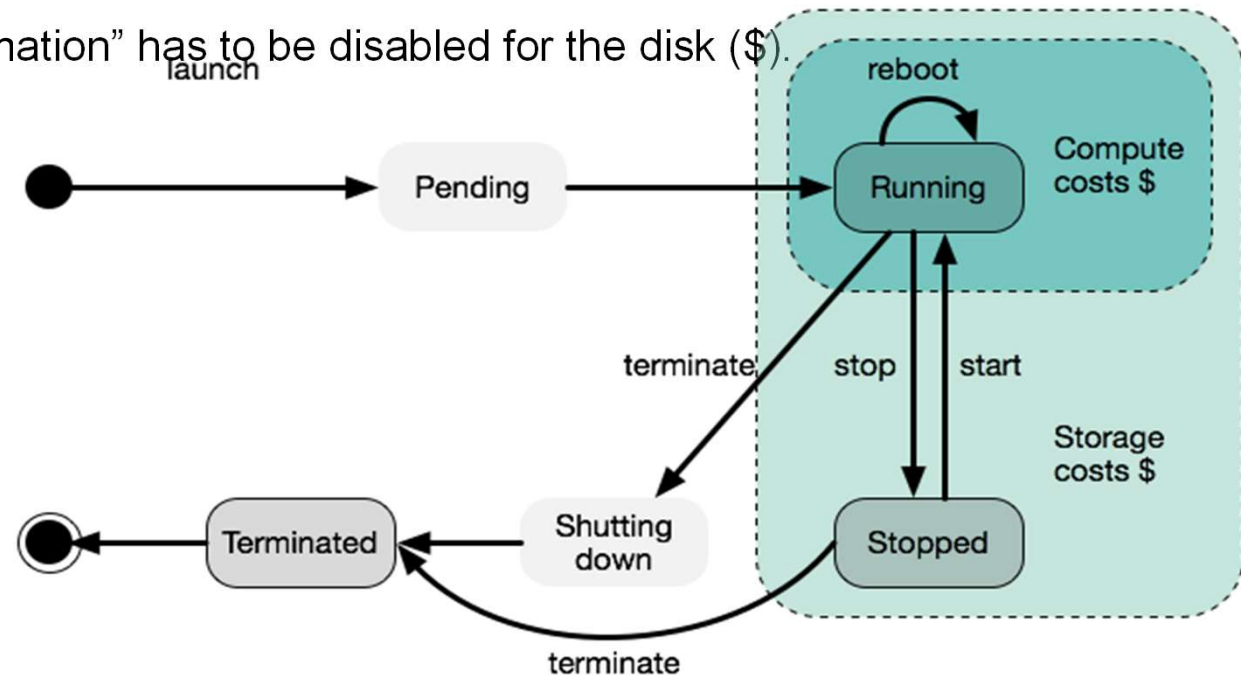
- **Spot Instances**

- Bid for unused EC2 capacity
- Mention your Spot Price and if the market rate is less than your Bid, you get your instance
- Instance automatically terminates if your Spot Price becomes less than the current market rate

- cPU Power
 - Measured in *Elastic Compute Unit* (ECU) – Defined by Amazon as the equivalent CPU capacity of a 1.0-1.2 GHz 2007 Opteron/Zeon processor
- Memory
 - Measured in GiB
- Disk performance
 - Possibility to provision a guaranteed number of Input/Output Operations per Second (IOPS)
- Network performance

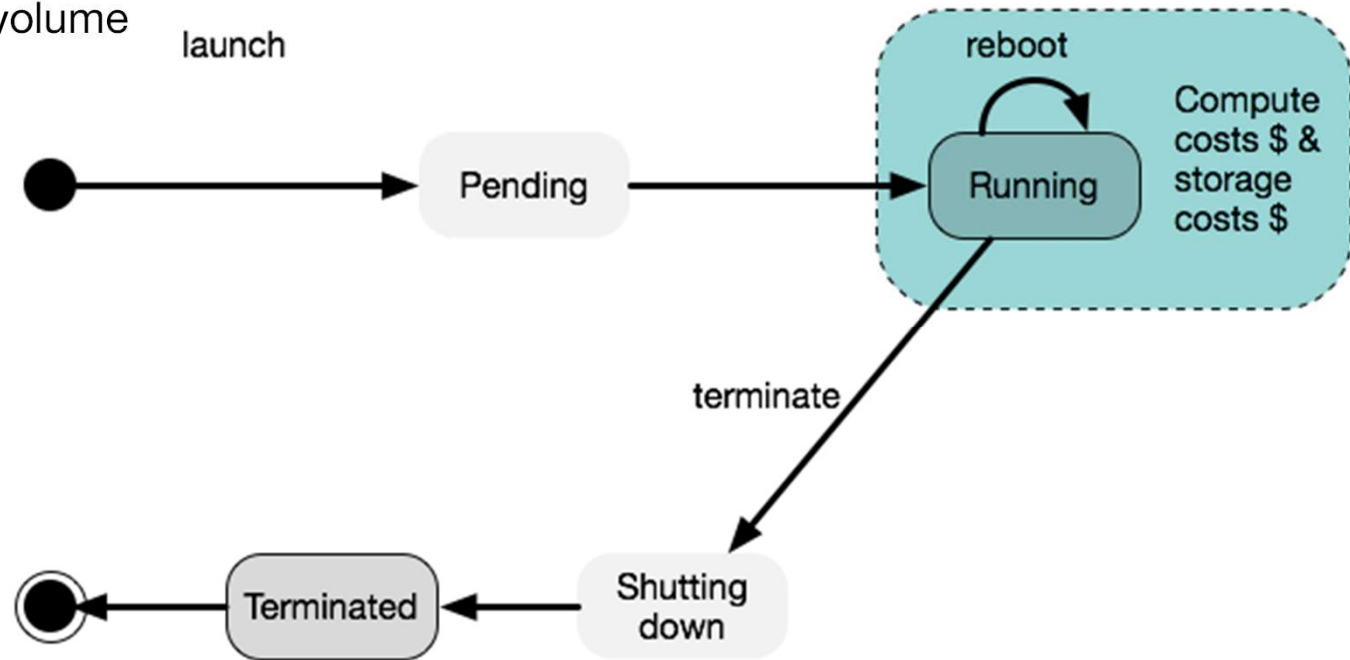
EC2 instance life cycle

- When the instance is terminated, the disk is deallocated as well and the data is lost.
- EC2 Instance with an EBS Volume with “Delete on Termination” enabled (this is the default case)
- To keep the data, “Delete on Termination” has to be disabled for the disk (\$).



EC2 instance life cycle

- EC2 Instance with an instance store volume



- When the instance is terminated, the disk is deallocated as well and the data is lost.
- In this case there is no way to keep the data beyond the lifetime of the instance.

Object Storage: Simple Storage Service (S3)

- Amazon Simple Storage Service (S3) is a high-performance, highly-available web-oriented storage service that supports very large files.
- You can write, read and delete *objects* (files) into S3 containing from 1 byte to 5 TB of data. The number of objects that can be stored is unlimited.

The screenshot displays the AWS S3 console interface. At the top, there's a navigation bar with 'Services', 'Edit', and user information 'Marcel Graf', 'Global', and 'Help'. Below this, there are buttons for 'Upload', 'Create Folder', and 'Actions'. The breadcrumb path is 'Buckets / heigvd-cld / assets'. A table lists the contents of the 'assets' folder:

	Name	Storage Class	Size
<input type="checkbox"/>	bg_logo.png	Standard	5.2 KB
<input type="checkbox"/>	bg_menu_sprite_standard.png	Standard	281 bytes
<input checked="" type="checkbox"/>	breadcrumbs-center-bg.png	Standard	3.3 KB
<input type="checkbox"/>	breadcrumbs-left-bg.png	Standard	3.1 KB
<input type="checkbox"/>	breadcrumbs-right-bg.png	Standard	2.9 KB

On the right, the properties for the selected object 'breadcrumbs-center-bg.png' are shown:

- Bucket: heigvd-cld
- Folder: assets
- Name: breadcrumbs-center-bg.png
- Link: <https://s3-eu-west-1.amazonaws.com/heigvd-cld/assets/breadcrumbs-center-bg.png>
- Size: 3.3 KB
- Last Modified: Mon Mar 04 21:50:46 GMT+100 2013
- Owner: Me
- ETag: 95f52e7e8b394b3121293305bba4348b
- Expiry Date: None
- Expiration Rule: N/A


- Each object is stored in a *bucket*
- A bucket is a container for objects. You can store any number of objects in a bucket.
- If the object named photos/puppy.jpg is stored in the DOC-EXAMPLE-BUCKET bucket in the US West (Oregon) Region, then it is addressable using the URL `https://DOC-EXAMPLE-BUCKET.s3.us-west-2.amazonaws.com/photos/puppy.jpg`.

- Data kept on object storage devices are accessed directly through APIs or http/https.
- Data : photos, videos, and log files.
- The object store guarantees that the data will not be lost.
- Object storage data can be replicated across different data centers and offer simple web services interfaces for access.

When to use the block storage?


- When Strong consistency is needed:
 - real-time systems such as transactional databases that are constantly being written to.
 - a read request must return the most updated version of the data.
- When scalability is not an issue
- When scalability becomes more difficult within a geographically distributed system.

When to use Object storage?

- Solving the increasing problem of data growth
- Solving the provisioning management issues: Web content, data backup, and archives are good use cases
- metadata functionality, facilitate this ease of use.
- Resiliency
 - at least three copies of every file are stored
 - 
 The distributed storage design for high availability allows less-expensive commodity hardware to be used
 - The downside is that there is no guarantee that a read request returns the most recent version of the data.

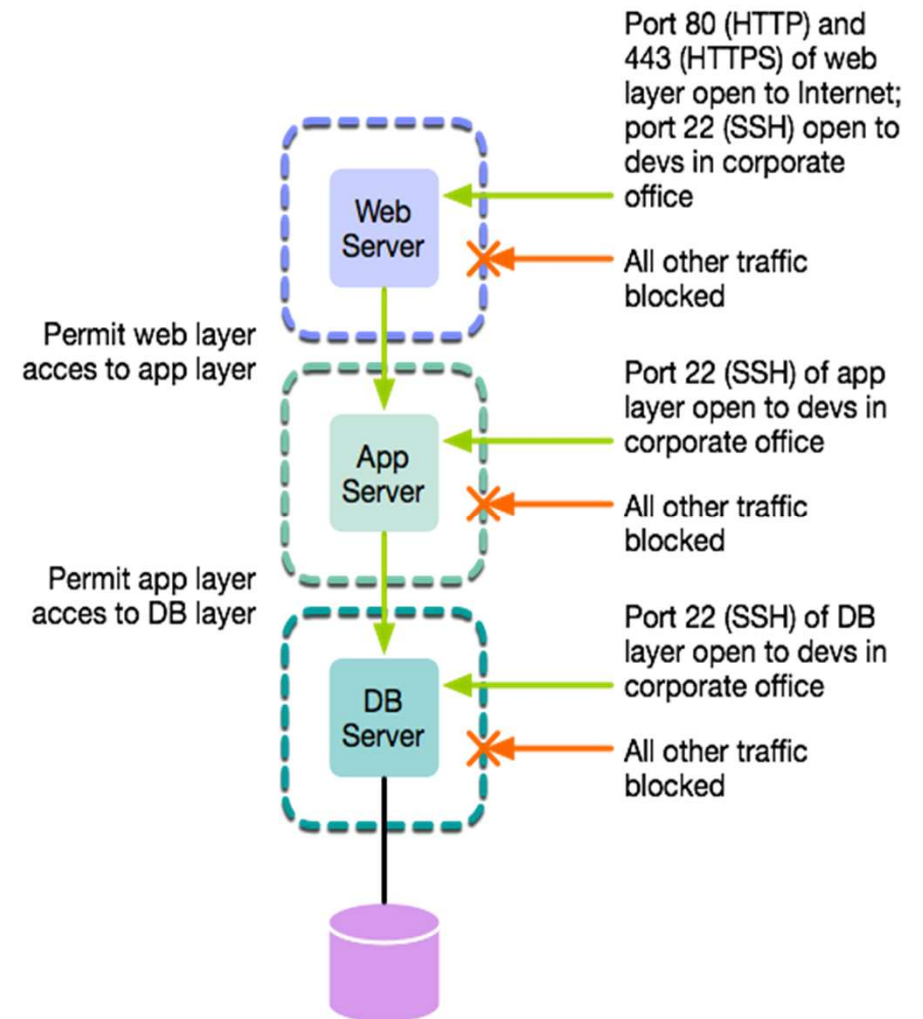
Security Group

- An EC2 instance has a public IP address and is therefore accessible from the Internet.
 - This is a security risk.
- Each virtual machine comes with a mandatory virtual firewall.
- The configuration of a firewall is called *Security Group*.
 - A configuration can be reused for several firewalls.
- Firewall performs inbound and outbound filtering.
 - Based on protocols / port numbers
 - Based on IP addresses
- To be able to log into a Linux instance, port 22 (SSH) has to be open.



Security Group Configuration

- Three-tier web application
- Restrictive configuration
- Distinguish access from
 - open Internet: HTTP and HTTPS open
 - corporate office (developers): SSH open
 - another layer of the web application: specific ports open



EC2 Elastic IP address

- An EC2 Elastic IP Address is a fixed IP address.
- It can be assigned to any instance in a region.
- Allows to keep the address
 - when stopping / starting an instance,
 - when changing the instance.

Allocate New Address



Are you sure you want to allocate a new IP address?

EIP used in: EC2

Cancel

Yes, Allocate

Associate Address

Select the instance to which you wish to associate this IP address (54.204.28.247)

Instance



Warning

If you associate an Elastic IP address with your instance, your current public IP addresses.

AWS Terminology

AWS term	Generic term	AWS term	Generic term
EC2	IaaS offering	Elastic Block Store (EBS) volume	Virtual disk on a SAN
EC2 Instance	Virtual machine	Instance store volume	Virtual disk co-located with virtual machine
Amazon Machine Image (AMI)	Virtual machine image	S3	Object storage service
Security Group	Firewall configuration	Cloud Watch	Monitoring service
Elastic IP address	Static external IP address		
Volume	Virtual disk		

Introduction to OpenStack



What is OpenStack ?

- Free and open-source cloud-computing software platform.
- Provides services for managing a Cloud environment on the fly.
- Consists of a group of interrelated projects that control pools of processing, storage, networking resources, authentication, storage, etc.
- Initially designed to provide services for an IaaS
- Today, OpenStack provides “high level” services



CANONICAL

- OpenStack foundation serves more than 60,000 Individual Members from over 180 countries around the world.

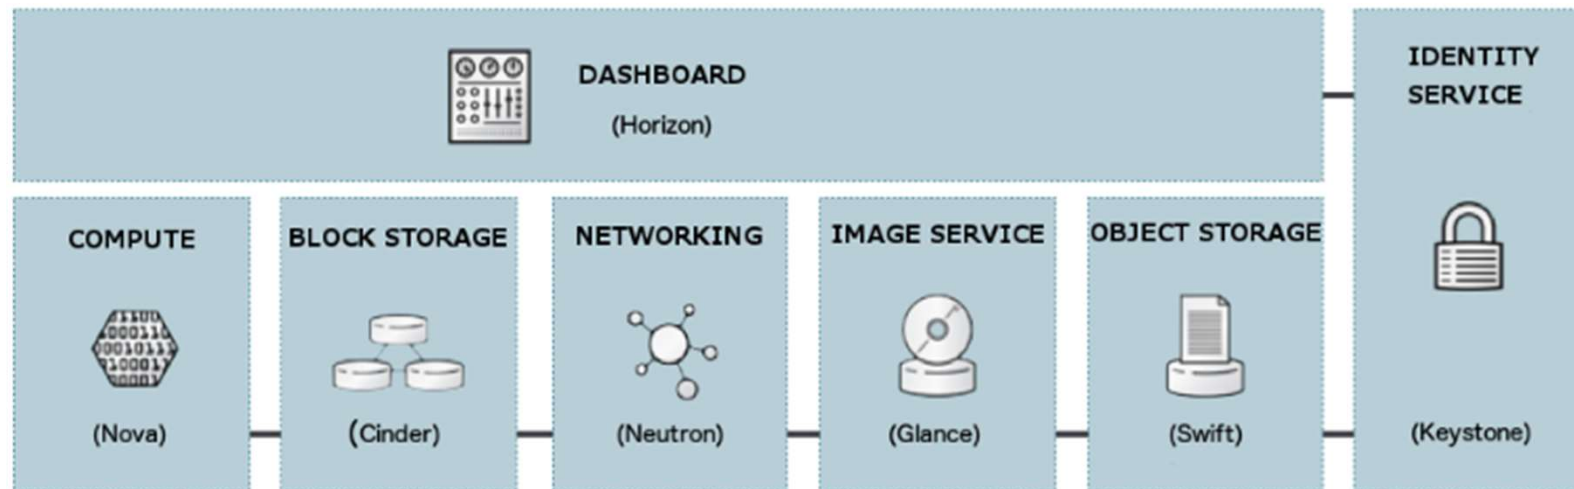


OpenStack releases

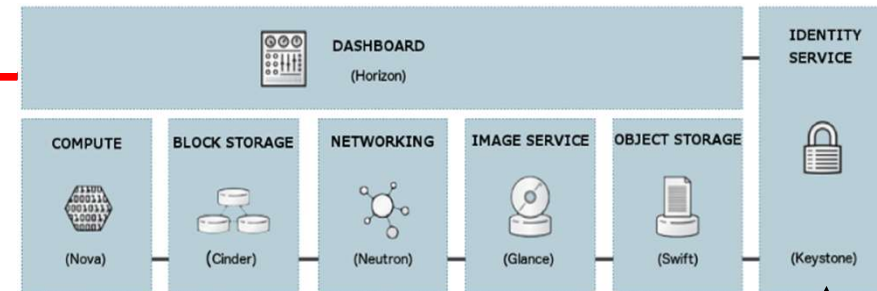
<https://releases.openstack.org/>

Series	Status	Initial Release Date	Next Phase	EOL Date
Yoga	Future	2022-03-30 <i>estimated (schedule)</i>	Development <i>estimated 2021-10-06</i>	
Xena	Development	2021-10-06 <i>estimated (schedule)</i>	Development <i>estimated 2021-04-14</i>	
Wallaby	Maintained	2021-04-14	Extended Maintenance <i>estimated 2022-10-14</i>	
Victoria	Maintained	2020-10-14	Extended Maintenance <i>estimated 2022-04-18</i>	
Ussuri	Maintained	2020-05-13	Extended Maintenance <i>estimated 2021-11-12</i>	
Train	Extended Maintenance (see note below)	2019-10-16	Unmaintained <i>TBD</i>	
Stein	Extended Maintenance (see note below)	2019-04-10	Unmaintained <i>TBD</i>	
Rocky	Extended Maintenance (see note below)	2018-08-30	Unmaintained <i>TBD</i>	
Queens	Extended Maintenance (see note below)	2018-02-28	Unmaintained <i>TBD</i>	
Pike	Extended Maintenance (see note below)	2017-08-30	Unmaintained <i>TBD</i>	
Ocata	Extended Maintenance (see note below)	2017-02-22	Unmaintained <i>estimated 2020-06-04</i>	
Newton	End Of Life	2016-10-06		2017-10-25
Mitaka	End Of Life	2016-04-07		2017-04-10
Liberty	End Of Life	2015-10-15		2016-11-17
Kilo	End Of Life	2015-04-30		2016-05-02
Juno	End Of Life	2014-10-16		2015-12-07
Icehouse	End Of Life	2014-04-17		2015-07-02
Havana	End Of Life	2013-10-17		2014-09-30
Grizzly	End Of Life	2013-04-04		2014-03-29
Folsom	End Of Life	2012-09-27		2013-11-19
Essex	End Of Life	2012-04-05		2013-05-06
Diablo	End Of Life	2011-09-22		2013-05-06
Cactus	End Of Life	2011-04-15		
Bexar	End Of Life	2011-02-03		
Austin	End Of Life	2010-10-21		

Basic services



Identity service : Keystone

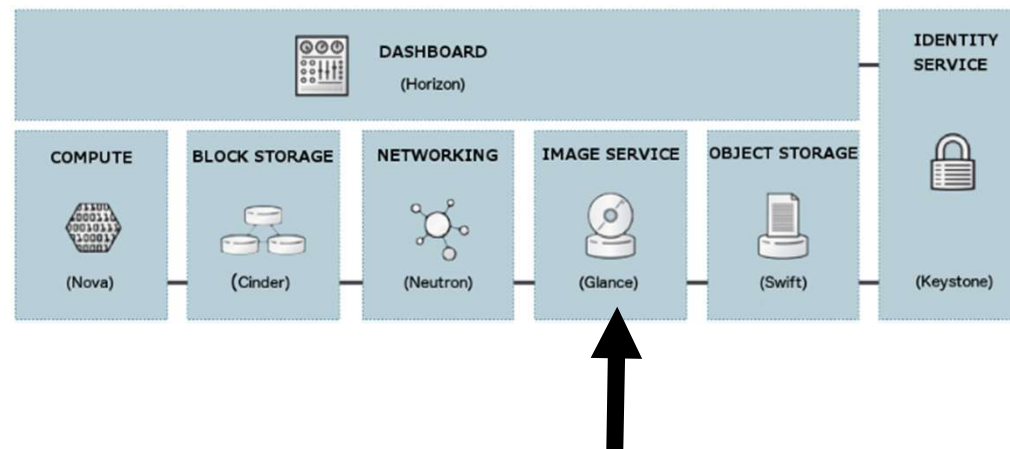


- **Keystone:** provides identity services for OpenStack.
- A central list of users/permissions mapped against OpenStack services.
- Keystone is organized as a group of internal services exposed on one or many endpoints.
- The Identity service generates authentication tokens that permit access to the OpenStack services REST APIs.

-
- Clients obtain this token and the URL endpoints for other service APIs by supplying their valid credentials to Keystone.
 - Each time you make a REST API request to an OpenStack service, you supply your authentication token in the X-Auth-Token request header.

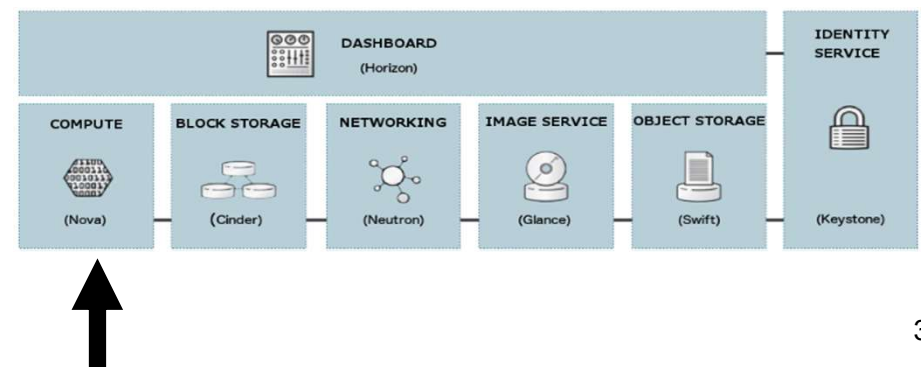
Image Service: Glance

- **Glance:** provides image services to OpenStack.
- "images" refers to images of hard disks. Used as templates for deploying new VMs.



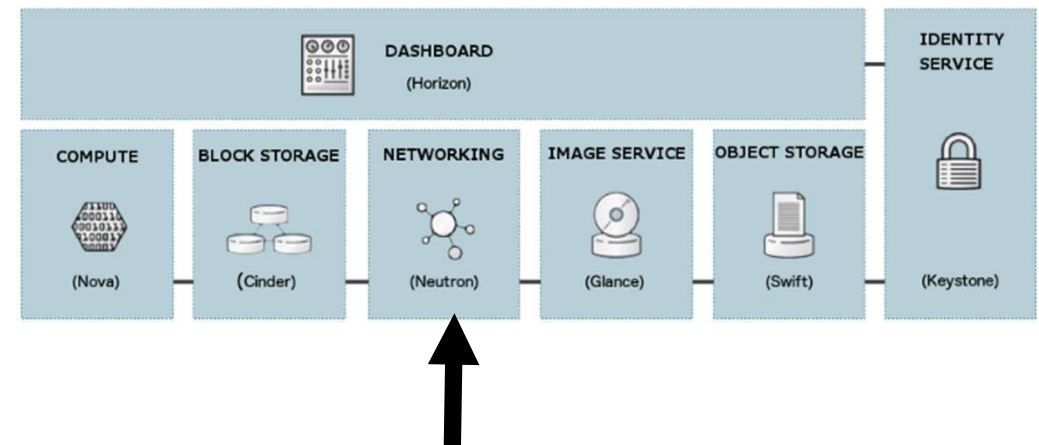
Compute service: OpenStack compute (nova)

- **OpenStack compute (nova):** Manages the lifecycle of compute instances in an OpenStack environment. Responsibilities include spawning, scheduling and decommissioning of virtual machines on demand.
- Tools using *nova*:
 - Horizon
 - OpenStack Client: Official CLI for OpenStack Projects. It includes not just nova commands but also commands for most of the projects in OpenStack.




Networking service: Neutron

- **Networking service (Neutron):** provides the networking capability for OpenStack. Enables Network-Connectivity-as-a-Service for other OpenStack services.
- OpenStack Networking enables projects to create virtual network topologies which may include services such as firewalls, load balancers and Virtual Private Network (VPN).
- Networking provides networks, subnets, and routers as object abstractions.

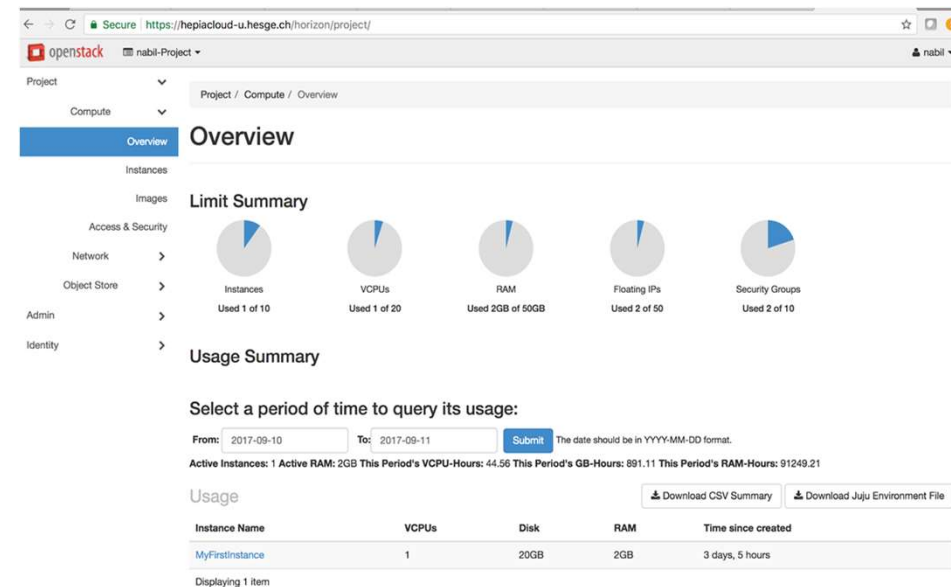


Networking service: Neutron



Dashboard: Horizon

- **Horizon:** The front office of OpenStack.
- The only native graphical interface to OpenStack.



- **Cinder:** block storage component, analogous to the traditional access on a disk drive.
- **Swift:** storage system for objects and files.
 - Users refer to a unique file identifiers: OpenStack decides where to store/back-up etc.

API Endpoints



MASTER OF SCIENCE
IN ENGINEERING

SWITCH [CH] | https://engines.switch.ch/horizon/project/api_access/

Anthony-Nabil • ZH ▾ nabil.abdennadher@hesge.ch ▾

View Credentials Download OpenStack RC File ▾

Network >

Orchestration > Displaying 14 items

Object Store >

Quickstart >

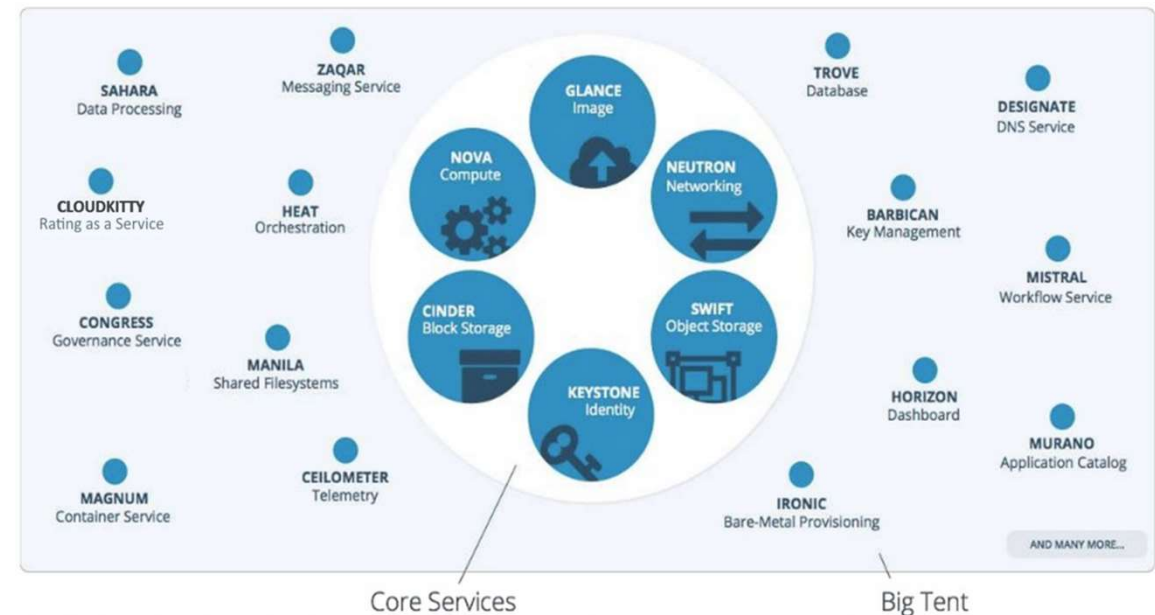
Identity >

Service	Service Endpoint
Cloudformation	https://api.zhdk.cloud.switch.ch:8000/v1
Compute	https://api.zhdk.cloud.switch.ch:8774/v2/cd8dd95f755744e79163571e2c5cfa3e
Dns	https://api.zhdk.cloud.switch.ch:9001
EC2	https://service-is-disabled
Identity	https://keystone.cloud.switch.ch:5000/v3
Image	https://api.zhdk.cloud.switch.ch:9292
Network	https://api.zhdk.cloud.switch.ch:9696
Object Store	https://os.zhdk.cloud.switch.ch/swift/v1
Orchestration	https://api.zhdk.cloud.switch.ch:8004/v1/cd8dd95f755744e79163571e2c5cfa3e
Placement	https://api.zhdk.cloud.switch.ch:8780/placement
S3	https://os.zhdk.cloud.switch.ch
Volume	https://api.zhdk.cloud.switch.ch:8776/v1/cd8dd95f755744e79163571e2c5cfa3e
Volumev2	https://api.zhdk.cloud.switch.ch:8776/v2/cd8dd95f755744e79163571e2c5cfa3e
Volumev3	https://api.zhdk.cloud.switch.ch:8776/v3/cd8dd95f755744e79163571e2c5cfa3e

Displaying 14 items

Services

- Compute
- Hardware Lifecycle
- Storage
- Networking
- Shared Services
- Orchestration
- Workload Provisioning
- Application Lifecycle
- API Proxies
- Web Frontend



<https://www.openstack.org/software/project-navigator/openstack-components#openstack-services>

Not all services are equally used or deployed!

