

# OPENSTACK: THE OPEN CLOUD

---

Anuj Sehgal (s.anuj@jacobs-university.de)

AIMS 2012 Labs

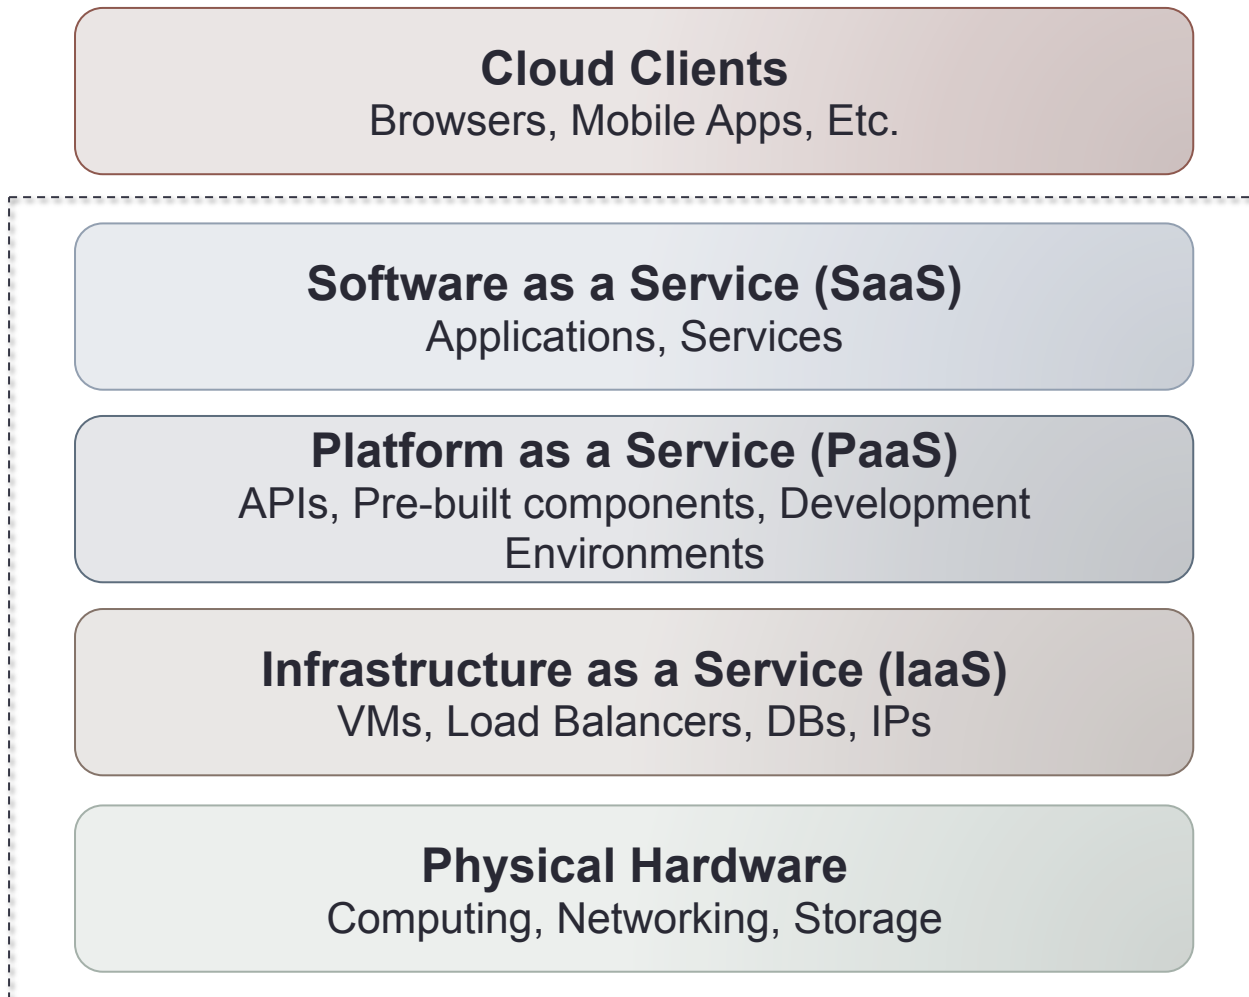
*04 June 2012*

# Outline

- What is the cloud?
- Background
- Architecture
- OpenStack Nova
- OpenStack Glance

# What is the Cloud?

- Cloud computing is the paradigm shift towards providing computing as a service.



# Background

- Amazon EC2 and S3 initial movers.
- Rackspace countered with their solution.
- Eucalyptus was the most popular open source IaaS platform.
  - NASA major supporter


# What went wrong?



- NASA Nebula Project
  - Required massive scalability (1 million machines and 60 million VMs)
- Scalability of Eucalyptus was insufficient
- Eucalyptus had closed source core
  - NASA could not contribute modules

# Birth of OpenStack

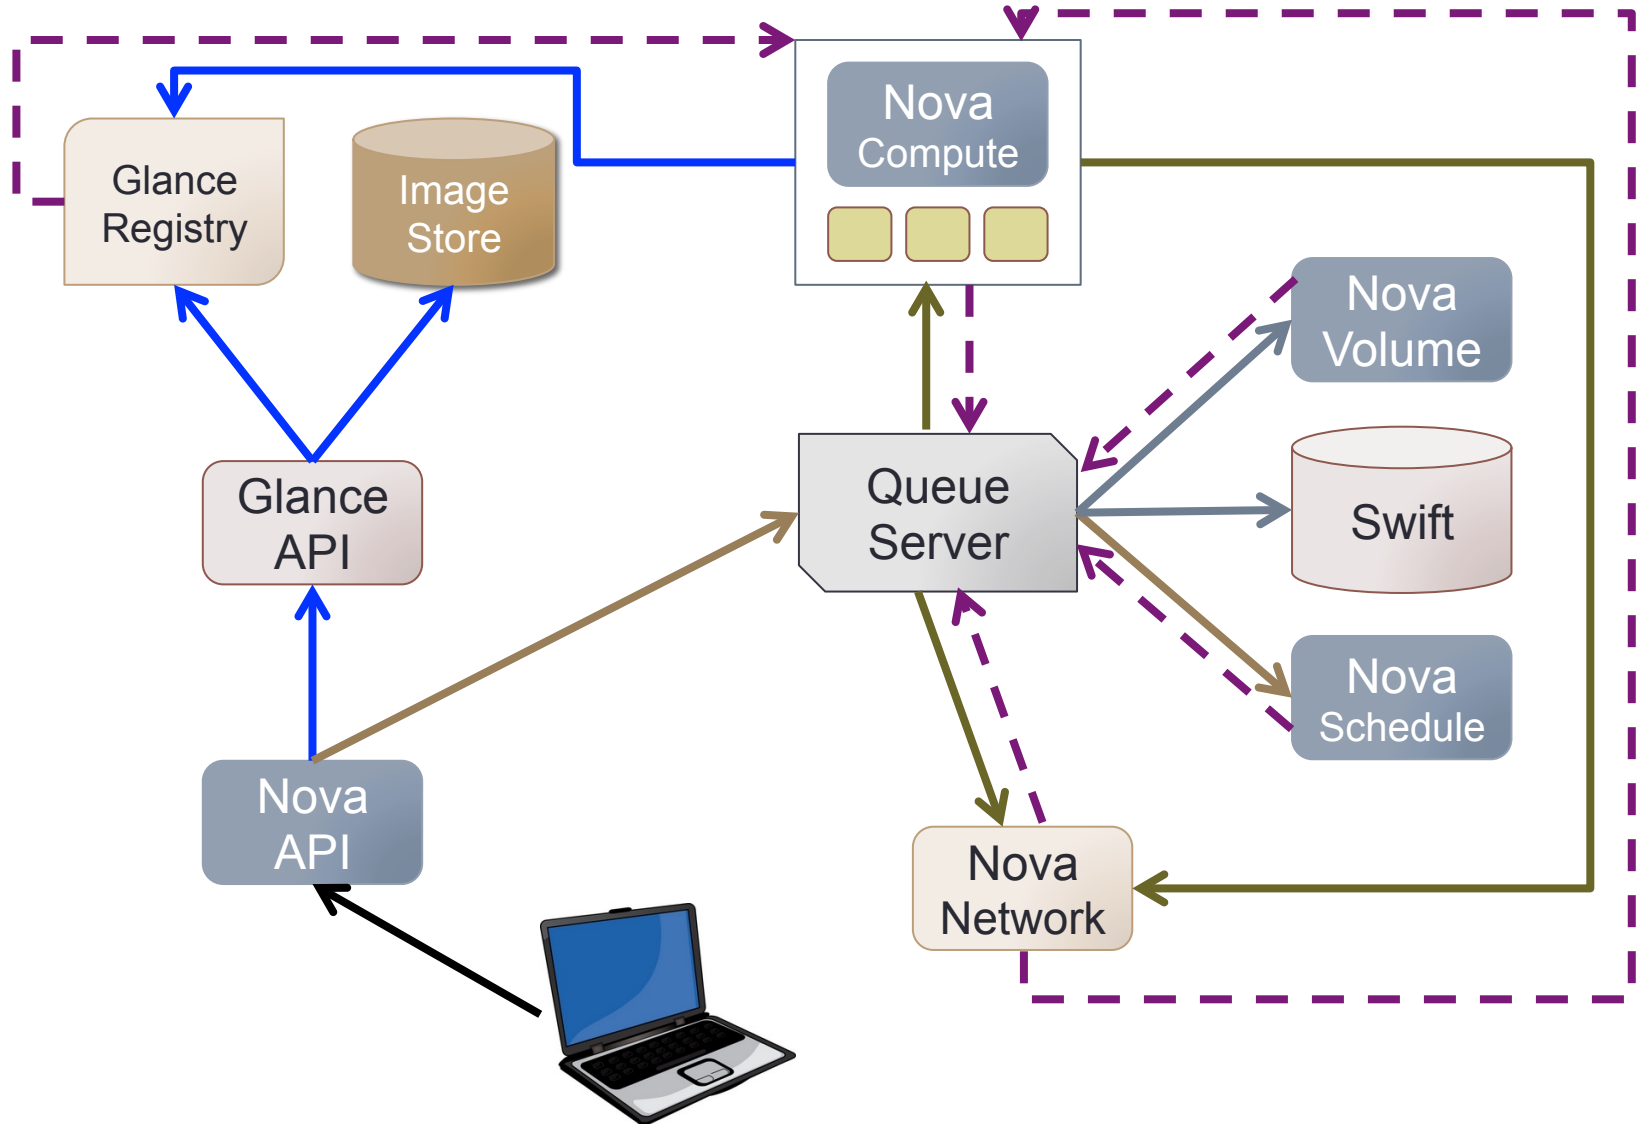
- OpenStack became a collaboration of Rackspace and NASA.
- Rackspace *Ozone* cloud controller.
- NASA *Nova* cloud computing from Nebula.
- Rackspace *Cloud Files* storage engine (*swift*).



# OpenStack Components

- Compute (Nova)
  - Used to orchestrate, manage and offer virtual machines upon many hypervisors.
  - Support for Xen, XenServer/XCP, KVM, UML, VMware vSphere and Hyper-V.
  - Analogous to the Amazon Elastic Compute Cloud (EC2).
- Object Storage (Swift)
  - Provides redundant storage for static objects.
  - Scalable to massive data sizes. Theoretically infinite storage.
- Image Service (Glance)
  - Provides storage for virtual disk, kernel and machine images.
  - Also provides image registration and querying services.
  - Accepts images in many formats, including AMI, AKI and ARI.

# OpenStack Architecture





# OpenStack Nova

- Is the heart of OpenStack
  - scheduler, networking, virtual machines, APIs and etc.
- **nova-api** provides an endpoint for management
  - Handles all API queries (Amazon EC2 and OpenStack)
  - Initiates all the orchestration of activities through message queue
  - Enforces policies
- **nova-schedule** decides where a new VM is instantiated
  - Maps API calls to appropriate OpenStack components
  - Picks servers from a pool
  - Decides based on load, memory, physical distance, CPU architecture and etc.
- **rabbitmq-server** provides the message queuing services
  - The API talks to the message queue
  - Facilitates an asynchronous request-response architecture via callbacks

# OpenStack Nova

- Three primary infrastructure modules
  - **nova-compute** creates and terminates VM instances, supports multiple hypervisors (but single one within a cloud)
  - **nova-volume** manages creation, attaching and detaching of persistent volumes to VMs (uses iSCSI)
  - **nova-network** manipulates the network by configuring VLANs, changing firewall rules, applying private/public IPs
- An SQL database (MySQL, PostgreSQL, SQLite) is used to keep track of runtime state
- Libvirt is used to manage the VM instances
- Written almost entirely in Python

# OpenStack Glance


- Provides the registration and management of VM images
- **glance** provides a list of available VM images to instantiate
  - images can be registered using EC2 APIs
  - uses filesystem or Swift for storage of images
- Requires kernel, ramdisk and filesystem images, which are provided to *nova-compute* for instantiation
- Multiple possible flavors for instantiation can be registered
  - CPUs, root-fs size, RAM and etc.
- Uses SQL to keep track of all state information

# Highly Distributed

- Single Node Architecture

## Cloud Controller

nova-api  
nova-network  
nova-scheduler



## Message Queue

rabbitmq-server

## Image Host

glance

## Compute Node

nova-compute

## Volume Node

nova-volume

# Highly Distributed

- Multi Node Architecture

## Cloud Controller



nova-api  
nova-network  
nova-scheduler

## Message Queue


rabbitmq-server

## Image Host

glance



**Compute Node**  
nova-compute



**Volume Node**  
nova-volume

# Highly Distributed

- Multi Node Architecture

## Cloud Controller

nova-api

nova-network

nova-scheduler

## Message Queue



rabbitmq-server

## Image Host

glance


## Volume Node

nova-volume



## Compute Node

nova-compute



## Compute Node

nova-compute

# Exercises

- Single Node Scenario
  - Ubuntu 11.10 Server base OS
  - RabbitMQ
  - Glance
  - Cloud Controller
  - Volume Host
  - Compute Host

# Exercises

- Multi Node Scenario (Time permitting/Homework/Demo)
  - Ubuntu 11.10 Server base OS
  - RabbitMQ
  - Glance
  - Cloud Controller
  - Volume Host
- Compute Host