CS 395 Special Topics: Cloud Computing

Overview of 

Kyoungho An

Vanderbilt University

2011/09/06

What is OpenStack?

- http://www.openstack.org/

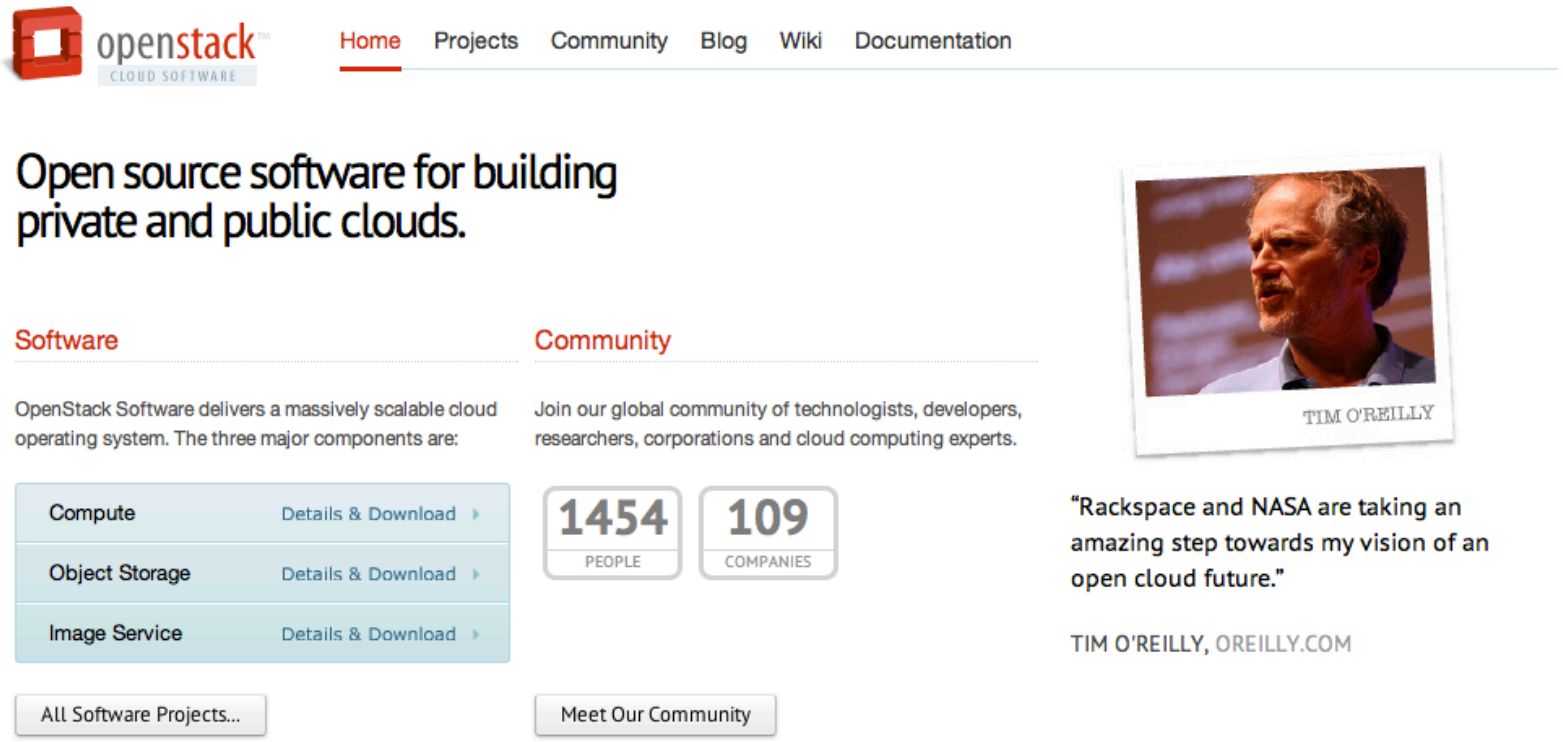

What projects are in OpenStack?

The OpenStack Core Projects

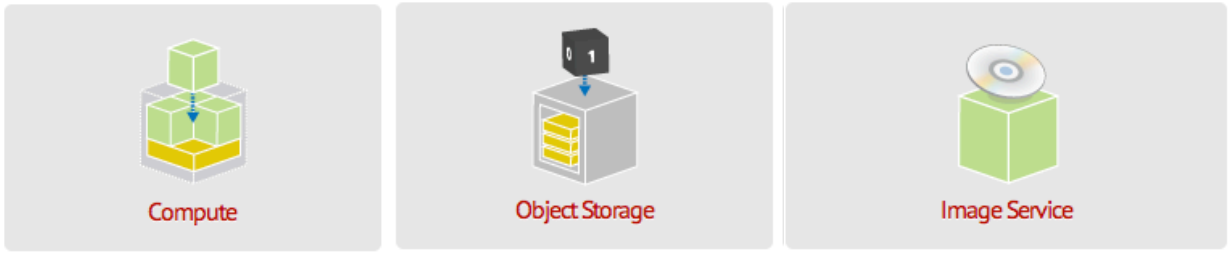

OPENSTACK COMPUTE: open source software and standards for large-scale deployments of automatically provisioned virtual compute instances.

OPENSTACK OBJECT STORAGE: open source software and standards for large-scale, redundant storage of static objects.

OPENSTACK IMAGE SERVICE: provides discovery, registration, and delivery services for virtual disk images.

Understanding *Swift*

- Scalable and redundant object storage service

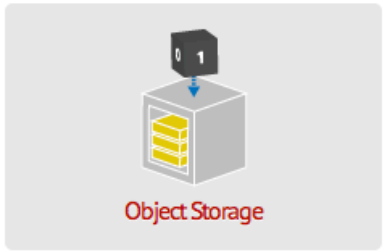

OPENSTACK OBJECT STORAGE: open source software and standards for large-scale, redundant storage of static objects.

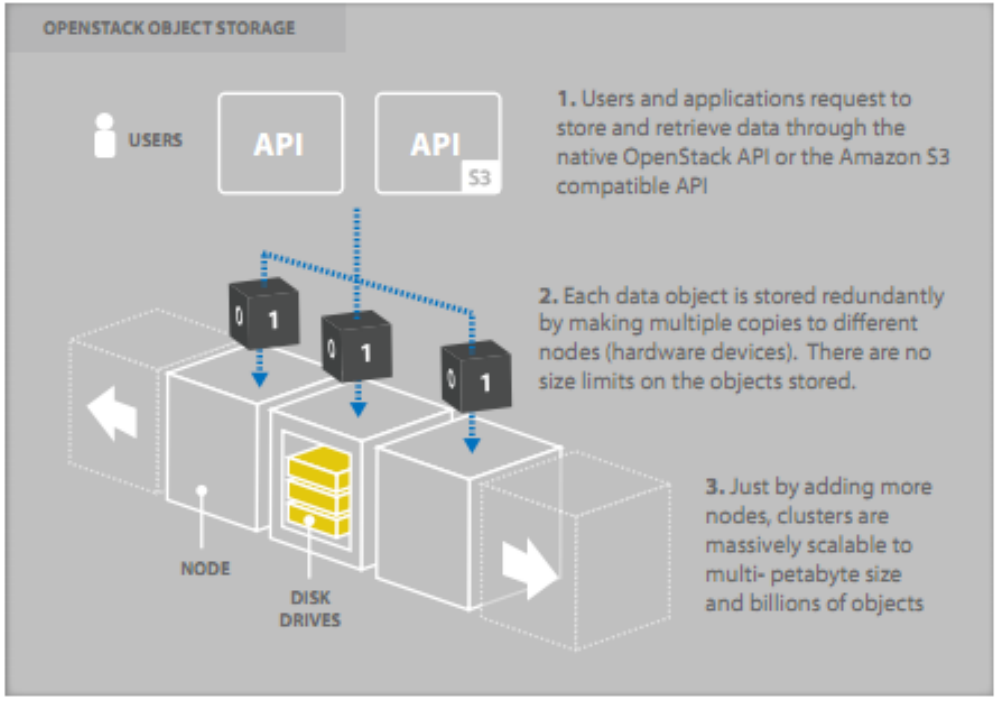

Understanding *Swift*

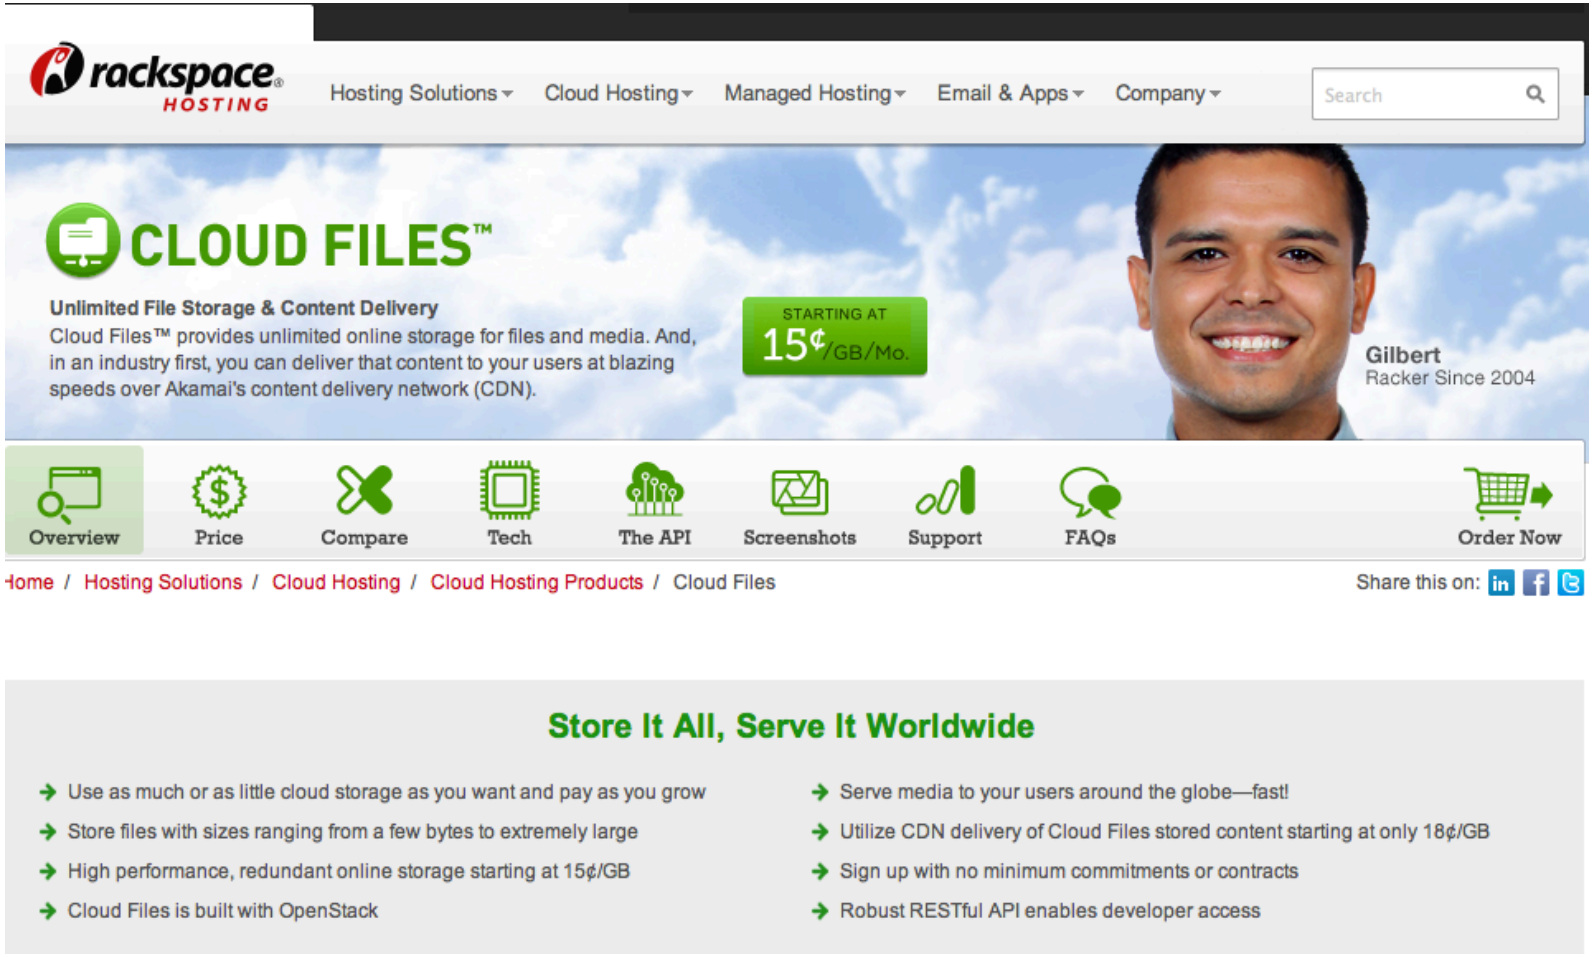

Understanding *Swift*

Amazon Simple Storage Service (Amazon S3)

Amazon S3 is storage for the Internet. It is designed to make web-scale computing easier for developers.

Amazon S3 provides a simple web services interface that can be used to store and retrieve any amount of data, at any time, from anywhere on the web. It gives any developer access to the same highly scalable, reliable, secure, fast, inexpensive infrastructure that Amazon uses to run its own global network of web sites. The service aims to maximize benefits of scale and to pass those benefits on to developers.

Easy to sign up,
pay only for what you use

Sign Up Now

This page contains the following categories of information. Click to jump down:

- ↓ **Amazon S3 Functionality**
- ↓ **Protecting Your Data**
- ↓ **Pricing**
- ↓ **Getting Started with Amazon S3**
- ↓ **Transferring Large Amounts of Data**
- ↓ **Common Use Cases**
- ↓ **Resources**
- ↓ **Amazon S3 Design Requirements**
- ↓ **Intended Usage and Restrictions**

Understanding *Swift*

- To provide scalable and redundant service
  - Multiple copies of each object to multiple storage servers within separate “zones”
  - Zones are logical grouping of storage servers (servers, racks, or data center)

Understanding *Swift*

- Configurable in terms of how many copies as well as how many zones
- Current best practices call for three replicas written across five zones

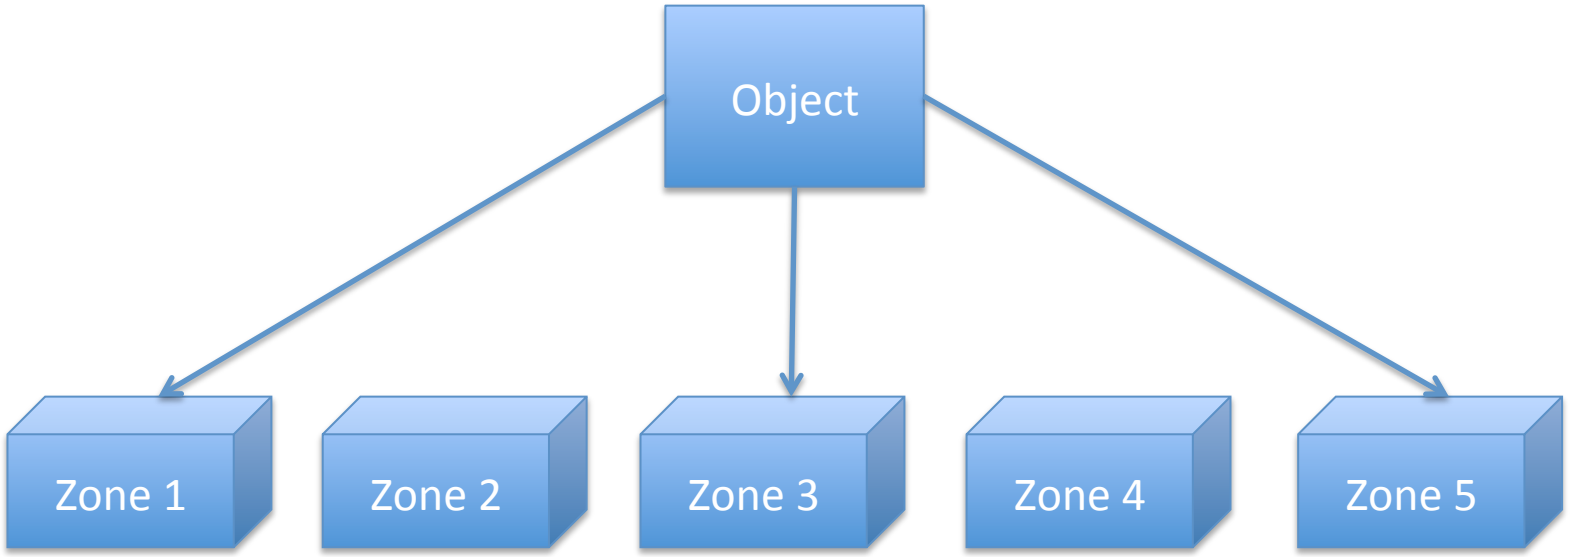

Architecture of *Swift*

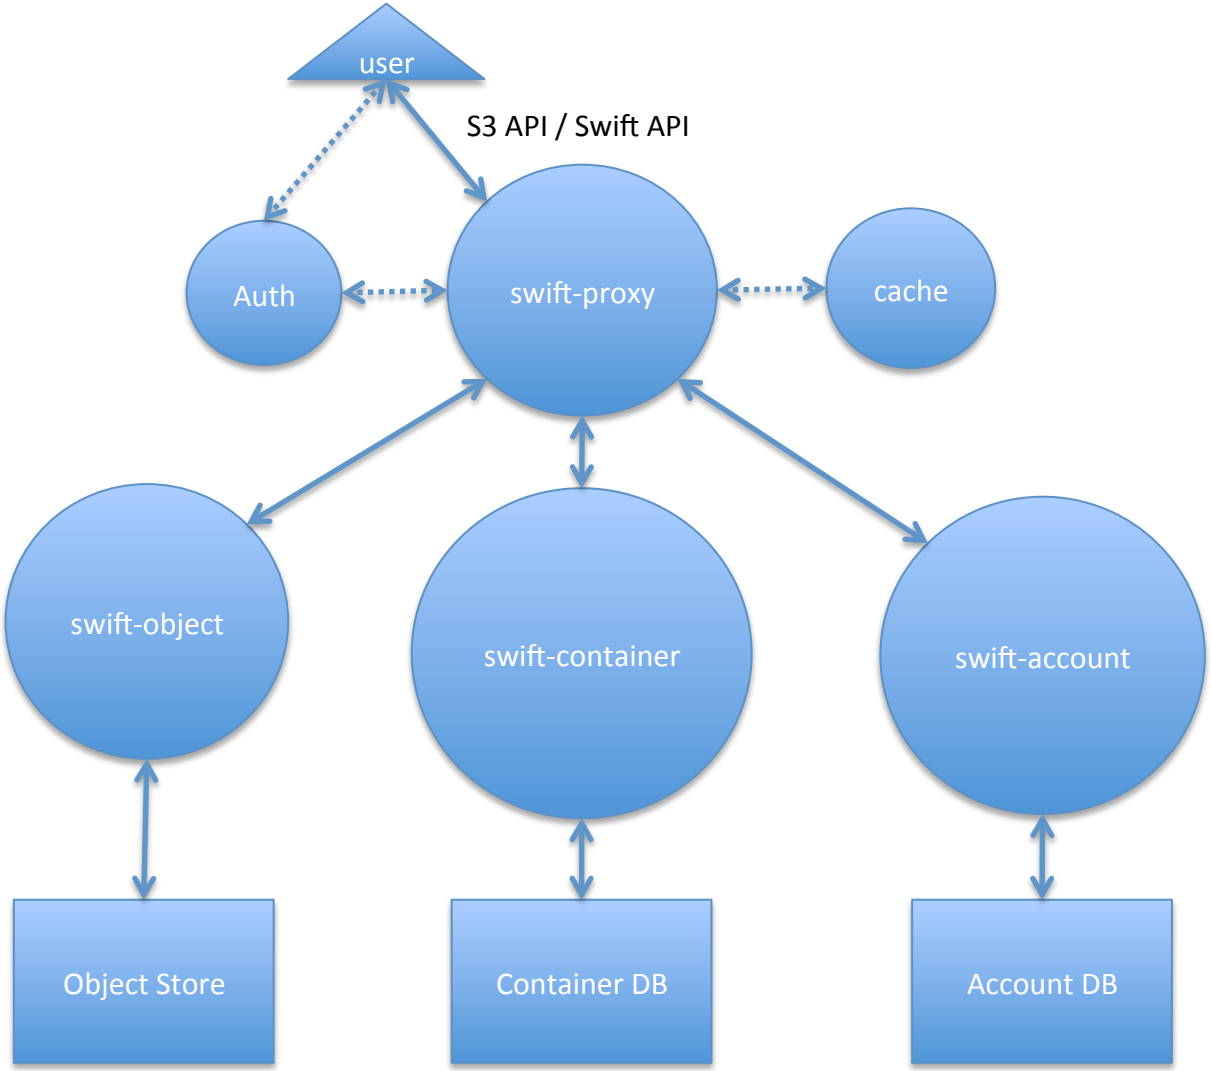

Understanding *Glance*

- A set of service for storing and querying virtual disk images
- Standalone service
- Provides REST interfaces
- Variety of back-end stores

OPENSTACK IMAGE SERVICE: provides discovery, registration, and delivery services for virtual disk images.

Architecture of *Glance*

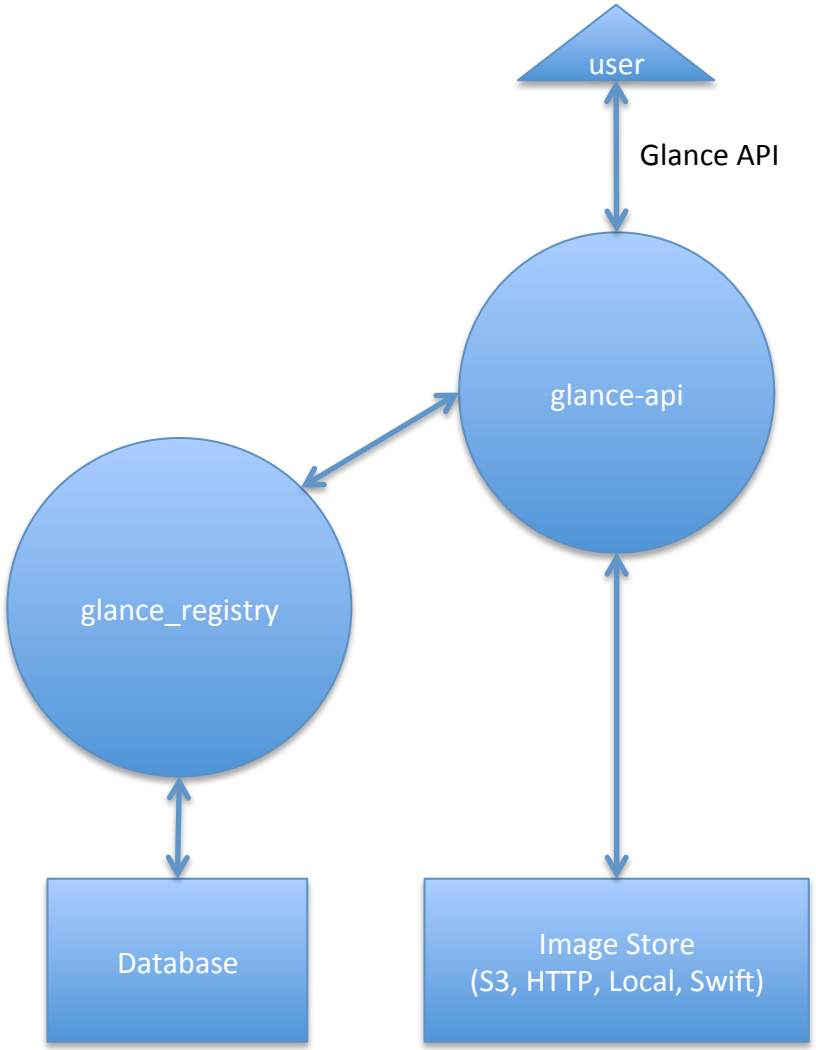

Understanding *Nova*

- Management of virtual compute instances

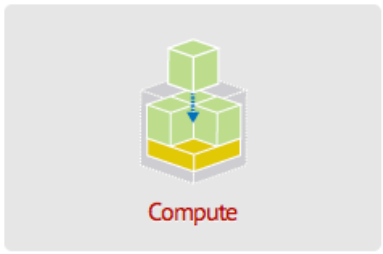

OPENSTACK COMPUTE: open source software and standards for large-scale deployments of automatically provisioned virtual compute instances.

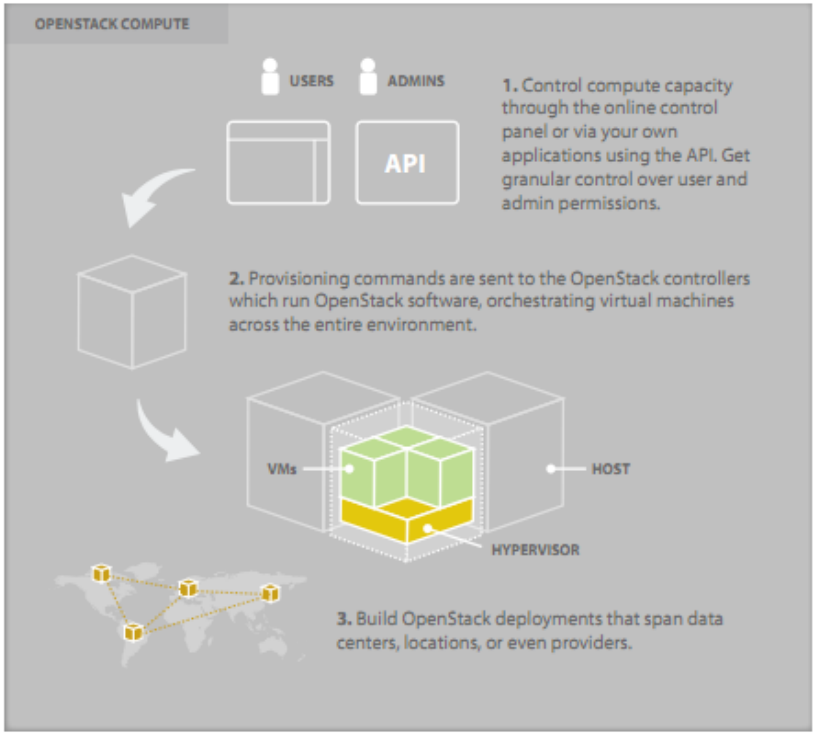

Understanding *Nova*

Amazon Elastic Compute Cloud (Amazon EC2)

Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides resizable compute capacity in the cloud. It is designed to make web-scale computing easier for developers.

Amazon EC2's simple web service interface allows you to obtain and configure capacity with minimal friction. It provides you with complete control of your computing resources and lets you run on Amazon's proven computing environment. Amazon EC2 reduces the time required to obtain and boot new server instances to minutes, allowing you to quickly scale capacity, both up and down, as your computing requirements change. Amazon EC2 changes the economics of computing by allowing you to pay only for capacity that you actually use. Amazon EC2 provides developers the tools to build failure resilient applications and isolate themselves from common failure scenarios.

Easy to sign up,
pay only for what you use

Sign Up Now

This page contains the following categories of information. Click to jump down:

- **Amazon EC2 Functionality**
- **Service Highlights**
- **Features**
- **Instance Types**
- **Operating Systems and Software**
- **Pricing**
- **Resources**
- **Detailed Description**
- **Intended Usage and Restrictions**

Architecture of *Nova*

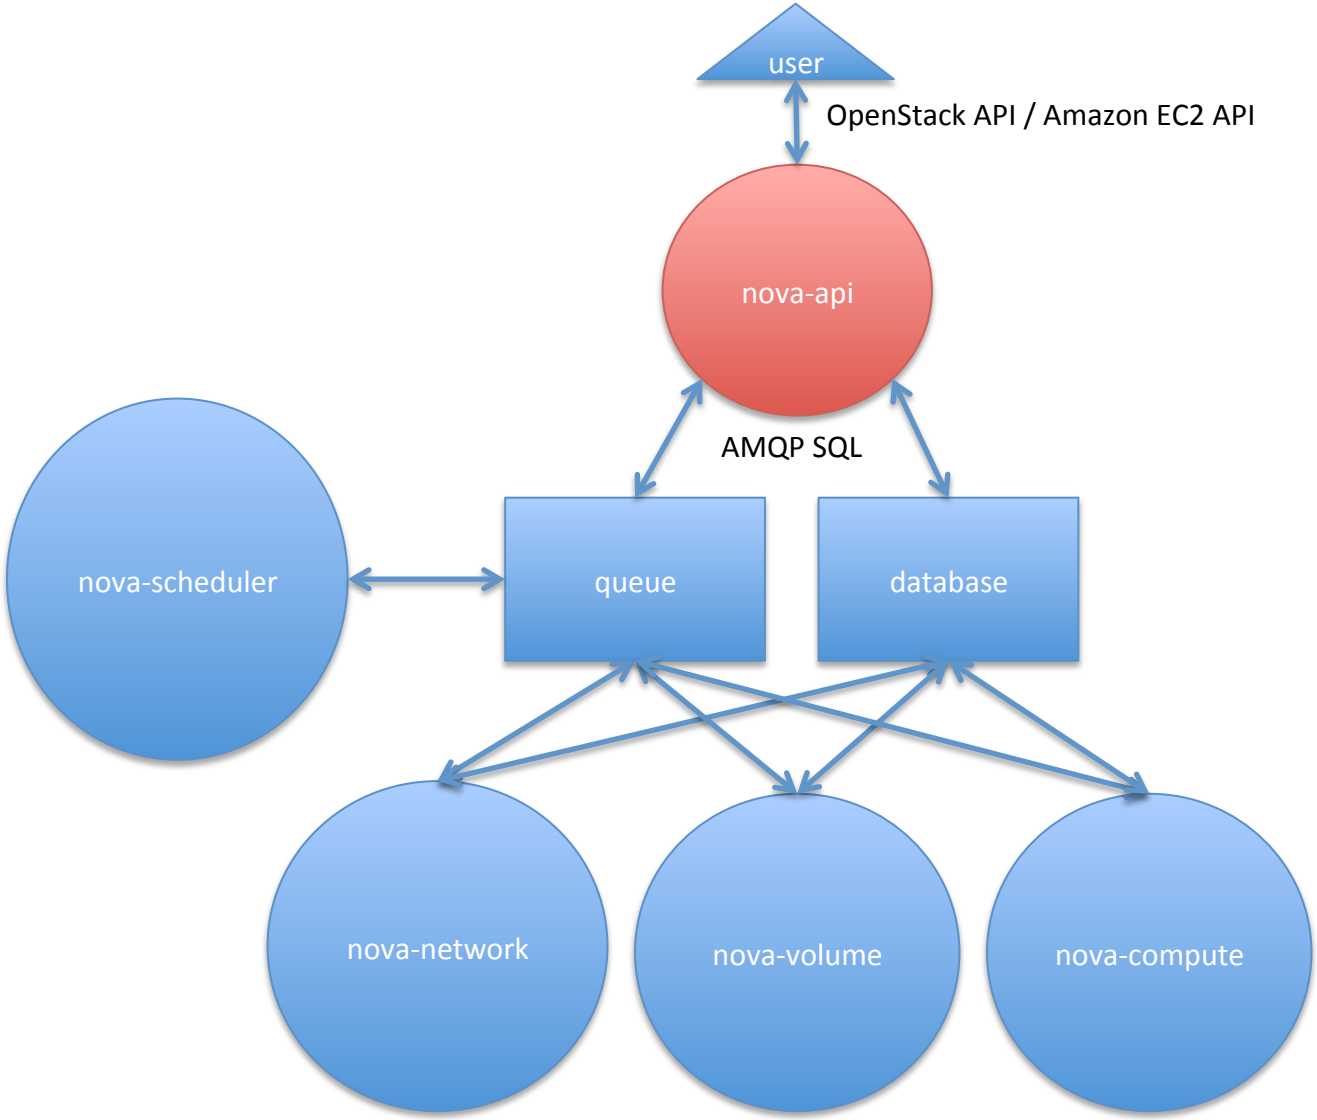

Nova API

- The heart of Nova
- Accept and fulfill incoming API requests
- For requests,
  - Query database and return the answer
  - Pass messages to other daemons through writing information to the database and adding messages to the queue

Architecture of *Nova*

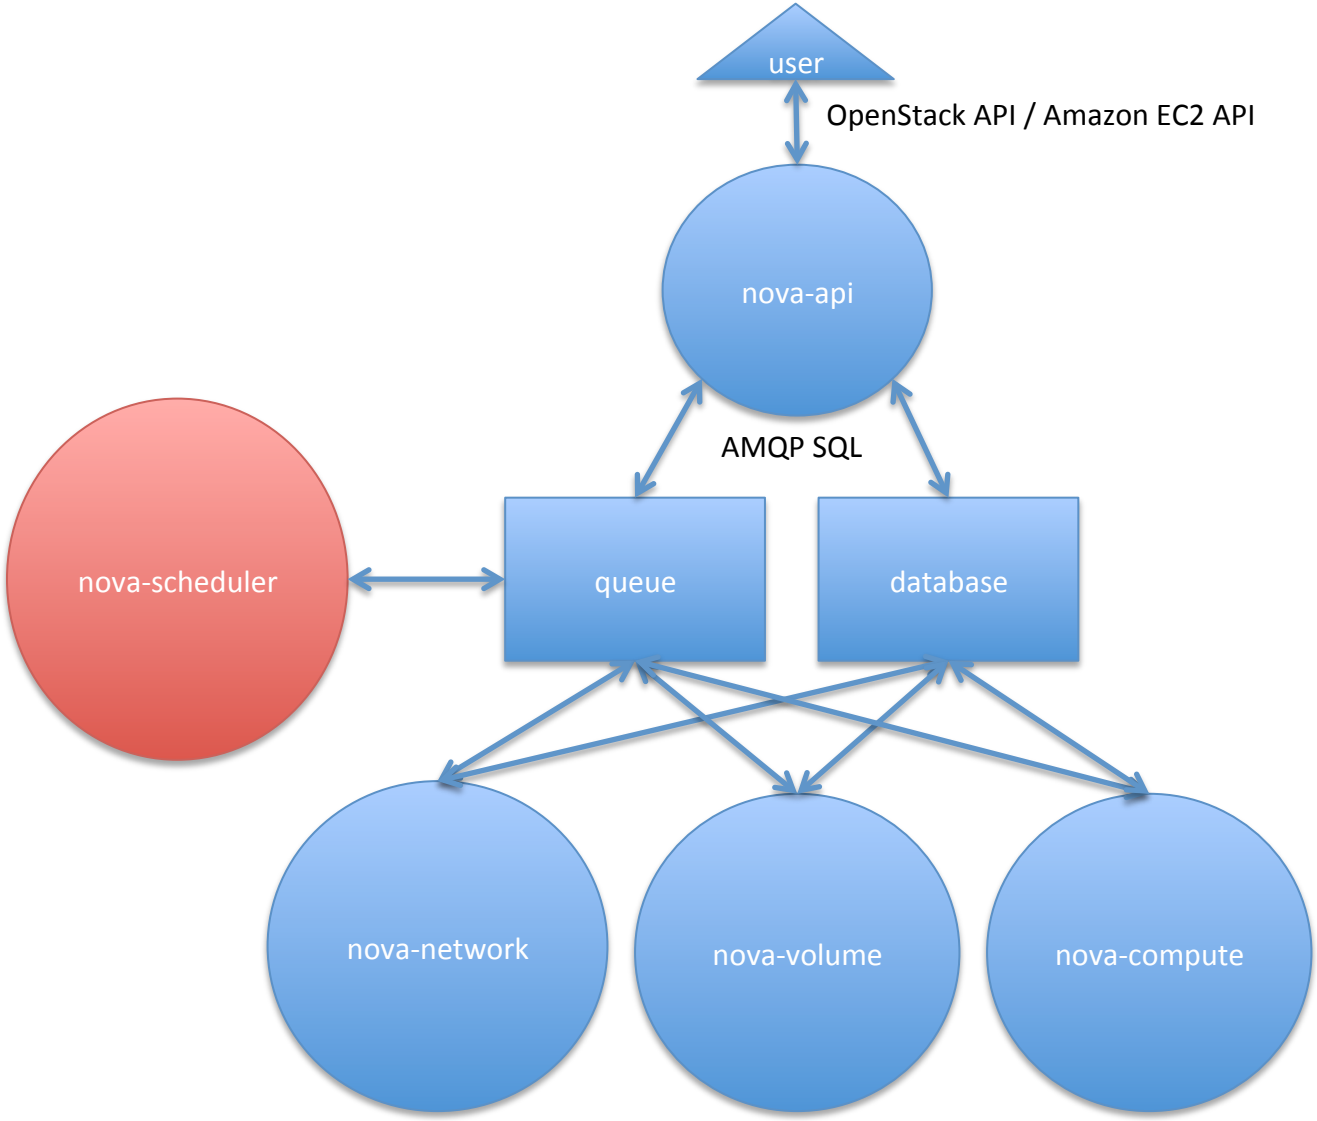

Nova Scheduler

- Pluggable architecture
- Current scheduler choices
  - Simple
  - Chance
  - Zone

Architecture of *Nova*

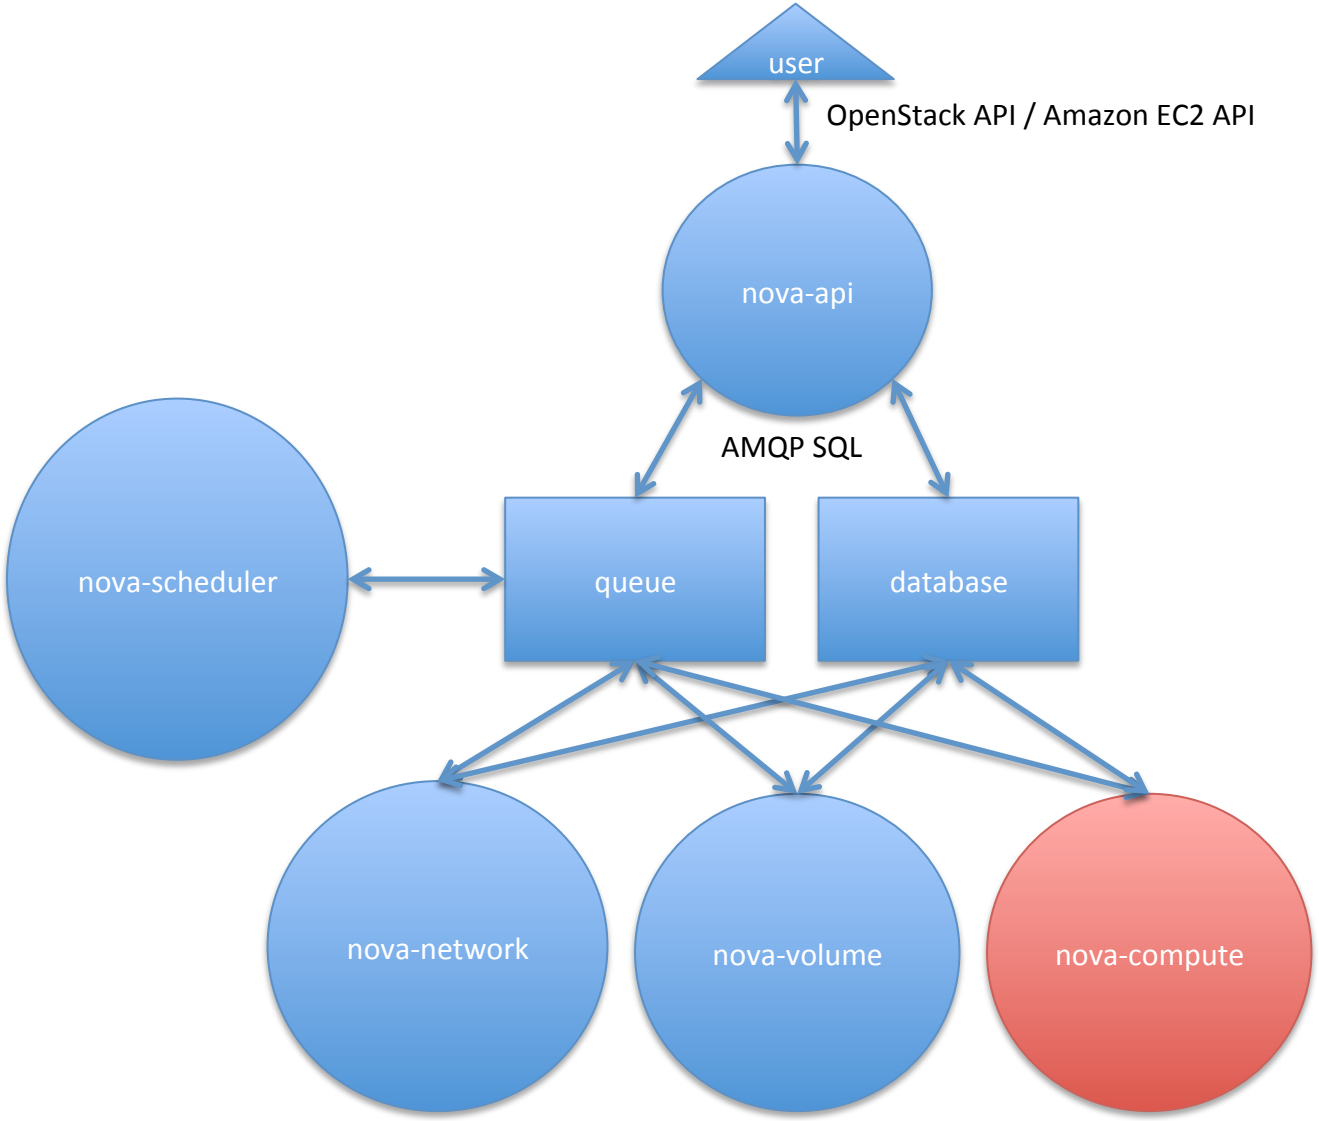

Nova Compute Worker

- Create and terminate virtual instances
- Process
  1. Accept actions from the queue
  2. Perform virtual machine API calls
  3. Update state in the database
- Wide support for virtualization technologies
  - KVM, Xen, VMware, UML, Hyper-V, Virtual Box…

Architecture of *Nova*

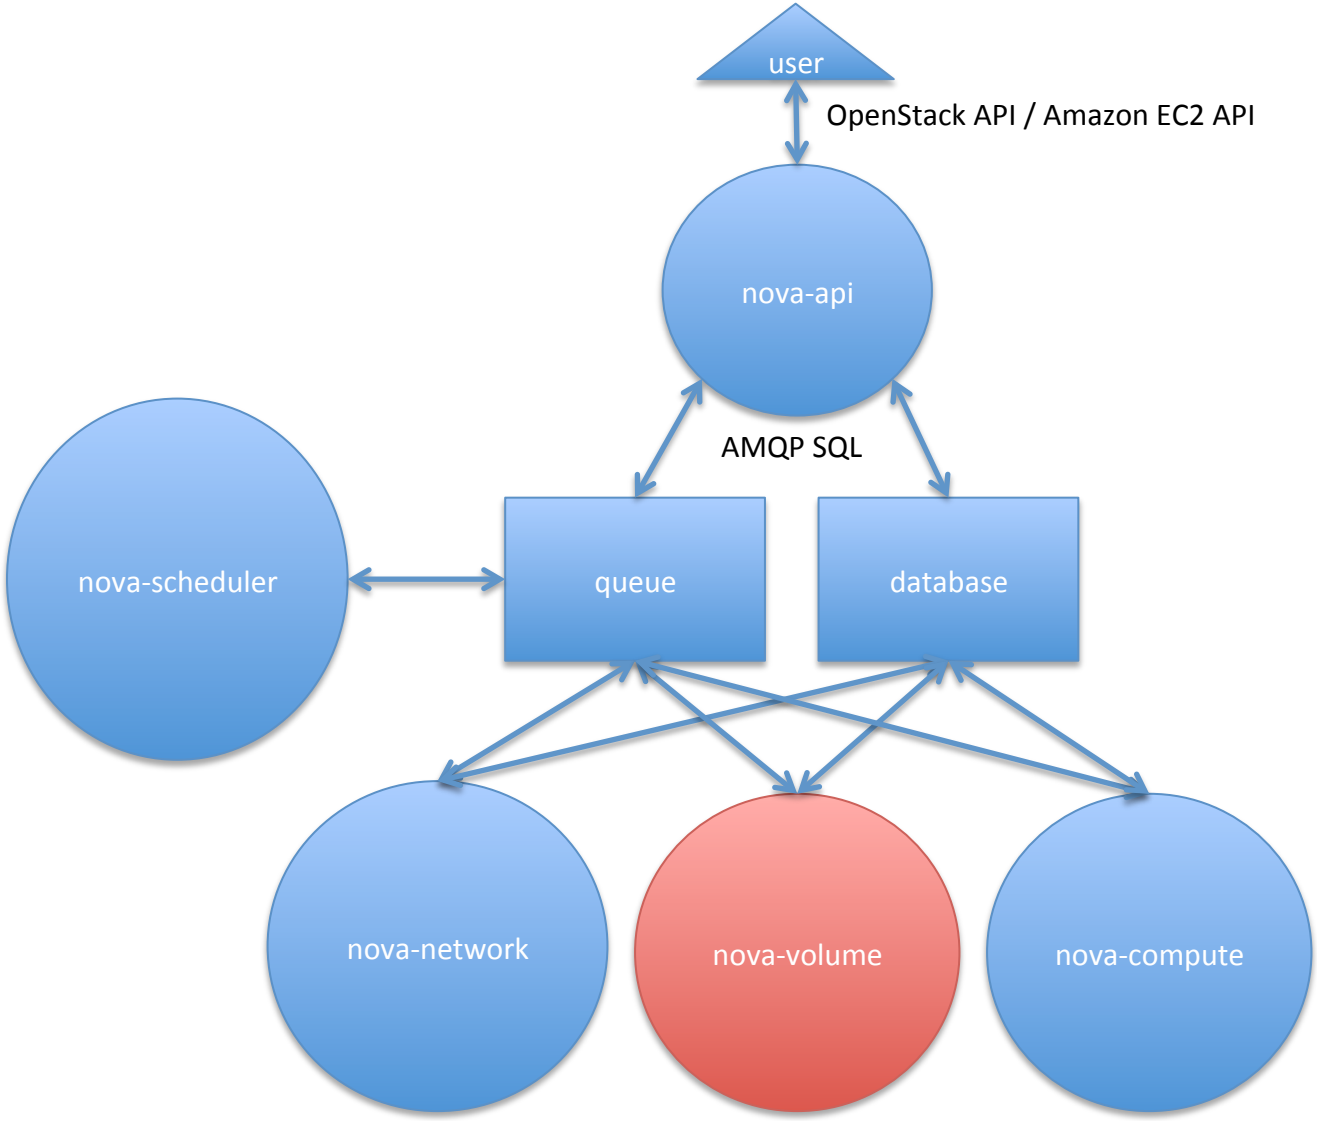

Nova Volume Worker

- Create, attach, and detach persistent volumes to compute instances
- Similar to Amazon's Elastic Block Stroage
- Volume provider options
  - AoE, iSCSI, Sheepdog, RBD, LeftHand…

Architecture of *Nova*

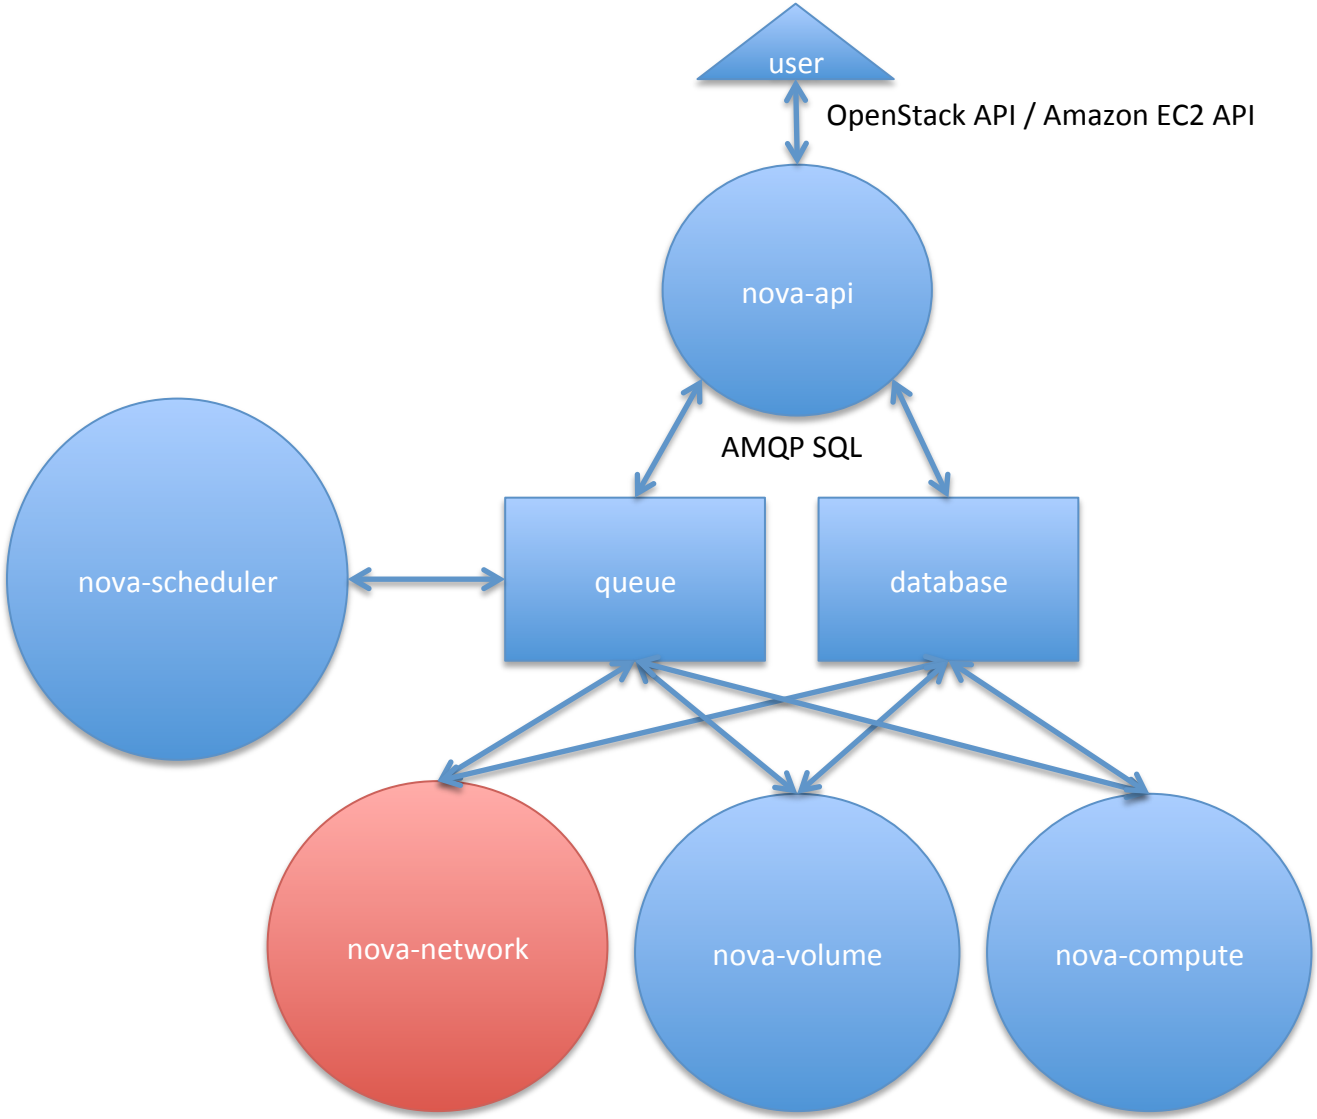

Nova Network Worker

- Accepts networking tasks from the queue, and performs system commands to manipulate the network
  - Setting up bridging interfaces
  - Changing *iptables* rules
- Types of IP addresses for an instance
  - Private IPs (Fixed)
  - Public IPs (Floating)

Architecture of *Nova*

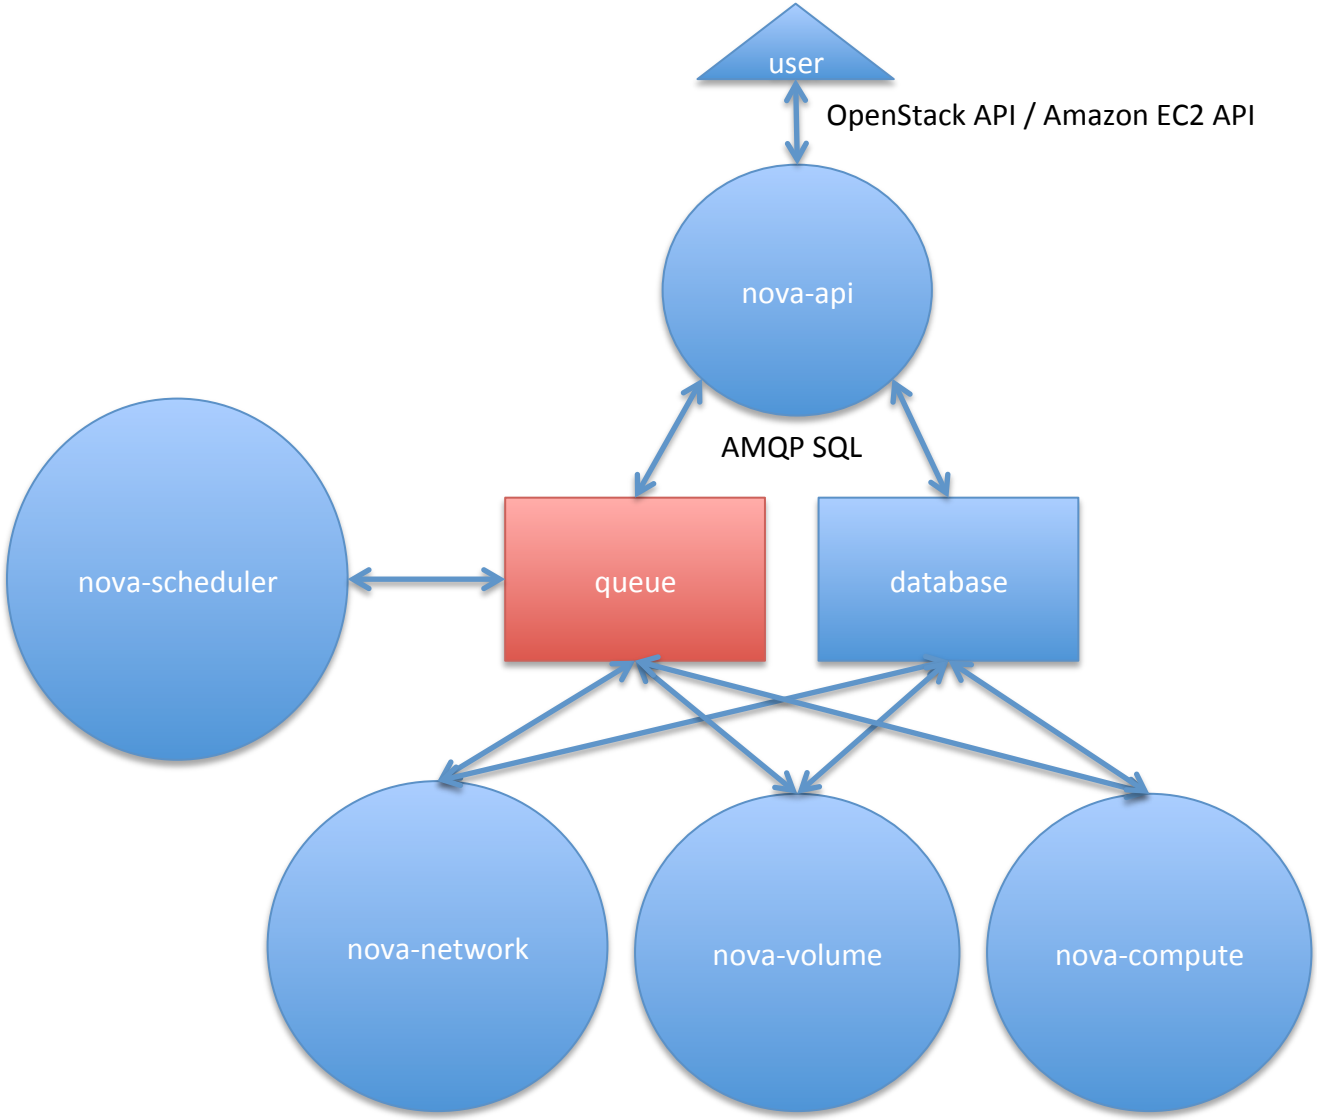

Nova Queue

- Provide a central hub for passing message between daemons
- Implemented with RabbitMQ
- Types of message queues
  - Topic queues (broadcast to worker daemons)
  - Fanout queues (nova-scheduler workers)
  - Host queues (specific host's compute worker)

Architecture of *Nova*

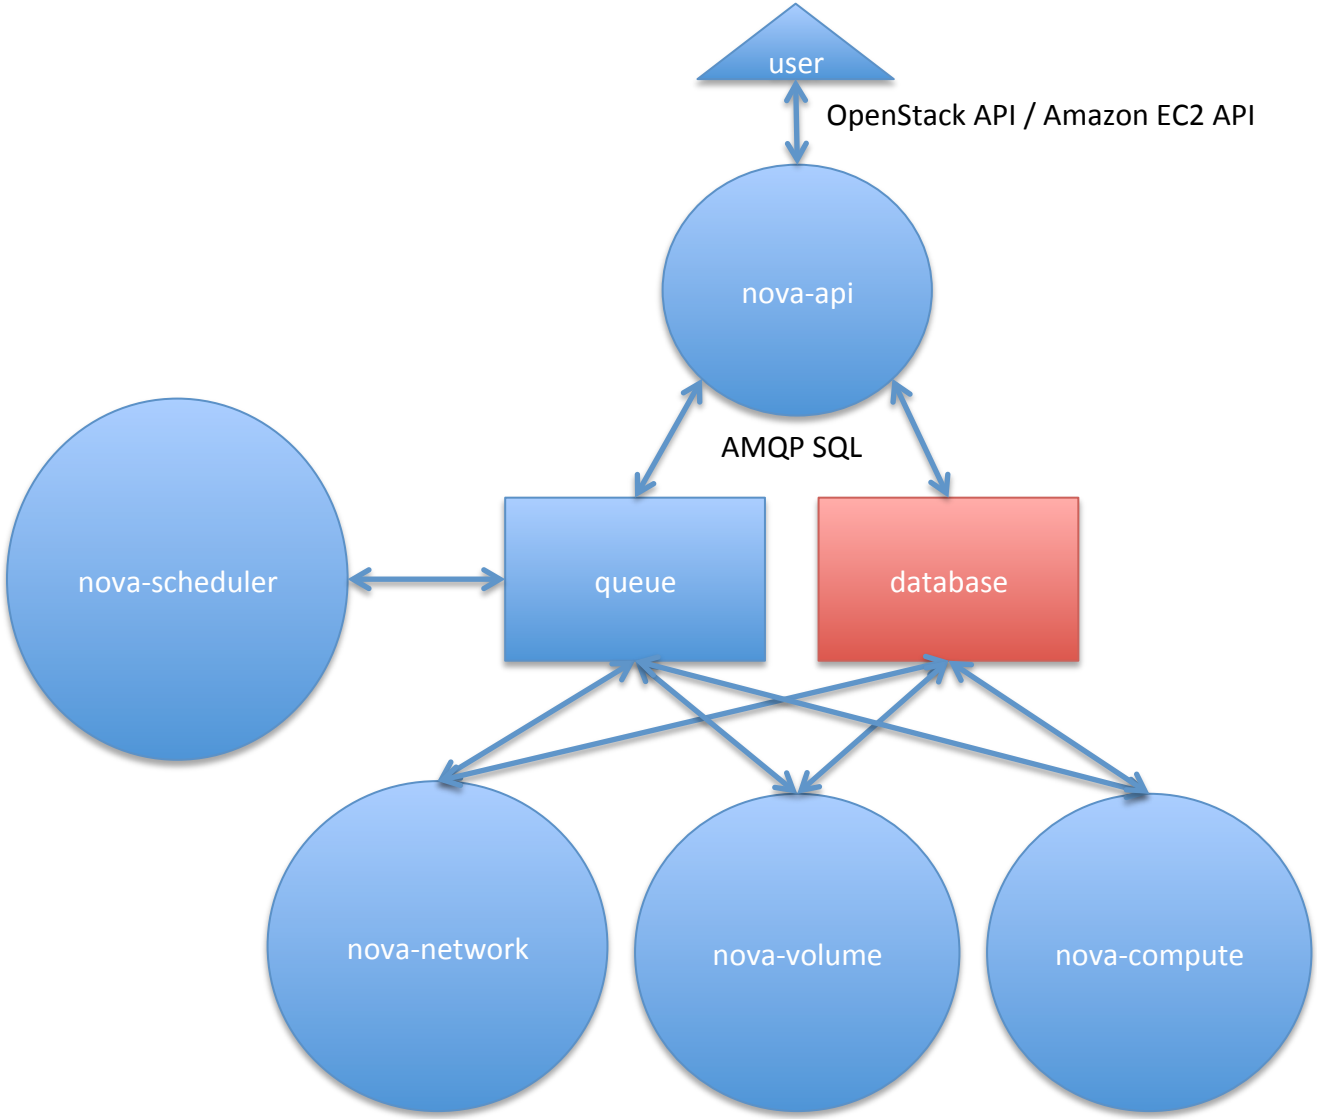

Nova Database

- Stores configuration and run-time state
- Supports a wide range of databases (MySQL and PostgreSQL)
- Nova database tables
  - Projects
  - Users
  - Certificates
  - Instances
  - Images
  - Volumes
  - Key Pairs

Questions?