OpenStack

Orgad Kimchi

Principal Software Engineer

Oracle ISV Engineering

Safe Harbor Statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Agenda

- OpenStack Overview
- OpenStack as a DevOps platform
- OpenStack Services
- OpenStack Use Cases

OpenStack Overview

- What is OpenStack?
  - Open Source Cloud Software
    - Foundation for IaaS, PaaS and SaaS
  - Combines compute, network and storage resources
    - Web portal for cloud admins and self-service users
    - Cloud services exposed through APIs

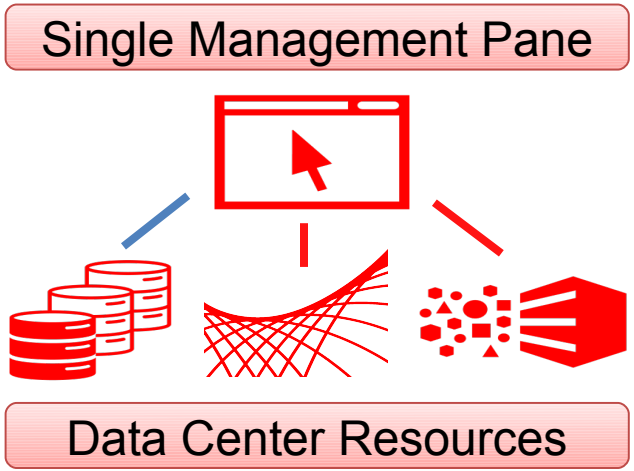

OpenStack Overview

- OpenStack is a cloud operating system that controls large pools of compute, storage, and networking resources throughout a data center
- All managed through a dashboard that gives administrators control while empowering users to provision resources through a web interface

Who's behind OpenStack?

- Founded by Rackspace Hosting and NASA
- OpenStack has grown to be a global software community of developers collaborating on a standard and massively scalable open source cloud operating system.
- The OpenStack Foundation promotes the development, distribution and adoption of the OpenStack cloud operating system.

OpenStack Releases

Release Name	Release Date	Included Components
Austin	21 October 2010	Nova, Swift
Bexar	3 February 2011	Nova, Glance, Swift
Cactus	15 April 2011	Nova, Glance, Swift
Diablo	22 September 2011	Nova, Glance, Swift
Essex	5 April 2012	Nova, Glance, Swift, Horizon, Keystone
Folsom	27 September 2012	Nova, Glance, Swift, Horizon, Keystone, Quantum, Cinder
Grizzly	4 April 2013	Nova, Glance, Swift, Horizon, Keystone, Quantum, Cinder
Havana	17 October 2013	Nova, Glance, Swift, Horizon, Keystone, Neutron, Cinder, Ceilometer, Heat
Icehouse	17 April 2014	Nova, Glance, Swift, Horizon, Keystone, Neutron, Cinder, Ceilometer, Heat, Trove
Juno	October 2014	Nova, Glance, Swift, Horizon, Keystone, Neutron, Cinder, Ceilometer, Heat, Trove (more to be added)

Enterprise OpenStack

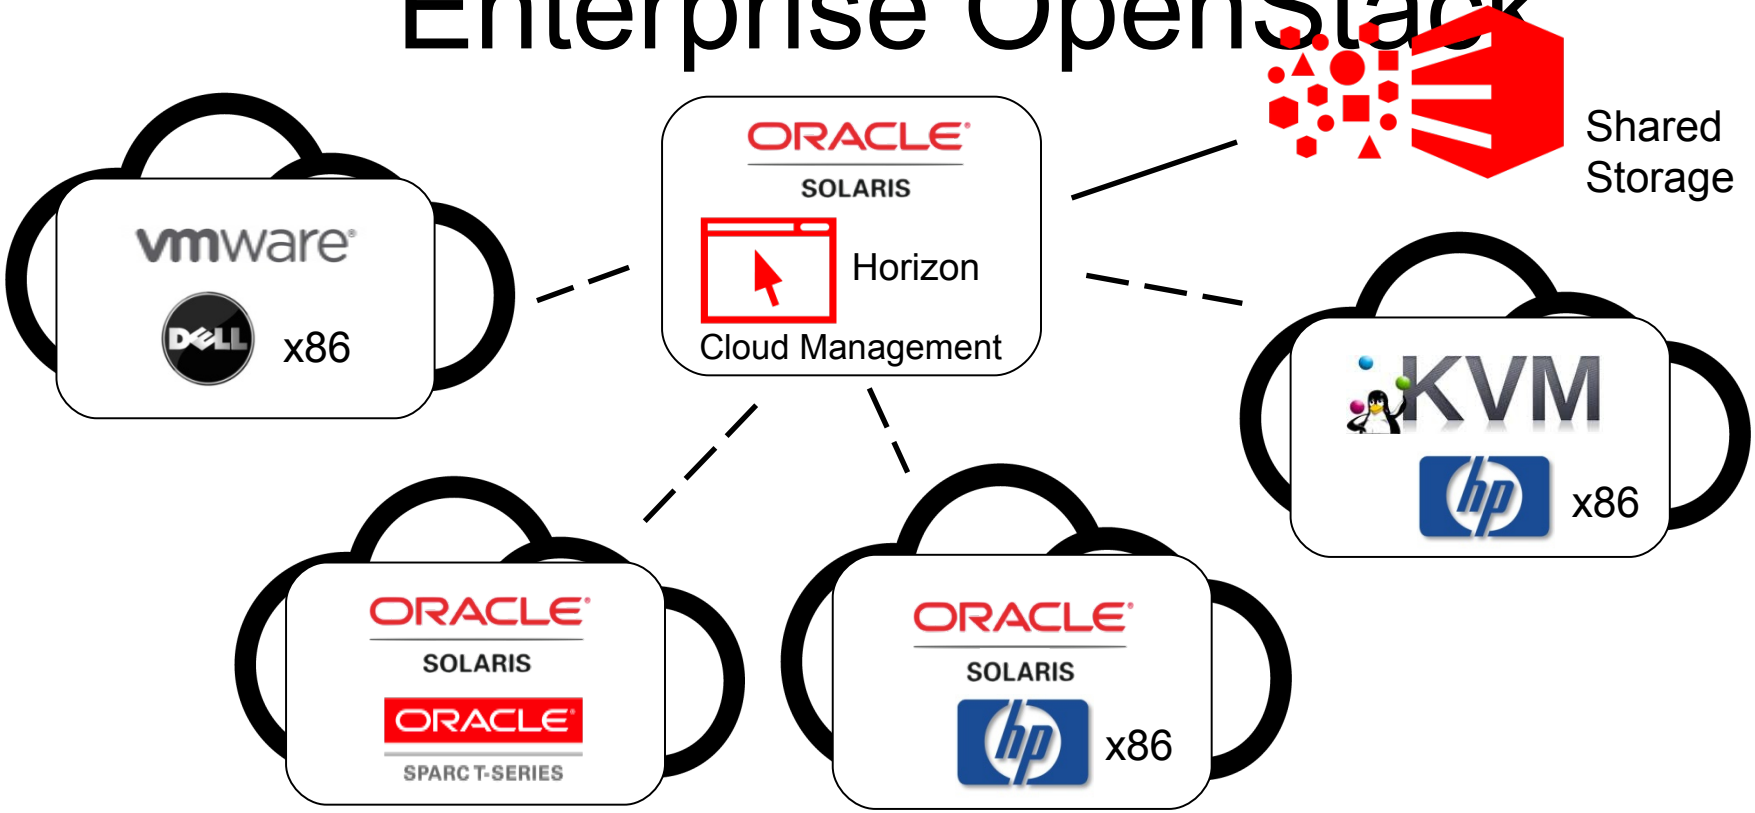

OpenStack Architecture

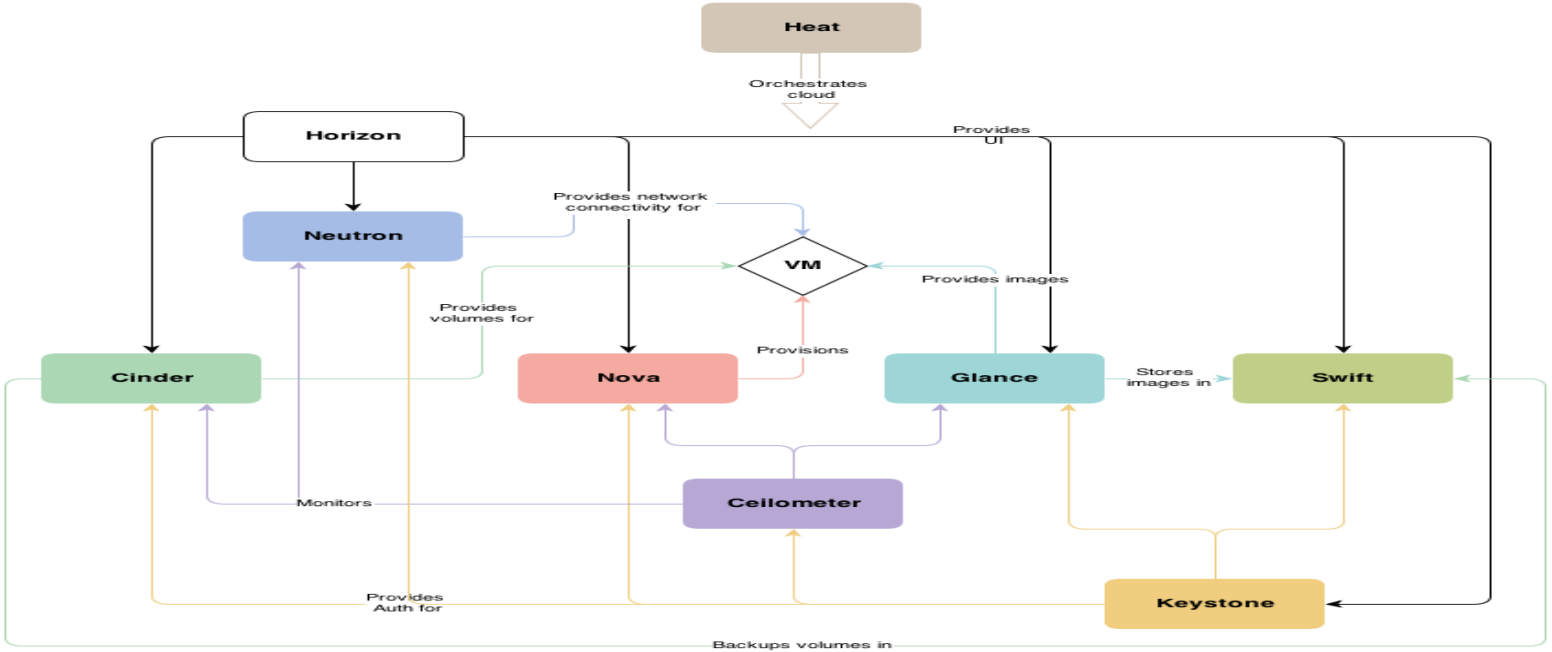

Source http://openstack.org

Image Service - Glance

- The OpenStack Image Service provides discovery, registration and delivery services for disk and server images.
- The ability to copy or snapshot a server image and immediately store it away is a powerful capability of the OpenStack cloud operating system.

Glance

- Stored images can be used as a template to get new servers up and running quickly—and more consistently
- If you are provisioning multiple servers—than installing a server operating system and individually configuring additional services.
- It can also be used to store and catalog an unlimited number of backups.

Glance Container Formats

- Glance also supports the concept of container formats:
- OVF An open standard for distributing one or more virtual machine images
- aki, ari, ami Amazon kernel, ramdisk, or machine image (respectively)
- Docker New container format to support Docker

Networking in OpenStack

- Networking in OpenStack is one of the most powerful and sophisticated feature sets.
- The OpenStack networking service, Neutron, offers a complete SDN solution along with various network services, out of the box.
- The network services Neutron can support include: routing, firewall, DNS, DHCP, load balance, VPN

User Isolation Multi Tenancy

- Allowing multiple users to share the same physical environment with complete separation between them is a key feature in OpenStack.
- OpenStack is designed in a way that many tenants can share the same physical resources, without being aware that they do so.
- OpenStack offer ways to share virtual resources between tenants, but maintains complete separation where needed.

OpenStack Compute

- Nova provides facilities to provision and manage virtual machine instances.
- Similar in functionality and scope to Amazon's EC2 service
- It allows you to create, manage, and destroy virtual servers based on machine images located in Glance through a programmable API.

OpenStack Compute

- OpenStack Compute provides virtual machines on demand for users.
- Virtual machines, or instances as they are called in Nova parlance, can be controlled via API calls to OpenStack Compute.
- Through these API calls, users can start, assign IP addresses, attach additional storage or access their instances consoles.

OpenStack Dashboard Horizon

- OpenStack Dashboard (or Horizon as it is codenamed) provides a web frontend for OpenStack services.
- This currently includes all of the core OpenStack services as well as some of the incubating projects.
- Depending on the user's credentials, it will display end user or operator specific screens and functionality.

Horizon

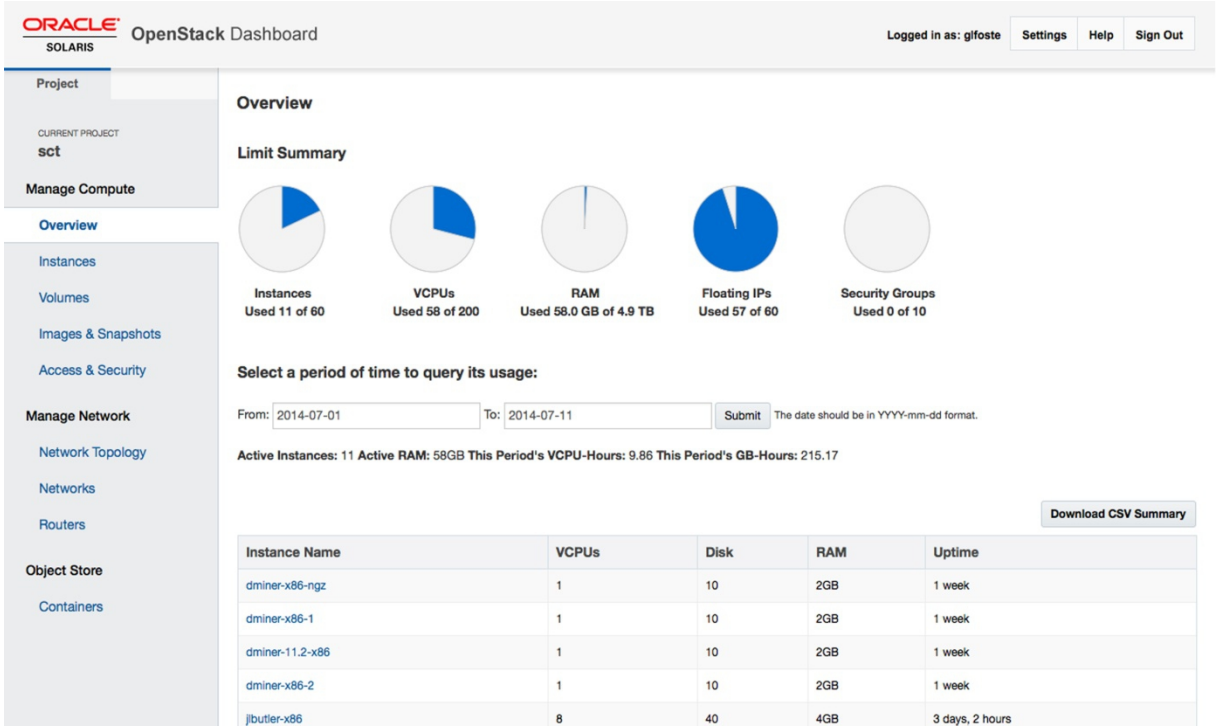

Object Storage - Swift

- OpenStack provides redundant, scalable object storage using clusters of standardized servers capable of storing petabytes of data
- Object Storage is not a traditional file system, but rather a distributed storage system for static data such as virtual machine images, photo storage, email storage, backups and archives.

Object Storage - Swift

- Global cluster capability: This allows replication and distribution of data around the world.
- This functionality helps with disaster recovery, distribution of hot data.
- Partial object retrieval: For example, if you want just a portion of a movie object or a TAR file.

OpenStack Metering

- Ceilometer is OpenStack's telemetry project.
- The project offers metering information about the resource consumption on OpenStack clouds.

Ceilometer

- Alarms form the basis of Ceilometer's monitoring support and OpenStack Heat integration.
- They are a threshold on a particular meter and resource that will create an event when they are exceeded.
- For example, an alarm can be set on an instance's CPU utilization with a threshold of 75%.

Orchestration module - Heat

- OpenStack Orchestration is a template-driven engine that allows application developers to describe and automate the deployment of infrastructure.

Orchestration module - Heat

- The flexible template language can specify compute, storage and networking configurations as well as detailed post-deployment activity
- To automate the full provisioning of infrastructure as well as services and applications.

Orchestration module - Heat

- Through integration with the Telemetry service, the Orchestration engine can also perform auto-scaling of certain infrastructure elements.

Heat

- heat command-line client
- A CLI that communicates with the heat-api to run AWS CloudFormation APIs.
- End developers can directly use the Orchestration REST API.
- heat-api componentAn OpenStack-native REST API that processes API requests by sending them to the heat-engine over Remote Procedure Call (RPC).

OpenStack Automation - PackStack

- PackStack is a utility that uses Puppet modules to deploy various parts of OpenStack on multiple pre-installed servers over SSH automatically

Database service - Trove

- The Database service provides scalable and reliable cloud provisioning functionality for both relational and non-relational database engines.
- Users can quickly and easily use database features without the burden of handling complex administrative tasks.
- Cloud users and database administrators can provision and manage multiple database instances as needed.

Trove

- Its aim is to provide a full-fledged database environment, while minimizing the administrative turmoil and pains of managing the surrounding infrastructure.

Data processing service -Sahara

- The Data processing service (Sahara) enables users to provide a scalable data processing stack and associated management interfaces.
- This includes provision and operation of data processing clusters as well as scheduling and operation of data processing jobs.

Sahara

- The solution addresses the following use cases:
- Fast provisioning of Hadoop clusters on OpenStack for development and QA.
- Utilization of unused compute power from general purpose OpenStack IaaS cloud.
- Analytics-as-a-Service for ad-hoc or bursty analytic workloads.

Popular Use Cases

- Bursting workloads from private to public OpenStack clouds
- High availability across clouds (for technical diversity)

Bursting workloads from private to public OpenStack clouds

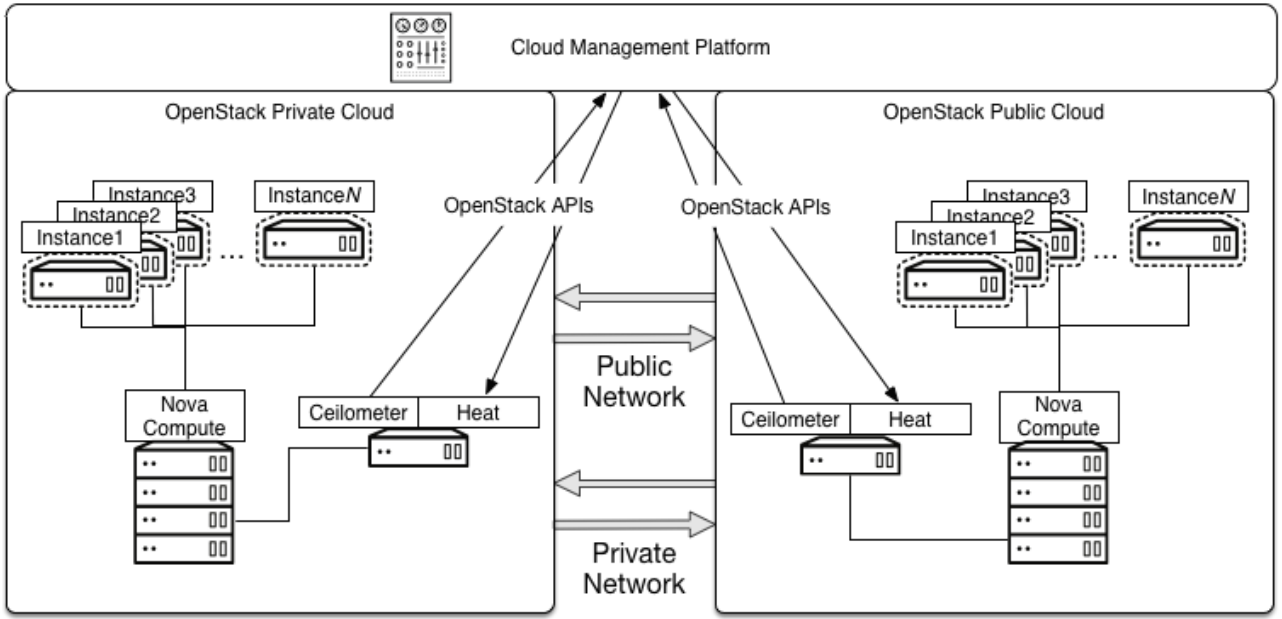

Source http://openstack.org

High availability/disaster recovery

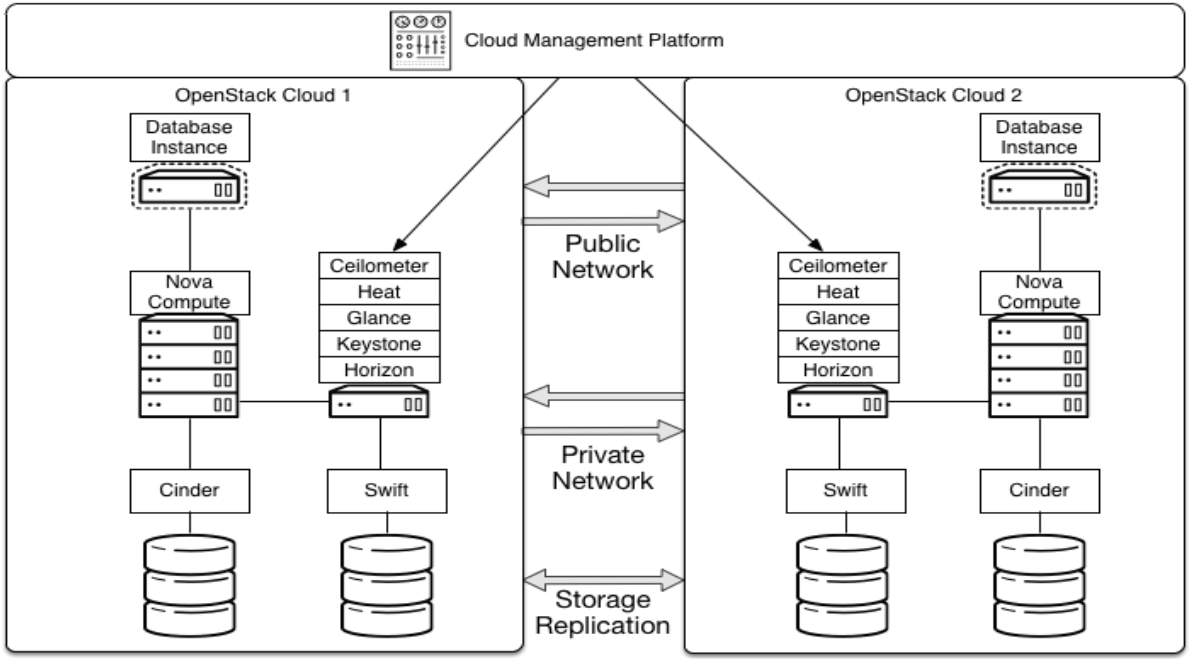

Source http://openstack.org