Objectives

- Understand the Architecture of scalable IT infrastructure
- Understand the concepts, architecture and projects of OpenStack
- Familiar with the basic operation of OpenStack

Part 1

At a Glance: Scalable IT Infrastructure

九州云
Cloud

IT Infrastructure for Online Services

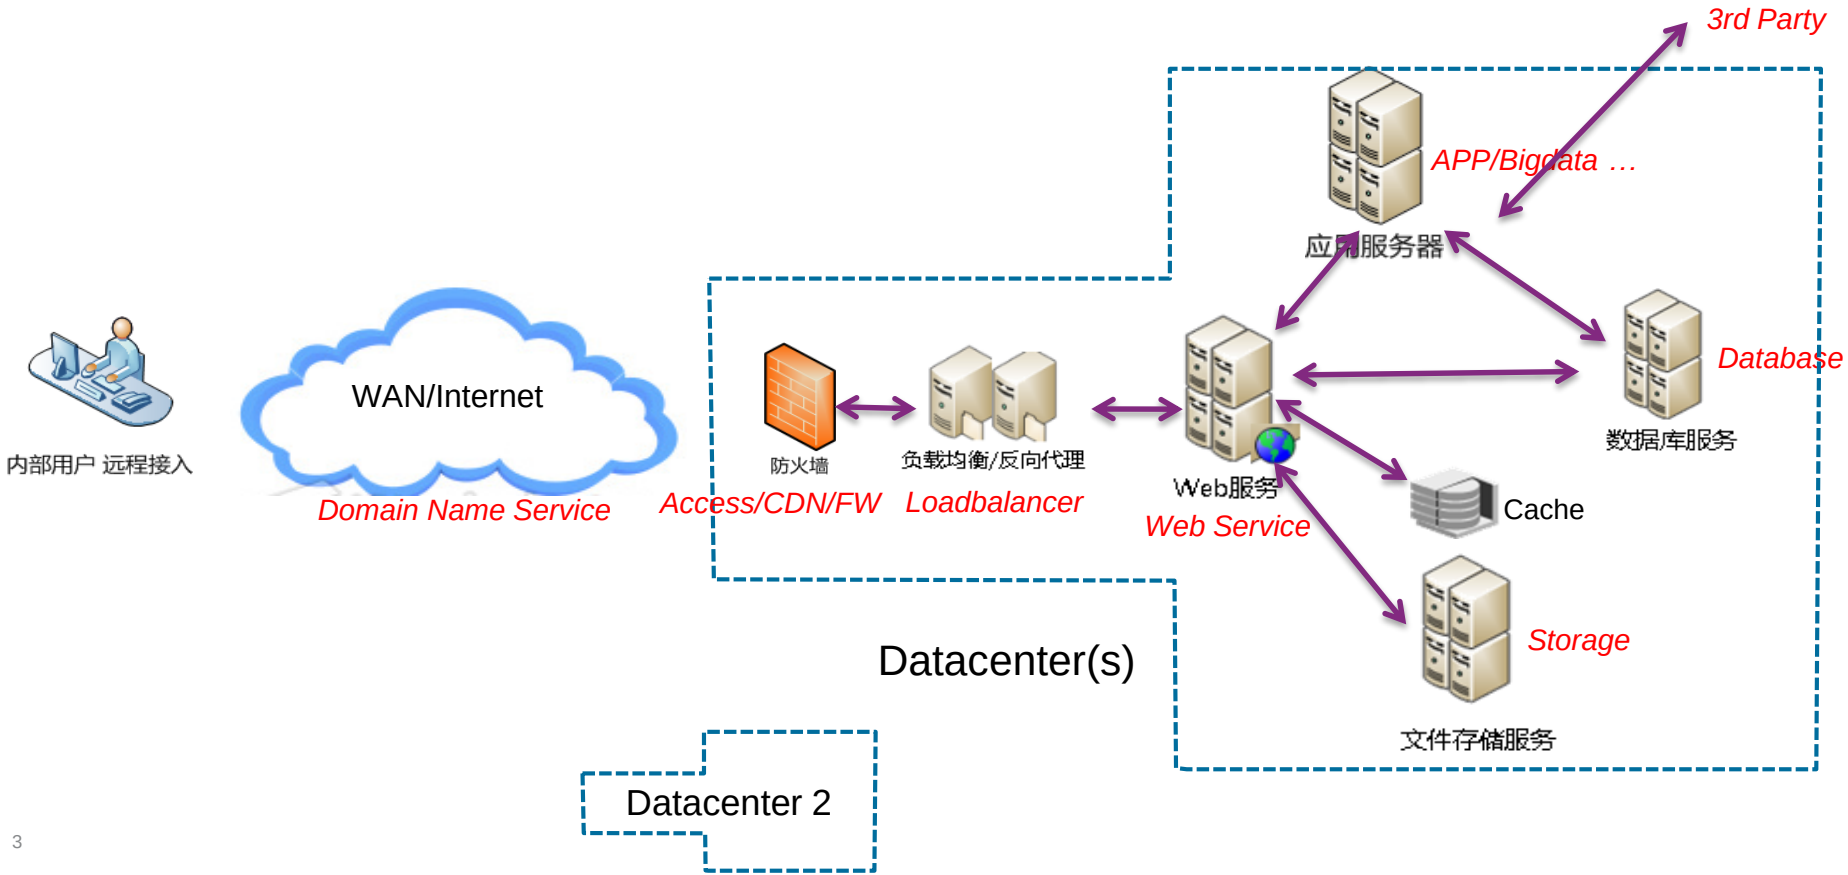

Network

- Acess, WAN, Datacenter
- Underlay network(physical)
- Overlay network(SDN part)

Physical network (Underlay) in a datacenter

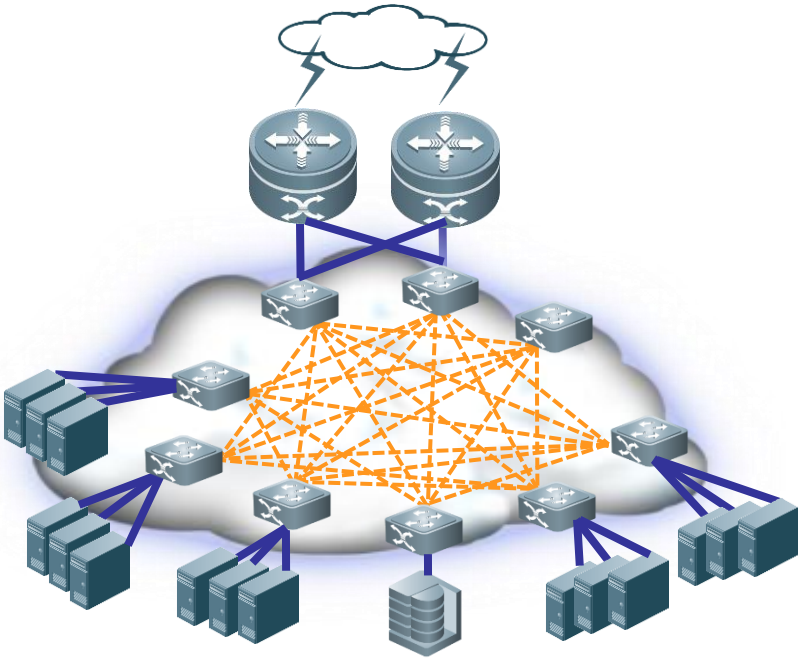

- Troug ut: bandwidth; time delay
- Availability: fail over; convergence
- Security: firewall; VPN; VLAN。。。

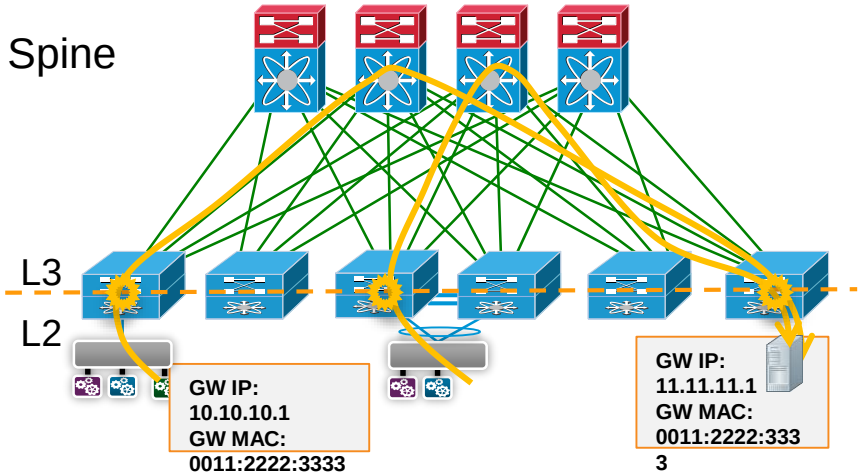

Overlay Network

- Tenant Network
- VXLAN,Virtual Router, Service Chain
- Orchestration

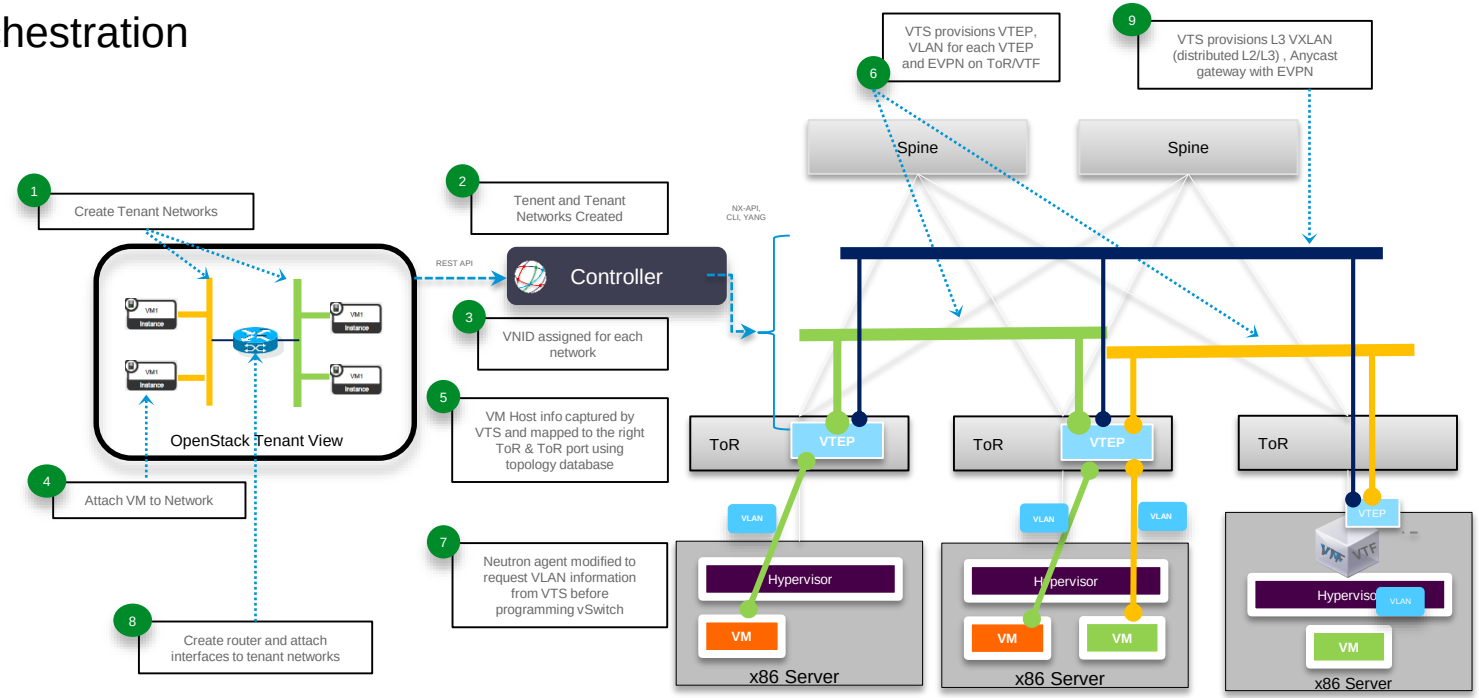

Storage Models

Storage types

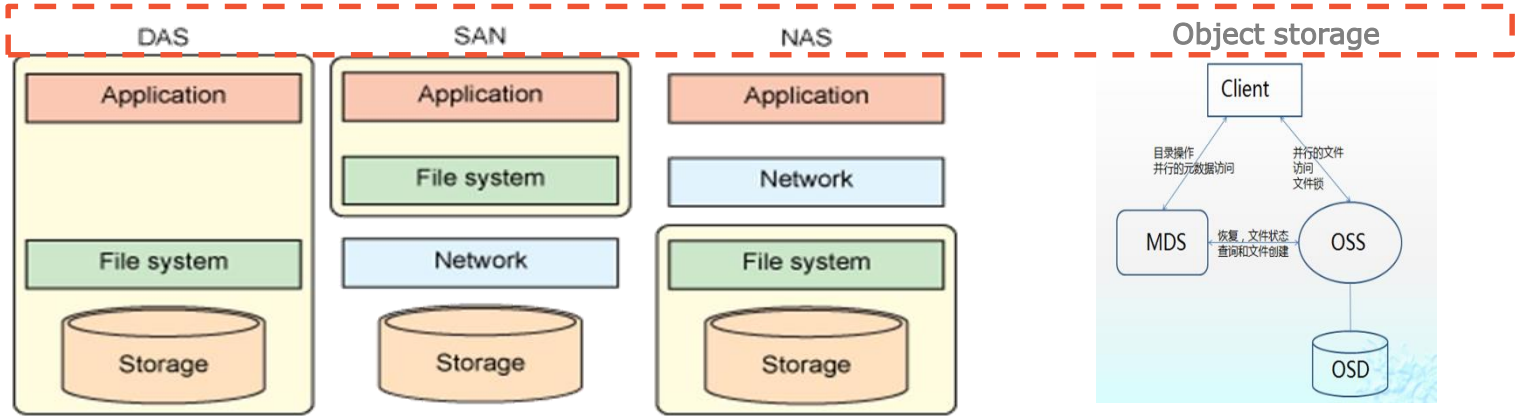

Distributed storage

Block	Cinder; Ceph; Sheepdog
File	HDFS; GlusterFS
Object	Swift ; Dynamo; Ceph
Database	Bigtable; Hbase; O/M

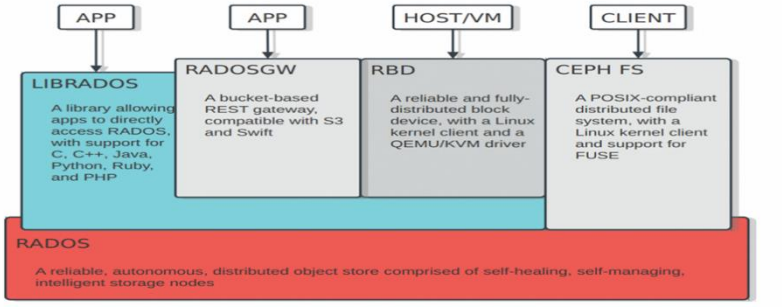

Distributed object storage

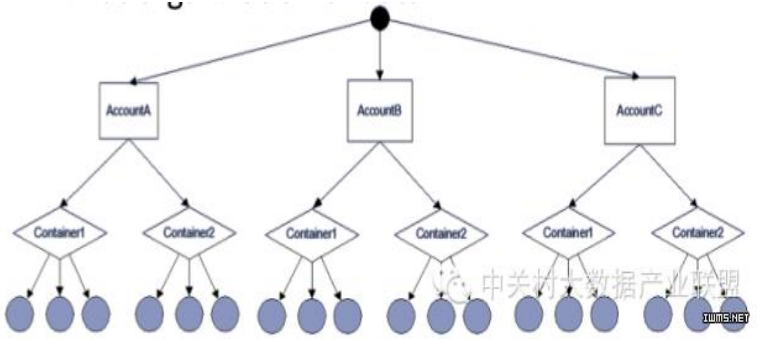

Web API, Restful

No directory, flat stucture

Distributed, high available, scalable, cost effective, multi-site

Computing

- Multi-hypervisor
- Virtualization
- Scale up vs. Scale out
- High availability

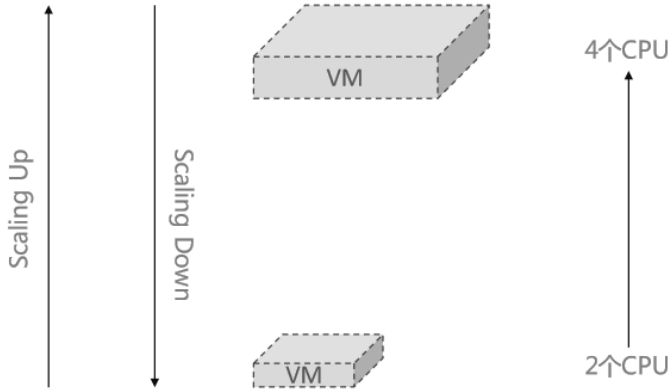

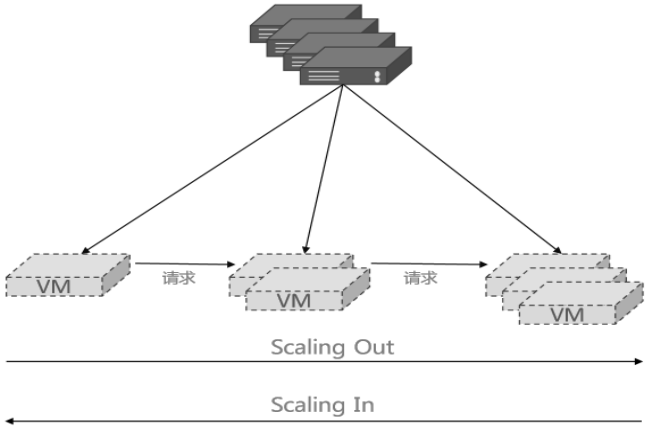

Part 2
Openstack Introduction

九州云
Cloud

What is Openstack

“Open source software for building private and public clouds”

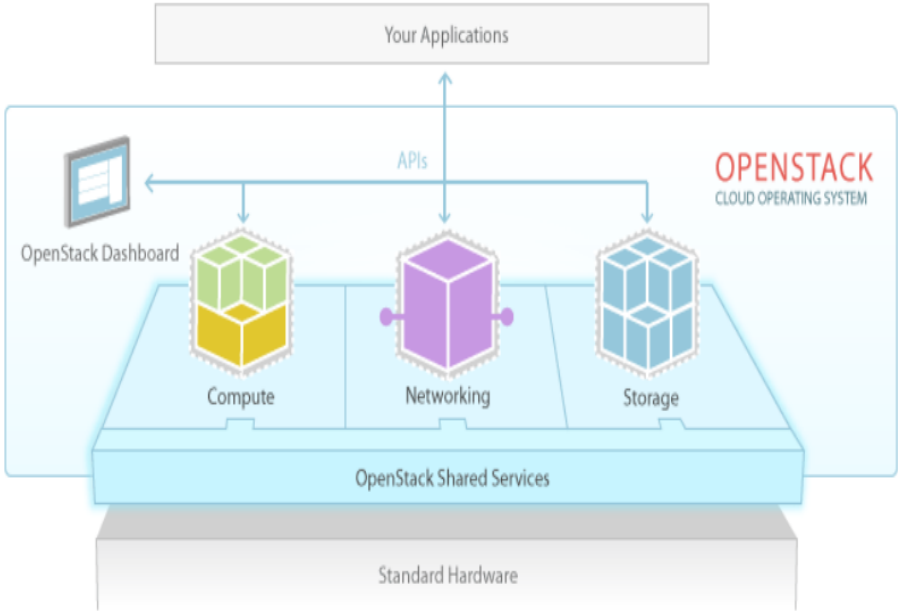

- **Virtual machine resource**
- **Storage resource**
- **Network resource**
- **Object and image store**
- **Multi-tenant:**
  - Tenant&user manage
  - Quota
  - Resource isolate

Openstack Sponsors

AT&T, Canonical, HP, IBM

Rackspace, Red Hat, Inc., SUSE, Nebula

500+ Sponsors

Contributors: 30000+

From 175 Countries

Aptira, CCAT, Cisco, Cloudscaling, Dell

DreamHost, eNovance, Ericsson, Hitachi, Huawei

Intel, Juniper Networks, Mirantis, Morphlabs, NEC

NetApp, Piston Cloud, VMware, Yahoo!

Nutanix, NxtGen, Oracle, Pactera, Parallels

PayPal, RiverMeadow Software, Scality, Seagate

Servergy, Servosity, Smartscale Systems, Solinea

StackLabs, SwiftStack Inc, Symantec, Tesora, Thales Group

Transcend Computing, UnitedStack Inc., VCE, Virtual Bridges

Western Digital Corporation, Wind River, Xemeti, ZTE Corporation

A10 Networks, Adaptive Computing, Alcatel Lucent, Anuta Networks

Arista Networks, Avaya, Blue Box, Brocade, Bull

Canadian Web Hosting, Cloudwatt, Comcast, Coraid

Cumulus Networks, Cyan Inc., EMC, Evident.io, F5 Networks

Flexiant, Fujitsu, Fusion-io, Gale Technologies, GoDaddy

Openstack Releases and Projects

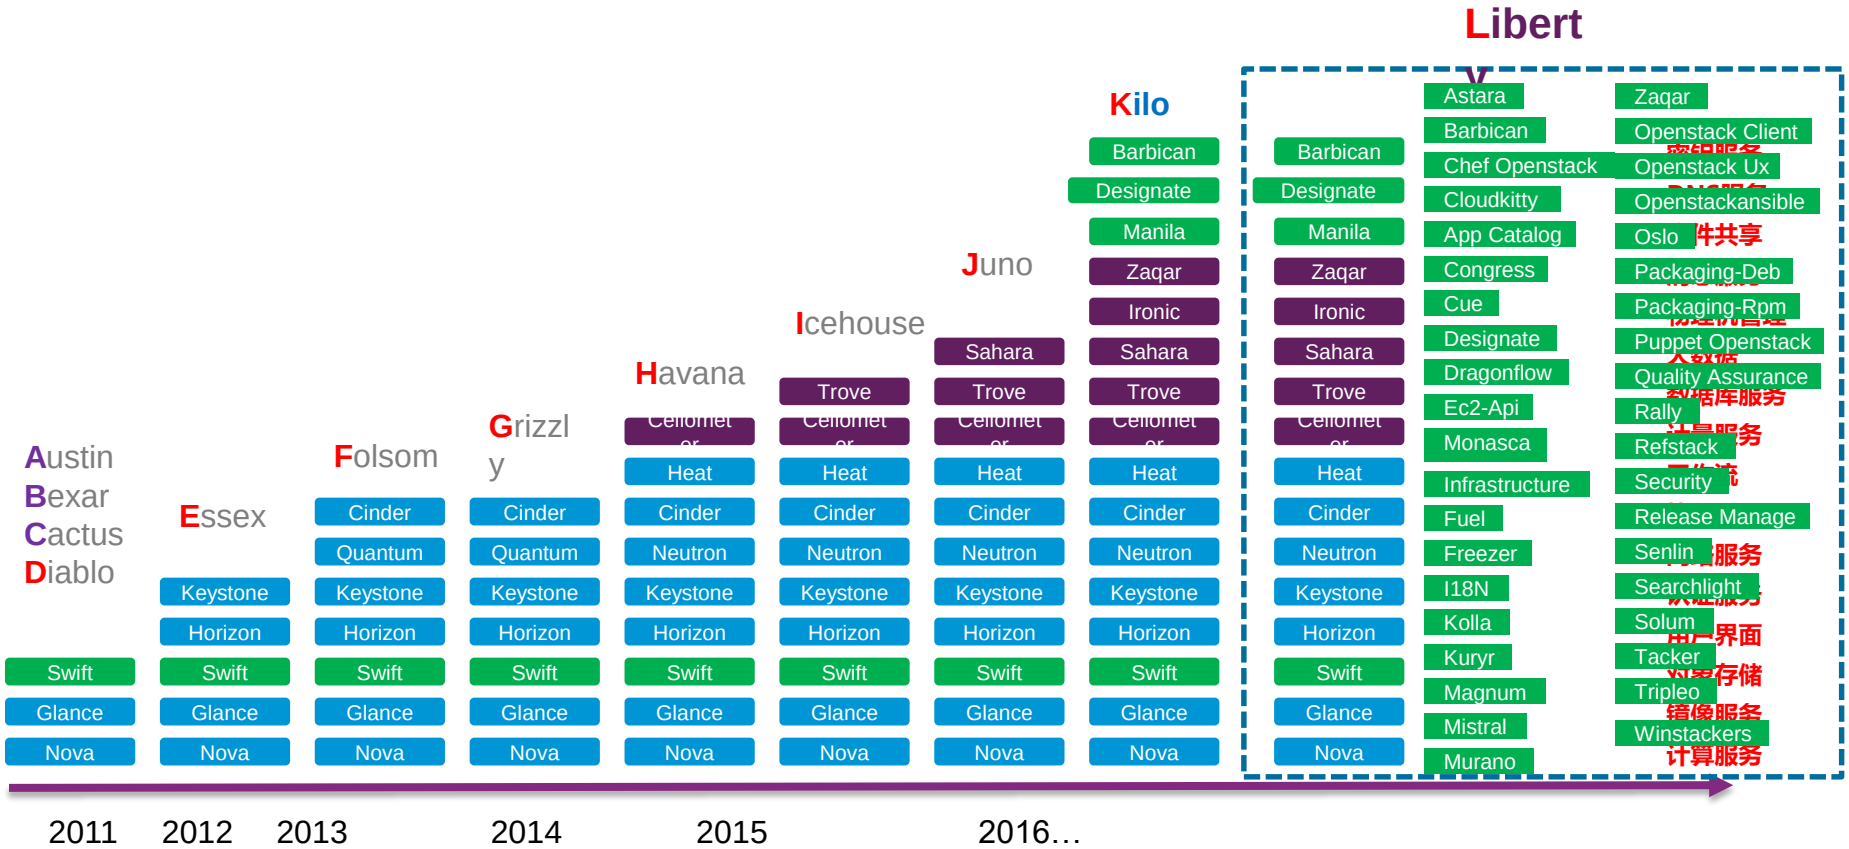

API Centric

Commercial Cloud Platform

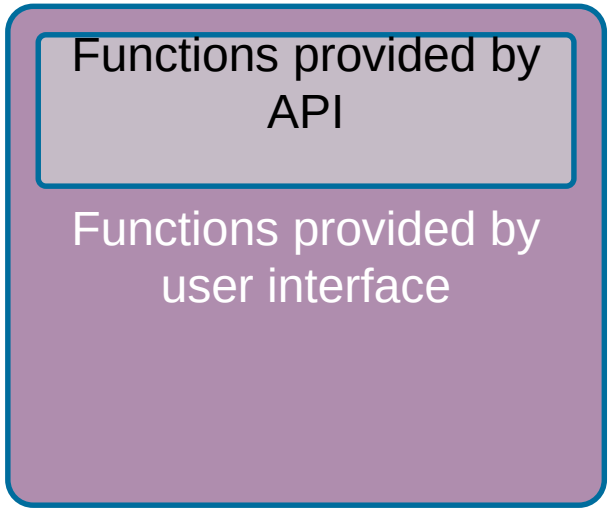

- Operators oriented
- User experience
- Manual operation

Openstack

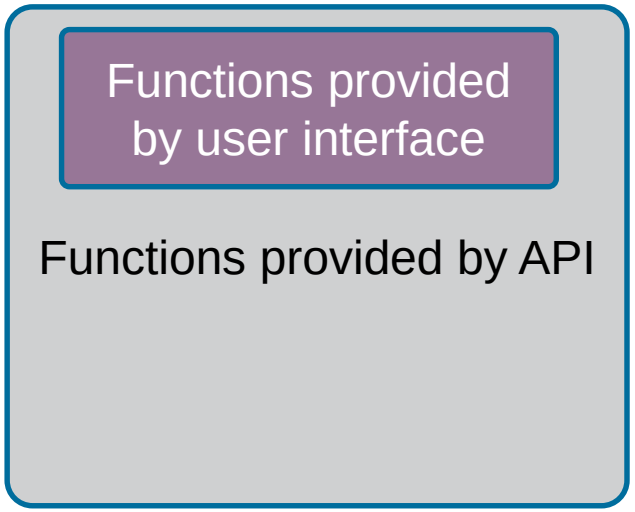

- API(object) oriented
- Distributed, scale
- Automate operation by machines

Openstack Core projects

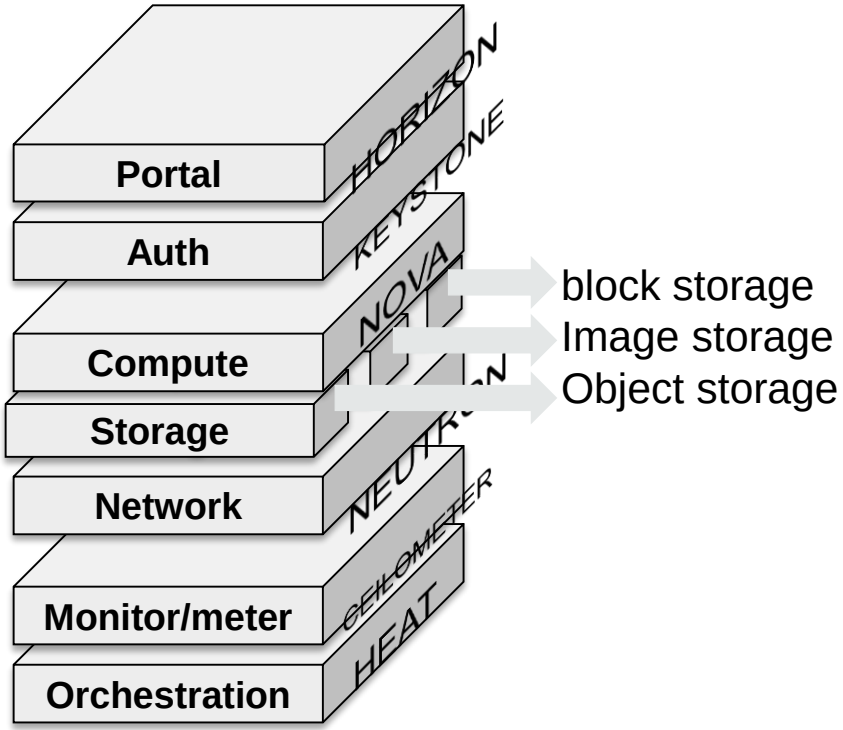

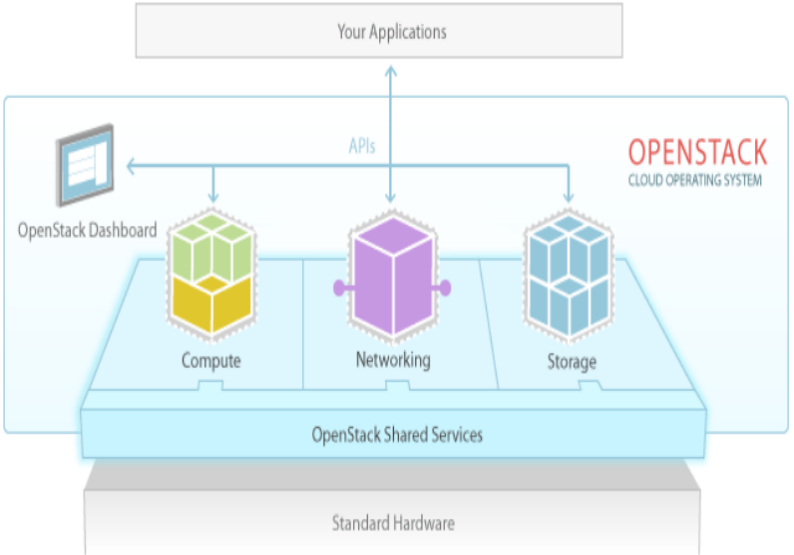

OpenStack Conceptual Architecture

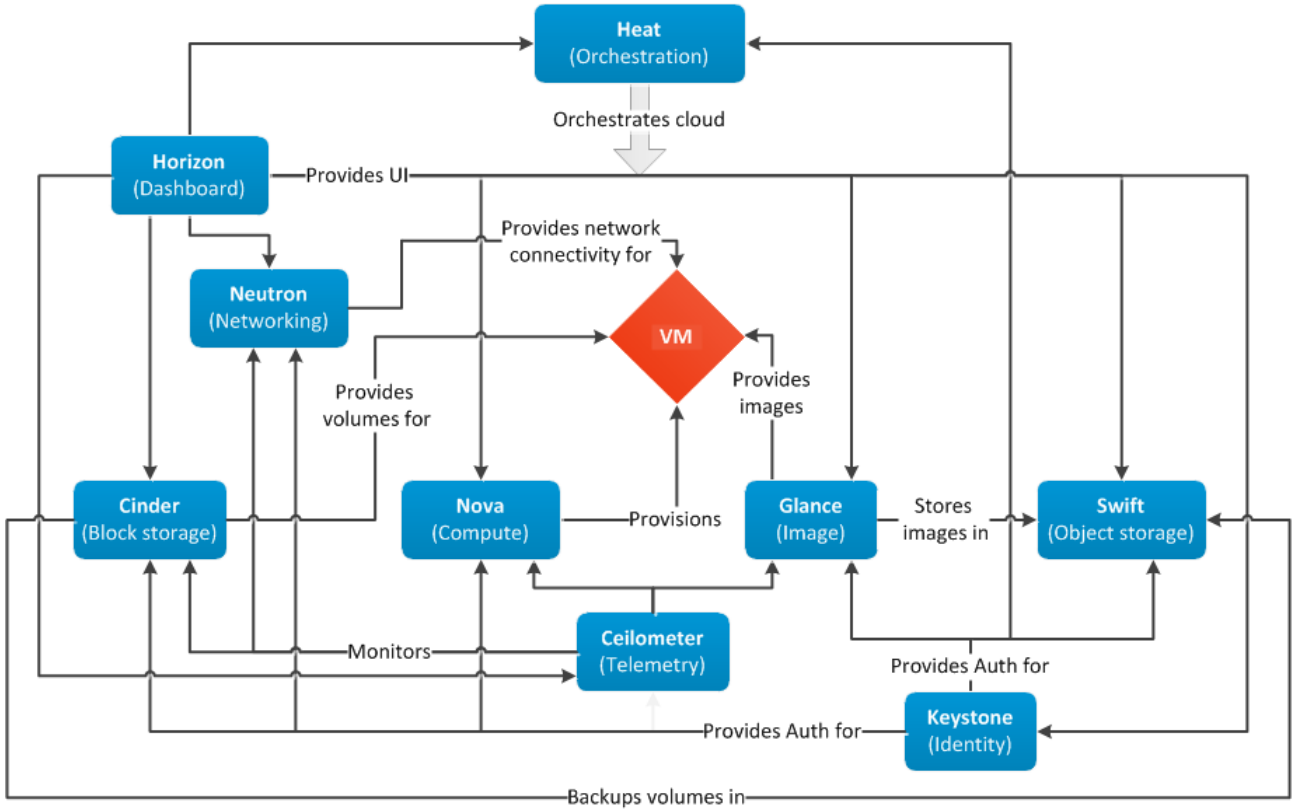

Basic OpenStack Physical Components

This sample three-node configuration shows the basic components of an OpenStack installation

- Cloud Controller node
- Network node
- One or more Compute nodes

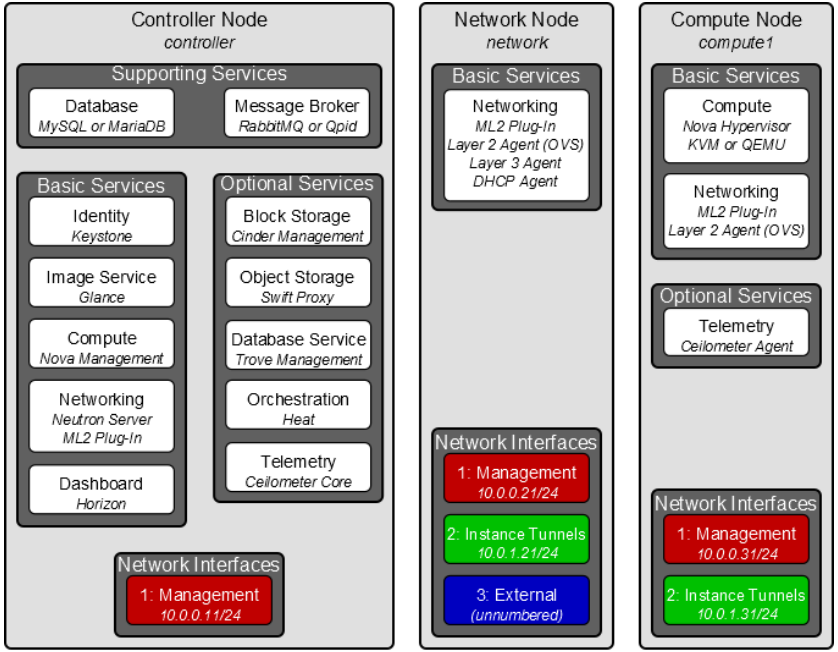

NOTICE: All of the basic components can exist on the same device or horizontally distributed, with with multiple network, compute, and controller nodes.

Keystone in OpenStack Conceptual Architecture

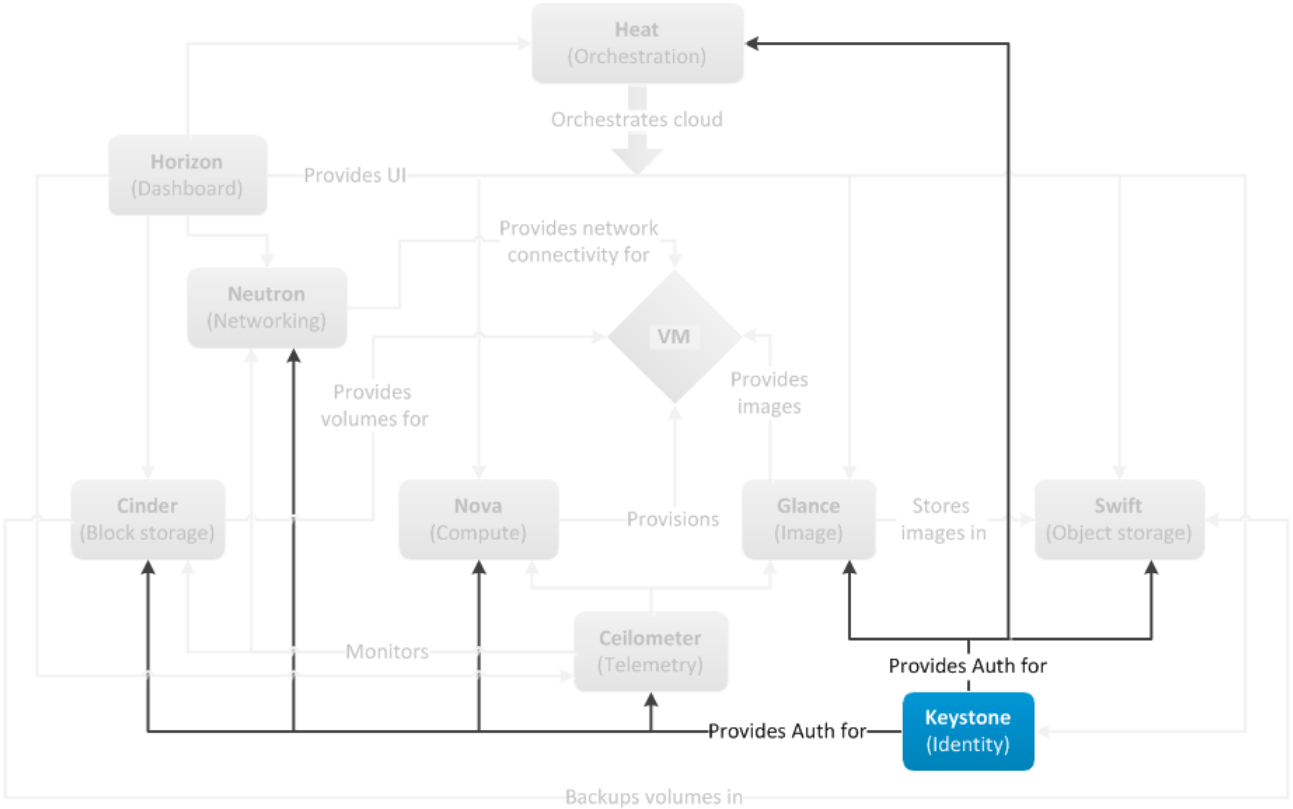

Tenant, User & Role

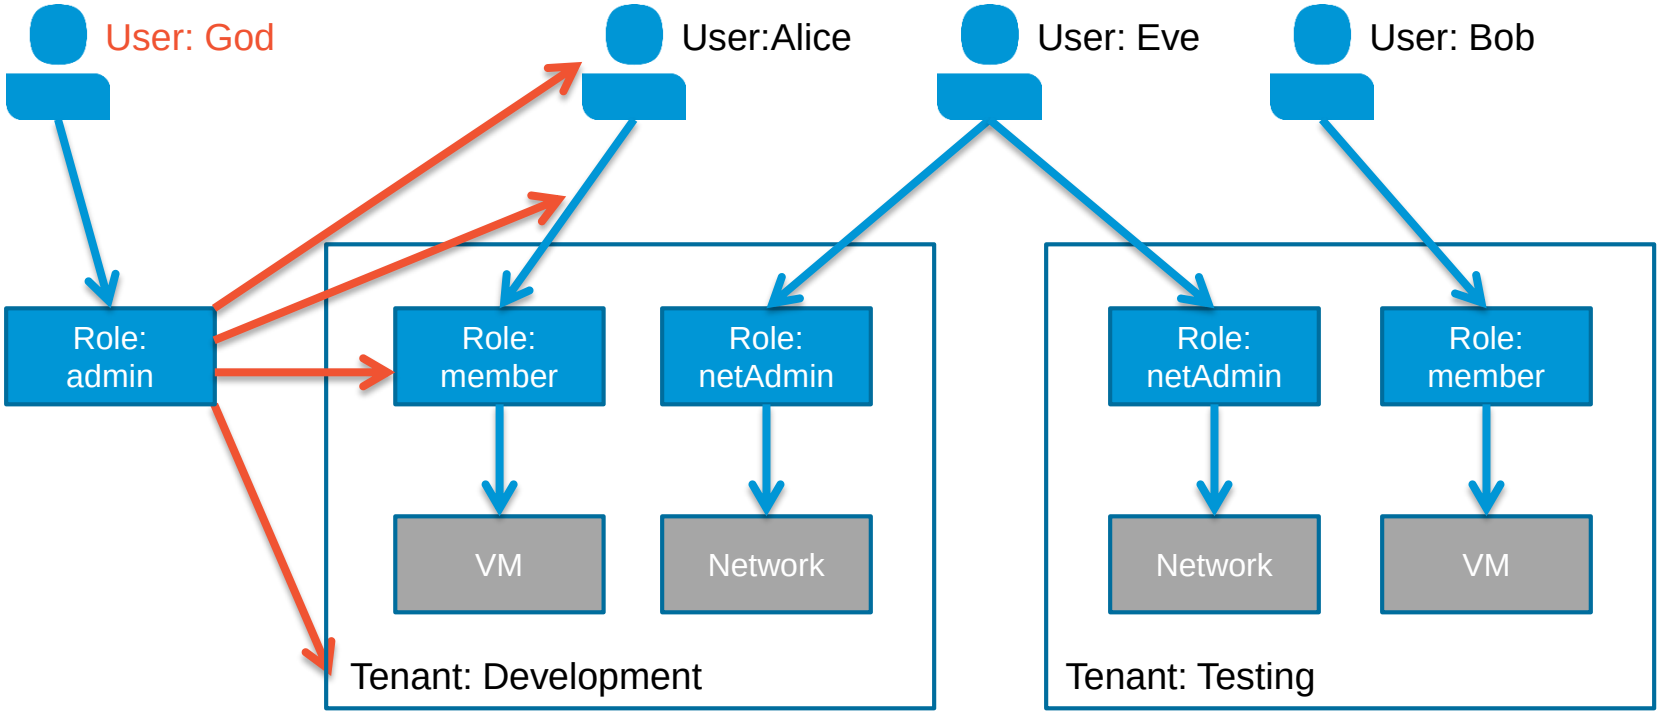

Keystone Backends

	SQL	LDAP	KVS	Memcach	Redis	File	Memory	Noop
Identity	√	√	√					
Token	√		√	√				
Cache				√	√		√	√
Catalog	√		√			√		
Policy	√					√		

NOTICE: The backend for each function is defined in the keystone.conf file.

TIP: IDM can be integrated with keystone

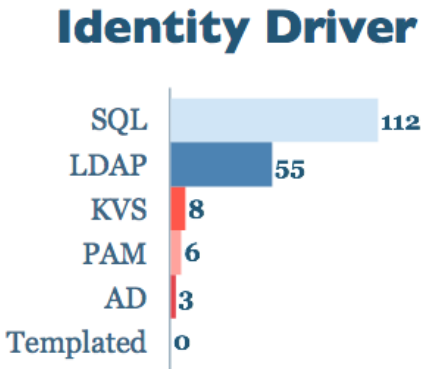

Glance
- Image Service

九州云
Cloud

Glance in OpenStack Conceptual Architecture

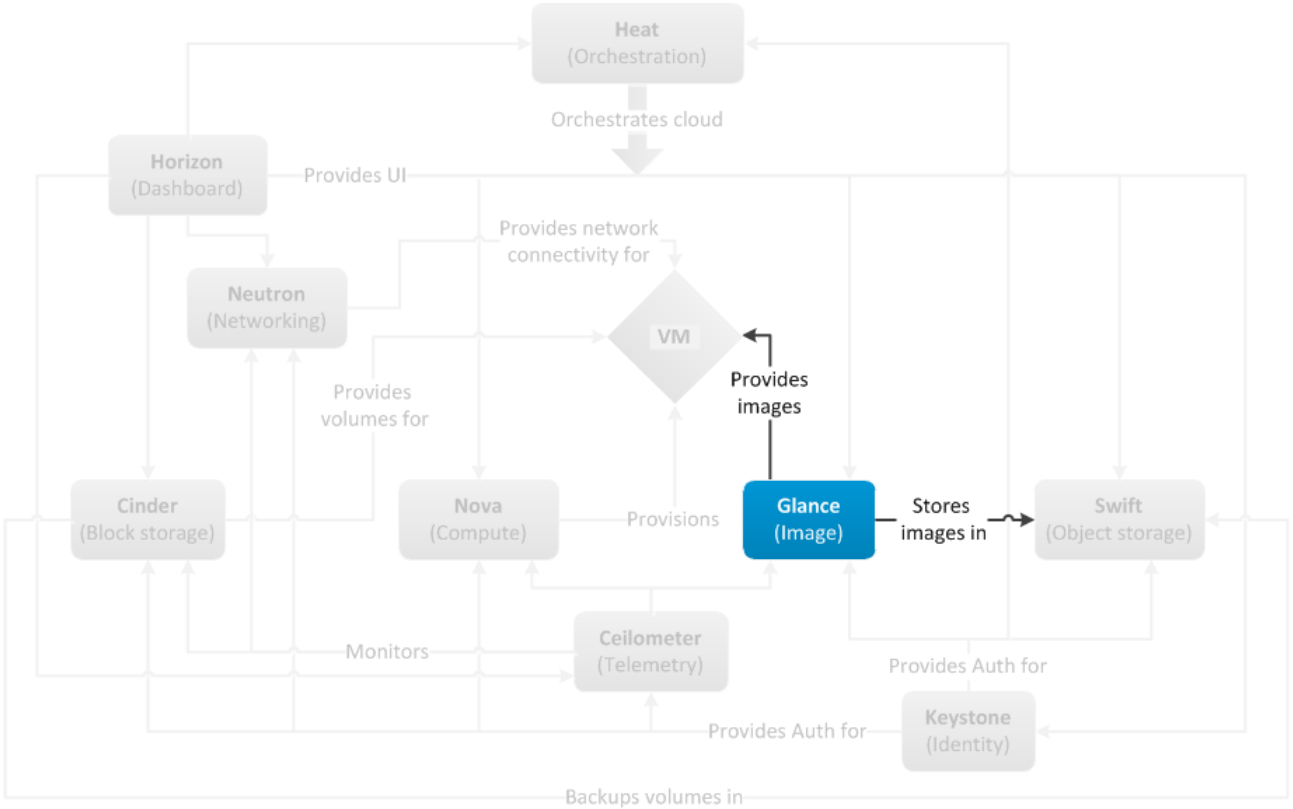

Glance Architecture

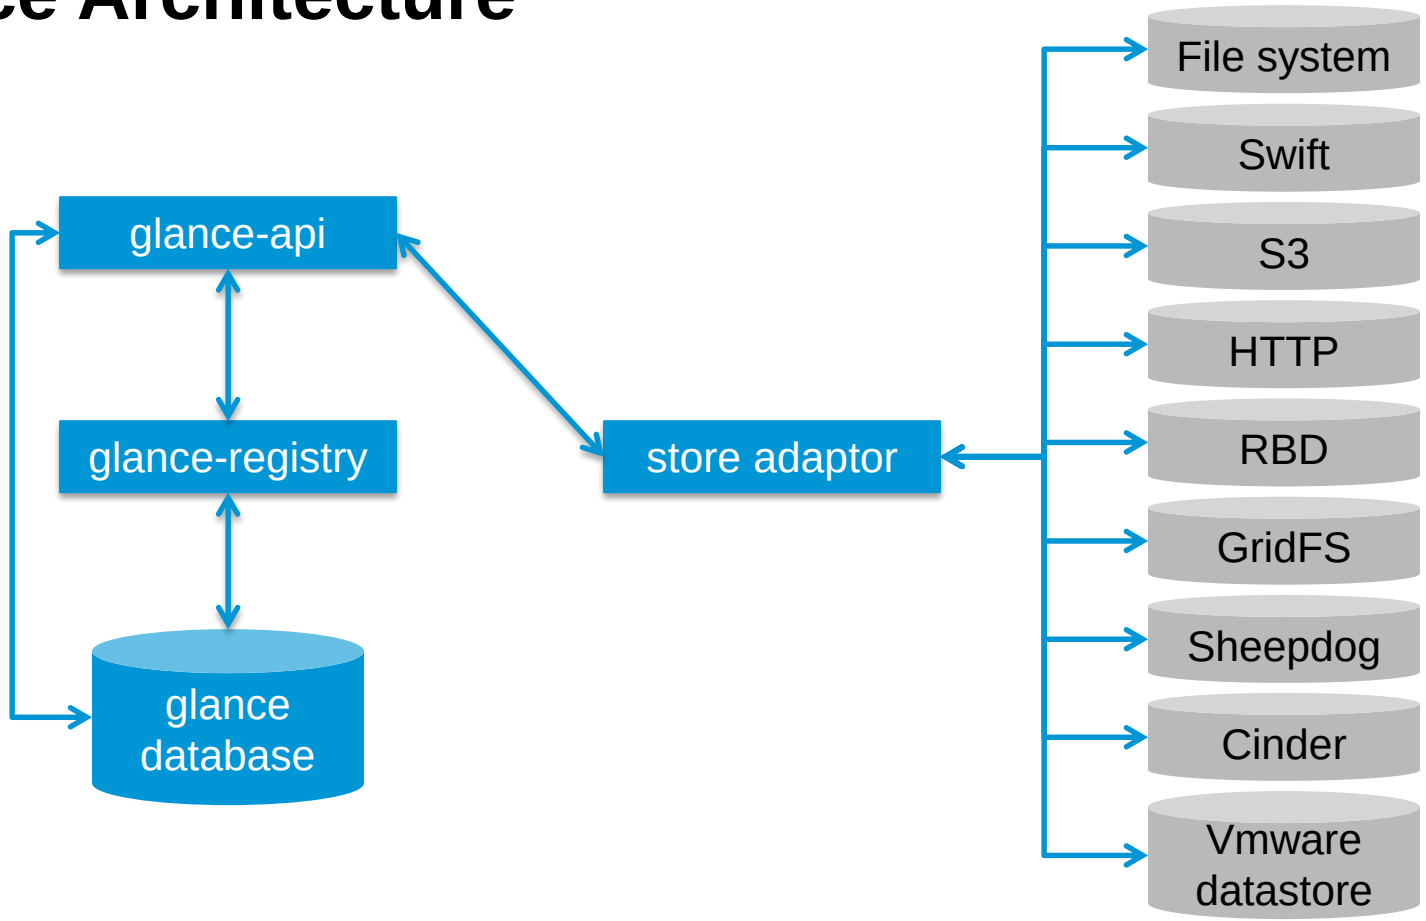

Glance supported Image disk format

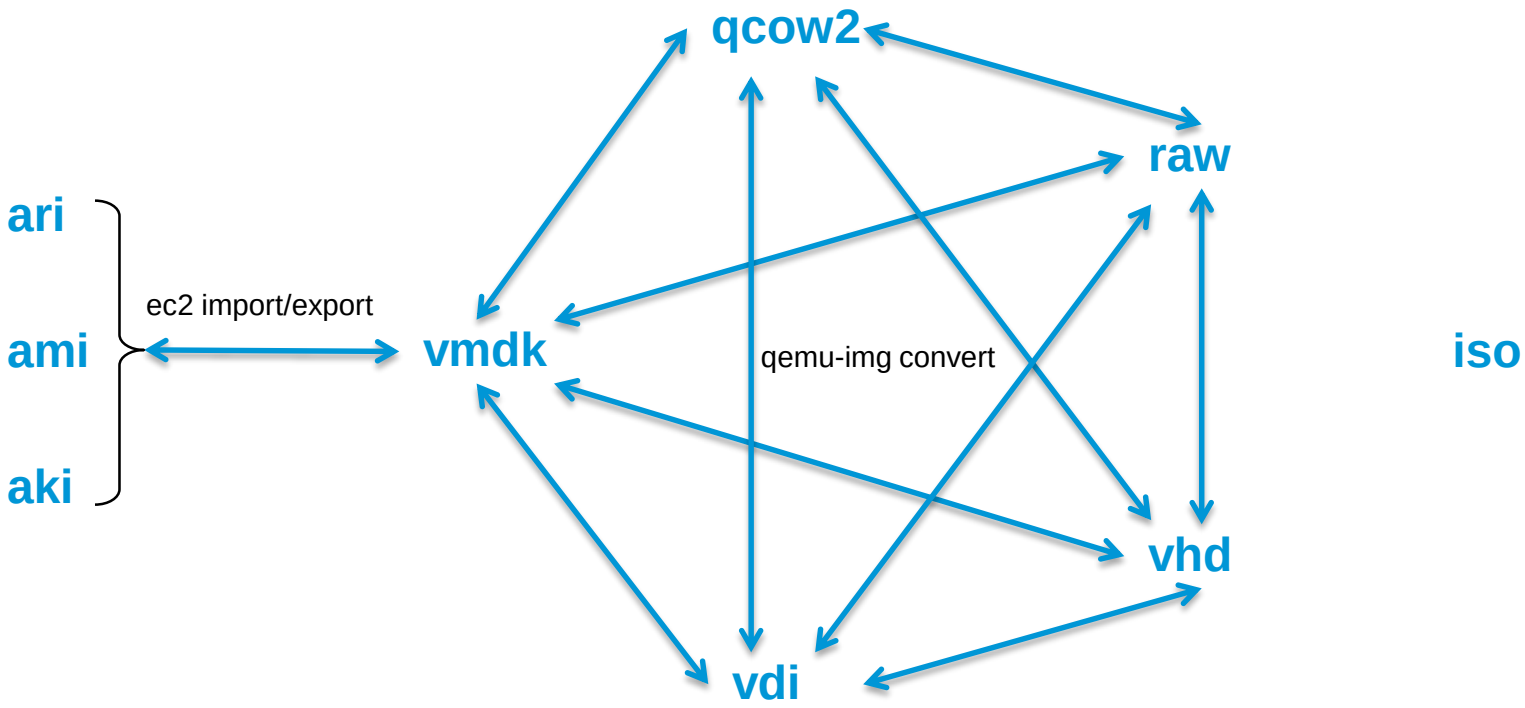

Nova
- Compute Service

Nova in OpenStack Conceptual Architecture

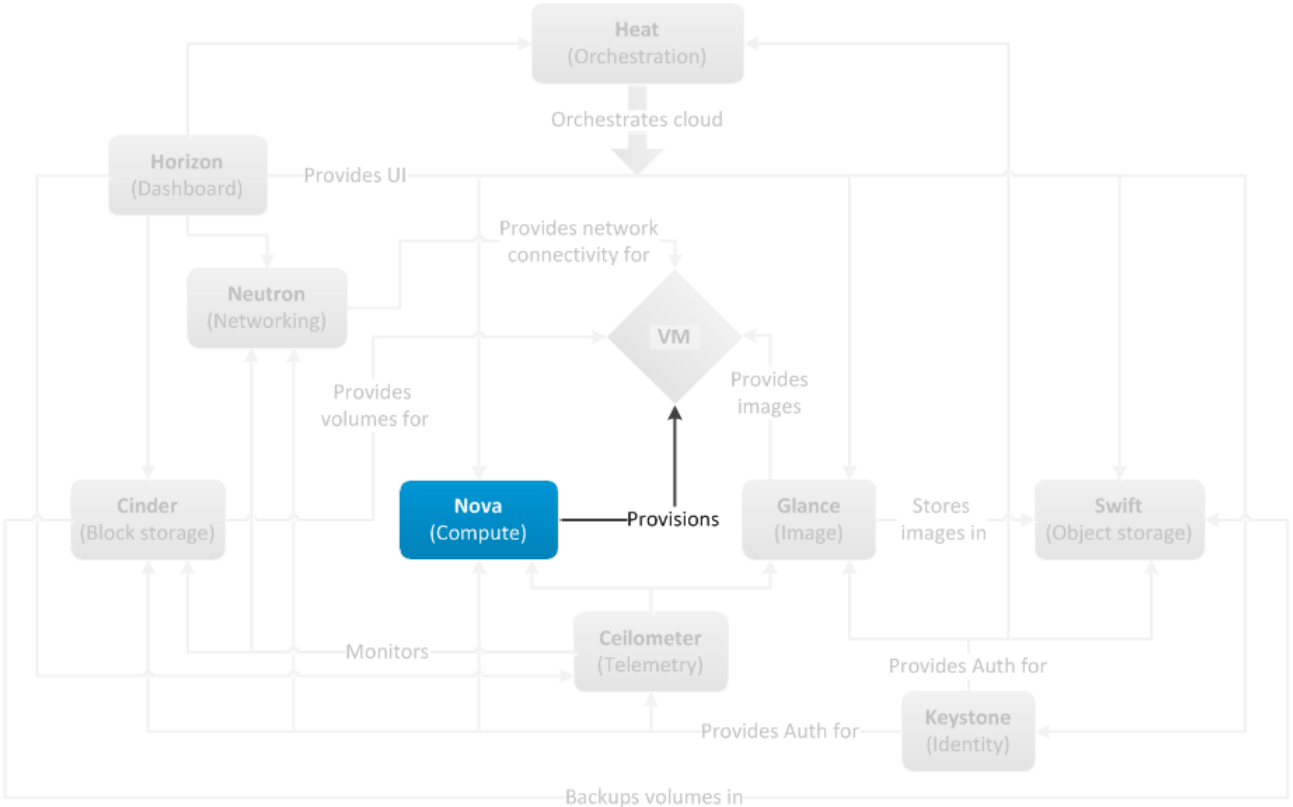

VM Provisioning

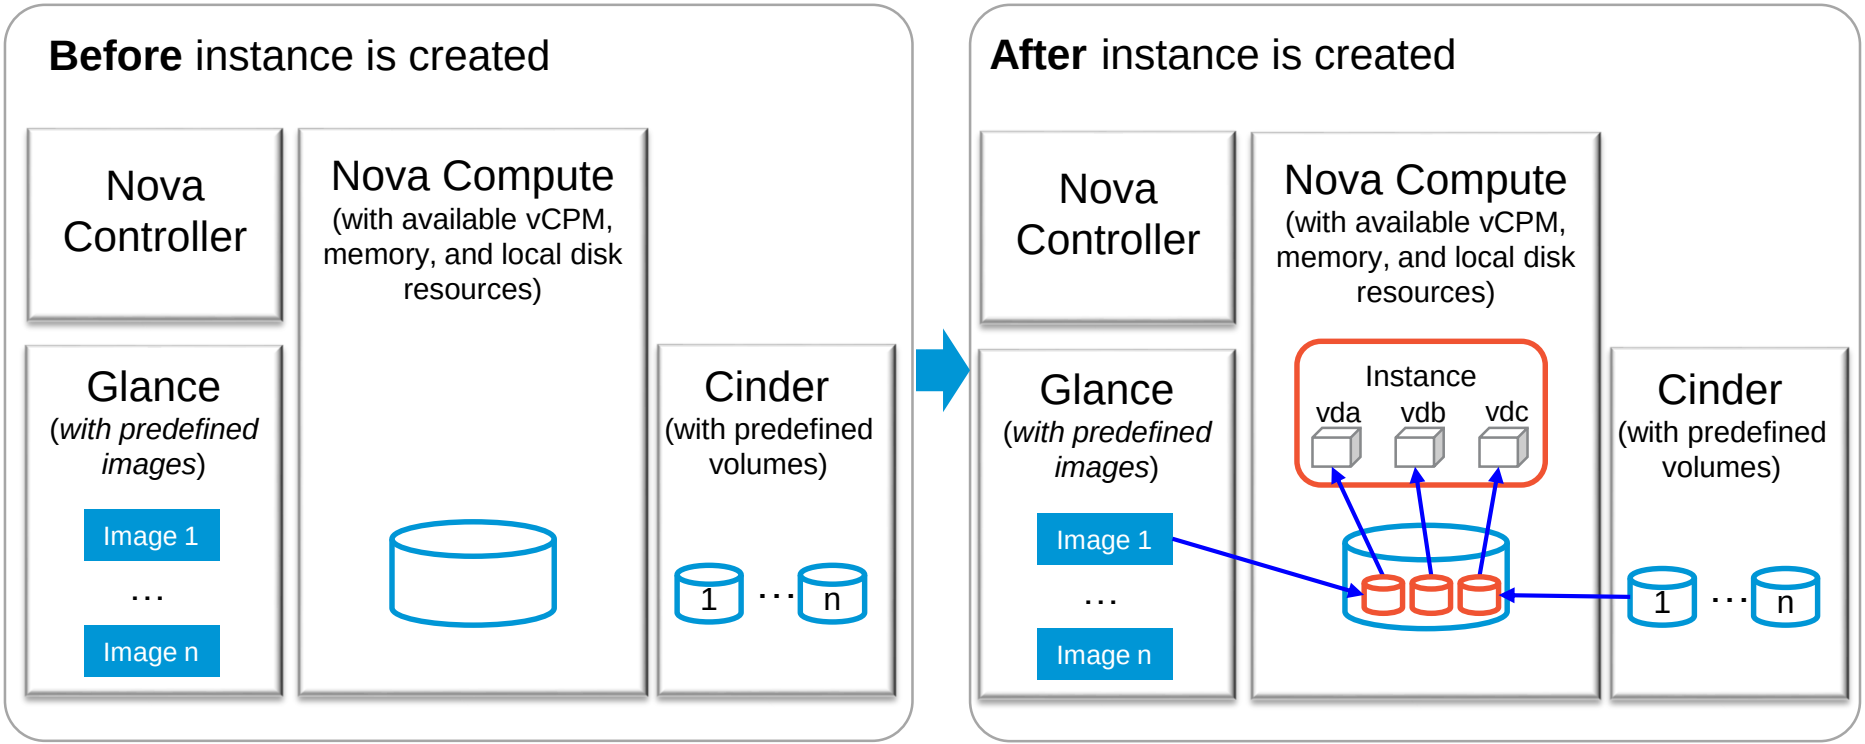

Nova Supported Hypervisors

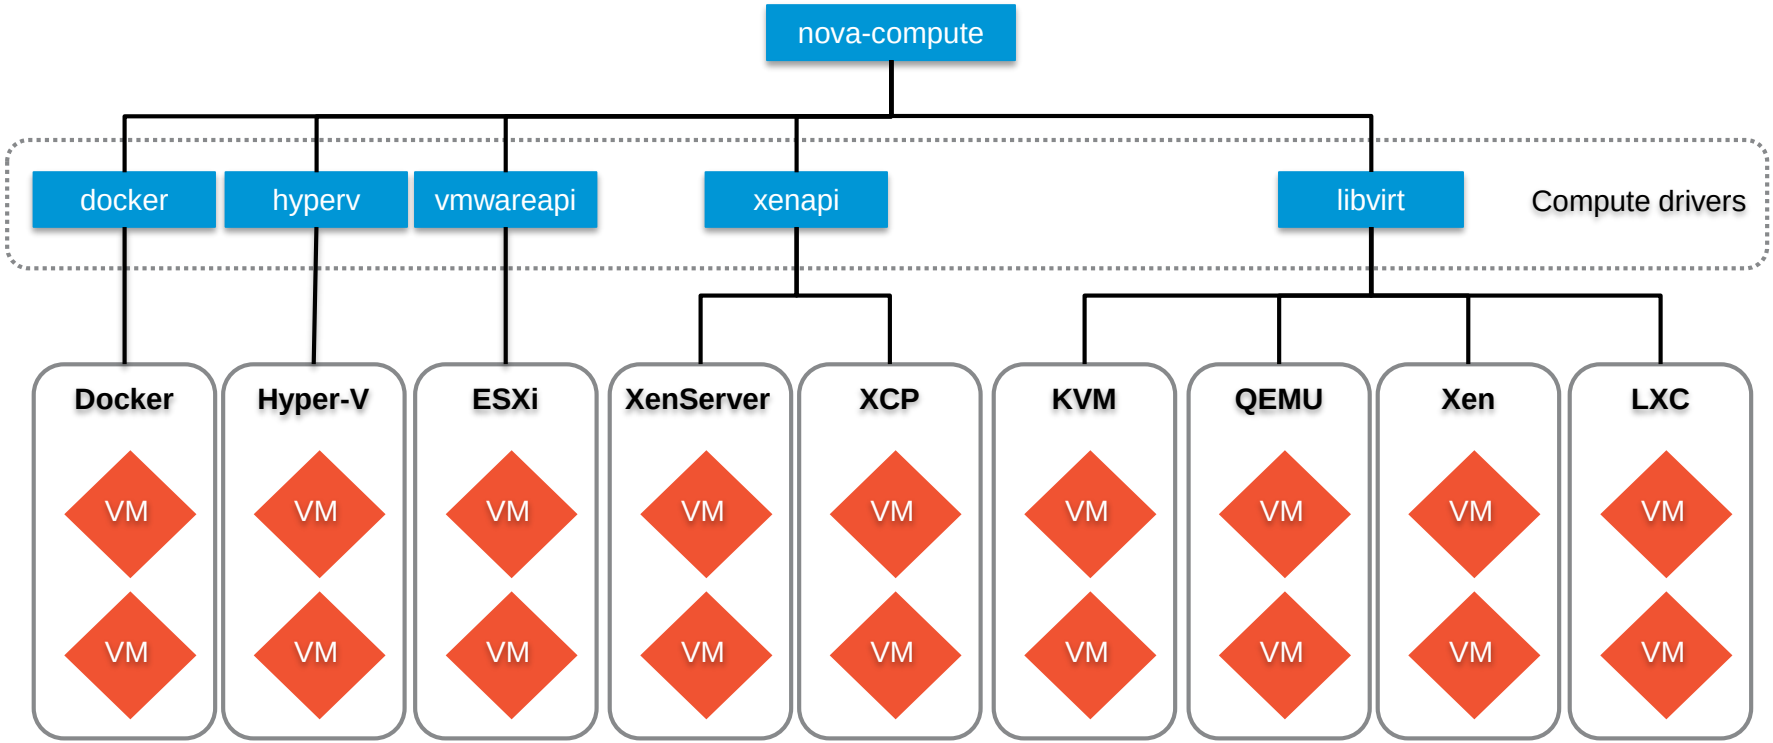

Horizon
- Dashboard Service

九州云
Cloud

Horizon in OpenStack Conceptual Architecture

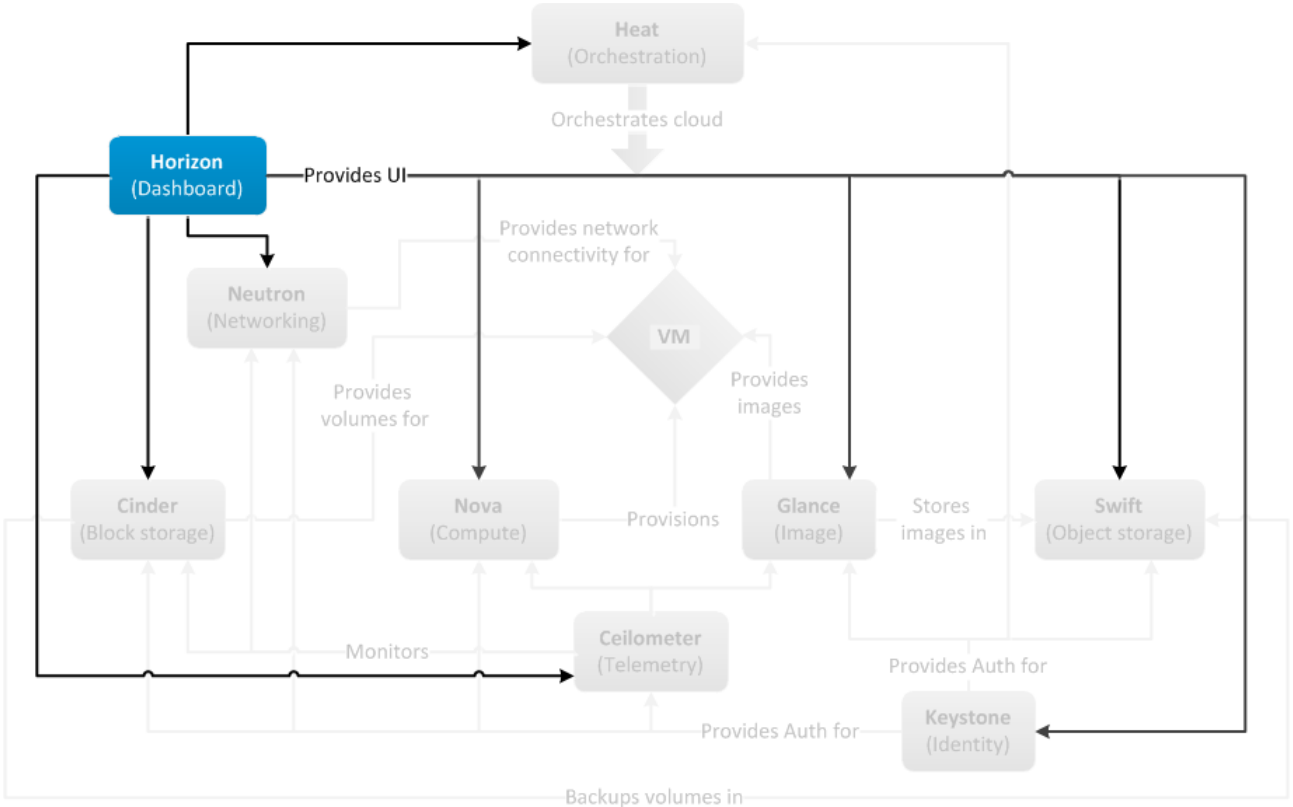

Horizon – The Face of OpenStack

Web-based graphical user interface to OpenStack services (e.g. Nova, Swift, Keystone, Glance, etc.)

- Stateless, can be scaled horizontally
- Supports a **subset** of OpenStack API, **not all**

Neutron
- Networking Service

九州云
Cloud

Neutron in OpenStack Conceptual Architecture

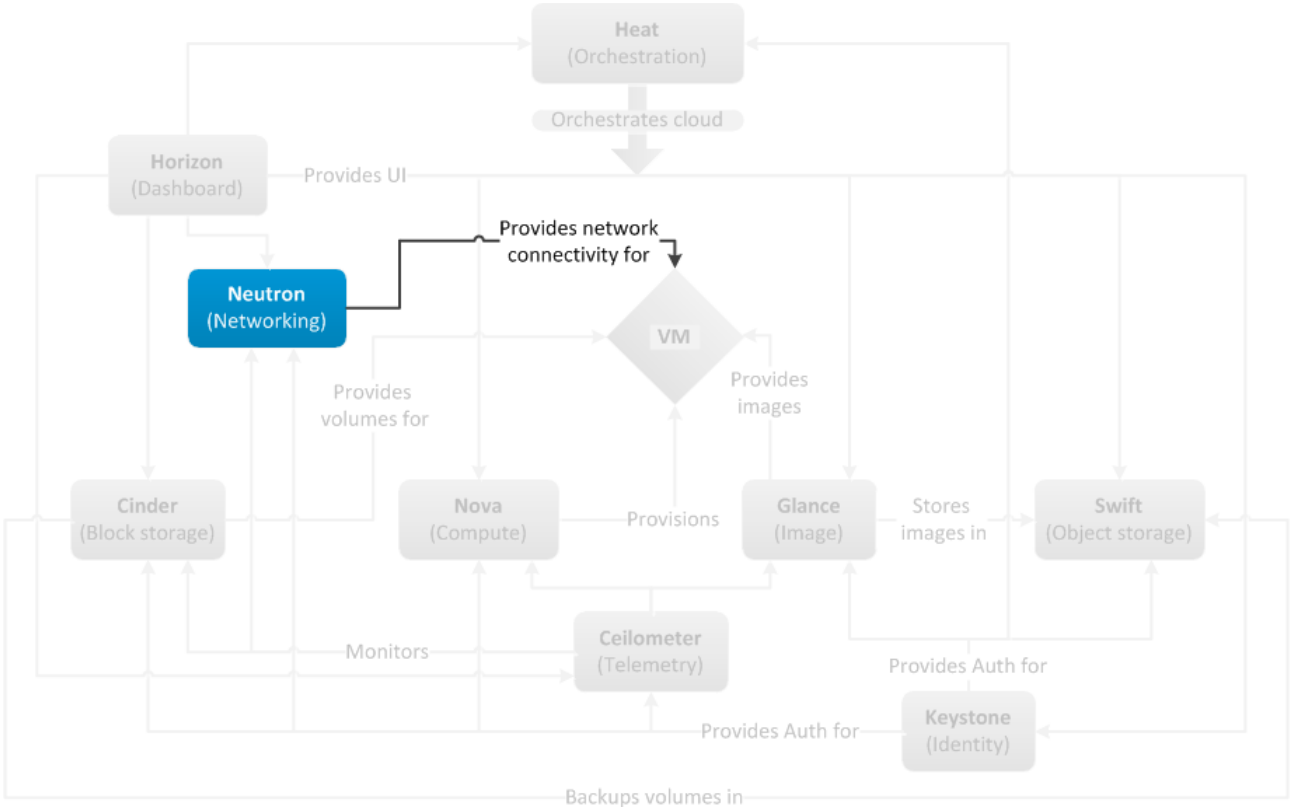

Neutron Overview

Network as a Service

- Allow tenants to control their own private networks
  - Isolated
  - Rich topologies
  - Choose their own IP addressing scheme, even if those IP addresses overlap with those used by other tenants
- Open APIs for operating logical networks
  - Technology agnostic. Separate logical operations and backend provisions
  - Backend technologies are provisioned/configured by plugins/drivers
- Modular design
  - API specifies service, vendor provides its implementation
  - Extensions for vendor-specific features
- Advance services support
  - Load balancing, VPN, Firewall
- Support emerging network technology
  - SDN/OpenFlow-based network
  - Overlay tunneling (VXLAN, NVGRE, STT, …)

Neutron Deployment Use Cases

Per-tenant Routers with Private Networks

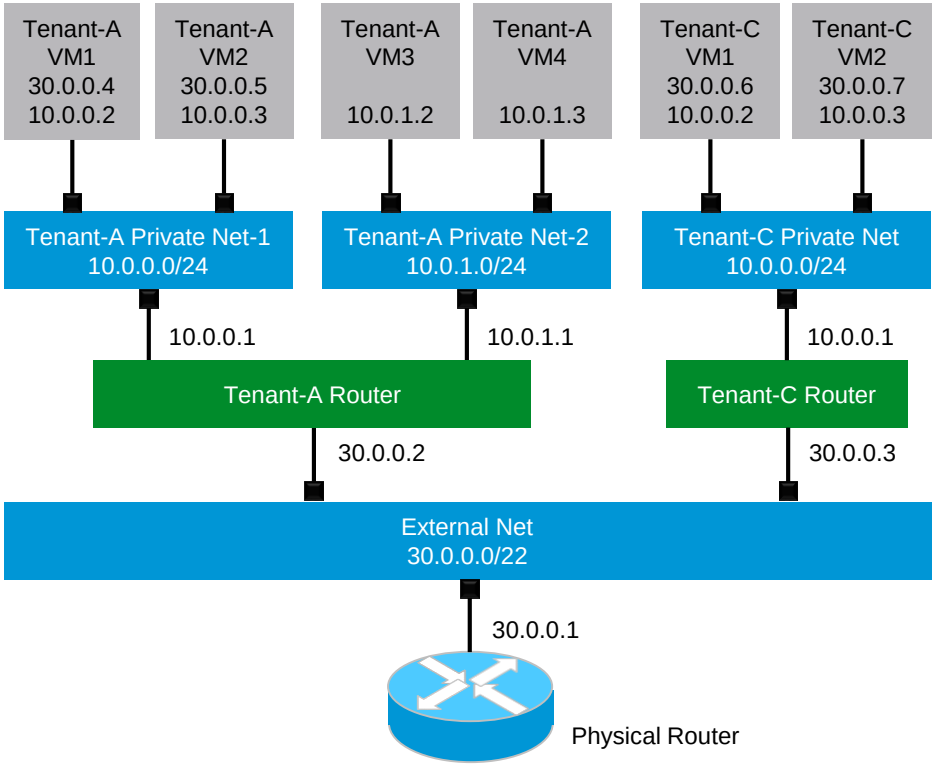

Advanced Services

Load balance

- HAProxy
- Embrane
- Netscaler
- Radware

VPN

- IPsec
- Cisco IPsec
- Cisco CSR

Firewall

- Iptables
- Varmour

Cinder
- Block Storage Service

九州云
Cloud

Cinder in OpenStack Conceptual Architecture

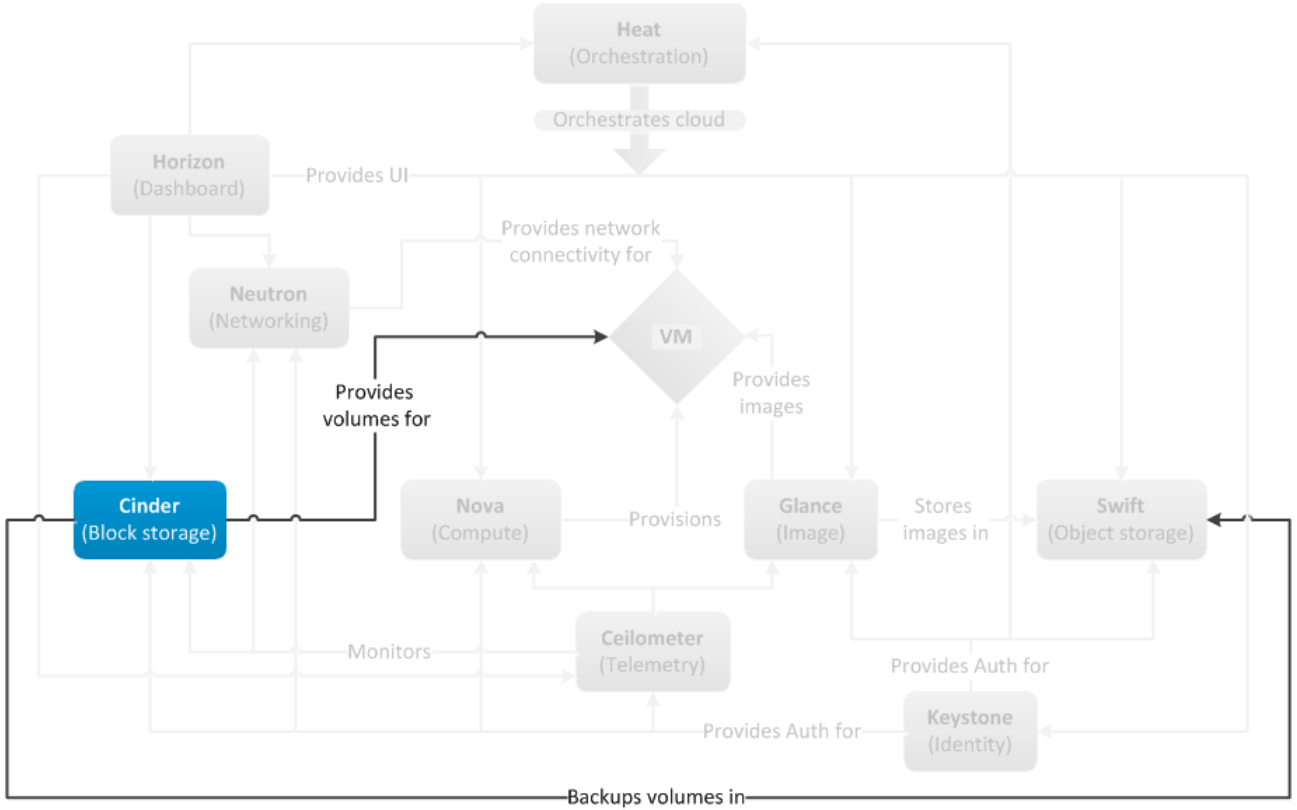

OpenStack Storage Overview

Comparison of standard boot storage, block storage, and object storage technologies in OpenStack

On-instance / ephemeral	Volumes block storage (Cinder)	Object Storage (Swift)
OS and scratch space	Persistent storage volume for VM	VM images and data
Persists until VM is terminated	Persists until deleted	Persists until deleted
Access associated with VM	Access associated with VM	Available from anywhere
Implemented as filesystem	Mounted via Cinder controlled protocol (iSCSI, FC, etc.)	REST API
Size settings based on flavors	Sizing based on need	Easily scalable for future growth
Example: 10GB first disk, 30GB/core second disk	Example: 1TB extra hard drive	Example: 10s of TBs of dataset storage

Cinder Overview

- Provides volumes to VMs provisioned by OpenStack
- Uses RESTful APIs to communicate with other OpenStack components
- Storage can SSD or HD and reside in the compute node or a dedicated storage device
- Uses iSCSI or FC to communicate between VM and dedicated storage device
- Supports NFS and CIFS file systems
- Volumes cannot be shared by servers
- Can be backed up to Swift

NOTE: By default, primary OS volume doesn't come from Cinder, but rather from ephemeral storage, however, you can create a Cinder boot volume

Cinder Components

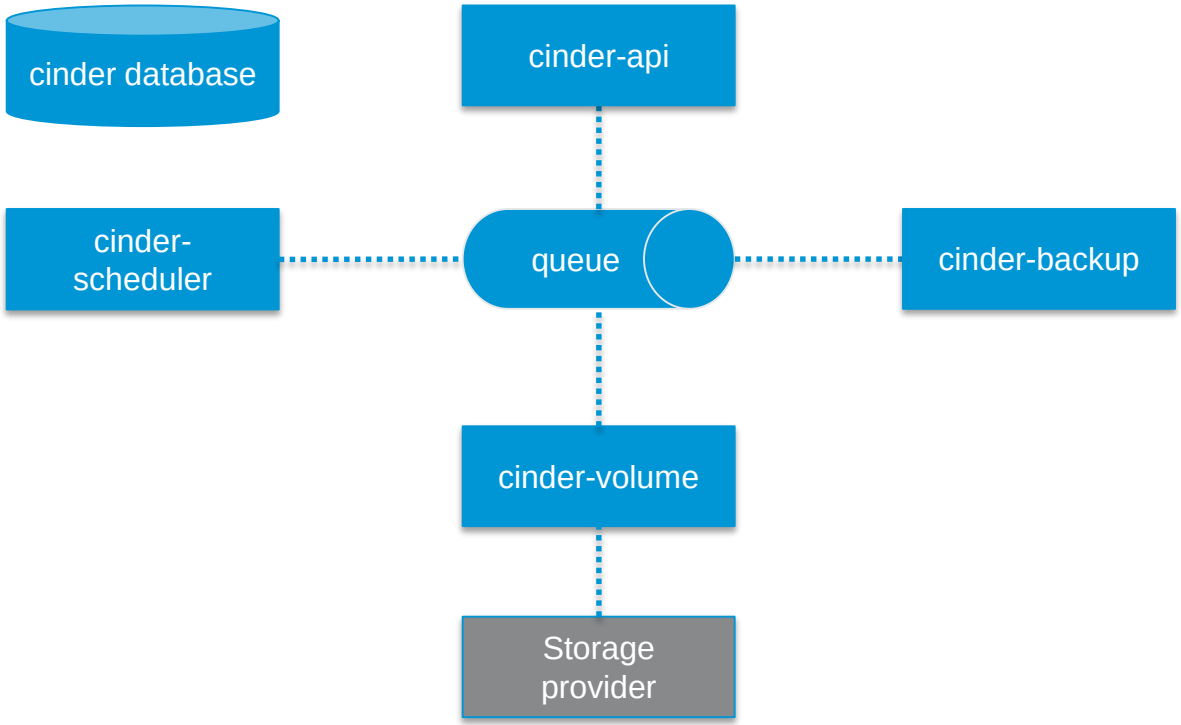

Snapshots

- A snapshot in cinder is a read-only point in time copy of a volume
- The snapshot can be created from a volume that is currently in use or in an available state
- The snapshot can be used to create a new volume via create from snapshot.

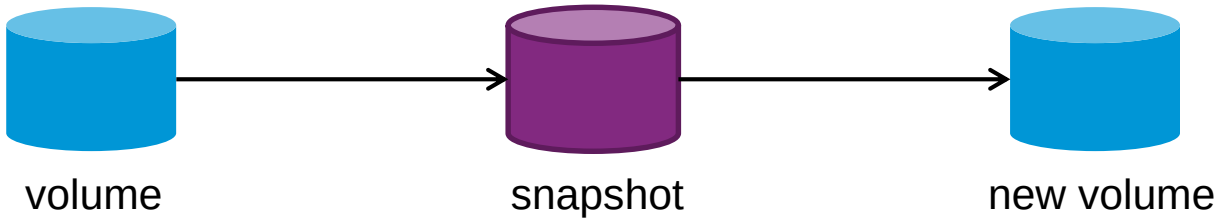

Swift
- Object Storage Service

九州云
Cloud

Swift in OpenStack Conceptual Architecture

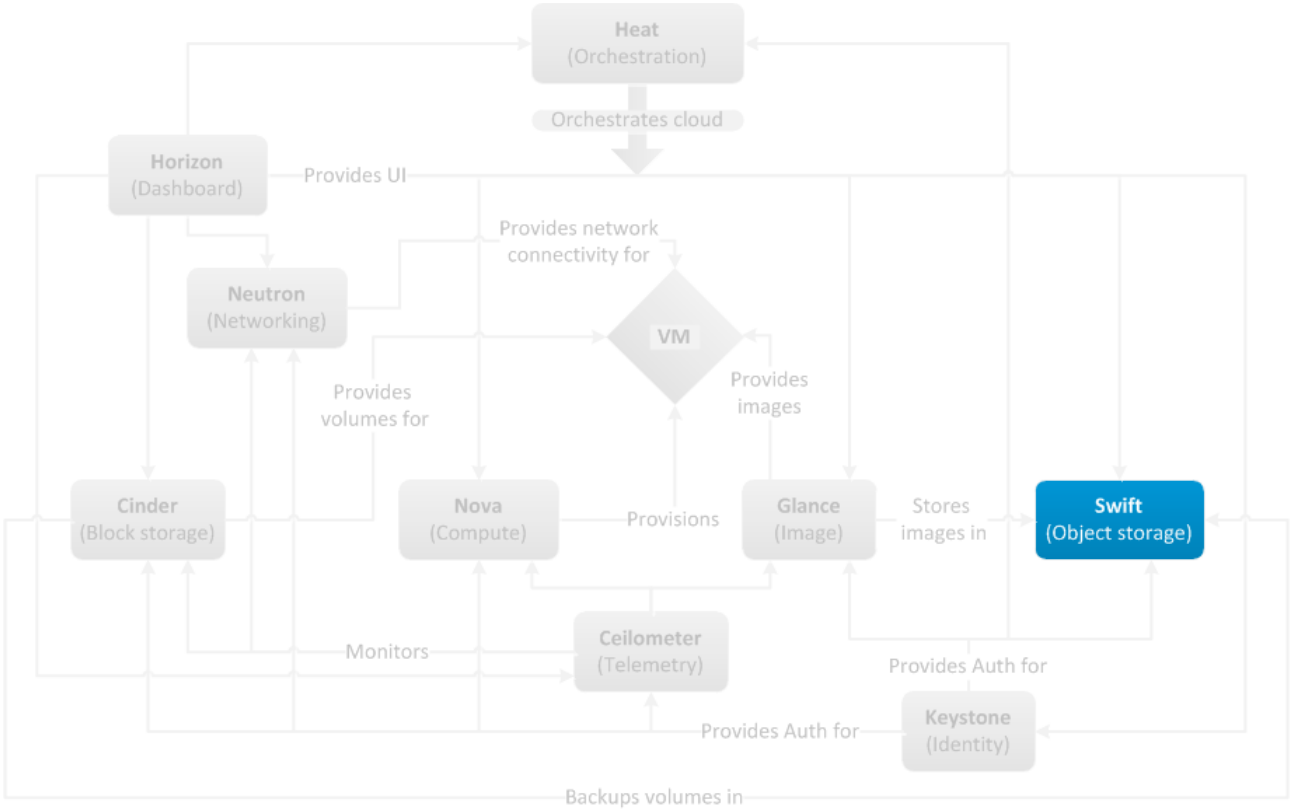

Object Storage

Swift is an OpenStack service that allows you to store and retrieve data objects

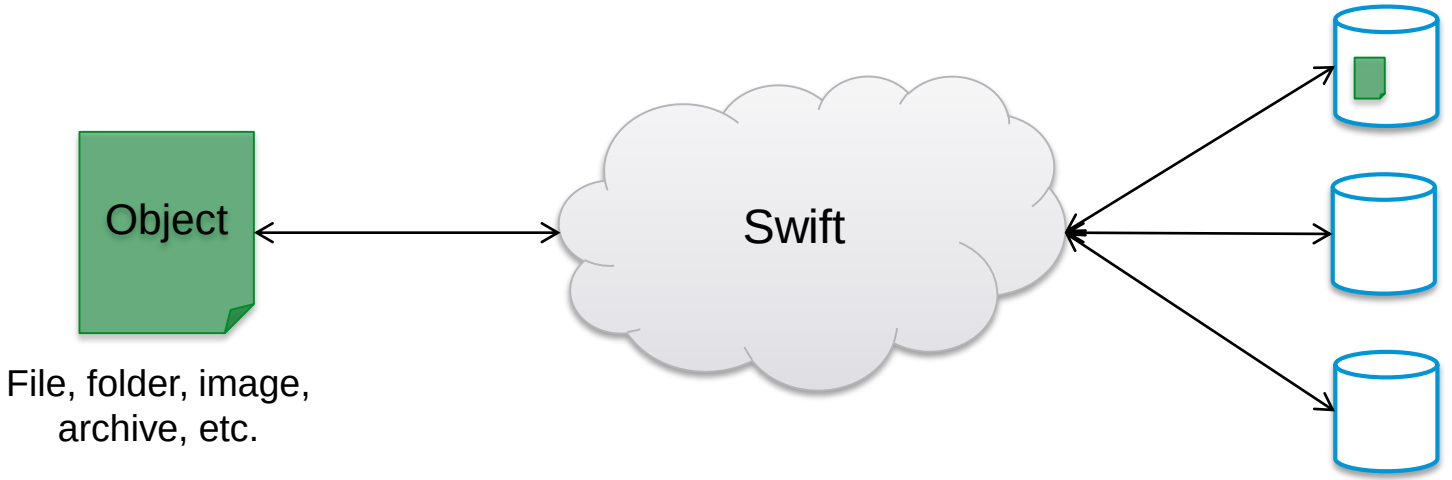

Swift Overview

Software to reliably store billions of objects distributed across standard hardware

Object storage

- Support multi-tenancy
- It’s not a file system
- Access with RESTful API
- Augments SAN/NAS/DAS
- Runs on commodity hardware
- Hardware agnostic (no need for RAID)
- Fully distributed evenly throughout the system
- No central database
- Provides data redundancy
- Eventually consistent
- Auditors to check staleness of data
- Optimized for scaling to zettabytes

User View: Accounts, Containers, Objects

User accesses an account

Account owns containers

Containers contain objects

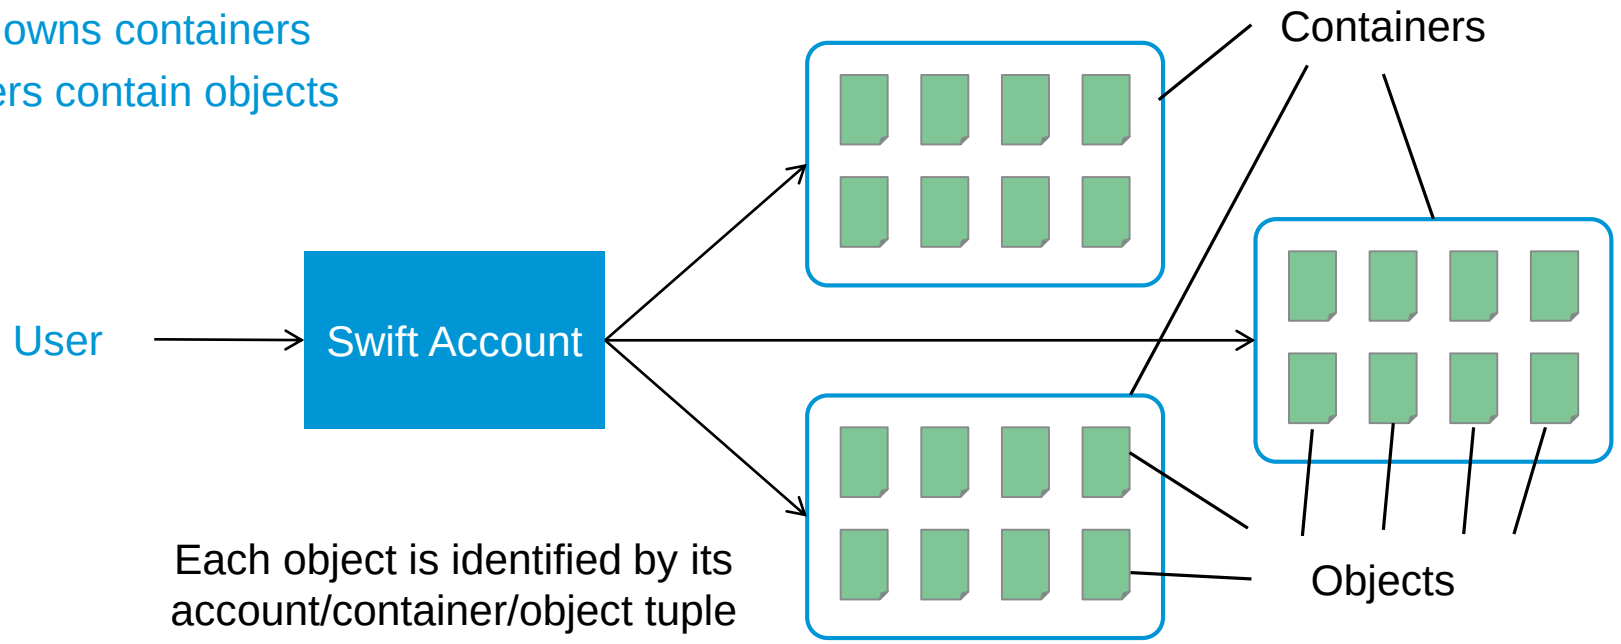

Partitions

Unit of Swift storage in which objects are placed

- Swift are divided into a power of 2 number of partitions
- Partitions are sections of address space, not disk space
- Directory on a disk with a hash table that describes the contents of a partition

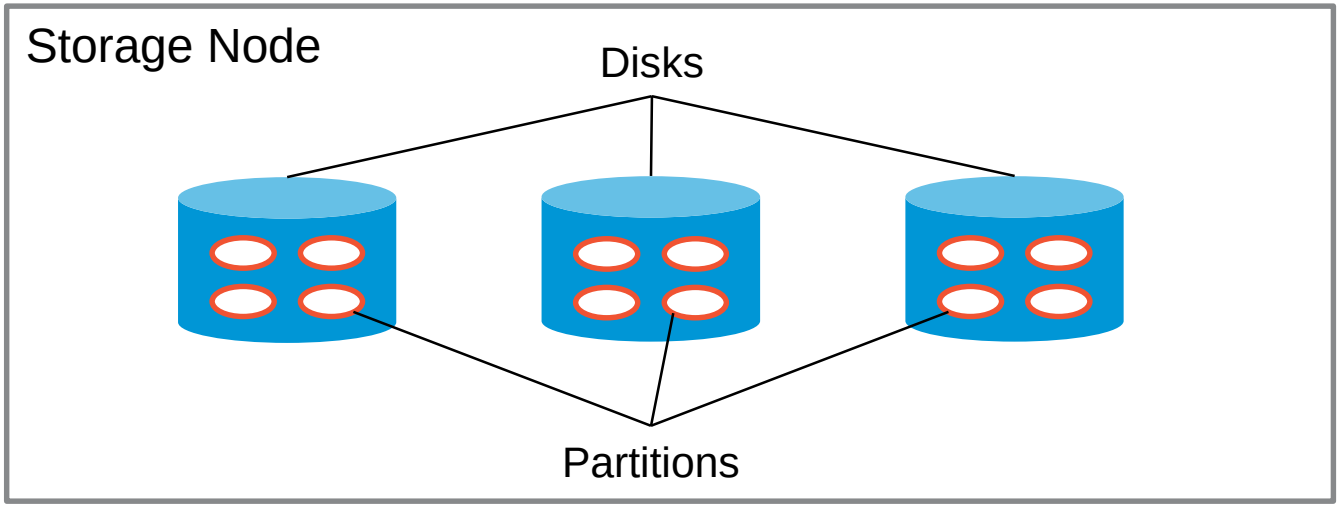

Zones

- Zones are user defined single points of failure within the storage cluster
- Partitions can be replicated across zones
- 3 partitions at least for high availability, 5 partitions is recommended
- Lowest latency placement algorithm used

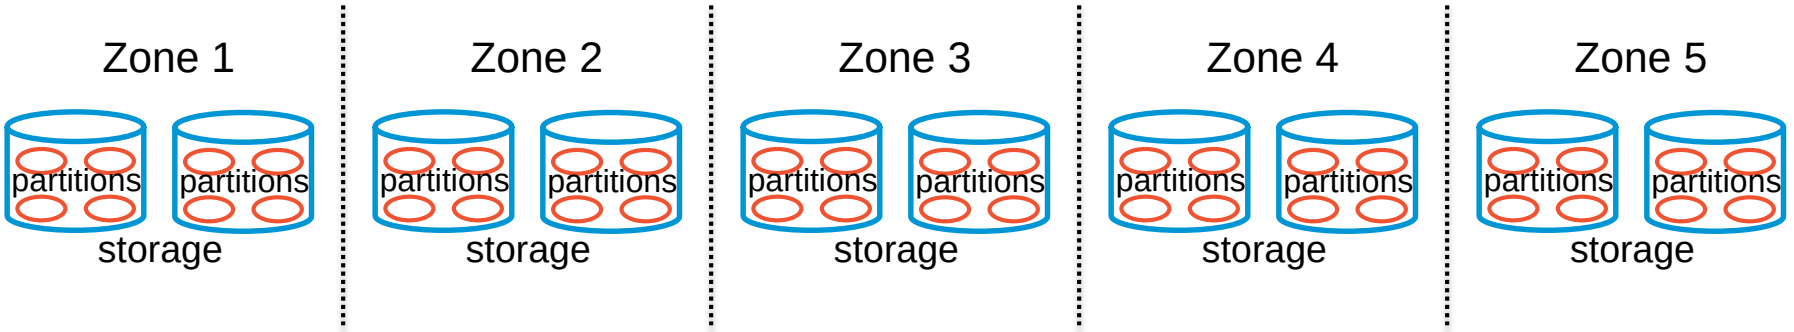

Ceilometer
- Metering Service

Ceilometer in OpenStack Conceptual Architecture

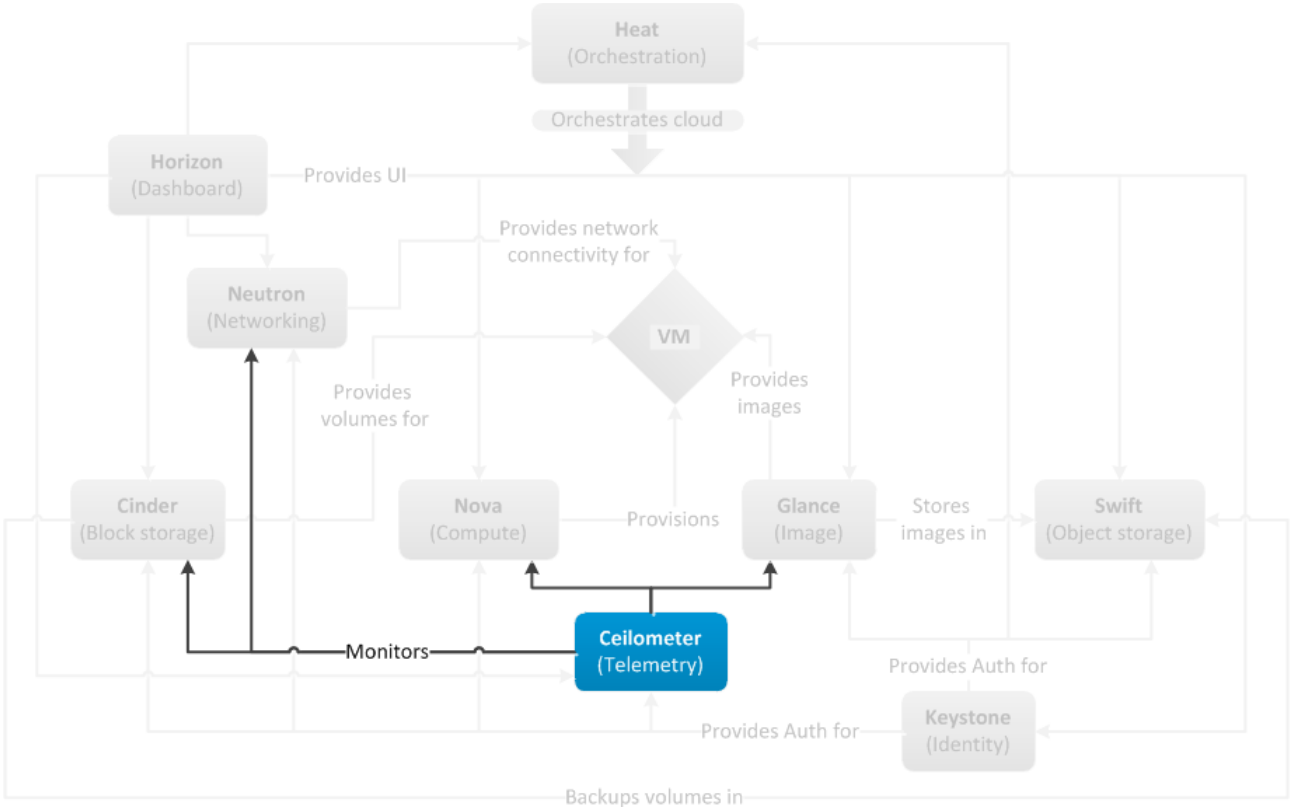

Ceilometer Overview

- Provide efficient collection of metering data, in terms of CPU and network costs.
- Allow deployers to integrate with the metering system directly or by replacing components.
- Data may be collected by monitoring notifications sent from existing services or by polling the infrastructure.
- Allow deployers to configure the type of data collected to meet their operating requirements.
- The data collected by the metering system is made visible to some users through a REST API.

Heat
- Orchestration Service

九州云
Cloud

Heat in OpenStack Conceptual Architecture

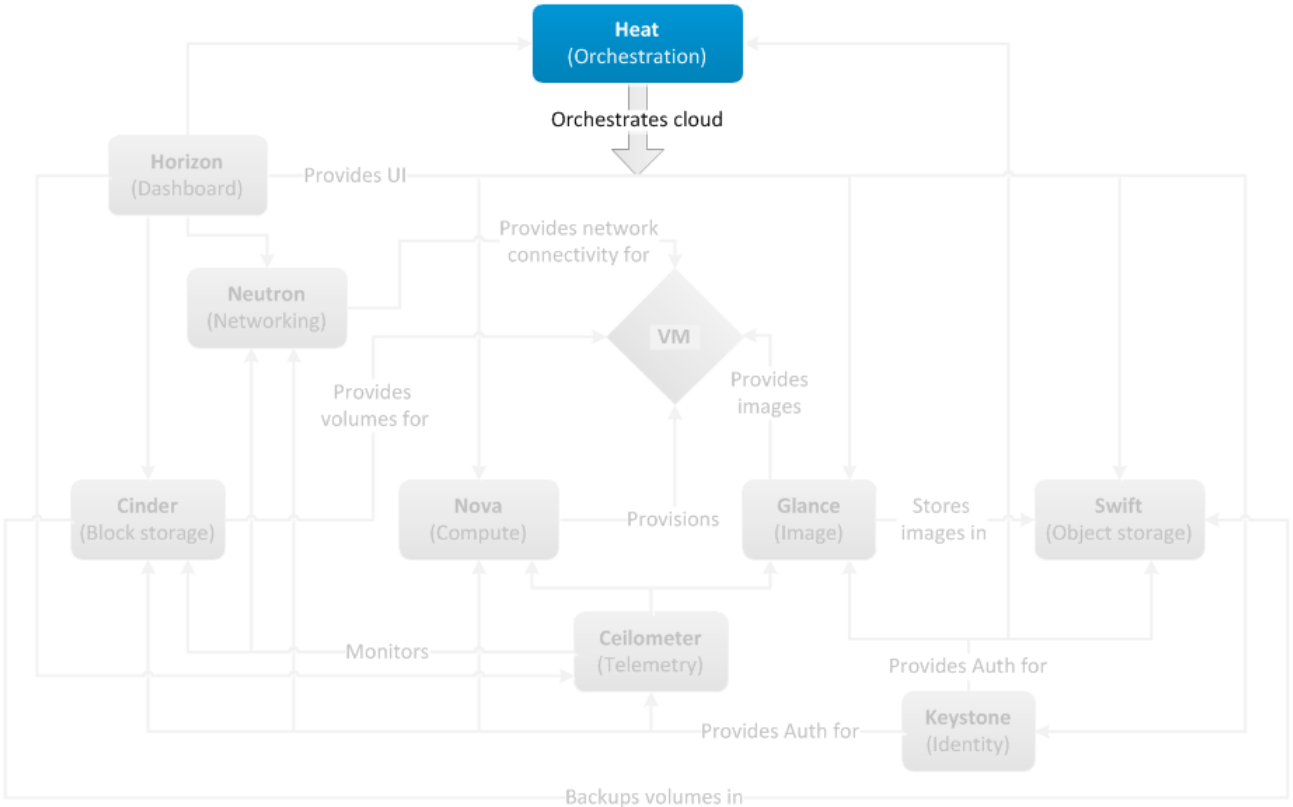

Heat Overview

- Orchestration service that executes OpenStack API calls to generate running cloud applications
- AWS CloudFormation implementation for OpenStack API
- Uses template mechanism
  - HOT
  - TOSCA
- Allow creation of resource types
  - Instances, networks, floating IPs, volumes, security groups, users, etc.
- Some advanced services
  - HA, auto-scaling
- Integrated well with Puppet and Chef

Part 3
Learn Openstack and get Certificated

九州云
Cloud

COA: Certified OpenStack Administrator

- Openstack Foundation Certified
- Openstack Administrator
- With Hands-on
- Installation, Operation, Troubleshooting
- 3 Days Class
- 150 Min exam

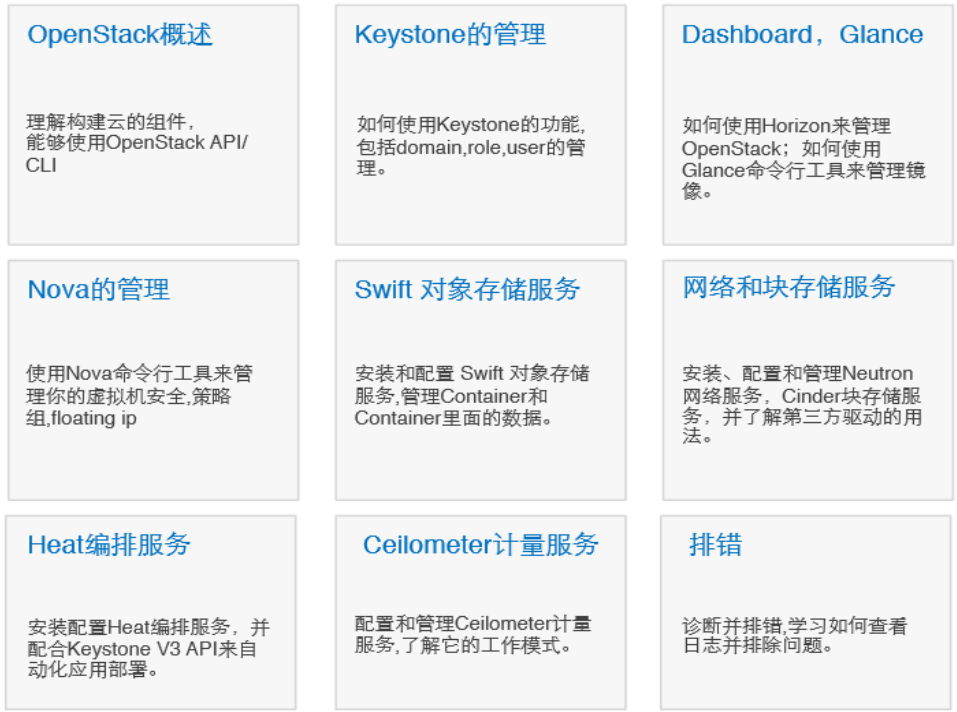

Thank you

九州云
Cloud