Openstack and Ceph

Oyedeji Shola

Seminar on Cluster, Gride and Cloud

OpenStack

- OpenStack is a set of software tools for building and managing cloud computing platforms for public and private clouds.
- OpenStack lets users deploy virtual machines and other instances which handle different tasks for managing a cloud environment on the fly.
- It makes horizontal scaling easy, which means that tasks which benefit from running concurrently can easily serve more or less users on the fly by just spinning up more instances.

What are the components of OpenStack?

- OpenStack is made up of many different moving parts because of its open nature, anyone can add additional components to OpenStack to help it to meet their needs.
- **Nova** is the primary computing engine behind OpenStack.
- **Swift** is a storage system for objects and files.
- **Cinder** is a block storage component, which is more analogous to the traditional notion of a computer being able to access specific locations on a disk drive.
- **Neutron** provides the networking capability for OpenStack.

How is OpenStack used in a cloud environment?

- **Horizon** is the dashboard behind OpenStack. It is the only graphical interface to OpenStack.
- **Keystone** provides identity services for OpenStack.
- **Glance** provides image services to OpenStack.
- **Ceilometer** provides telemetry services, which allow the cloud to provide billing services to individual users of the cloud.
- **Heat** is the orchestration component of OpenStack, which allows developers to store the requirements of a cloud application in a file that defines what resources are necessary for that application.

Ceph

- Distributed object store and file storage designed to provide excellent performance, reliability and scalability.
- Open Source Software, commercial support

Why care about new storage system?

Requirements

- Diverse storage needs
- Object storage
- Block devices (for VMs) with snapshots, cloning
- Shared file system with POSIX, coherent caches
- Structured data… files, block devices, or objects?
- **Scale**
- Terabytes, petabytes, exabytes
- Heterogeneous hardware
- Reliability and fault tolerance

Why care about new storage system?

Time

- Ease of administration
- No manual data migration, load balancing
- Painless scaling
- Expansion and Contraction
- Seamless migration

Why care about new storage system?

Cost

- Linear function of size or performance
- Incremental expansion
- No vendor lock-in
- Choice of hardware
- Choice of software
- Open

Ceph Architecture

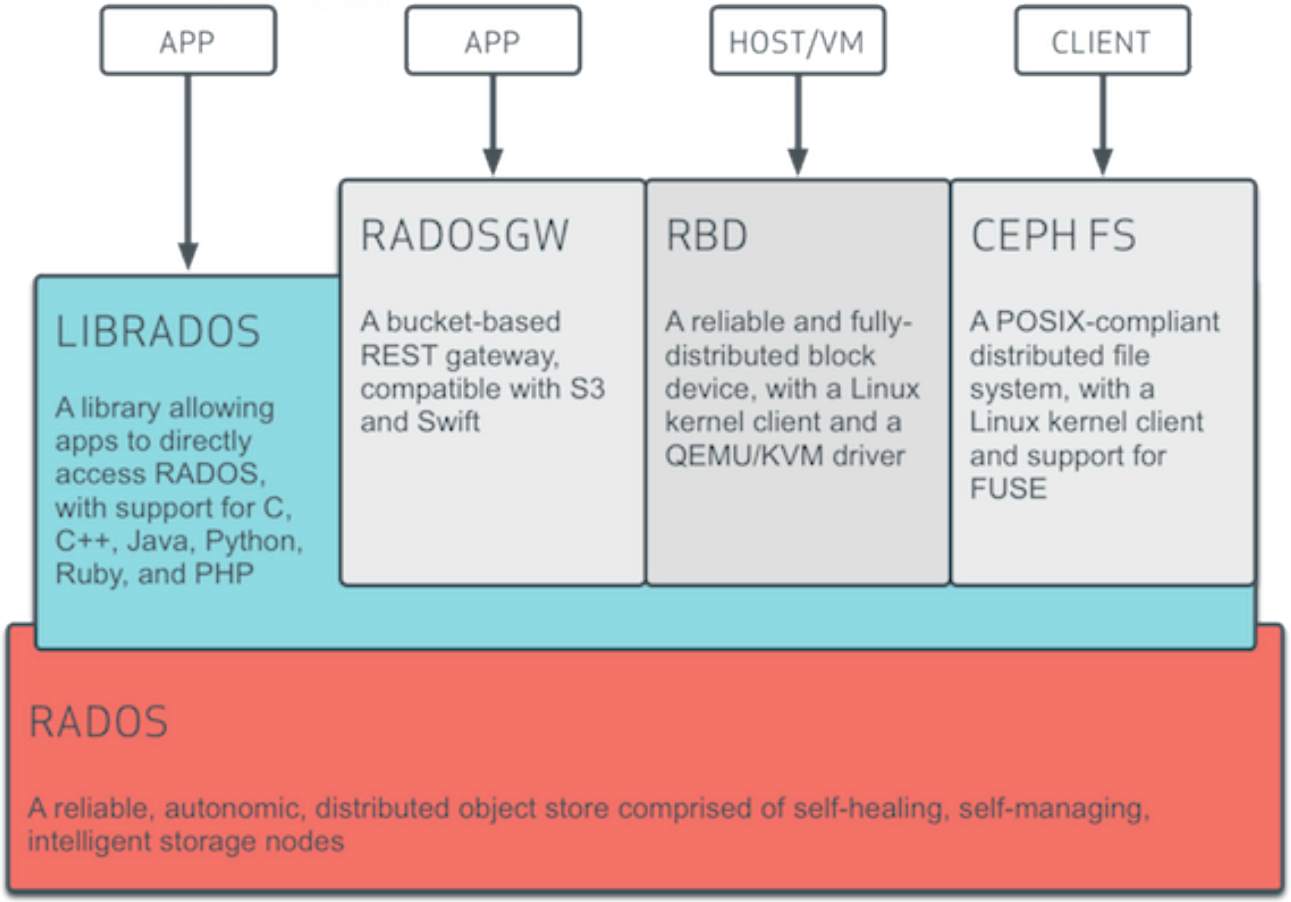

Components of Ceph

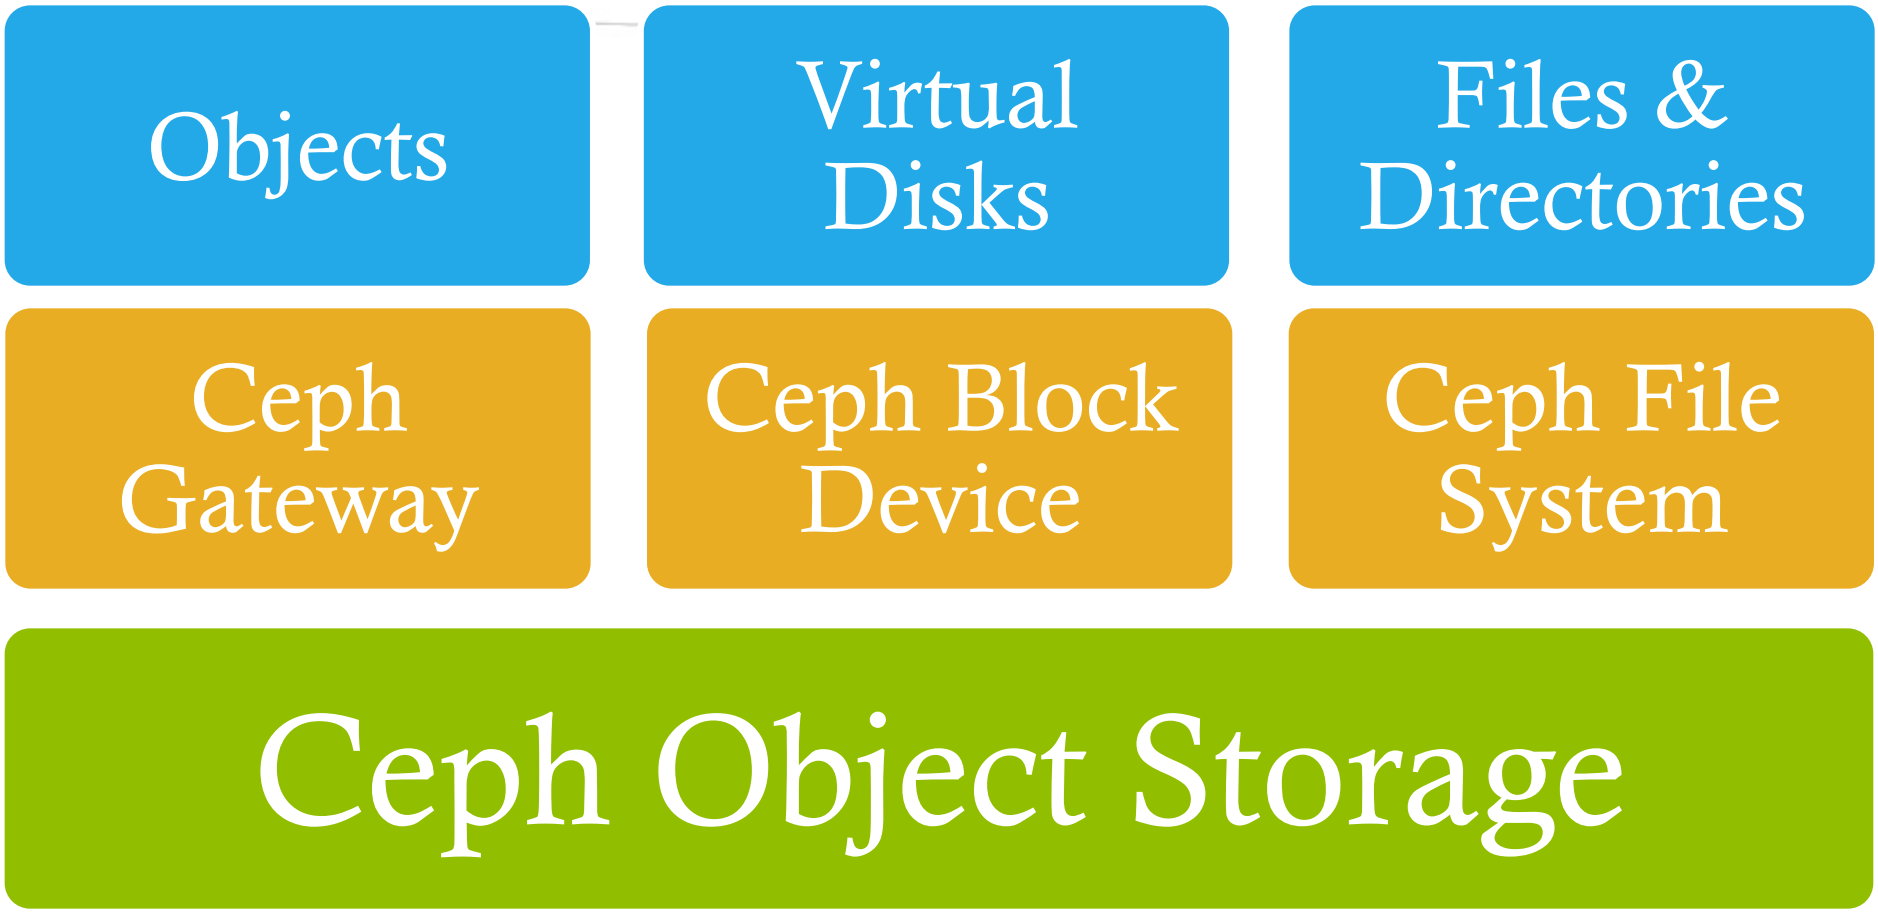

Ceph Design Goals

- Every component must scale
- There can be no single point of failure
- Software based, not an appliance
- Open Source
- Run on commodity hardware
- Everything must self-manage wherever possible

Key Features

- Decoupled data and metadata
  - CRUSH
    - Files striped onto predictably named objects
    - CRUSH maps objects to storage devices
- Dynamic Distributed Metadata Management
  - Dynamic subtree partitioning
    - Distributes metadata amongst MDSs
- Object-based storage
  - OSDs handle migration, replication, failure detection and recovery

Client Access Example

1. Client sends *open* request to MDS
2. MDS returns capability, file node, file size and stripe information
3. Client read/write directly from/to OSDs
4. MDS manages the capability
5. Client sends *close* request, relinquishes capability, provides details to MDS

Openstack and Ceph

- Used for *Glance* Images, *Cinder* Volumes and *Nova* ephemeral disk.
- Ceph is the most popular network block storage backend for OpenStack
- Ceph + OpenStack offers compelling features:
  - clones, layered volumes, snapshots, boot from volume, live migration
  - Cost effective with Thin Provisioning

Ceph Storage in OpenStack

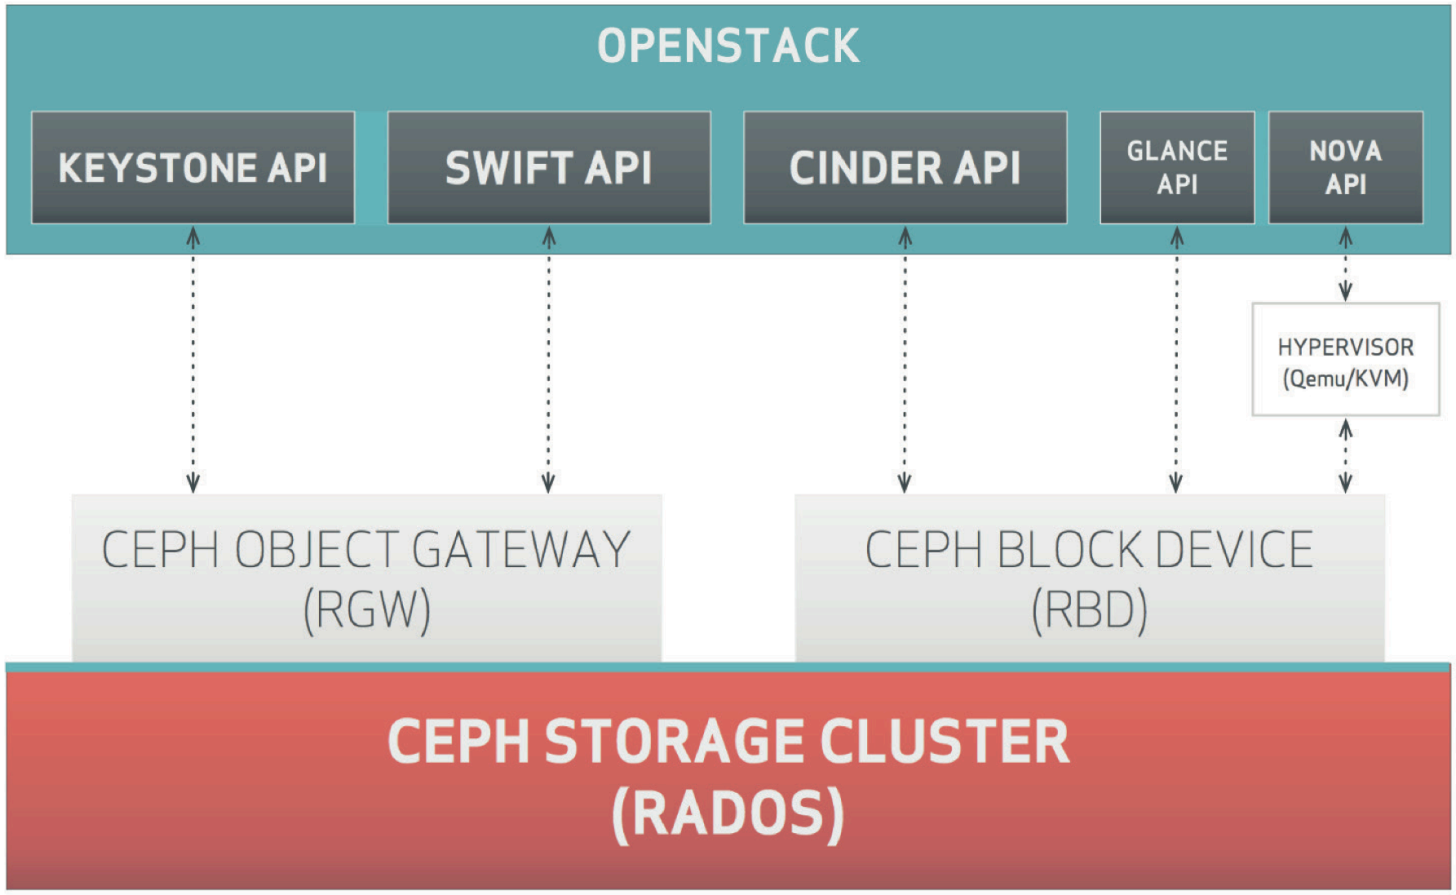

Questions

- Thank You