OPENSTACK & DISK STORAGE

Name – Sumeet Thombre

Email – Sumeet.Thombre@student.lut.fi

Professor – Andrey Shevel

Date- 02/06/2015

Place- ITMO university, St. Petersburg, Russia

Course- Computing clusters, clouds, grids

CONTENTS

- Cloud storage
- Cloud storage architecture
- Ephemeral and Persistent storage
- Object storage
- Swift
- Block storage
- Cinder
- File-level storage
- Overview & OpenStack Flavours
- Choosing storage back-ends

CLOUD STORAGE

- Digital data in logical pools

- Physical storage spans multiple servers

- OpenStack – lets you build
IaaS cloud that runs on
Commodity hardware

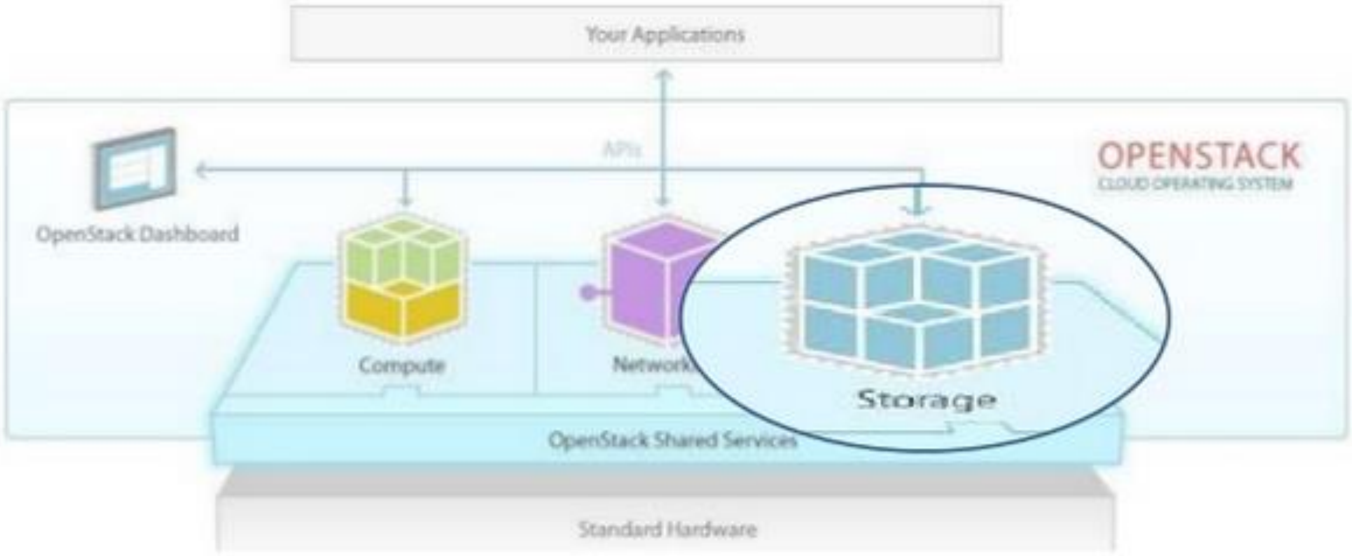

- OpenStack – “The Linux of the cloud”

CLOUD STORAGE - ARCHITECTURE

- Based on highly virtualised infrastructure
- Many distributed resources, still acts as one
- Highly fault tolerant through redundancy & distribution of data
- Highly durable through the creation of versioned copies
- Eventually consistent with regard to data replicas

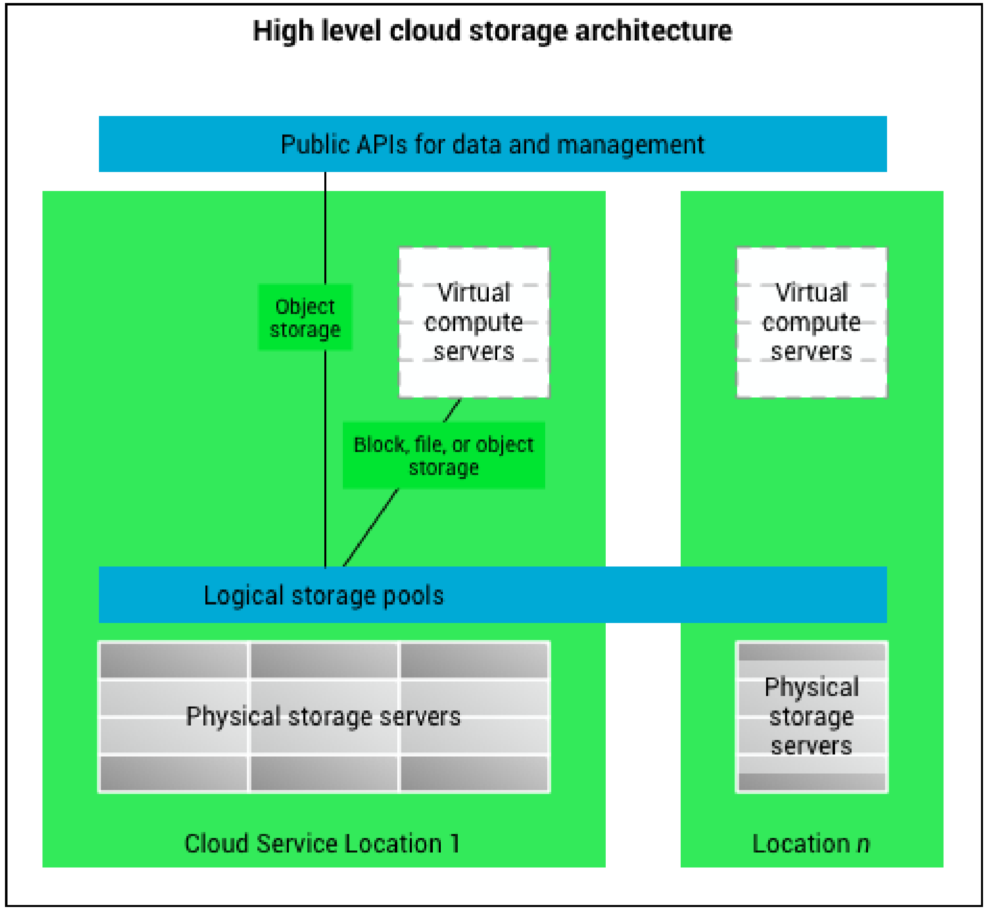

EPHEMERAL & PERSISTENT STORAGE

- Ephemeral storage
  - Disks effectively disappear when the VM is terminated
  - Comes with OpenStack Compute Service (Nova)
- Persistent storage
  - Storage resource outlives any other resource and is always available
  - OpenStack supports two forms of persistent storage,
    - Object storage (Swift)
    - Block storage (Cinder)

OBJECT STORAGE IN OPENSTACK

- Binary objects through a REST API
- Data, metadata & identifier
- Three main traffic flows in object network patterns,

  - among object, container & account servers

  - between servers & proxies

  - between proxies & users
- Object access path –

  Account > containers > objects

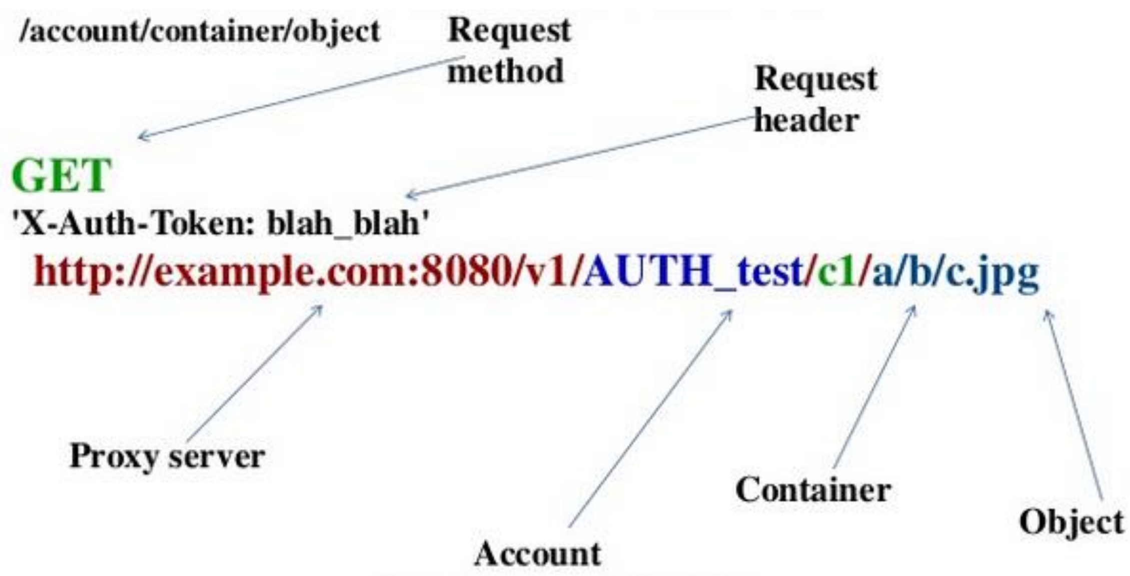

OPENSTACK SWIFT

- Highly available, distributed, eventually consistent object store
- Scales horizontally
- Involves the Ring - represents a mapping between the names of entities stored on disk and their physical location
- Decoupled/independent from OpenStack
- Unstructured data, alternative to Amazon S3

- But you can't
  - Mount it
  - Have file hierarchies

BLOCK STORAGE IN OPENSTACK

- Physical record, sequence of bits/bytes, max block size
- Low level I/O
- Tracks & sectors
- No metadata
- User interaction by attaching volumes to their running VM instances

OPENSTACK CINDER

- Cinder volumes, attached to VMs
- Backend devices opaque to users
- Components: API, scheduler, volume
- Provides APIs such as create/delete, back-up/restore snapshot/ clone volumes

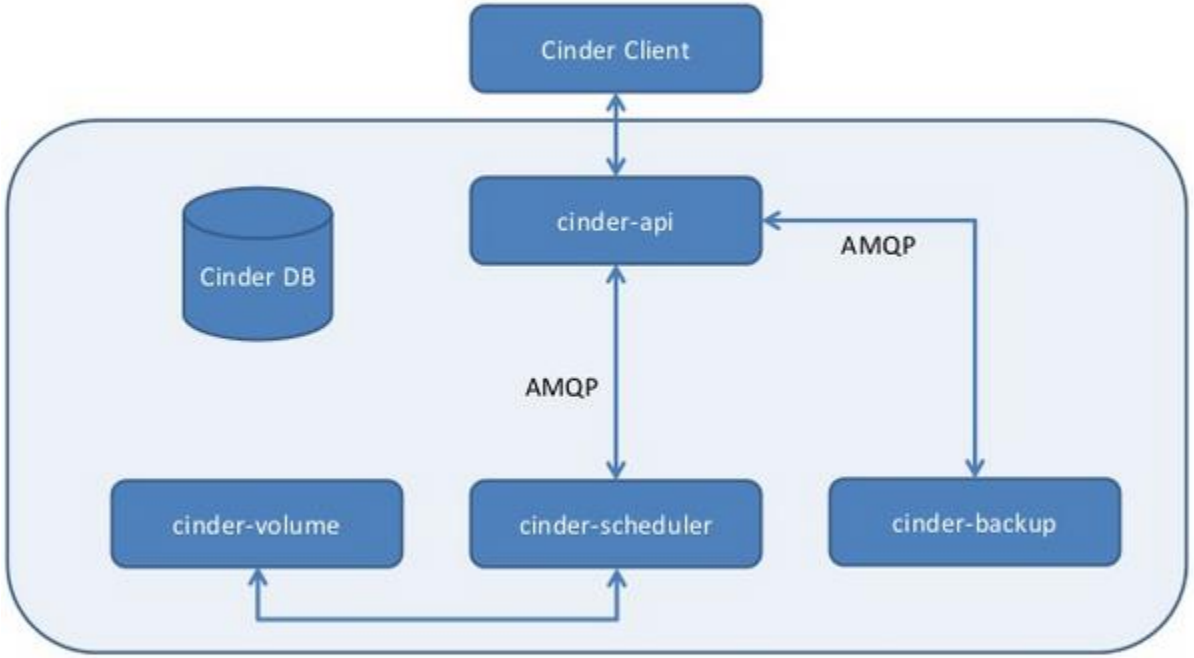

FILE LEVEL STORAGE

- Users access stored data using the operating system's file system interface, files & directories
- Simple access
- NFS in UNIX, CIFS in Windows (most common form)
- Not supported in OpenStack!
- Live Migration

OVERVIEW OF STORAGE TECHNIQUES

	Ephemeral storage	Block storage	Object storage
Used to...	Run operating system and scratch space	Add additional persistent storage to a virtual machine (VM)	Store data, including VM images
Accessed through...	A file system	A block device that can be partitioned, formatted, and mounted (such as, /dev/vdc)	The REST API
Accessible from...	Within a VM	Within a VM	Anywhere
Managed by...	OpenStack Compute (nova)	OpenStack Block Storage (cinder)	OpenStack Object Storage (swift)
Persists until...	VM is terminated	Deleted by user	Deleted by user
Sizing determined by...	Administrator configuration of size settings, known as *flavors*	User specification in initial request	Amount of available physical storage
Example of typical usage...	10 GB first disk, 30 GB second disk	1 TB disk	10s of TBs of dataset storage

OPENSTACK FLAVORS

- Virtual hardware templates
- Define sizes for memory in RAM, root disk size, amount of ephemeral data disk space available, and number of cores
- Storage performance (spindles/core), memory availability (RAM/core), network bandwidth (Gbps/core) and overall CPU performance (CPU/core)

Name	Virtual cores	Memory	Disk	Ephemeral
m1.tiny	1	512 MB	1 GB	0 GB
m1.small	1	2 GB	10 GB	20 GB
m1.medium	2	4 GB	10 GB	40 GB
m1.large	4	8 GB	10 GB	80 GB
m1.xlarge	8	16 GB	10 GB	160 GB

CHOOSING STORAGE BACKENDS

- Best fit!
- Object vs Block vs File-level storage
- Cost vs performance
- Redundancy & distribution of storage
- Backend drivers – LVM, Ceph, GlusterFS

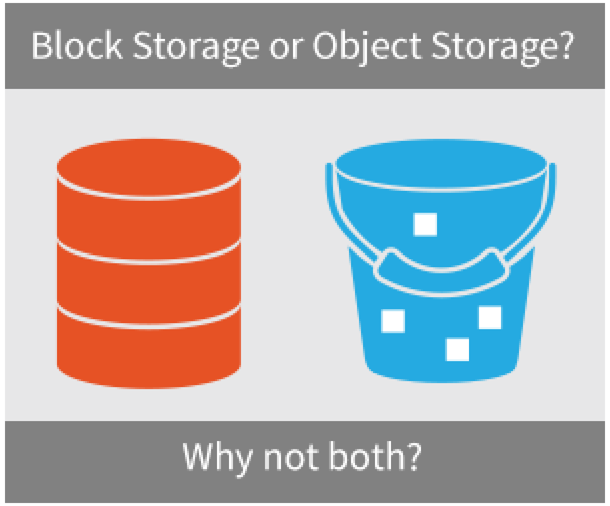

SNAPSHOT FROM THE TOOL

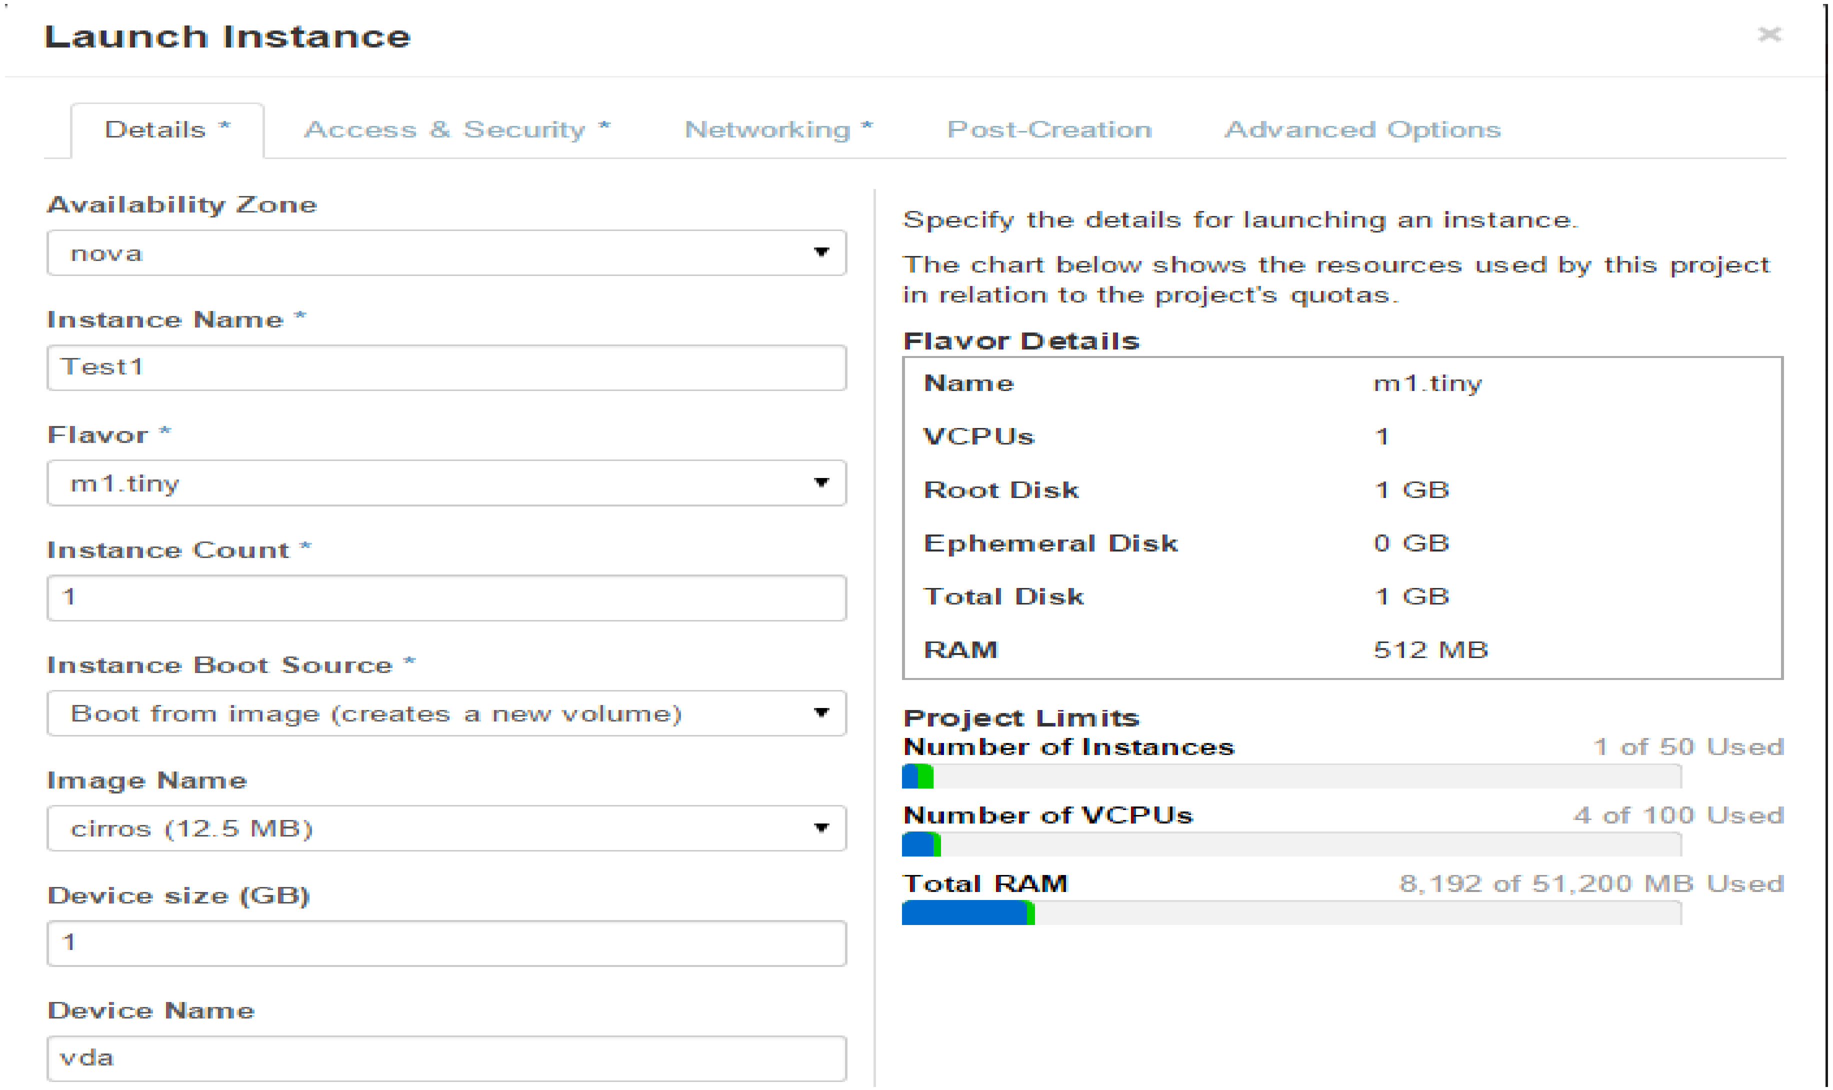

THANKS!

ANY QUERIES?