ceph

CEPH

CEPHALOCON APAC 2018

THE FUTURE OF STORAGE

22-23 March 2018 | BEIJING

Basic and Advanced Analysis of Ceph Volume Backend Driver in Cinder

Netmarble is a Great Game Culture Company

Netmarble Games, the #1 Online Game publisher in Korea, is teaming up with various top development studios to bring quality mobile entertainment to the world.

2017 GLOBAL PUBLISHERS TOP 3

Rank	Publisher	HQ Country
1	Tencent	China
2	NetEase	China
3	Netmarble	South Korea
4	Activision Blizzard	United States
5	Supercell	Finland
6	mixi	Japan
7	BANDAI NAMCO	Japan
8	Sony	Japan
9	MZ	United States
10	LINE	Japan

CEPH in Netmarble

How Many? 8 Clusters

How Much? around 4 PB

Details?

KRBD Block Storage	3 Clusters
Block Storage as a Cinder Backend	4 Clusters
Archiving Storage via S3 interface	1 Cluster

Plan? Origin CDN Storage
Data Processing Storage

Today's Journey

Default Cinder Features

Complex Cinder Feature

Who's Speker

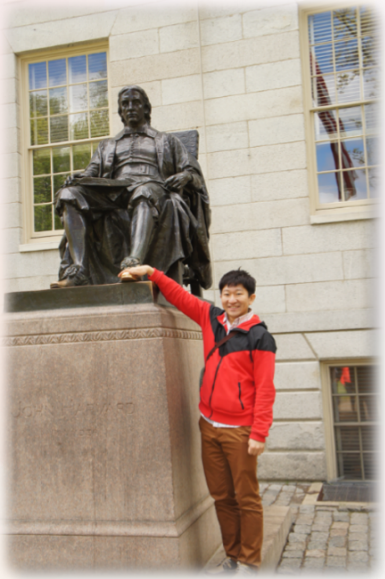

John Haan (Seungjin Han)

2011 – 2016 (Samsung SDS)
Cloud Solution Engineer
OpenStack Development

2016 – present (Netmarble)
Cloud Platform Engineer
Responsible for Ceph Storage

Basic

Default Cinder Features

- Concept of Cinder Volume
  - RBD Copy on Write
- Cinder Snapshot
  - RBD Snapshot / Flatten
- Cinder Backup
  - RBD export-diff / import-diff

OpenStack Cinder & CEPH

OpenStack Cinder

- Block storage service for OpenStack
- Provides storage resources to end users
- Dominent users use Ceph as a cinder backend

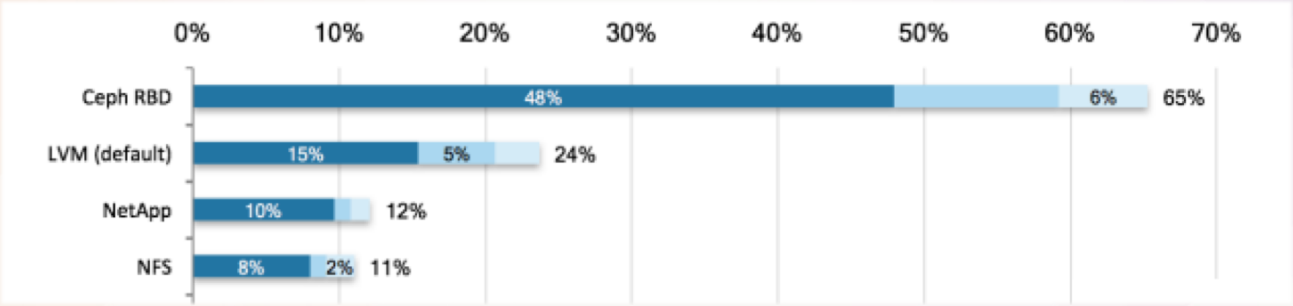

OpenStack Cinder & CEPH

How to Integrate Ceph as a Cinder Backend

- http://docs.ceph.com/docs/master/rbd/rbd-openstack/

OpenStack Cinder & CEPH

OpenStack Configurations

```
### glance
[glance_store]
rbd_store_pool = images

### cinder
[rbd]
rbd_pool = volumes

### nova
[libvirt]
images_rbd_pool = vms
```

Create Ceph Pools

```
ceph osd pool create volumes …
ceph osd pool create images …
ceph osd pool create backups …
ceph osd pool create vms …
```

OpenStack Cinder & CEPH

```
POOLS:
    NAME        ID     USED       %USED     MAX AVAIL     OBJECTS
    volumes     58      5454G     40.71         6814G      702361
    images      59     30924M      0.23         6814G        3879
    vms         60       152G      1.14         6814G       39209
    backups     61          0         0         6814G           0
```

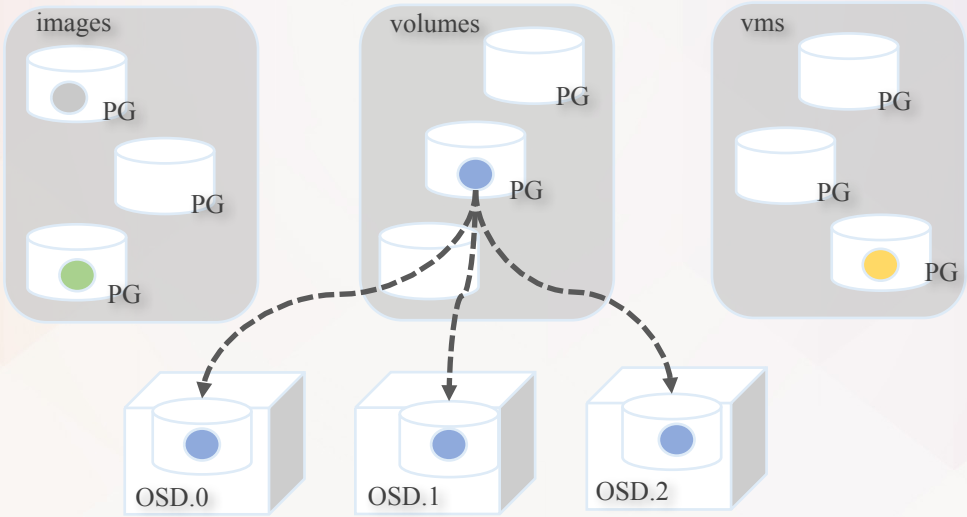

Anatomy of Cinder Volume

- 20GB of cinder volume

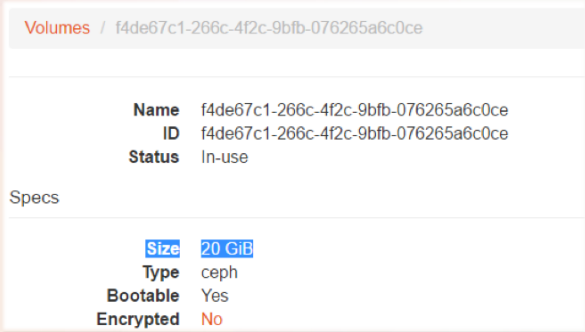

- could change rados default size

```
### cinder
[rbd]
rbd_store_chunk_size = {chunk size in MB}
# default value is 4MB
```

- 4MB chuk of rbd volume

```
rbd image '087aa587-4abc-48e0-be52-3397c705333c':
        size 44032 MB in 11008 objects
        order 22 (4096 kB objects)
        block_name_prefix: rbd_data.5c47933fc8e78f
        format: 2
```

- 8MB chunk of rbd volumes

```
rbd image 'volume-f4de67c1-266c-4f2c-9bfb-076265a6c0ce':
        size 20480 MB in 2560 objects
        order 23 (8192 kB objects)
        block_name_prefix: rbd_data.11ab362117b302
        format: 2
```

Where is my data stored in Filestore

cinder volume

Where is my data stored in Filestore


```
rbd image 'volume-da19afe9-3f6f-4fcc-9093-bc69fb913e86':
        size 44032 MB in 11008 objects
        order 22 (4096 kB objects)
        block_name_prefix: rbd_data.03ba561c5863d0
        format: 2
```

Where is my data stored in Filestore


```
[root@Pff1-1ctrl01 ~]# rados -p volumes-4 ls | grep rbd_data.03ba561c5863d0
rbd_data.03ba561c5863d0.0000000000000d64
rbd_data.03ba561c5863d0.0000000000001649
rbd_data.03ba561c5863d0.0000000000000416
rbd_data.03ba561c5863d0.0000000000000406
rbd_data.03ba561c5863d0.00000000000003d6
rbd_data.03ba561c5863d0.0000000000001f66
rbd_data.03ba561c5863d0.0000000000000d07
```

Where is my data stored in Filestore

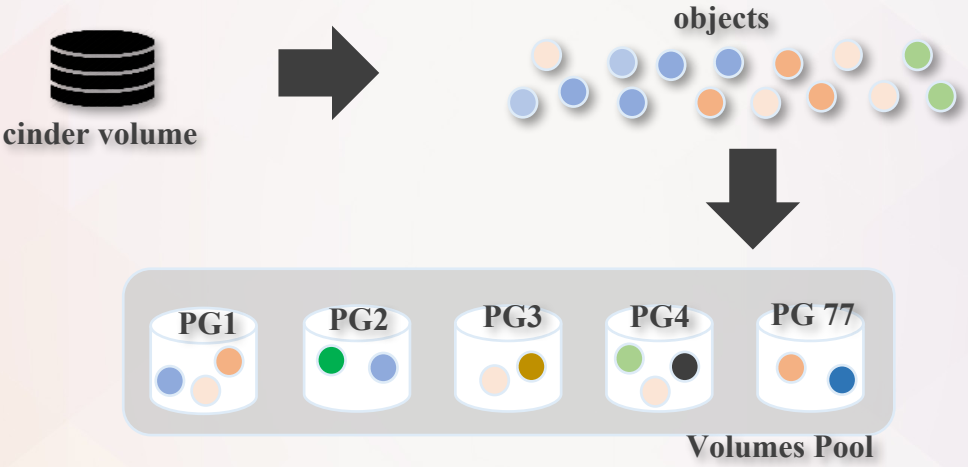

Where is my data stored in Filestore

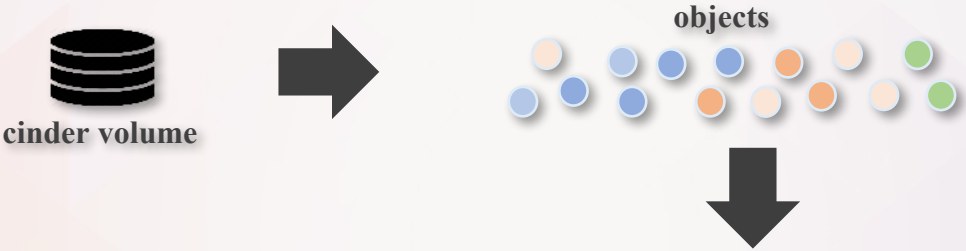


```
[root@Pff1-1ctrl01 ~]# ceph osd map volumes-4 rbd_data.03ba561c5
63d0.0000000000000d64
osdmap e19235 pool 'volumes-4' (77) object 'rbd_data.03ba561c586
3d0.0000000000000d64' -> pg 77.e75d9500 (77.500) -> up ([345,323
,264], p345) acting ([345,323,264], p345)
```

Where is my data stored in Filestore

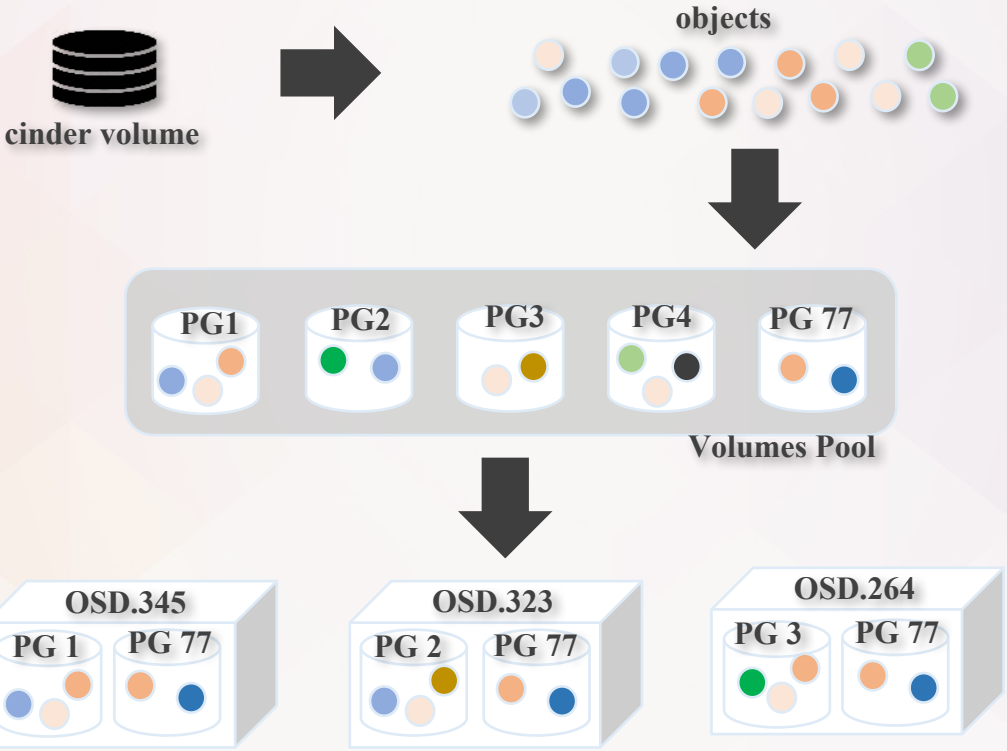

Where is my data stored in Filestore

```
[root@Pff1-4cephosd04 current]# lls 77.500_
77.500_head/ 77.500_TEMP/
[root@Pff1-4cephosd04 current]# ls 77.500_head/
__head_00000500__4d                                                rbd\udata.
462313693a40b.0000000000001e74__head_3852C500__4d
rbd\udata.03ba561c5863d0.000000000000d64__head_E75D9500__4d  rbd\udata.
462313693a40b.00000000000022a1__head_CEB0A500__4d
rbd\udata.03df911573c784.000000000000d6d__head_ED2E1500__4d  rbd\udata.
462313693a40b.0000000000002671__head_AEE97500__4d
```

Volumes Pool

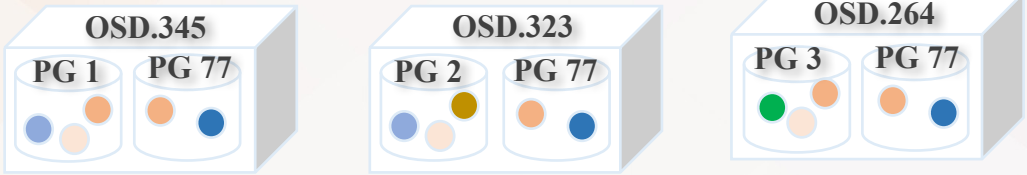

Where is my data stored in Bluestore

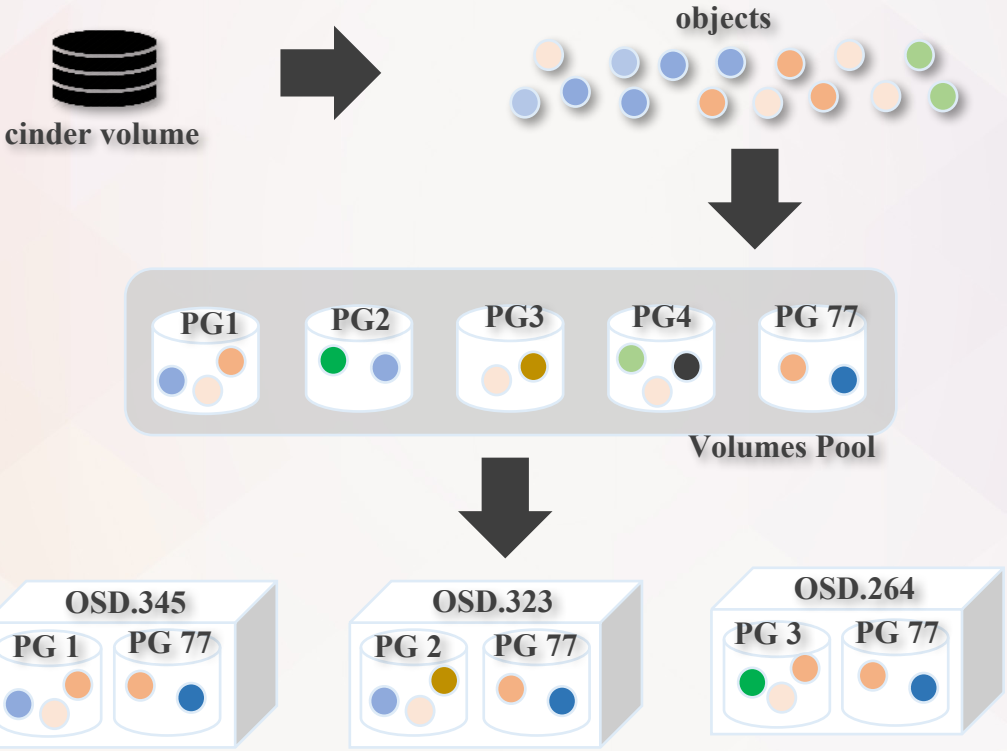

Where is my data stored in Bluestore

via ceph-objectstore-tool

```
root@ceph01:/mnt/tmp/12.0_head/all/#12:05e22483:::rbd_data.4d574c73a9c6.0000000000000000:head## ll
total 512
drwx------ 0 root root       0 Jan  1  1970 attr/
-rwx------ 0 root root       9 Jan  1  1970 bitwise_hash*
-rwx------ 1 root root 4194304 Jan  1  1970 data*
drwx------ 0 root root       0 Jan  1  1970 omap/
-rwx------ 0 root root       0 Jan  1  1970 omap_header*
```

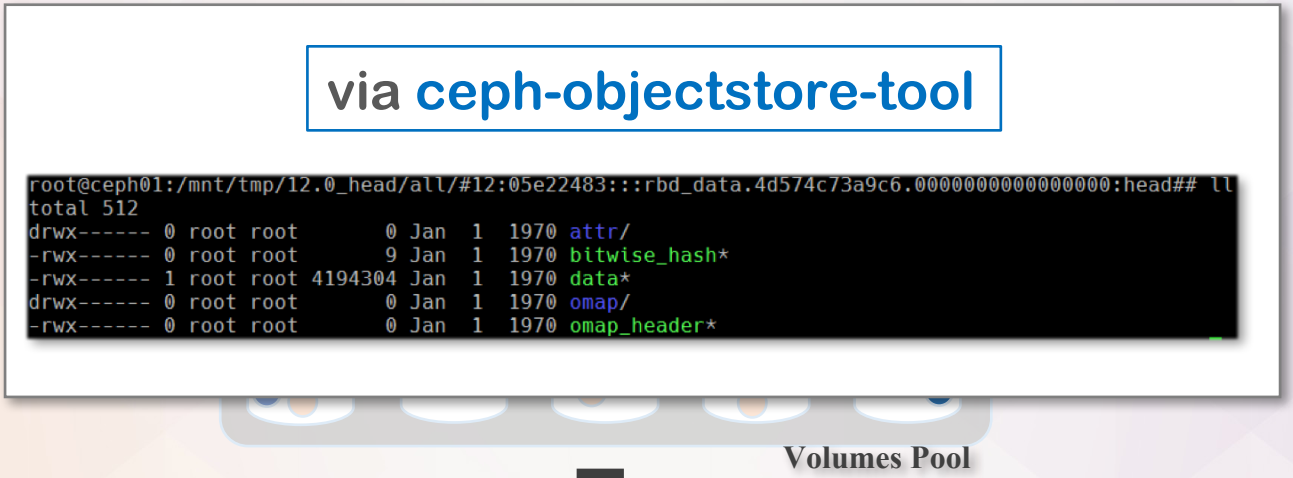

Copy on Write

- original Image(parent) → protected snapshot → cloned(child) image
- fast volume provisioning than legacy storage

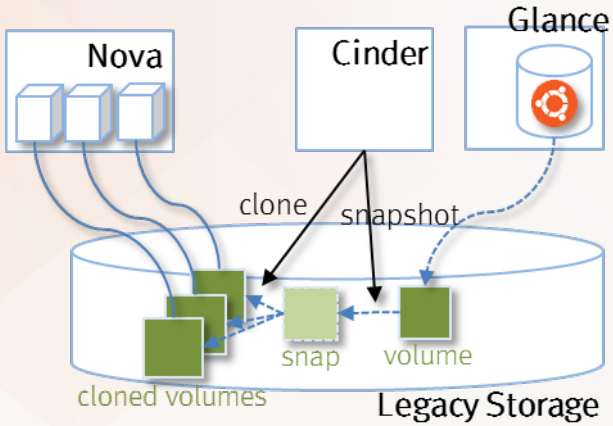

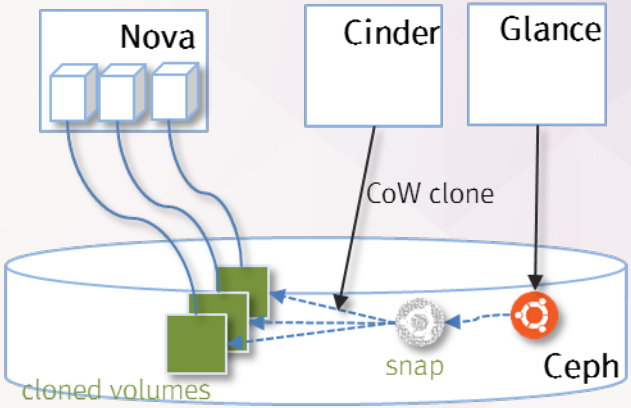

Copy on Write

- must set show_image_direct_url to True in glance configuration
- children volumes related to snapshot

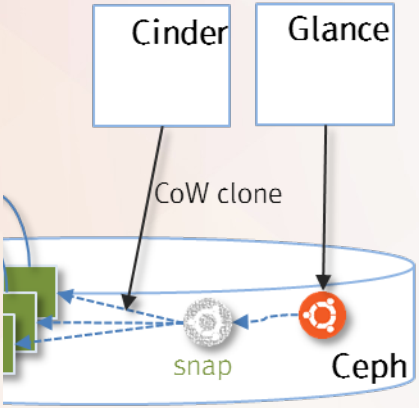

<protected snapshot from the image>

```
[root@Pff1-1ctrl01 ~]# rbd -p images info d962d3bd-0463-4f78-a02c-dab4af38886e@snap
rbd image 'd962d3bd-0463-4f78-a02c-dab4af38886e':
        size 40960 MB in 10240 objects
        order 22 (4096 kB objects)
        block_name_prefix: rbd_data.57c95e5fe68266
        format: 2
        features: layering, exclusive-lock, object-map, fast-diff, deep-flatten
        flags:
        protected: True
```

<children volume from the snapshot>

```
[root@Pff1-1ctrl01 ~]# rbd -p volumes-1 info volume-013d1af4-38af-4316-9736-32913dcc125b
rbd image 'volume-013d1af4-38af-4316-9736-32913dcc125b':
        size 44032 MB in 11008 objects
        order 22 (4096 kB objects)
        block_name_prefix: rbd_data.8eb24954da257d
        format: 2
        features: layering, exclusive-lock, object-map, fast-diff, deep-flatten
        flags:
        parent: images/d962d3bd-0463-4f78-a02c-dab4af38886e@snap
        overlap: 40960 MB
```

Copy on Write

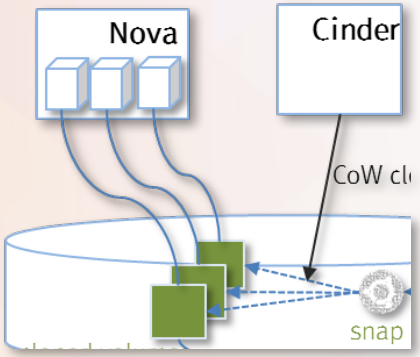

- The volume connected with compute node

```
[root@Pff1-1ctrl01 ~]# rbd -p volumes-1 status volume-013c
913dcc125b
Watchers:
        watcher=           :0/2389303156 client.21944604
```

- libvirt xml of instance

```
    <disk type='network' device='disk'>
      <driver name='qemu' type='raw' cache='writeback'/>
      <auth username='cinder'>
        <secret type='ceph' uuid='457eb676-33da-42ec-9a8c-92
      </auth>
      <source protocol='rbd' name='volumes-1/volume-9d0acbc3
6cdb33aa'>
```

Snapshot of Cinder

Cinder Snapshot Analysis

- capture of what a volume like at a particular moment
- cinder snapshot can be restored into a cinder volume
- data of cinder snapshot can be transferred into a cinder backup volume

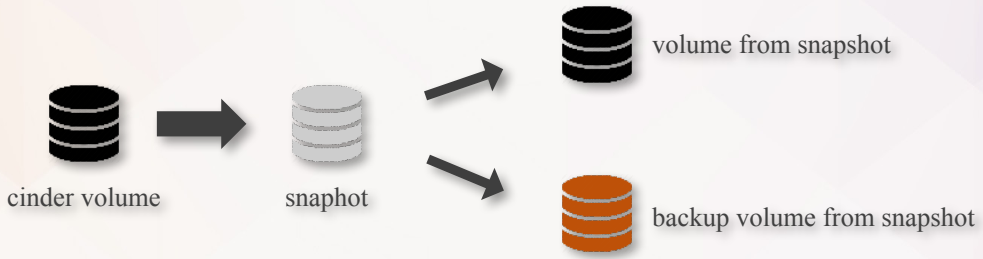

Snapshot of Ceph

Ceph Snapshot Analysis

- snapshot concept is the same as cinder
- volume can be flatten or cloned when creating from snapshot

```
### cinder
[rbd]
rbd_flatten_volume_from_snapshot = {false or true}
```

```
### CEPH
rbd –p volumes flatten volume-xxxxxx
Image flatten: 100% complete...done.
```

Snapshot of Cinder & Ceph

cinder snapshot from a cinder volume

```
[root@Pff1-1ctrl01 ~]# cinder snapshot-list
+--------------------------------------+--------------------------------------
----------+-----------+---------------+------+
| ID                                   | Volume ID
          | Status    | Name          | Size |
+--------------------------------------+--------------------------------------
----------+-----------+---------------+------+
| f5ee3a0f-f874-4b87-bd26-0dfe08a2bcbc | 7a84bb5c-ae8a-4ada-88d2-90
7483fcbd3f | available | sjvol-snap-01 | 1    |
+--------------------------------------+--------------------------------------
----------+-----------+---------------+------+
```

snapshot image from the rbd volume

```
[root@Pff1-1ctrl01 ~]# rbd -p volumes-1 snap ls volume-7a84bb5c-ae8
a-4ada-88d2-907483fcbd3f
SNAPID NAME                                                 SIZE TIMEST
AMP
    26 snapshot-f5ee3a0f-f874-4b87-bd26-0dfe08a2bcbc 1024 MB Fri Ma
r  2 15:36:41 2018
```

Volume from a Ceph Snapshot

Non flatten volume

```
        format: 2
        features: layering, exclusive-lock, object-map, fast-diff,
deep-flatten
        flags:
        create_timestamp: Fri Mar  2 15:51:23 2018
        parent: volumes-1/volume-7a84bb5c-ae8a-4ada-88d2-907483fcbd
3f@snapshot-f5ee3a0f-f874-4b87-bd26-0dfe08a2bcbc
        overlap: 1024 MB
```

Flatten volume

```
[root@Pff1-1ctrl01 ~]# rbd -p volumes-1 info volume-82269f01-063c-4
48-bac1-e1dfc1ff1c30
rbd image 'volume-82269f01-063c-4548-bac1-e1dfc1ff1c30':
        size 1024 MB in 256 objects
        order 22 (4096 kB objects)
        block_name_prefix: rbd_data.b29ac9720a33c0
        format: 2
        features: layering, exclusive-lock, object-map, fast-diff,
deep-flatten
        flags:
        create_timestamp: Fri Mar  2 15:51:23 2018
```

Cinder Backup with CEPH Backend

- Utilize rbd export-diff & import-diff
- Support full backup & incremental backup

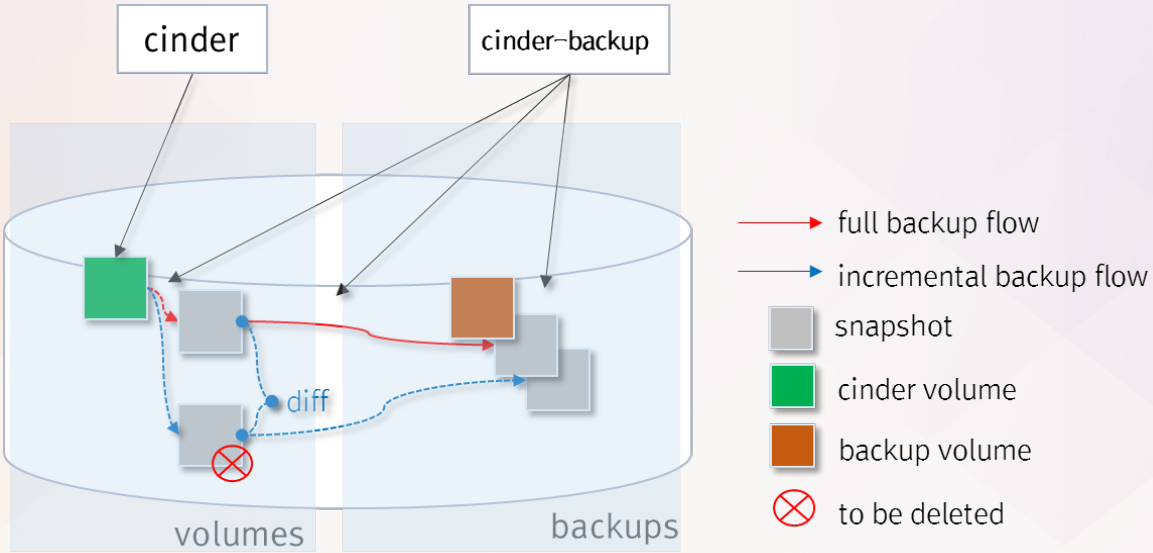

Full Backup from a Cinder Volume

- full backup volumes

```
stack@devstack02:~$ cinder backup-create backup-test-1 --name fullbackup
+-----------+--------------------------------------+
| Property  | Value                                |
+-----------+--------------------------------------+
| id        | c445d420-9408-4aef-b733-ee39dca67ba0 |
| name      | fullbackup                           |
| volume_id | b73547b0-74dc-4be5-a958-368f7ea70128 |
+-----------+--------------------------------------+
```

```
stack@devstack02:~$ cinder backup-list
+--------------------------------------+--------------------------
-----------+-----------+------------+------+--------------+--------
---+
| ID                                   | Volume ID
          | Status    | Name       | Size | Object Count | Contain
er |
+--------------------------------------+--------------------------
-----------+-----------+------------+------+--------------+--------
---+
| c445d420-9408-4aef-b733-ee39dca67ba0 | b73547b0-74dc-4be5-a958-36
8f7ea70128 | available | fullbackup | 1    | 0            | backups
   |
+--------------------------------------+--------------------------
-----------+-----------+------------+------+--------------+--------
---+
```

Full Backup from a Cinder Volume

- ceph volume image from backups pool

```
root@ceph01:~# rbd -p backups info volume-b73547b0-74dc-4be5-a958-3
68f7ea70128.backup.base
rbd image 'volume-b73547b0-74dc-4be5-a958-368f7ea70128.backup.base'
:
        size 1024 MB in 256 objects
        order 22 (4096 kB objects)
        block_name_prefix: rbd_data.170e42963e5a
        format: 2
        features: layering
        flags:
        create_timestamp: Fri Feb  9 04:42:27 2018
```

- snapshot volume image from the backup volume

```
root@ceph01:~# rbd -p backups snap ls volume-b73547b0-74dc-4be5-a95
8-368f7ea70128.backup.base
SNAPID NAME
      SIZE TIMESTAMP
    20 backup.c445d420-9408-4aef-b733-ee39dca67ba0.snap.1520003729.
13 1024 MB Fri Feb  9 04:42:29 2018
    21 backup.4e7f9c0f-50f2-4833-92ea-ef7853ff3912.snap.1520004177.
55 1024 MB Fri Feb  9 04:49:58 2018
    22 backup.83ffce61-7606-49bf-8b44-5a0c35dd01f8.snap.1520004389.
52 1024 MB Fri Feb  9 04:53:30 2018
```

Incremental Backups

- Incremental backup volumes

```
stack@devstack02:~$ cinder backup-create backup-test-1 --name inc-b
ackup-1 --incremental
+-----------+--------------------------------------+
| Property  | Value                                |
+-----------+--------------------------------------+
| id        | 4e7f9c0f-50f2-4833-92ea-ef7853ff3912 |
| name      | inc-backup-1                         |
| volume_id | b73547b0-74dc-4be5-a958-368f7ea70128 |
+-----------+--------------------------------------+
```

```
stack@devstack02:~$ cinder backup-list | grep inc
| 4e7f9c0f-50f2-4833-92ea-ef7853ff3912 | b73547b0-74dc-4be5-a958-36
8f7ea70128 | available | inc-backup-1 | 1    | 0            | backu
ps   |
| 83ffce61-7606-49bf-8b44-5a0c35dd01f8 | b73547b0-74dc-4be5-a958-36
8f7ea70128 | available | inc-backup-2 | 1    | 0            | backu
ps   |
```

Incremental Backups

- Incremental Snapshot with RBD

```
$ rbd export-diff –-pool volumes –-from-snap backup.4e volumes/volume-b7..@backup.4e7.. - | rbd import-diff –backups – backups/volume-4e7..
```

- snapshot volumes from the backup volume

```
root@ceph01:~# rbd -p backups snap ls volume-b73547b0-74dc-4be5-a958-368f7ea70128.backup.base
SNAPID NAME
      SIZE TIMESTAMP
    20 backup.c445d420-9408-4aef-b733-ee39dca67ba0.snap.1520003729.13 1024 MB Fri Feb  9 04:42:29 2018
    21 backup.4e7f9c0f-50f2-4833-92ea-ef7853ff3912.snap.1520004177.55 1024 MB Fri Feb  9 04:49:58 2018
    22 backup.83ffcc61-7606-49bf-8b44-5a0c35dd01f8.snap.1520004389.52 1024 MB Fri Feb  9 04:53:30 2018
```

Advanced

Complex Cinder Features

- Image Cached Volume
- Cinder Replication
  - RBD Mirroring

Concern about CoW

```
root@flosctrl01:~# rbd -p images children
volumes-1/volume-5811fae1-e0f0-4d2f-b3d4-
volumes-1/volume-59882a1e-70c1-4a5c-ad89-
volumes-1/volume-839c8e5a-ce8d-4565-8514-
volumes-1/volume-c64b3707-5e2d-46e5-96e9-
volumes-2/volume-38130b19-8dfe-4641-beed-
volumes-2/volume-54929465-1d3b-49b7-9614-
volumes-2/volume-8858fad1-2523-4798-a7ea-
volumes-2/volume-bdb13b98-7dd2-4c24-98e2-
```

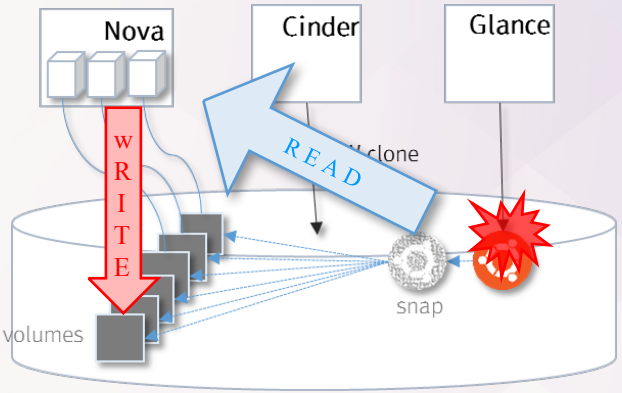

Image-Volume Cache

Image-Volume Cache Defines

- improve the performance of creating a volume from an image
- volume is first created from an image
  → a new cached image-volume will be created

Configuration

```
[DEFAULT]
cinder_internal_tenant_project_id = {project id}
cinder_internal_tenant_user_id = {user id}
[ceph]
image_volume_cache_enabled = True
image_volume_cache_max_size_gb = 10000
image_volume_cache_max_count = 1000
```

Image-Volume Cache

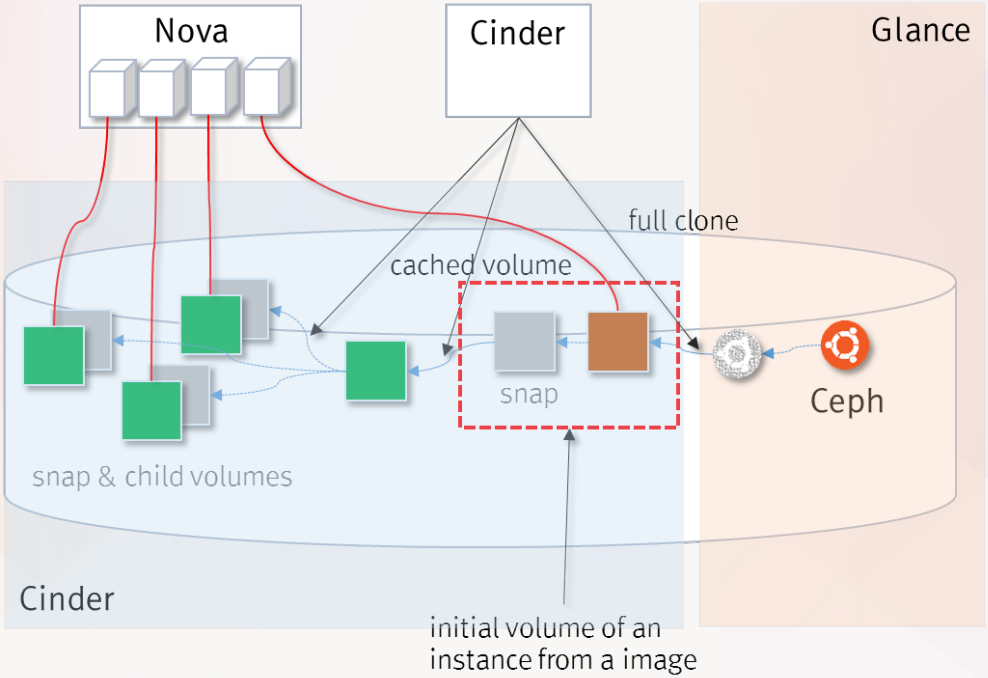

Image-Volume Cache in Cinder

- cached volume in cinder

```
| 02d6b027-59ae-43c0-b22b-4dd08faca7ee | available      | image-965170a8-a646-4f56-88f1-f39fb4f518bb | 41   | ceph
  | false    |                                      |
| 070df76d-3de2-4404-bacf-0fa9410bd68c | available      | image-0e3436ae-137f-4a6e-b88e-70405e5b0095 | 41   | ceph
  | false    |                                      |
| 15c6e8b3-a18e-4035-bbac-38bbfcb7d871 | available      | image-cfdbbbdf-6b00-4ca6-bb3e-71e224803126 | 40   | ceph
  | false    |                                      |
| 17fc911c-c343-4aa5-aad6-636948c3914e | available      | image-0e3436ae-137f-4a6e-b88e-70405e5b0095 | 41   | ceph
  | false    |                                      |
| 1e1f4de0-f0fa-4587-b26b-816c62f7cf9b | available      | image-0e3436ae-137f-4a6e-b88e-70405e5b0095 | 41   | ceph
  | false    |                                      |
| 4a6e7f4a-97c5-4f31-8c96-cc28a2996699 | available      | image-452c4abf-cad3-4373-acaa-4629596d0af1 | 40   | ceph
  | false    |                                      |
| 5346eacb-cafa-4e55-ab6c-112b25db0512 | available      | image-cfdbbbdf-6b00-4ca6-bb3e-71e224803126 | 40   | ceph
  | false    |                                      |
| 68d4d89a-9c6f-410e-bfb9-a71c1f79f1cd | available      | image-72550f64-2fa1-4fda-b0db-20a7b7cf74c0 | 40   | ceph
  | false    |                                      |
```

How Ceph provides image-cache

volume from an instance initially created.

```
root@f1cephmon01:~# rbd -p volumes-5 info volume-db5
rbd image 'volume-db5d7b9d-9759-4de6-b0eb-a900902fbe
        size 40960 MB in 10240 objects
```

snapshot from the the volume

```
root@f1cephmon01:~# rbd -p volumes-5 info volu
0eb-a900902fbe3b.deleted@volume-6ff9d1a7-1996-
clone_snap
rbd image 'volume-db5d7b9d-9759-4de6-b0eb-a900
        size 40960 MB in 10240 objects
```

cached image volume

```
root@f1cephmon01:~# rbd -p volumes-5 info volume-6ff9
92a-4831297cfe72
rbd image 'volume-6ff9d1a7-1996-4f86-892a-4831297cfe7
        parent: volumes-5/volume-db5d7b9d-9759-4de6-b0
.deleted@volume-6ff9d1a7-1996-4f86-892a-4831297cfe72.c
```

Concern about Disaster Recovery

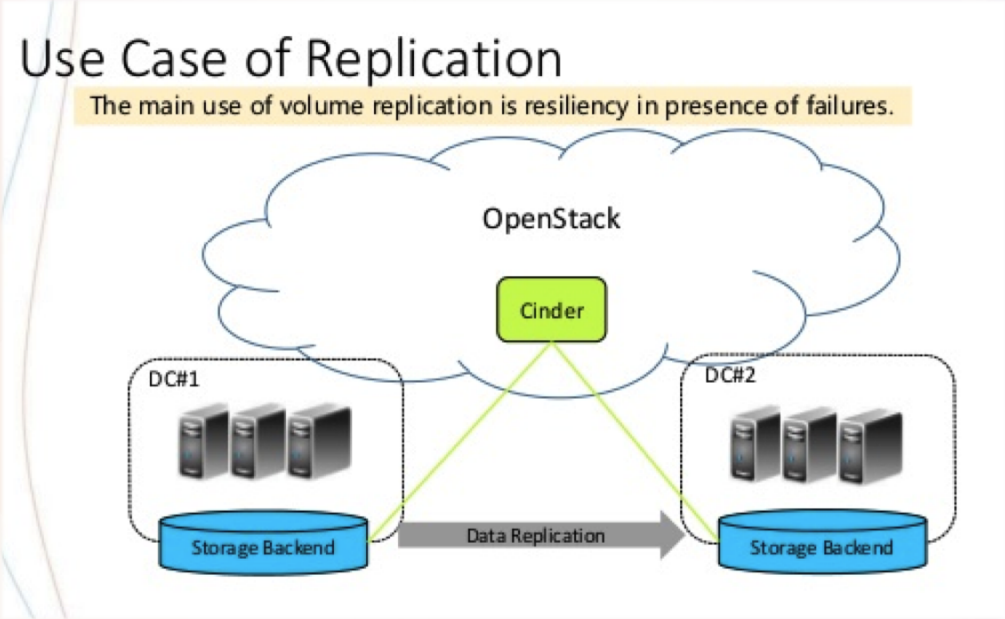

Replication in Cinder

Replication in Cinder Analysis

- depends on the driver's implementation
- There is no automatic failover since the use case is Disaster Recovery, and it must be done manually when the primary backend is out.

Ceph Replication in Cinder

Step	To Do
1	prepare different ceph cluster
2	configure ceph clusters in mirrored mode and to mirror the pool used by cinder.
3	copy cluster keys to the cinder volume node.
4	configure ceph driver in cinder to use replication.

RBD Mirroring

RBD Mirroring Defines

- asynchronously mirrored between two ceph clusters
- uses the RBD journaling image feature to ensure crash-consistent replication
- 2 ways : per pool(automatically mirrored), specific subset of images
- Use Cases : disaster recovery, global block device distribution

Cinder Volume Replication

Peering

- primary

```
root@cluster001:~ rbd mirror pool peer add volumes client.admin@ceph-secondary
5d4fbcdb-c7e5-4966-9c24-fdfcf4413b28
root@cluster001:~# rbd mirror pool status volumes
health: OK
images: 0 total
```

- secondary

```
root@cluster002:~# rbd mirror pool peer add volumes client.admin@ceph-primary
d6ec5046-becd-4a06-9ad2-8f18cb396e08
root@cluster002:~# rbd mirror pool status volumes
health: OK
images: 0 total
```

Cinder Volume Replication

Replication Status

- cluster 1

```
root@cluster001:~# rbd -p volumes info volume-aab62739-b544-454a-a219-12d9b4006372
rbd image 'volume-aab62739-b544-454a-a219-12d9b4006372':
size 1024 MB in 256 objects
order 22 (4096 kB objects)
block_name_prefix: rbd_data.3a9c3a810770
format: 2
features: layering, exclusive-lock, object-map, fast-diff, deep-flatten, journaling
journal: 3a9c3a810770
mirroring state: enabled
mirroring global id: 82eaa5f2-be3d-4954-a5fe-d14477fb5fed
mirroring primary: true
```

```
root@cluster001:~# rbd -p volumes mirror image status volume-aab62739-b544-454a-a219-12d9b4006372
  global_id:   82eaa5f2-be3d-4954-a5fe-d14477fb5fed
  state:       up+stopped
  description: remote image is non-primary or local image is primary
  last_update: 2017-05-01 00:57:19
```

- cluster 2

```
root@cluster001:~# rbd -p volumes info volume-aab62739-b544-454a-a219-12d9b4006372
rbd image 'volume-aab62739-b544-454a-a219-12d9b4006372':
size 1024 MB in 256 objects
order 22 (4096 kB objects)
block_name_prefix: rbd_data.3a9c3a810770
format: 2
features: layering, exclusive-lock, object-map, fast-diff, deep-flatten, journaling
journal: 3a9c3a810770
mirroring state: enabled
mirroring global id: 82eaa5f2-be3d-4954-a5fe-d14477fb5fed
mirroring primary: true
```

```
root@cluster2001:/etc/ceph# rbd -p volumes mirror image status volume-aab62739-b544-454a-a219-12d9b4006372
volume-aab62739-b544-454a-a219-12d9b4006372:
  global_id:   82eaa5f2-be3d-4954-a5fe-d14477fb5fed
  state:       up+replaying
  description: replaying, master_position=[object_number=3, tag_tid=1, entry_tid=3], mirror_position=[object_number=3, tag_tid=1, entry_tid=3], entries_behind_master=0
```

Cinder Volume Replication

Process of Fail Over

- **Before Fail-Over**

```
stack@openstack001:~$ cinder service-list --binary cinder-volume --withreplication
+--------------+-------------------------+------+---------+-------+----------------------------+-------------------+---
| Binary       | Host                    | Zone | Status  | State | Updated_at                 | Replication Status | Active Bac
+--------------+-------------------------+------+---------+-------+----------------------------+-------------------+---
| cinder-volume | openstack001@ceph-1     | nova | enabled | up    | 2017-04-30T16:12:52.000000 | enabled
+--------------+-------------------------+------+---------+-------+----------------------------+-------------------+---
```

- **After Fail-Over**

```
stack@openstack001:~$ cinder failover-host openstack001@ceph-1
stack@openstack001:~$ cinder service-list --binary cinder-volume --withreplication
+--------------+-------------------------+------+---------+-------+----------------------------+-------------------+-------
| Binary       | Host                    | Zone | Status  | State | Updated_at                 | Replication Status | Active Backer
+--------------+-------------------------+------+---------+-------+----------------------------+-------------------+-------
| cinder-volume | openstack001@ceph-1     | nova | enabled | up    | 2017-04-30T16:16:53.000000 | failing-over       |
+--------------+-------------------------+------+---------+-------+----------------------------+-------------------+-------
```

Cinder Volume Replication

Result of Fail Over

- cinder list

```
stack@openstack001:~$ cinder list
+--------------------------------------+-----------+-----------------+------+-------------+----------+-------------+
| ID                                   | Status    | Name            | Size | Volume Type | Bootable | Attached to |
+--------------------------------------+-----------+-----------------+------+-------------+----------+-------------+
| 7d886f63-3d00-4919-aa40-c89ce78b76e2 | error     | normal-ceph     | 1    | ceph        | false    |             |
| aab62739-b544-454a-a219-12d9b4006372 | available | replicated-ceph4 | 1    | replicated  | false    |             |
+--------------------------------------+-----------+-----------------+------+-------------+----------+-------------+
```

- cluster 1 & 2

```
root@cluster001:~# rbd -p volumes mirror image status volume-aab62739-b544-454a-a219-12d9b4006372
volume-aab62739-b544-454a-a219-12d9b4006372:
  global_id:   82eaa5f2-be3d-4954-a5fe-d14477fb5fed
  state:       up+replaying
  description: replaying, master_position=[], mirror_position=[], entries_behind_master=0
  last_update: 2017-05-01 01:17:49
```

```
root@cluster002:~# rbd -p volumes mirror image status volume-aab62739-b544-454a-a219-12d9b4006372
volume-aab62739-b544-454a-a219-12d9b4006372:
  global_id:   82eaa5f2-be3d-4954-a5fe-d14477fb5fed
  state:       up+stopped
  description: remote image is non-primary or local image is primary
  last_update: 2017-05-01 01:17:41
```


Straight roads do not make skillful drivers.
- Paulo Coelho -
PETRONOZ
BUZZ
Mobil 2

Q & A

Thanks