

..... 3

GlusterFS 3

GFS2 3

Ceph 3


 Ceph 1.3 (Hammer v0.94) 5

 Ceph 3 (Luminous v12.x) 5

 Ceph 4 (Nautilus 14.x) 5

 Ceph 5 (Pacific 16.x) 5

 Ceph 5



•

GlusterFS

- [Gluster Storage and Virtualization Compatibility](#)
- [GlusterFS 가](#)
- [GlusterFS Quick Start](#)
- [GlusterFS - Commands](#)
- [Gluster With Heketi REST API](#)


GFS2

- [GFS2 MetaData Unsync WithDraw](#)
- [GFS2 journal lock](#)
- [GFS2 lock dump](#)
- [multipath error getting device](#)


Ceph

Ceph Network Architecture .

1. Openstack



2. Ceph Storage



- [Ceph](#)
- [Ceph Codename - Version](#)
- [Ceph Compatibility Matrix](#)
- [Ceph RedHat - Community version](#)
- [Ceph Quick Start](#)
- [Ceph Storage PG reduce](#)
- [Linux](#)

- Redhat Ceph Storage Lifecycle
- - Ceph RBD with Samba HA (Pacemaker,LVM)
- Ceph Storage install in RHEL/CentOS
- Ceph HEALTH_WARN with 'mons are allowing insecure global_id reclaim'
- PG auto-scaling
- remove all objects from a pool
- Pool delete

Ceph 1.3 (Hammer v0.94)

- Ceph /
- Ceph Storage CRUSH map
- Ceph OSD 가 /
- Ceph Configuration
- Ceph OSD Failed Disk
- Ceph rbd lvm volume

Ceph 3 (Luminous v12.x)

- Ceph Journal disk
- Ceph OSD node
- CEPH pg,object repair - ceph status
- pg down peering warning

Ceph 4 (Nautilus 14.x)

- 1 daemons have recently crashed

Ceph 5 (Pacific 16.x)

Ceph

- Kraken 11.0

OSD	
OSD	
OSD	/etc/init.d/ceph stop osd.28 /sbin/stop ceph-osd id=28

- Luminous 12.0

systemd	
	systemctl stop ceph-[DAEMON_TYPE@\$ID]
OSD	systemctl stop ceph-osd.target
	systemctl stop ceph.target
	systemctl start ceph-[DAEMON_TYPE@\$ID]
OSD	systemctl start ceph-osd.target
	systemctl start ceph.target
	systemctl restart ceph-[DAEMON_TYPE@\$ID]
OSD	systemctl restart ceph-osd.target
	systemctl restart ceph.target

From:
<https://atl.kr/dokuwiki/> - AllThatLinux!

Permanent link:
<https://atl.kr/dokuwiki/doku.php/%EC%8A%A4%ED%86%A0%EB%A6%AC%EC%A7%80?rev=1655276047>

Last update: **2022/06/15 06:54**

