

**→ Convenciones:**

```
# En todos los nodos como 'sudo su'.  
[root@srv1 ~]# Solo en servidor 'srv1' → como 'sudo su'.  
[root@srv2 ~]# Solo en servidor 'srv2' → como 'sudo su'.
```

**363.2 Ceph Storage Clusters (weight: 8)****Weight** 8**Description** Candidates should be able to manage and maintain a Ceph Cluster. This includes the configuration of RGW, RDB devices and CephFS.**Key Knowledge Areas:**

- Understand the architecture and components of Ceph
- Manage OSD, MGR, MON and MDS
- Understand and manage placement groups and pools
- Understand storage backends (FileStore and BlueStore)
- Initialize a Ceph cluster
- Create and manage Rados Block Devices
- Create and manage CephFS volumes, including snapshots
- Mount and use an existing CephFS
- Understand and adjust CRUSH maps
- Configure high availability aspects of Ceph
- Scale up a Ceph cluster
- Restore and verify the integrity of a Ceph cluster after an outage
- Understand key concepts of Ceph updates, including update order, tunables and features

**Partial list of the used files, terms and utilities:**

- ceph-deploy (including relevant subcommands)
- ceph.conf
- ceph (including relevant subcommands)
- rados (including relevant subcommands)
- rbd (including relevant subcommands)
- cephfs (including relevant subcommands)
- ceph-volume (including relevant subcommands)
- ceph-authtool
- ceph-bluestore-tool
- crushtool

## → Conceptos Teóricos Preliminares

Red Hat Ceph Storage (RHCS) es una plataforma de almacenamiento escalable, abierta y definida por software que combina la versión más estable del sistema de almacenamiento Ceph con una plataforma de gestión Ceph, utilidades de implantación y servicios de soporte.

Red Hat Ceph Storage está diseñado para la infraestructura de nube y el almacenamiento de objetos a escala web. Los clústeres de Red Hat Ceph Storage constan de los siguientes tipos de nodos:

### Nodo de administración de Red Hat Ceph Storage Ansible

Este tipo de nodo actúa como lo hacía el nodo tradicional de administración de Ceph en las versiones anteriores de Red Hat Ceph Storage. Este tipo de nodo proporciona las siguientes funciones:

- Gestión centralizada del clúster de almacenamiento
- Los archivos de configuración y las claves de Ceph
- Opcionalmente, repositorios locales para instalar Ceph en nodos que no pueden acceder a Internet por razones de seguridad

#### Monitorear los nodos → [ceph-mon](#)

Cada nodo monitor ejecuta el demonio monitor (ceph-mon), que mantiene una copia maestra del mapa del cluster. El mapa del cluster incluye la topología del cluster. Un cliente que se conecta al clúster Ceph recupera la copia actual del mapa del clúster desde el monitor, lo que permite al cliente leer y escribir datos en el clúster.

#### Importante:

Ceph puede funcionar con un solo monitor; sin embargo, para asegurar una alta disponibilidad en un cluster de producción, Red Hat sólo soportará implementaciones con al menos tres nodos de monitorización. Red Hat recomienda desplegar un total de 5 monitores Ceph para clusters de almacenamiento que superen los 750 OSD.

### Nodos OSD → [ceph-osd](#)

Cada nodo de Dispositivo de Almacenamiento de Objetos (OSD) ejecuta el demonio Ceph OSD (ceph-osd), que interactúa con los discos lógicos conectados al nodo. Ceph almacena los datos en estos nodos OSD.

Ceph puede funcionar con muy pocos nodos OSD, que por defecto son tres, pero los clusters de producción obtienen un mejor rendimiento a partir de escalas modestas, por ejemplo 50 OSD en un cluster de almacenamiento. Lo ideal es que un clúster Ceph tenga varios nodos OSD, lo que permite crear dominios de fallo aislados mediante la creación del mapa CRUSH.

### Nodos MDS → [ceph-mds](#)

Cada nodo del Servidor de Metadatos (MDS) ejecuta el demonio MDS (ceph-mds), que gestiona los metadatos relacionados con los archivos almacenados en el Sistema de Archivos Ceph (CephFS). El demonio MDS también coordina el acceso al clúster compartido.

### **Nodo de la pasarela de objetos → [ceph-radosgw](#)**

El nodo Ceph Object Gateway ejecuta el demonio Ceph RADOS Gateway (ceph-radosgw), y es una interfaz de almacenamiento de objetos construida sobre **Librados** para proporcionar a las aplicaciones una puerta de enlace RESTful a los clusters de almacenamiento Ceph. El Ceph Object Gateway soporta dos interfaces:

#### **S3**

Proporciona la funcionalidad de almacenamiento de objetos con una interfaz que es compatible con un gran subconjunto de la API RESTful de Amazon S3.

#### **Swift**

Proporciona funcionalidad de almacenamiento de objetos con una interfaz que es compatible con un gran subconjunto de la API de OpenStack Swift.

→ **Configuración** → [ceph-deploy](#) ==> **Para CentOS 7** → [nautilus](#)

```
ceph-nodo1 ~ # python -V
```

```
Python 3.9.2
```

```
ceph-nodo1 ~ # pip-3.9 install ceph-deploy
```

```

-^-
/ \
|O o| ceph-deploy v2.0.1
)-.-(
'///\
|' |
|' |
|' |

```

→ **Topología:**

```
# vim /etc/hosts
```

```
# Cluster ceph.
```

```
192.168.10.160 ceph-mon.cadilinea.lan ceph-mon
```

```
192.168.10.161 ceph-osd1.cadilinea.lan ceph-osd1
```

```
192.168.10.162 ceph-osd2.cadilinea.lan ceph-osd2
```

```
192.168.10.163 ceph-osd3.cadilinea.lan ceph-osd3
```

→ **Sincronización Horaria** → [chronyd.service](#)

→ **Permisos sudo y ssh:**

```
# vim /etc/sudoers.d/ceph
```

```
ceph ALL=(root) NOPASSWD:ALL
```

```
# id ceph
```

```
uid=167(ceph) gid=167(ceph) grupos=167(ceph)
```

```
ceph-mon ~ # ssh-keygen
```

```
ceph-mon ~ # ssh-copy-id ceph-osd1
```

```
ceph-mon ~ # ssh-copy-id ceph-osd2
```

```
ceph-mon ~ # ssh-copy-id ceph-osd3
```

```
ceph-mon ~ # ssh-copy-id ceph-mon
```

```
ceph-mon ~ # vim ~/.ssh/config
```

```
Host ceph-mon
```

```
  Hostname ceph-mon.cadilinea.lan
```

```
  User root
```

```
Host ceph-osd1
```

```
  Hostname ceph-osd1.cadilinea.lan
```

```
  User root
```

```
Host ceph-osd2
```

```
  Hostname ceph-osd2.cadilinea.lan
```

```
  User root
```

```
Host ceph-osd3
```

```
  Hostname ceph-osd3.cadilinea.lan
```

```
  User root
```

```
ceph-mon ~ # chmod 0600 ~/.ssh/config
```

### → Instalación del software

```
# yum install epel-release
```

```
# rpm --import 'https://download.ceph.com/keys/release.asc'
```

```
# vim /etc/yum.repos.d/ceph.repo
```

```
[ceph]
```

```
name=Ceph packages for $basearch
```

```
baseurl=https://download.ceph.com/rpm-nautilus/el7/$basearch
```

```
enabled=1
```

```
priority=2
```

```
gpgcheck=1
```

```
gpgkey=https://download.ceph.com/keys/release.asc
```

```
[ceph-noarch]
```

```
name=Ceph noarch packages
```

```
baseurl=https://download.ceph.com/rpm-nautilus/el7//noarch
```

```
enabled=1
```

```
priority=2
gpgcheck=1
gpgkey=https://download.ceph.com/keys/release.asc
```

```
[ceph-source]
name=Ceph source packages
baseurl=https://download.ceph.com/rpm-nautilus/el7/SRPMS
enabled=0
priority=2
gpgcheck=1
gpgkey=https://download.ceph.com/keys/release.asc
```

```
# yum update
# yum install ceph ceph-deploy
```

→ → [Despliegue ceph](#) → [nautilus](#)

→ [Construimos el cluster](#)

```
ceph-mon ~ # mkdir /etc/ceph
ceph-mon ~ # cd /etc/ceph
```

```
ceph-mon ceph # ls /etc/ceph/
rbdmap
```

```
ceph-mon ceph # cat /etc/ceph/rbdmap
# RbdDevice      Parameters
#poolname/imagename      id=client,keyring=/etc/ceph/ceph.client.keyring
```

**\*\*  [Toda la configuración del despliegue se albergará en este directorio](#)  →  [/etc/ceph](#)  \*\***

```
ceph-mon ceph # ceph-deploy
usage: ceph-deploy [-h] [-v | -q] [--version] [--username USERNAME]
      [--overwrite-conf] [--ceph-conf CEPH_CONF]
      COMMAND ...
```

Easy Ceph deployment

```
  _^_
 /  \
|O o| ceph-deploy v2.0.1
) .-(
'///\`
|  |  |
  |  |
```

Full documentation can be found at: <http://ceph.com/ceph-deploy/docs>

optional arguments:

- h, --help show this help message and exit
- v, --verbose be more verbose
- q, --quiet be less verbose
- version the current installed version of ceph-deploy
- username USERNAME the username to connect to the remote host
- overwrite-conf overwrite an existing conf file on remote host (if present)
- ceph-conf CEPH\_CONF use (or reuse) a given ceph.conf file

commands:

COMMAND	description
<b>new</b>	Start deploying a new cluster, and write a CLUSTER.conf and keyring for it.
<b>install</b>	Install Ceph packages on remote hosts.
<b>rgw</b>	Ceph RGW daemon management
<b>mgr</b>	Ceph MGR daemon management
<b>mds</b>	Ceph MDS daemon management
<b>mon</b>	Ceph MON Daemon management
<b>gatherkeys</b>	Gather authentication keys for provisioning new nodes.
<b>disk</b>	Manage disks on a remote host.
<b>osd</b>	Prepare a data disk on remote host.
<b>repo</b>	Repo definition management
<b>admin</b>	Push configuration and client.admin key to a remote host.
<b>config</b>	Copy ceph.conf to/from remote host(s)
<b>uninstall</b>	Remove Ceph packages from remote hosts.
<b>purgedata</b>	Purge (delete, destroy, discard, shred) any Ceph data from /var/lib/ceph
<b>purge</b>	Remove Ceph packages from remote hosts and purge all data.
<b>forgetkeys</b>	Remove authentication keys from the local directory.
<b>pkg</b>	Manage packages on remote hosts.
<b>calamari</b>	Install and configure Calamari nodes. Assumes that a repository with Calamari packages is already configured. Refer to the docs for examples ( <a href="http://ceph.com/ceph-deploy/docs/conf.html">http://ceph.com/ceph-deploy/docs/conf.html</a> )

**ceph-mon ceph-cluster # ceph --version**

ceph version 14.2.21 (5ef401921d7a88aea18ec7558f7f9374ebd8f5a6) **nautilus (stable)**

→ **Desplegamos el monitoring** → **ceph-mon**

**ceph-mon ceph # ceph-deploy new ceph-mon --public-network 192.168.10.0/24**

```
[ceph_deploy.conf][DEBUG ] found configuration file at: /root/.cephdeploy.conf
[ceph_deploy.cli][INFO ] Invoked (2.0.1): /usr/bin/ceph-deploy new ceph-mon --public-network
192.168.10.0/24
[ceph_deploy.cli][INFO ] ceph-deploy options:
[ceph_deploy.cli][INFO ] username           : None
[ceph_deploy.cli][INFO ] func             : <function new at 0x7fdc192aa2a8>
[ceph_deploy.cli][INFO ] verbose          : False
[ceph_deploy.cli][INFO ] overwrite_conf   : False
[ceph_deploy.cli][INFO ] quiet            : False
[ceph_deploy.cli][INFO ] cd_conf          : <ceph_deploy.conf.cephdeploy.Conf instance at
0x7fdc18c285a8>
[ceph_deploy.cli][INFO ] cluster          : ceph
[ceph_deploy.cli][INFO ] ssh_copykey      : True
[ceph_deploy.cli][INFO ] mon              : ['ceph-mon']
[ceph_deploy.cli][INFO ] public_network   : 192.168.10.0/24
[ceph_deploy.cli][INFO ] ceph_conf        : None
[ceph_deploy.cli][INFO ] cluster_network  : None
[ceph_deploy.cli][INFO ] default_release  : False
[ceph_deploy.cli][INFO ] fsid             : None
[ceph_deploy.new][DEBUG ] Creating new cluster named ceph
[ceph_deploy.new][INFO ] making sure passwordless SSH succeeds
[ceph-mon][DEBUG ] connected to host: ceph-mon
[ceph-mon][DEBUG ] detect platform information from remote host
[ceph-mon][DEBUG ] detect machine type
[ceph-mon][DEBUG ] find the location of an executable
[ceph-mon][INFO ] Running command: /usr/sbin/ip link show
[ceph-mon][INFO ] Running command: /usr/sbin/ip addr show
[ceph-mon][DEBUG ] IP addresses found: [u'192.168.10.160']
[ceph_deploy.new][DEBUG ] Resolving host ceph-mon
[ceph_deploy.new][DEBUG ] Monitor ceph-mon at 192.168.10.160
[ceph_deploy.new][DEBUG ] Monitor initial members are ['ceph-mon']
[ceph_deploy.new][DEBUG ] Monitor addrs are [u'192.168.10.160']
[ceph_deploy.new][DEBUG ] Creating a random mon key...
[ceph_deploy.new][DEBUG ] Writing monitor keyring to ceph.mon.keyring...
[ceph_deploy.new][DEBUG ] Writing initial config to ceph.conf...
```

### ceph-mon ceph # ls

```
ceph.conf ceph-deploy-ceph.log ceph.mon.keyring rbdmap
```

### ceph-mon ceph # cat ceph.conf

```
[global]
fsid = 2137b4f3-ea3e-4721-bf88-98bfe6419420
public_network = 192.168.10.0/24
mon_initial_members = ceph-mon
mon_host = 192.168.10.160
auth_cluster_required = cephx
auth_service_required = cephx
```

```
auth_client_required = cephx
```

### ceph-mon ceph # cat ceph.mon.keyring

```
[mon.]
```

```
key = AQDdxahgAAAAABAAQsciS4Dh56TkwtWG5JhIfw==
```

```
caps mon = allow *
```

### ceph-mon ceph # cat ceph-deploy-ceph.log

```
[2021-05-22 10:50:36,901][ceph_deploy.conf][DEBUG ] found configuration file at:
/root/.cephdeploy.conf
[2021-05-22 10:50:36,902][ceph_deploy.cli][INFO ] Invoked (2.0.1): /usr/bin/ceph-deploy new
ceph-mon --public-network 192.168.10.0/24
[2021-05-22 10:50:36,902][ceph_deploy.cli][INFO ] ceph-deploy options:
[2021-05-22 10:50:36,902][ceph_deploy.cli][INFO ] username           : None
[2021-05-22 10:50:36,902][ceph_deploy.cli][INFO ] func                   : <function new at
0x7fdc192aa2a8>
[2021-05-22 10:50:36,902][ceph_deploy.cli][INFO ] verbose                : False
[2021-05-22 10:50:36,902][ceph_deploy.cli][INFO ] overwrite_conf        : False
[2021-05-22 10:50:36,902][ceph_deploy.cli][INFO ] quiet                  : False
[2021-05-22 10:50:36,902][ceph_deploy.cli][INFO ] cd_conf                :
<ceph_deploy.conf.cephdeploy.Conf instance at 0x7fdc18c285a8>
[2021-05-22 10:50:36,902][ceph_deploy.cli][INFO ] cluster                : ceph
[2021-05-22 10:50:36,902][ceph_deploy.cli][INFO ] ssh_copykey            : True
[2021-05-22 10:50:36,903][ceph_deploy.cli][INFO ] mon                    : ['ceph-mon']
[2021-05-22 10:50:36,903][ceph_deploy.cli][INFO ] public_network         : 192.168.10.0/24
[2021-05-22 10:50:36,903][ceph_deploy.cli][INFO ] ceph_conf              : None
[2021-05-22 10:50:36,903][ceph_deploy.cli][INFO ] cluster_network        : None
[2021-05-22 10:50:36,903][ceph_deploy.cli][INFO ] default_release        : False
[2021-05-22 10:50:36,903][ceph_deploy.cli][INFO ] fsid                   : None
[2021-05-22 10:50:36,903][ceph_deploy.new][DEBUG ] Creating new cluster named ceph
[2021-05-22 10:50:36,903][ceph_deploy.new][INFO ] making sure passwordless SSH succeeds
[2021-05-22 10:50:36,967][ceph-mon][DEBUG ] connected to host: ceph-mon
[2021-05-22 10:50:36,978][ceph-mon][DEBUG ] detect platform information from remote host
[2021-05-22 10:50:36,996][ceph-mon][DEBUG ] detect machine type
[2021-05-22 10:50:37,000][ceph-mon][DEBUG ] find the location of an executable
[2021-05-22 10:50:37,001][ceph-mon][INFO ] Running command: /usr/sbin/ip link show
[2021-05-22 10:50:37,026][ceph-mon][INFO ] Running command: /usr/sbin/ip addr show
[2021-05-22 10:50:37,031][ceph-mon][DEBUG ] IP addresses found: [u'192.168.10.160']
[2021-05-22 10:50:37,031][ceph_deploy.new][DEBUG ] Resolving host ceph-mon
[2021-05-22 10:50:37,031][ceph_deploy.new][DEBUG ] Monitor ceph-mon at 192.168.10.160
[2021-05-22 10:50:37,031][ceph_deploy.new][DEBUG ] Monitor initial members are ['ceph-mon']
[2021-05-22 10:50:37,031][ceph_deploy.new][DEBUG ] Monitor addrs are [u'192.168.10.160']
[2021-05-22 10:50:37,031][ceph_deploy.new][DEBUG ] Creating a random mon key...
[2021-05-22 10:50:37,031][ceph_deploy.new][DEBUG ] Writing monitor keyring to
ceph.mon.keyring...
[2021-05-22 10:50:37,031][ceph_deploy.new][DEBUG ] Writing initial config to ceph.conf...
```



→ **Instalar paquetería de despliegue en el cluster** → **ceph** → **Todos los nodos !.**

**ceph-mon ceph # ceph-deploy install ceph-mon ceph-osd1 ceph-osd2 ceph-osd3**

```
[ceph_deploy.conf][DEBUG ] found configuration file at: /root/.cephdeploy.conf
[ceph_deploy.cli][INFO ] Invoked (2.0.1): /usr/bin/ceph-deploy install ceph-mon ceph-osd1 ceph-osd2 ceph-osd3
[ceph_deploy.cli][INFO ] ceph-deploy options:
[ceph_deploy.cli][INFO ] verbose           : False
[ceph_deploy.cli][INFO ] testing          : None
[ceph_deploy.cli][INFO ] cd_conf         : <ceph_deploy.conf.cephdeploy.Conf instance at 0x7f19c348c7e8>
[ceph_deploy.cli][INFO ] cluster         : ceph
[ceph_deploy.cli][INFO ] dev_commit      : None
[ceph_deploy.cli][INFO ] install_mds     : False
[ceph_deploy.cli][INFO ] stable          : None
[ceph_deploy.cli][INFO ] default_release : False
[ceph_deploy.cli][INFO ] username        : None
[ceph_deploy.cli][INFO ] adjust_repos    : True
[ceph_deploy.cli][INFO ] func            : <function install at 0x7f19c4369a28>
[ceph_deploy.cli][INFO ] install_mgr     : False
[ceph_deploy.cli][INFO ] install_all     : False
[ceph_deploy.cli][INFO ] repo            : False
[ceph_deploy.cli][INFO ] host            : ['ceph-mon', 'ceph-osd1', 'ceph-osd2', 'ceph-osd3']
...
[ceph-osd3][DEBUG ] Running transaction test
[ceph-osd3][DEBUG ] Transaction test succeeded
[ceph-osd3][DEBUG ] Running transaction
[ceph-osd3][DEBUG ] Instalando   : mailcap-2.1.41-2.el7.noarch           1/2
[ceph-osd3][DEBUG ] Instalando   : 2:ceph-radosgw-14.2.21-0.el7.x86_64 2/2
[ceph-osd3][DEBUG ] Comprobando  : mailcap-2.1.41-2.el7.noarch           1/2
[ceph-osd3][DEBUG ] Comprobando  : 2:ceph-radosgw-14.2.21-0.el7.x86_64 2/2
[ceph-osd3][DEBUG ]
[ceph-osd3][DEBUG ] Instalado:
[ceph-osd3][DEBUG ] ceph-radosgw.x86_64 2:14.2.21-0.el7
[ceph-osd3][DEBUG ]
[ceph-osd3][DEBUG ] Dependencia(s) instalada(s):
[ceph-osd3][DEBUG ] mailcap.noarch 0:2.1.41-2.el7
[ceph-osd3][DEBUG ]
[ceph-osd3][DEBUG ] ¡Listo!
[ceph-osd3][INFO ] Running command: ceph --version
[ceph-osd3][DEBUG ] ceph version 14.2.21 (5ef401921d7a88aea18ec7558f7f9374ebd8f5a6)
nautilus (stable)
```

→ **Desplegar monitorización inicial y** → **gatherkeys**

**ceph-mon ceph # ceph-deploy mon create-initial**

```
[ceph_deploy.conf][DEBUG ] found configuration file at: /root/.cephdeploy.conf
[ceph_deploy.cli][INFO ] Invoked (2.0.1): /usr/bin/ceph-deploy mon create-initial
[ceph_deploy.cli][INFO ] ceph-deploy options:
[ceph_deploy.cli][INFO ] username           : None
[ceph_deploy.cli][INFO ] verbose           : False
[ceph_deploy.cli][INFO ] overwrite_conf    : False
[ceph_deploy.cli][INFO ] subcommand        : create-initial
[ceph_deploy.cli][INFO ] quiet             : False
[ceph_deploy.cli][INFO ] cd_conf           : <ceph_deploy.conf.cephdeploy.Conf instance at
0x7fa102a6b200>
[ceph_deploy.cli][INFO ] cluster           : ceph
[ceph_deploy.cli][INFO ] func              : <function mon at 0x7fa102cc5848>
[ceph_deploy.cli][INFO ] ceph_conf         : None
[ceph_deploy.cli][INFO ] default_release   : False
[ceph_deploy.cli][INFO ] keyrings          : None
[ceph_deploy.mon][DEBUG ] Deploying mon, cluster ceph hosts ceph-mon
[ceph_deploy.mon][DEBUG ] detecting platform for host ceph-mon ...
...
[ceph_deploy.gatherkeys][INFO ] Storing ceph.client.admin.keyring
[ceph_deploy.gatherkeys][INFO ] Storing ceph.bootstrap-mds.keyring
[ceph_deploy.gatherkeys][INFO ] Storing ceph.bootstrap-mgr.keyring
[ceph_deploy.gatherkeys][INFO ] keyring 'ceph.mon.keyring' already exists
[ceph_deploy.gatherkeys][INFO ] Storing ceph.bootstrap-osd.keyring
[ceph_deploy.gatherkeys][INFO ] Storing ceph.bootstrap-rgw.keyring
[ceph_deploy.gatherkeys][INFO ] Destroy temp directory /tmp/tmp3SexnG
```

**ceph-mon ceph # ls**

```
ceph.bootstrap-mds.keyring ceph.bootstrap-osd.keyring ceph.client.admin.keyring ceph-deploy-
ceph.log rbdmap
ceph.bootstrap-mgr.keyring ceph.bootstrap-rgw.keyring ceph.conf ceph.mon.keyring
```

**ceph-mon ceph # cat ceph.bootstrap-mds.keyring**

```
[client.bootstrap-mds]
key = AQD/z6hgINp8GRAASAqDJHAQW9qSjgxOBKo2DA==
caps mon = "allow profile bootstrap-mds"
```

**ceph-mon ceph # cat ceph.bootstrap-osd.keyring**

```
[client.bootstrap-osd]
key = AQD/z6hgGfR8GRAAlfePyYlgvK/naPYsQED+KQ==
caps mon = "allow profile bootstrap-osd"
```

**ceph-mon ceph # cat ceph.client.admin.keyring**

```
[client.admin]
key = AQD/z6hgxsd8GRAAmq/f7TS6XtFq6lC8AC6DpA==
caps mds = "allow *"
caps mgr = "allow *"
```

```
caps mon = "allow *"
caps osd = "allow *"
```

### ceph-mon ceph # cat ceph.bootstrap-mgr.keyring

```
[client.bootstrap-mgr]
key = AQD/z6hgheed8GRAAJf7Va/G8YCXe0UH01Sw4pQ==
caps mon = "allow profile bootstrap-mgr"
```

### ceph-mon ceph # cat ceph.bootstrap-rgw.keyring

```
[client.bootstrap-rgw]
key = AQD/z6hgdx9GRAAdQpBq+R0zWsR5Z+Oaedsrw==
caps mon = "allow profile bootstrap-rgw"
```

→ **Desplegamos la configuración para admin y para administrar todos los nodos, y de forma autónoma.**

### ceph-mon ceph # ceph-deploy admin ceph-mon ceph-osd1 ceph-osd2 ceph-osd3

```
[ceph_deploy.conf][DEBUG ] found configuration file at: /root/.cephdeploy.conf
[ceph_deploy.cli][INFO ] Invoked (2.0.1): /usr/bin/ceph-deploy admin ceph-mon ceph-osd1 ceph-
osd2 ceph-osd3
[ceph_deploy.cli][INFO ] ceph-deploy options:
[ceph_deploy.cli][INFO ] username           : None
[ceph_deploy.cli][INFO ] verbose           : False
[ceph_deploy.cli][INFO ] overwrite_conf    : False
[ceph_deploy.cli][INFO ] quiet             : False
[ceph_deploy.cli][INFO ] cd_conf           : <ceph_deploy.conf.cephdeploy.Conf instance at
0x7f292cd5f758>
[ceph_deploy.cli][INFO ] cluster           : ceph
[ceph_deploy.cli][INFO ] client            : ['ceph-mon', 'ceph-osd1', 'ceph-osd2', 'ceph-osd3']
[ceph_deploy.cli][INFO ] func              : <function admin at 0x7f292d5f4668>
[ceph_deploy.cli][INFO ] ceph_conf         : None
[ceph_deploy.cli][INFO ] default_release   : False
[ceph_deploy.admin][DEBUG ] Pushing admin keys and conf to ceph-mon
[ceph-mon][DEBUG ] connected to host: ceph-mon
[ceph-mon][DEBUG ] detect platform information from remote host
[ceph-mon][DEBUG ] detect machine type
[ceph-mon][DEBUG ] write cluster configuration to /etc/ceph/{cluster}.conf
[ceph_deploy.admin][DEBUG ] Pushing admin keys and conf to ceph-osd1
[ceph-osd1][DEBUG ] connected to host: ceph-osd1
[ceph-osd1][DEBUG ] detect platform information from remote host
[ceph-osd1][DEBUG ] detect machine type
[ceph-osd1][DEBUG ] write cluster configuration to /etc/ceph/{cluster}.conf
[ceph_deploy.admin][DEBUG ] Pushing admin keys and conf to ceph-osd2
[ceph-osd2][DEBUG ] connected to host: ceph-osd2
[ceph-osd2][DEBUG ] detect platform information from remote host
[ceph-osd2][DEBUG ] detect machine type
[ceph-osd2][DEBUG ] write cluster configuration to /etc/ceph/{cluster}.conf
```

```
[ceph_deploy.admin][DEBUG ] Pushing admin keys and conf to ceph-osd3
[ceph-osd3][DEBUG ] connected to host: ceph-osd3
[ceph-osd3][DEBUG ] detect platform information from remote host
[ceph-osd3][DEBUG ] detect machine type
[ceph-osd3][DEBUG ] write cluster configuration to /etc/ceph/{cluster}.conf
```

```
ceph-mon ceph # ssh ceph-osd3
```

```
ceph-osd3 ~ # tree /etc/ceph/
```

```
/etc/ceph/
├── ceph.client.admin.keyring
├── ceph.conf
├── rbdmap
└── tmpMWUJcy
```

```
ceph-osd3 ~ # cat /etc/ceph/ceph.client.admin.keyring
```

```
[client.admin]
    key = AQD/z6hgxsd8GRAAmq/f7TS6XtFq6lC8AC6DpA==
    caps mds = "allow *"
    caps mgr = "allow *"
    caps mon = "allow *"
    caps osd = "allow *"
```

```
ceph-osd3 ~ # cat /etc/ceph/ceph.conf
```

```
[global]
fsid = 2137b4f3-ea3e-4721-bf88-98bfe6419420
public_network = 192.168.10.0/24
mon_initial_members = ceph-mon
mon_host = 192.168.10.160
auth_cluster_required = cephx
auth_service_required = cephx
auth_client_required = cephx
```

```
ceph-osd3 ~ # cat /etc/ceph/rbdmap
```

```
# RbdDevice      Parameters
#poolname/imagename    id=client,keyring=/etc/ceph/ceph.client.keyring
```

```
ceph-osd3 ~ # cat /etc/ceph/tmpMWUJcy
```

→ Estado actual del cluster

```
ceph-mon ceph # ceph -s
```

```
cluster:
id: 2137b4f3-ea3e-4721-bf88-98bfe6419420
health: HEALTH_WARN
    mon is allowing insecure global_id reclaim
```

```
services:
  mon: 1 daemons, quorum ceph-mon (age 35m)
  mgr: no daemons active
  osd: 0 osds: 0 up, 0 in
```

```
data:
  pools: 0 pools, 0 pgs
  objects: 0 objects, 0 B
  usage: 0 B used, 0 B / 0 B avail
  pgs:
```

```
ceph-mon ceph # ceph config set mon auth_allow_insecure_global_id_reclaim false
```

```
ceph-mon ~ # ceph config dump
```

```
WHO MASK LEVEL OPTION VALUE RO
  mon advanced auth_allow_insecure_global_id_reclaim false
```

```
ceph-mon ceph # ceph -w
```

```
cluster:
  id: 2137b4f3-ea3e-4721-bf88-98bfe6419420
  health: HEALTH_OK
```

```
services:
  mon: 1 daemons, quorum ceph-mon (age 39m)
  mgr: no daemons active
  osd: 0 osds: 0 up, 0 in
```

```
data:
  pools: 0 pools, 0 pgs
  objects: 0 objects, 0 B
  usage: 0 B used, 0 B / 0 B avail
  pgs:
```

```
2021-05-22 12:13:43.386908 mon.ceph-mon [INF] Health check cleared:
AUTH_INSECURE_GLOBAL_ID_RECLAIM_ALLOWED (was: mon is allowing insecure
global_id reclaim)
```

```
2021-05-22 12:13:43.386933 mon.ceph-mon [INF] Cluster is now healthy
```

```
<CTRL-C>
```

→ **Desplegamos los gatherkeys para aprovisionar los nodos.**

```
ceph-mon ceph # ceph-deploy gatherkeys ceph-mon
```

```
[ceph_deploy.conf][DEBUG ] found configuration file at: /root/.cephdeploy.conf
[ceph_deploy.cli][INFO ] Invoked (2.0.1): /usr/bin/ceph-deploy gatherkeys ceph-mon
[ceph_deploy.cli][INFO ] ceph-deploy options:
[ceph_deploy.cli][INFO ] username           : None
[ceph_deploy.cli][INFO ] verbose           : False
[ceph_deploy.cli][INFO ] overwrite_conf    : False
```

```

[ceph_deploy.cli][INFO ] quiet                : False
[ceph_deploy.cli][INFO ] cd_conf             : <ceph_deploy.conf.cephdeploy.Conf instance at
0x7fcf8768bf80>
[ceph_deploy.cli][INFO ] cluster            : ceph
[ceph_deploy.cli][INFO ] mon                : ['ceph-mon']
[ceph_deploy.cli][INFO ] func               : <function gatherkeys at 0x7fcf878d2de8>
[ceph_deploy.cli][INFO ] ceph_conf          : None
[ceph_deploy.cli][INFO ] default_release    : False
[ceph_deploy.gatherkeys][INFO ] Storing keys in temp directory /tmp/tmpwID6sj
[ceph-mon][DEBUG ] connected to host: ceph-mon
[ceph-mon][DEBUG ] detect platform information from remote host
[ceph-mon][DEBUG ] detect machine type
[ceph-mon][DEBUG ] get remote short hostname
[ceph-mon][DEBUG ] fetch remote file
[ceph-mon][INFO ] Running command: /usr/bin/ceph --connect-timeout=25 --cluster=ceph --
admin-daemon=/var/run/ceph/ceph-mon.ceph-mon.asok mon_status
[ceph-mon][INFO ] Running command: /usr/bin/ceph --connect-timeout=25 --cluster=ceph --name
mon. --keyring=/var/lib/ceph/mon/ceph-ceph-mon/keyring auth get client.admin
[ceph-mon][INFO ] Running command: /usr/bin/ceph --connect-timeout=25 --cluster=ceph --name
mon. --keyring=/var/lib/ceph/mon/ceph-ceph-mon/keyring auth get client.bootstrap-mds
[ceph-mon][INFO ] Running command: /usr/bin/ceph --connect-timeout=25 --cluster=ceph --name
mon. --keyring=/var/lib/ceph/mon/ceph-ceph-mon/keyring auth get client.bootstrap-mgr
[ceph-mon][INFO ] Running command: /usr/bin/ceph --connect-timeout=25 --cluster=ceph --name
mon. --keyring=/var/lib/ceph/mon/ceph-ceph-mon/keyring auth get client.bootstrap-osd
[ceph-mon][INFO ] Running command: /usr/bin/ceph --connect-timeout=25 --cluster=ceph --name
mon. --keyring=/var/lib/ceph/mon/ceph-ceph-mon/keyring auth get client.bootstrap-rgw
[ceph_deploy.gatherkeys][INFO ] keyring 'ceph.client.admin.keyring' already exists
[ceph_deploy.gatherkeys][INFO ] keyring 'ceph.bootstrap-mds.keyring' already exists
[ceph_deploy.gatherkeys][INFO ] keyring 'ceph.bootstrap-mgr.keyring' already exists
[ceph_deploy.gatherkeys][INFO ] keyring 'ceph.mon.keyring' already exists
[ceph_deploy.gatherkeys][INFO ] keyring 'ceph.bootstrap-osd.keyring' already exists
[ceph_deploy.gatherkeys][INFO ] keyring 'ceph.bootstrap-rgw.keyring' already exists
[ceph_deploy.gatherkeys][INFO ] Destroy temp directory /tmp/tmpwID6sj

```

### → Despliegue/Creación de **ceph-mgr** daemon

#### **ceph-mon** **ceph** # **ceph-deploy mgr create ceph-mon**

```

[ceph_deploy.conf][DEBUG ] found configuration file at: /root/.cephdeploy.conf
[ceph_deploy.cli][INFO ] Invoked (2.0.1): /usr/bin/ceph-deploy mgr create ceph-mon
[ceph_deploy.cli][INFO ] ceph-deploy options:
[ceph_deploy.cli][INFO ] username                : None
[ceph_deploy.cli][INFO ] verbose                 : False
[ceph_deploy.cli][INFO ] mgr                     : [('ceph-mon', 'ceph-mon')]
[ceph_deploy.cli][INFO ] overwrite_conf          : False
[ceph_deploy.cli][INFO ] subcommand              : create
[ceph_deploy.cli][INFO ] quiet                   : False
[ceph_deploy.cli][INFO ] cd_conf                 : <ceph_deploy.conf.cephdeploy.Conf instance at

```

```

0x7fb606ab3a70>
[ceph_deploy.cli][INFO ] cluster           : ceph
[ceph_deploy.cli][INFO ] func           : <function mgr at 0x7fb607117578>
[ceph_deploy.cli][INFO ] ceph_conf      : None
[ceph_deploy.cli][INFO ] default_release : False
[ceph_deploy.mgr][DEBUG ] Deploying mgr, cluster ceph hosts ceph-mon:ceph-mon
[ceph-mon][DEBUG ] connected to host: ceph-mon
[ceph-mon][DEBUG ] detect platform information from remote host
[ceph-mon][DEBUG ] detect machine type
[ceph_deploy.mgr][INFO ] Distro info: CentOS Linux 7.9.2009 Core
[ceph_deploy.mgr][DEBUG ] remote host will use systemd
[ceph_deploy.mgr][DEBUG ] deploying mgr bootstrap to ceph-mon
[ceph-mon][DEBUG ] write cluster configuration to /etc/ceph/{cluster}.conf
[ceph-mon][WARNIN] mgr keyring does not exist yet, creating one
[ceph-mon][DEBUG ] create a keyring file
[ceph-mon][DEBUG ] create path recursively if it doesn't exist
[ceph-mon][INFO ] Running command: ceph --cluster ceph --name client.bootstrap-mgr --
keyring /var/lib/ceph/bootstrap-mgr/ceph.keyring auth get-or-create mgr.ceph-mon mon allow
profile mgr osd allow * mds allow * -o /var/lib/ceph/mgr/ceph-ceph-mon/keyring
[ceph-mon][INFO ] Running command: systemctl enable ceph-mgr@ceph-mon
[ceph-mon][WARNIN] Created symlink from /etc/systemd/system/ceph-mgr.target.wants/ceph-
mgr@ceph-mon.service to /usr/lib/systemd/system/ceph-mgr@.service.
[ceph-mon][INFO ] Running command: systemctl start ceph-mgr@ceph-mon
[ceph-mon][INFO ] Running command: systemctl enable ceph.target
    
```

### ceph-mon ceph # systemctl status ceph

```

ceph-crash.service      ceph-mgr@ceph-mon.service ceph-mon@ceph-mon.service ceph-
osd.target             ceph.target
ceph-mds.target         ceph-mgr.target          ceph-mon.target          ceph-radosgw.target
    
```

### ceph-mon ceph # systemctl status ceph.target

- ceph.target - ceph target allowing to start/stop all ceph\*@.service instances at once  
 Loaded: loaded (/usr/lib/systemd/system/ceph.target; enabled; vendor preset: enabled)  
 Active: active since sáb 2021-05-22 08:29:58 CEST; 4h 34min ago

```

may 22 08:29:58 ceph-mon.cadilinea.lan systemd[1]: Reached target ceph target allowing to
start/stop all ceph*@.service instanc... once.
    
```

Hint: Some lines were ellipsized, use -l to show in full.

### ceph-mon ceph # ceph -s

```

cluster:
  id: 2137b4f3-ea3e-4721-bf88-98bfe6419420
  health: HEALTH_WARN
         OSD count 0 < osd_pool_default_size 3

services:
  mon: 1 daemons, quorum ceph-mon (age 95m)
    
```

```
mgr: ceph-mon(active, since 8m)
osd: 0 osds: 0 up, 0 in
```

```
data:
pools: 0 pools, 0 pgs
objects: 0 objects, 0 B
usage: 0 B used, 0 B / 0 B avail
pgs:
```

→ **Habilitamos Firewall** → **ceph**

```
ceph-mon ceph # firewall-cmd --permanent --add-service={ceph,ceph-mon}
ceph-mon ceph # firewall-cmd --reload
```

→ **Preparamos discos (como blocks) en los hosts remotos** → **ceph-osd1 ceph-osd2 ceph-osd3**

```
# lsblk /dev/vd[b-z]
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vdb 252:16 0 20G 0 disk
vdc 252:32 0 20G 0 disk
vdd 252:48 0 20G 0 disk
```

→ **Creación de servidores ceph-osd**

```
ceph-mon ceph # ceph-deploy osd create ceph-osd1 --data /dev/vdb
ceph-mon ceph # ceph-deploy osd create ceph-osd2 --data /dev/vdb
ceph-mon ceph # ceph-deploy osd create ceph-osd3 --data /dev/vdb
```

→ **Ejemplo:**

```
ceph-mon ceph # ceph-deploy osd create ceph-osd3 --data /dev/vdb
[ceph_deploy.conf][DEBUG ] found configuration file at: /root/.cephdeploy.conf
[ceph_deploy.cli][INFO ] Invoked (2.0.1): /usr/bin/ceph-deploy osd create ceph-osd3 --data
/dev/vdb
[ceph_deploy.cli][INFO ] ceph-deploy options:
[ceph_deploy.cli][INFO ] verbose : False
[ceph_deploy.cli][INFO ] bluestore : None
[ceph_deploy.cli][INFO ] cd_conf : <ceph_deploy.conf.cephdeploy.Conf instance at
0x7f7645854758>
[ceph_deploy.cli][INFO ] cluster : ceph
[ceph_deploy.cli][INFO ] fs_type : xfs
...
```

```
[ceph-osd3][WARNIN] --> ceph-volume lvm prepare successful for: /dev/vdb
[ceph-osd3][WARNIN] Running command: /bin/chown -R ceph:ceph /var/lib/ceph/osd/ceph-2
[ceph-osd3][WARNIN] Running command: /bin/ceph-bluestore-tool --cluster=ceph prime-osd-dir
--dev /dev/ceph-4cc2b8e2-bbdf-45ec-9f84-613a9d26616b/osd-block-9d3f2eca-2f78-491f-ac3e-
640f03d264ce --path /var/lib/ceph/osd/ceph-2 --no-mon-config
```



```
[ceph-osd3][WARNIN] Running command: /bin/ln -snf /dev/ceph-4cc2b8e2-bbdf-45ec-9f84-613a9d26616b/osd-block-9d3f2eca-2f78-491f-ac3e-640f03d264ce /var/lib/ceph/osd/ceph-2/block
[ceph-osd3][WARNIN] Running command: /bin/chown -h ceph:ceph /var/lib/ceph/osd/ceph-2/block
[ceph-osd3][WARNIN] Running command: /bin/chown -R ceph:ceph /dev/dm-6
[ceph-osd3][WARNIN] Running command: /bin/chown -R ceph:ceph /var/lib/ceph/osd/ceph-2
[ceph-osd3][WARNIN] Running command: /bin/systemctl enable ceph-volume@lvm-2-9d3f2eca-2f78-491f-ac3e-640f03d264ce
[ceph-osd3][WARNIN] stderr: Created symlink from /etc/systemd/system/multi-user.target.wants/ceph-volume@lvm-2-9d3f2eca-2f78-491f-ac3e-640f03d264ce.service to /usr/lib/systemd/system/ceph-volume@.service.
[ceph-osd3][WARNIN] Running command: /bin/systemctl enable --runtime ceph-osd@2
[ceph-osd3][WARNIN] stderr: Created symlink from /run/systemd/system/ceph-osd.target.wants/ceph-osd@2.service to /usr/lib/systemd/system/ceph-osd@.service.
[ceph-osd3][WARNIN] Running command: /bin/systemctl start ceph-osd@2
[ceph-osd3][WARNIN] --> ceph-volume lvm activate successful for osd ID: 2
[ceph-osd3][WARNIN] --> ceph-volume lvm create successful for: /dev/vdb
[ceph-osd3][INFO ] checking OSD status...
[ceph-osd3][DEBUG ] find the location of an executable
[ceph-osd3][INFO ] Running command: /bin/ceph --cluster=ceph osd stat --format=json
[ceph_deploy.osd][DEBUG ] Host ceph-osd3 is now ready for osd use.
```

### ceph-mon ceph # ceph -s

```
cluster:
  id: 2137b4f3-ea3e-4721-bf88-98bfe6419420
  health: HEALTH_WARN
        2 osds down
        2 hosts (2 osds) down
```

```
services:
  mon: 1 daemons, quorum ceph-mon (age 19m)
  mgr: ceph-mon(active, since 18m)
  osd: 3 osds: 1 up (since 56s), 3 in (since 2m)
```

```
data:
  pools: 0 pools, 0 pgs
  objects: 0 objects, 0 B
  usage: 3.0 GiB used, 57 GiB / 60 GiB avail
  pgs:
```

### ceph-mon ceph # ssh ceph-osd1 lsblk /dev/vdb

```
NAME                                                                 MAJ:MIN RM SIZE RO
TYPE MOUNTPOINT
vdb                                                                 252:16  0 20G  0 disk
└─ceph--4744afc0--6372--44b5--a74e--7bd0d30d5ad7-osd--block--8acc47ed--3737--4f21--8f80--922a176b9c28 253:6  0 20G  0 lvm
```

**ceph-mon ceph # ceph osd tree**

ID	CLASS	WEIGHT	TYPE	NAME	STATUS	REWEIGHT	PRI-AFF
-1		0.05846	root	default			
-3		0.01949	host	ceph-osd1			
0	hdd	0.01949	osd.0	down	1.00000	1.00000	
-5		0.01949	host	ceph-osd2			
1	hdd	0.01949	osd.1	down	1.00000	1.00000	
-7		0.01949	host	ceph-osd3			
2	hdd	0.01949	osd.2	up	1.00000	1.00000	

**ceph-mon ceph # ceph-deploy disk list ceph-osd1**

```
...
[ceph-osd1][INFO ] Disk /dev/mapper/ceph--4744afc0--6372--44b5--a74e--7bd0d30d5ad7-osd--
block--8acc47ed--3737--4f21--8f80--922a176b9c28: 21.5 GB, 21470642176 bytes, 41934848
sectors
```

**ceph-mon ceph # ceph -s**

cluster:

id: 2137b4f3-ea3e-4721-bf88-98bfe6419420

health: HEALTH\_OK

services:

mon: 1 daemons, quorum ceph-mon (age 8m)

mgr: ceph-mon(active, since 7m)

osd: 3 osds: 3 up (since 17s), 3 in (since 11m)

data:

pools: 0 pools, 0 pgs

objects: 0 objects, 0 B

usage: 3.0 GiB used, 57 GiB / 60 GiB avail

pgs:

→ Creación de servidores de metadatos → ceph-mds

**ceph-mon ceph # ceph-deploy mds create ceph-osd1 ceph-osd2 ceph-osd3**

[ceph\_deploy.conf][DEBUG ] found configuration file at: /root/.cephdeploy.conf

[ceph\_deploy.cli][INFO ] Invoked (2.0.1): /usr/bin/ceph-deploy mds create ceph-osd1 ceph-osd2 ceph-osd3

[ceph\_deploy.cli][INFO ] ceph-deploy options:

[ceph\_deploy.cli][INFO ] username : None

[ceph\_deploy.cli][INFO ] verbose : False

[ceph\_deploy.cli][INFO ] overwrite\_conf : False

[ceph\_deploy.cli][INFO ] subcommand : create

[ceph\_deploy.cli][INFO ] quiet : False

[ceph\_deploy.cli][INFO ] cd\_conf : &lt;ceph\_deploy.conf.cephdeploy.Conf instance at

```
0x7f963bcd8638>
[ceph_deploy.cli][INFO ] cluster           : ceph
[ceph_deploy.cli][INFO ] func           : <function mds at 0x7f963bf2b398>
[ceph_deploy.cli][INFO ] ceph_conf      : None
[ceph_deploy.cli][INFO ] mds            : [('ceph-osd1', 'ceph-osd1'), ('ceph-osd2', 'ceph-
osd2'), ('ceph-osd3', 'ceph-osd3')]
...
[ceph-osd3][WARNIN] mds keyring does not exist yet, creating one
[ceph-osd3][DEBUG ] create a keyring file
[ceph-osd3][DEBUG ] create path if it doesn't exist
[ceph-osd3][INFO ] Running command: ceph --cluster ceph --name client.bootstrap-mds --
keyring /var/lib/ceph/bootstrap-mds/ceph.keyring auth get-or-create mds.ceph-osd3 osd allow rwx
mds allow mon allow profile mds -o /var/lib/ceph/mds/ceph-ceph-osd3/keyring
[ceph-osd3][INFO ] Running command: systemctl enable ceph-mds@ceph-osd3
[ceph-osd3][WARNIN] Created symlink from /etc/systemd/system/ceph-mds.target.wants/ceph-
mds@ceph-osd3.service to /usr/lib/systemd/system/ceph-mds@.service.
[ceph-osd3][INFO ] Running command: systemctl start ceph-mds@ceph-osd3
[ceph-osd3][INFO ] Running command: systemctl enable ceph.target
```

→ **Monitorización de todos los nodos** → [ceph-osd1 ceph-osd2 ceph-osd3](#)

**ceph-mon ceph # ceph-deploy mon create ceph-osd1 ceph-osd2 ceph-osd3**

```
[ceph_deploy.conf][DEBUG ] found configuration file at: /root/.cephdeploy.conf
[ceph_deploy.cli][INFO ] Invoked (2.0.1): /usr/bin/ceph-deploy mon create ceph-osd1 ceph-osd2
ceph-osd3
[ceph_deploy.cli][INFO ] ceph-deploy options:
[ceph_deploy.cli][INFO ] username           : None
[ceph_deploy.cli][INFO ] verbose            : False
[ceph_deploy.cli][INFO ] overwrite_conf     : False
[ceph_deploy.cli][INFO ] subcommand         : create
[ceph_deploy.cli][INFO ] quiet              : False
[ceph_deploy.cli][INFO ] cd_conf            : <ceph_deploy.conf.cephdeploy.Conf instance at
0x7fe73c6a8200>
[ceph_deploy.cli][INFO ] cluster           : ceph
[ceph_deploy.cli][INFO ] mon               : ['ceph-osd1', 'ceph-osd2', 'ceph-osd3']
[ceph_deploy.cli][INFO ] func              : <function mon at 0x7fe73c902848>
[ceph_deploy.cli][INFO ] ceph_conf         : None
[ceph_deploy.cli][INFO ] default_release   : False
[ceph_deploy.cli][INFO ] keyrings          : None
[ceph_deploy.mon][DEBUG ] Deploying mon, cluster ceph hosts ceph-osd1 ceph-osd2 ceph-osd3
[ceph_deploy.mon][DEBUG ] detecting platform for host ceph-osd1 ...
...
[ceph-osd3][DEBUG ] }
[ceph-osd3][DEBUG ]
*****
[ceph-osd3][INFO ] monitor: mon.ceph-osd3 is currently at the state of synchronizing
```

```
[ceph-osd3][INFO ] Running command: ceph --cluster=ceph --admin-daemon /var/run/ceph/ceph-
mon.ceph-osd3.asok mon_status
[ceph-osd3][WARNIN] ceph-osd3 is not defined in `mon initial members`
[ceph-osd3][WARNIN] monitor ceph-osd3 does not exist in monmap
```

```
# firewall-cmd --permanent --add-service={ceph,ceph-mon}
# firewall-cmd --reload
```

### ceph-mon ceph # ceph -s

```
cluster:
  id: 2137b4f3-ea3e-4721-bf88-98bfe6419420
  health: HEALTH_OK

services:
  mon: 4 daemons, quorum ceph-mon,ceph-osd1,ceph-osd2,ceph-osd3 (age 3m)
  mgr: ceph-mon(active, since 3m)
  mds: 3 up:standby
  osd: 3 osds: 3 up (since 3m), 3 in (since 57m)

data:
  pools: 0 pools, 0 pgs
  objects: 0 objects, 0 B
  usage: 3.0 GiB used, 57 GiB / 60 GiB avail
  pgs:
```

→ **Desplegamos ceph-mgr para todos los nodos** → **ceph-osd1 ceph-osd2 ceph-osd3**

### ceph-mon ceph # ceph-deploy mgr create ceph-osd1 ceph-osd2 ceph-osd3

```
[ceph_deploy.conf][DEBUG ] found configuration file at: /root/.cephdeploy.conf
[ceph_deploy.cli][INFO ] Invoked (2.0.1): /usr/bin/ceph-deploy mgr create ceph-osd1 ceph-osd2
ceph-osd3
[ceph_deploy.cli][INFO ] ceph-deploy options:
[ceph_deploy.cli][INFO ] username           : None
[ceph_deploy.cli][INFO ] verbose           : False
[ceph_deploy.cli][INFO ] mgr              : [('ceph-osd1', 'ceph-osd1'), ('ceph-osd2', 'ceph-
osd2'), ('ceph-osd3', 'ceph-osd3')]
[ceph_deploy.cli][INFO ] overwrite_conf   : False
[ceph_deploy.cli][INFO ] subcommand       : create
[ceph_deploy.cli][INFO ] quiet            : False
[ceph_deploy.cli][INFO ] cd_conf          : <ceph_deploy.conf.cephdeploy.Conf instance at
0x7fa971639a70>
[ceph_deploy.cli][INFO ] cluster         : ceph
...
```

```
[ceph-osd3][INFO ] Running command: ceph --cluster ceph --name client.bootstrap-mgr --
keyring /var/lib/ceph/bootstrap-mgr/ceph.keyring auth get-or-create mgr.ceph-osd3 mon allow
```

```
profile mgr osd allow * mds allow * -o /var/lib/ceph/mgr/ceph-ceph-osd3/keyring
[ceph-osd3][INFO ] Running command: systemctl enable ceph-mgr@ceph-osd3
[ceph-osd3][WARNIN] Created symlink from /etc/systemd/system/ceph-mgr.target.wants/ceph-
mgr@ceph-osd3.service to /usr/lib/systemd/system/ceph-mgr@.service.
[ceph-osd3][INFO ] Running command: systemctl start ceph-mgr@ceph-osd3
[ceph-osd3][INFO ] Running command: systemctl enable ceph.target
```

### ceph-mon ceph # ceph -s

```
cluster:
  id: 2137b4f3-ea3e-4721-bf88-98bfe6419420
  health: HEALTH_OK

services:
  mon: 4 daemons, quorum ceph-mon,ceph-osd1,ceph-osd2,ceph-osd3 (age 18m)
  mgr: ceph-mon(active, since 17m), standbys: ceph-osd1, ceph-osd2, ceph-osd3
  mds: 3 up:standby
  osd: 3 osds: 3 up (since 17m), 3 in (since 72m)

data:
  pools: 0 pools, 0 pgs
  objects: 0 objects, 0 B
  usage: 3.0 GiB used, 57 GiB / 60 GiB avail
  pgs:
```

→ Definimos un pool de datos → pool-test

```
ceph-mon ceph # ceph osd pool create pool-test 10
pool 'pool-Replicated' created
```

→ Permitimos el borrado de pool's

```
ceph-mon ~ # ceph tell mon.* injectargs '--mon-allow-pool-delete=true'
injectargs:mon_allow_pool_delete = 'true'
mon.ceph-mon: injectargs:mon_allow_pool_delete = 'true'
mon.ceph-osd1: injectargs:mon_allow_pool_delete = 'true'
mon.ceph-osd2: injectargs:mon_allow_pool_delete = 'true'
mon.ceph-osd3: injectargs:mon_allow_pool_delete = 'true'
```

→ También:

```
ceph-mon ceph # vim /etc/ceph/ceph.conf
...
mon_allow_pool_delete = true
ceph-mon ceph # systemctl restart ceph.target
```

→ Borrarnos como comprobación y creamos de nuevo.

```
ceph-mon ceph # ceph osd pool rm pool-test pool-test --yes-i-really-really-mean-it
```

```
pool 'pool-test' removed
```

```
ceph-mon ceph # ceph osd pool create pool-test 10 10
```

```
pool 'pool-test' created
```

```
ceph-mon ceph # ceph osd tree
```

ID	CLASS	WEIGHT	TYPE	NAME	STATUS	REWEIGHT	PRI	AFF
-1		0.05846	root	default				
-3		0.01949	host	ceph-osd1				
0	hdd	0.01949	osd	osd.0	up	1.00000	1.00000	
-5		0.01949	host	ceph-osd2				
1	hdd	0.01949	osd	osd.1	up	1.00000	1.00000	
-7		0.01949	host	ceph-osd3				
2	hdd	0.01949	osd	osd.2	up	1.00000	1.00000	

```
ceph-mon ceph # ceph osd df
```

ID	CLASS	WEIGHT	REWEIGHT	SIZE	RAW	USE	DATA	OMAP	META	AVAIL	%USE		
0	hdd	0.01949	1.00000	20 GiB	1.0 GiB	6.5 MiB	20 KiB	1024 MiB	19 GiB	5.03	1.00	0	up
1	hdd	0.01949	1.00000	20 GiB	1.0 GiB	6.5 MiB	20 KiB	1024 MiB	19 GiB	5.03	1.00	0	up
2	hdd	0.01949	1.00000	20 GiB	1.0 GiB	6.5 MiB	20 KiB	1024 MiB	19 GiB	5.03	1.00	0	up
				TOTAL	60 GiB	3.0 GiB	20 MiB	60 KiB	3.0 GiB	57 GiB	5.03		

MIN/MAX VAR: 1.00/1.00 STDDEV: 0

```
ceph-mon ceph # ceph osd metadata
```

```
...
{
  "id": 2,
  "arch": "x86_64",
  "back_addr": "[v2:192.168.10.163:6804/1180,v1:192.168.10.163:6805/1180]",
  "back_iface": "eth0",
  "bluefs": "1",
  "bluefs_single_shared_device": "1",
  "bluestore_bdev_access_mode": "blk",
  "bluestore_bdev_block_size": "4096",
  "bluestore_bdev_dev_node": "/dev/dm-5",
  "bluestore_bdev_driver": "KernelDevice",
  "bluestore_bdev_partition_path": "/dev/dm-5",
  "bluestore_bdev_rotational": "1",
  "bluestore_bdev_size": "21470642176",
  "bluestore_bdev_support_discard": "0",
  "bluestore_bdev_type": "hdd",
  "ceph_release": "nautilus",
  "ceph_version": "ceph version 14.2.21 (5ef401921d7a88aea18ec7558f7f9374ebd8f5a6)
nautilus (stable)",
  "ceph_version_short": "14.2.21",
  "cpu": "Intel Core Processor (Haswell, no TSX, IBRS)",
```

```

"default_device_class": "hdd",
"device_ids": "",
"device_paths": "vdb=/dev/disk/by-path/pci-0000:05:00.0",
"devices": "vdb",
"distro": "centos",
"distro_description": "CentOS Linux 7 (Core)",
"distro_version": "7",
"front_addr": "[v2:192.168.10.163:6802/1180,v1:192.168.10.163:6803/1180]",
"front_iface": "eth0",
"hb_back_addr": "[v2:192.168.10.163:6808/1180,v1:192.168.10.163:6809/1180]",
"hb_front_addr": "[v2:192.168.10.163:6806/1180,v1:192.168.10.163:6807/1180]",
"hostname": "ceph-osd3.cadilinea.lan",
"journal_rotational": "1",
"kernel_description": "#1 SMP Wed Apr 28 21:49:45 UTC 2021",
"kernel_version": "3.10.0-1160.25.1.el7.x86_64",
"mem_swap_kb": "2621436",
"mem_total_kb": "2847936",
"network_numa_unknown_ifaces": "eth0",
"objectstore_numa_unknown_devices": "vdb",
"os": "Linux",
"osd_data": "/var/lib/ceph/osd/ceph-2",
"osd_objectstore": "bluestore",
"rotational": "1"
}
]

```

→ **Instalar y crear** → [ceph-radosgw](#)

```

ceph-mon ceph # ceph-deploy install --rgw ceph-mon ceph-osd1 ceph-osd2 ceph-osd3
ceph-mon ceph # ceph-deploy rgw create ceph-mon ceph-osd1 ceph-osd2 ceph-osd3

```

```

...
[ceph-osd3][DEBUG ] connected to host: ceph-osd3
[ceph-osd3][DEBUG ] detect platform information from remote host
[ceph-osd3][DEBUG ] detect machine type
[ceph_deploy.rgw][INFO ] Distro info: CentOS Linux 7.9.2009 Core
[ceph_deploy.rgw][DEBUG ] remote host will use systemd
[ceph_deploy.rgw][DEBUG ] deploying rgw bootstrap to ceph-osd3
[ceph-osd3][DEBUG ] write cluster configuration to /etc/ceph/{cluster}.conf
[ceph-osd3][WARNIN] rgw keyring does not exist yet, creating one
[ceph-osd3][DEBUG ] create a keyring file
[ceph-osd3][DEBUG ] create path recursively if it doesn't exist
[ceph-osd3][INFO ] Running command: ceph --cluster ceph --name client.bootstrap-rgw --
keyring /var/lib/ceph/bootstrap-rgw/ceph.keyring auth get-or-create client.rgw.ceph-osd3 osd allow
rwx mon allow rw -o /var/lib/ceph/radosgw/ceph-rgw.ceph-osd3/keyring
[ceph-osd3][INFO ] Running command: systemctl enable ceph-radosgw@rgw.ceph-osd3
[ceph-osd3][WARNIN] Created symlink from
/etc/systemd/system/ceph-radosgw.target.wants/ceph-radosgw@rgw.ceph-osd3.service to

```

```
/usr/lib/systemd/system/ceph-radosgw@.service.  
[ceph-osd3][INFO ] Running command: systemctl start ceph-radosgw@rgw.ceph-osd3  
[ceph-osd3][INFO ] Running command: systemctl enable ceph.target  
[ceph_deploy.rgw][INFO ] The Ceph Object Gateway (RGW) is now running on host ceph-osd3  
and default port 7480
```

### ceph-mon ceph # ceph -s

cluster:

```
id: 2137b4f3-ea3e-4721-bf88-98bfe6419420  
health: HEALTH_OK
```

services:

```
mon: 4 daemons, quorum ceph-mon,ceph-osd1,ceph-osd2,ceph-osd3 (age 39m)  
mgr: ceph-mon(active, since 39m), standbys: ceph-osd1, ceph-osd3, ceph-osd2  
mds: 3 up:standby  
osd: 3 osds: 3 up (since 39m), 3 in (since 8h)  
rgw: 4 daemons active (ceph-mon, ceph-osd1, ceph-osd2, ceph-osd3)
```

task status:

data:

```
pools: 4 pools, 128 pgs  
objects: 189 objects, 2.3 KiB  
usage: 3.0 GiB used, 57 GiB / 60 GiB avail  
pgs: 128 active+clean
```

```
# firewall-cmd --permanent --add-port=7480/tcp ; firewall-cmd --reload
```

```
→ Creamos un pool → replicated → pool-test
```

```
ceph-mon ceph # ceph osd pool rm pool-test pool-test --yes-i-really-really-mean-it
```

```
pool 'pool-test' removed
```

```
ceph-mon ceph # ceph osd pool create pool-test 10 10 replicated
```

```
pool 'pool-test' created
```

```
ceph-mon ceph # ceph osd pool stats pool-test
```

```
pool pool-test id 7  
nothing is going on
```

### ceph-mon ceph # ceph -s

cluster:

```
id: 2137b4f3-ea3e-4721-bf88-98bfe6419420  
health: HEALTH_WARN  
1 pool(s) have non-power-of-two pg_num  
no active mgr
```



```
services:
  mon: 4 daemons, quorum ceph-mon,ceph-osd1,ceph-osd3 (age 0.854576s), out of quorum: ceph-
osd2
  mgr: no daemons active (since 6s)
  mds: 3 up:standby
  osd: 3 osds: 3 up (since 2h), 3 in (since 9h)
  rgw: 4 daemons active (ceph-mon, ceph-osd1, ceph-osd2, ceph-osd3)
```

task status:

```
data:
  pools: 5 pools, 138 pgs
  objects: 189 objects, 2.3 KiB
  usage: 3.0 GiB used, 57 GiB / 60 GiB avail
  pgs: 138 active+clean
```

```
ceph-mon ceph # ceph config set mon.* mon_warn_on_pool_pg_num_not_power_of_two false
```

```
ceph-mon ceph # vim ceph.conf
```

...

```
mon_warn_on_pool_pg_num_not_power_of_two = false
osd pool default pg num = 1024
osd pool default pgp num = 1024
```

```
ceph-mon ceph # ceph -s
```

```
cluster:
  id: 2137b4f3-ea3e-4721-bf88-98bfe6419420
  health: HEALTH_OK
```

```
services:
  mon: 4 daemons, quorum ceph-mon,ceph-osd1,ceph-osd2,ceph-osd3 (age 4m)
  mgr: ceph-osd1(active, since 5m), standbys: ceph-mon, ceph-osd2, ceph-osd3
  mds: 3 up:standby
  osd: 3 osds: 3 up (since 4m), 3 in (since 9h)
  rgw: 4 daemons active (ceph-mon, ceph-osd1, ceph-osd2, ceph-osd3)
```

task status:

```
data:
  pools: 5 pools, 138 pgs
  objects: 189 objects, 2.3 KiB
  usage: 3.0 GiB used, 57 GiB / 60 GiB avail
  pgs: 138 active+clean
```

```
→ Cargamos el módulo → rbd [todos los nodos]
```

```
# modprobe rbd
```

```
# lsmod | grep rbd
```

```
rbd                83733 0
libceph            306750 1 rbd
```

```
# touch /etc/modules-load.d/rbd
```

→ **Habilitamos la aplicación del pool**

```
ceph-mon ceph # ceph osd pool application enable pool-test rbd
enabled application 'rbd' on pool 'pool-test'
```

→ **Asignación de Discos → disco-01**

```
ceph-mon ceph # rbd create disco-01 --size 1 -p pool-test
ceph-mon ceph # rbd list -p pool-test
disco-01
```

```
ceph-mon ceph # rbd list --pool pool-test -l
NAME    SIZE  PARENT  FMT  PROT  LOCK
disco-01 1 MiB    2
```

→ **Mapeo del disco → disco-01**

```
ceph-mon ceph # rbd map disco-01 -p pool-test
rbd: sysfs write failed
RBD image feature set mismatch. You can disable features unsupported by the kernel with "rbd
feature disable pool-test/disco-01 object-map fast-diff deep-flatten".
In some cases useful info is found in syslog - try "dmesg | tail".
rbd: map failed: (6) No such device or address
```

```
ceph-mon ceph # dmesg | tail
[ 1740.108159] Key type dns_resolver registered
[ 1740.125663] Key type ceph registered
[ 1740.126294] libceph: loaded (mon/osd proto 15/24)
[ 1740.132144] rbd: loaded (major 251)
[ 1740.136877] libceph: mon0 192.168.10.160:6789 session established
[ 1740.137761] libceph: mon0 192.168.10.160:6789 socket closed (con state OPEN)
[ 1740.138645] libceph: mon0 192.168.10.160:6789 session lost, hunting for new mon
[ 1740.142242] libceph: mon1 192.168.10.161:6789 session established
[ 1740.143460] libceph: client234908 fsid 2137b4f3-ea3e-4721-bf88-98bfe6419420
[ 1740.230561] rbd: image disco-01: image uses unsupported features: 0x38
```

→ **Deshabilitamos característica del núcleo no soportada !**

```
ceph-mon ceph # rbd feature disable pool-test/disco-01 object-map fast-diff deep-flatten
```

→ **Mapeamos de nuevo el disco → disco-01**

```
ceph-mon ceph # rbd map disco-01 -p pool-test
/dev/rbd0
```

```
ceph-mon ceph # ls /dev/rbd/pool-test/disco-01
/dev/rbd/pool-test/disco-01
ceph-mon ceph # ls /dev/rbd0
/dev/rbd0
```

→ Configuración actual → `/etc/ceph/ceph.conf`

```
ceph-mon ceph # cat /etc/ceph/ceph.conf
[global]
fsid = 2137b4f3-ea3e-4721-bf88-98bfe6419420
public_network = 192.168.10.0/24
mon_initial_members = ceph-mon
mon_host = 192.168.10.160
auth_cluster_required = cephx
auth_service_required = cephx
auth_client_required = cephx
```

```
mon_allow_pool_delete = true
mon_warn_on_pool_pg_num_not_power_of_two = false
```

```
osd_pool_default_pg_num = 1024
osd_pool_default_pgp_num = 1024
```

→ Creación de un Filesystem|ext4

(Elegimos un FS ext4 porque puede efectuarse una reducción posterior. No sería así si fuese un FS xfs).

```
ceph-mon ceph # mkfs.ext4 /dev/rbd0
mke2fs 1.42.9 (28-Dec-2013)
```

El sistema de ficheros es demasiado pequeño para un fichero de transacciones

Discarding device blocks: hecho

Etiqueta del sistema de ficheros=

OS type: Linux

Tamaño del bloque=1024 (bitácora=0)

Tamaño del fragmento=1024 (bitácora=0)

Stride=4096 blocks, Stripe width=4096 blocks

128 inodes, 1024 blocks

51 blocks (4.98%) reserved for the super user

Primer bloque de datos=1

Número máximo de bloques del sistema de ficheros=1048576

1 bloque de grupo

8192 bloques por grupo, 8192 fragmentos por grupo  
128 nodos-i por grupo

Allocating group tables: hecho

Escribiendo las tablas de nodos-i: hecho

Escribiendo superbloques y la información contable del sistema de ficheros: hecho

```
ceph-mon ceph # mkdir /mnt/rbd-test
```

```
ceph-mon ceph # mount /dev/rbd0 /mnt/rbd-test/
```

```
ceph-mon ceph # df -hT /dev/rbd0
```

```
S.ficheros  Tipo Tamaño Usados  Disp Uso% Montado en
/dev/rbd0  ext4 1003K  21K 911K  3% /mnt/rbd-test
```

```
ceph-mon ceph # touch /mnt/rbd-test/Desde-ceph-mon.txt
```

```
ceph-mon ceph # rbd ls -l -p pool-test
```

```
NAME  SIZE PARENT FMT PROT LOCK
disco-01 1 MiB  2  excl
```

→ **Ampliación del FileSystem** → **ext4**

**\*\* Debe hacerse con el Stma. Montado. \*\***

```
ceph-mon ceph # rbd resize --image disco-01 --size 500M -p pool-test
```

Resizing image: 100% complete...done.

```
ceph-mon ceph # rbd ls -l -p pool-test
```

```
NAME  SIZE PARENT FMT PROT LOCK
disco-01 500 MiB  2
```

```
ceph-mon ceph # ls /mnt/rbd-test/
```

Desde-ceph-mon.txt lost+found

→ **Reducción del FileSystem** → **ext4** (**\*\* No permitido para xfs|FileSystem \*\***).

**\*\* Debe hacerse con el Stma. Montado \*\***

```
ceph-mon ceph # rbd resize --image disco-01 --size 250M --allow-shrink -p pool-test
```

Resizing image: 100% complete...done.

```
ceph-mon ceph # rbd ls -l -p pool-test
```

```
NAME  SIZE PARENT FMT PROT LOCK
disco-01 250 MiB  2
```

```
ceph-mon ceph # ls /mnt/rbd-test/
```

Desde-ceph-mon.txt lost+found

→ **Renombrar imagen RBD** → **disco-01 ==> rbd\_disco-01**

```
ceph-mon ceph # rbd mv disco-01 rbd_disco-01 -p pool-test
```

```
ceph-mon ceph # rbd ls -l -p pool-test
NAME      SIZE  PARENT FMT PROT LOCK
rbd_disco-01 250 MiB      2
ceph-mon ceph # ls /mnt/rbd-test/
Desde-ceph-mon.txt  lost+found
```

→ **Habilitar** → **dashboard** → **ceph-mon** → **8443/tcp** | **grafana** → **3000/tcp**

```
ceph-mon ceph # firewall-cmd --permanent --add-port=8443/tcp
ceph-mon ceph # firewall-cmd --permanent --add-port=3000/tcp
ceph-mon ceph # firewall-cmd --reload
```

```
ceph-mon ceph # yum install ceph-mgr python-jwt python-routes
```

```
ceph-mon ceph # vim /etc/yum.repos.d/grafana.repo
[grafana]
name=grafana
baseurl=https://packages.grafana.com/oss/rpm
repo_gpgcheck=1
enabled=1
gpgcheck=1
gpgkey=https://packages.grafana.com/gpg.key
sslverify=1
sslcacert=/etc/pki/tls/certs/ca-bundle.crt
```

```
ceph-mon ceph # yum install grafana
ceph-mon ceph # yum install fontconfig freetype* urw-fonts
```

```
ceph-mon ceph # yum install https://download.ceph.com/rpm-nautilus/el7/noarch/ceph-grafana-dashboards-14.2.21-0.17.noarch.rpm
```

```
ceph-mon ceph # yum install https://download.ceph.com/rpm-nautilus/el7/noarch/ceph-mgr-dashboard-14.2.21-0.el7.narch.rpm
```

```
ceph-mon ceph # systemctl enable --now grafana-server
```

```
ceph-mon ceph # ceph mgr module enable dashboard --force
ceph-mon ceph # ceph mgr module ls | less
```

```
...
  "enabled_modules": [
    "dashboard",
    "iostat",
    "restful"
  ],
...
```

```
ceph-mon ceph # ceph dashboard create-self-signed-cert
```

Self-signed certificate created

```
ceph-mon ceph # ceph mgr services
{
  "dashboard": "https://ceph-mon.cadilinea.lan:8443/"
}
```

→ Creamos usuario/password de acceso → administrator

```
ceph-mon ceph # vim /etc/ceph/secret.admin
123456
```

```
ceph-mon ceph # ceph dashboard ac-user-create carlos -i /etc/ceph/secret.admin administrator
{"username": "carlos", "lastUpdate": 1622491815, "name": null, "roles": ["administrator"],
"password":
"$2b$12$IyO7SEr8XZqHzLhJODjLf.U1WMVXhDPCL4of9UkWzHXz0PJYVCvVG", "email":
null}
```

```
ceph-mon ceph # ceph dashboard ac-user-show
["carlos"]
```

→ Acceso Entorno Gráfico → <https://ceph-mon.cadilinea.lan:8443/> → carlos:123456

Mostramos varias imágenes !.

→ Activación de Meta Datos → ceph-mds → ceph-osd1 ceph-osd2 ceph-osd3  
==> Necesario para almacenar objetos → cephfs.

```
ceph-mon ~ # ceph -s
cluster:
  id: 2137b4f3-ea3e-4721-bf88-98bfe6419420
  health: HEALTH_OK

services:
  mon: 4 daemons, quorum ceph-mon,ceph-osd1,ceph-osd2,ceph-osd3 (age 28m)
  mgr: ceph-osd1(active, since 28m), standbys: ceph-mon, ceph-osd3, ceph-osd2
  mds: 3 up:standby
  osd: 3 osds: 3 up (since 27m), 3 in (since 14h)
  rgw: 4 daemons active (ceph-mon, ceph-osd1, ceph-osd2, ceph-osd3)

task status:

data:
  pools: 5 pools, 138 pgs
  objects: 226 objects, 1.0 MiB
  usage: 3.0 GiB used, 57 GiB / 60 GiB avail
```

pgs: 138 active+clean

```
ceph-mon ceph # ceph-deploy --overwrite-conf mds create ceph-osd1:ceph-mds1
```

...

```
[ceph-osd1][INFO ] Running command: systemctl enable ceph-mds@ceph-mds1
```

```
[ceph-osd1][INFO ] Running command: systemctl start ceph-mds@ceph-mds1
```

```
[ceph-osd1][INFO ] Running command: systemctl enable ceph.target
```

```
ceph-mon ceph # ceph-deploy --overwrite-conf mds create ceph-osd2:ceph-mds2
```

...

```
[ceph-osd2][INFO ] Running command: systemctl enable ceph-mds@ceph-mds2
```

```
[ceph-osd2][INFO ] Running command: systemctl start ceph-mds@ceph-mds2
```

```
[ceph-osd2][INFO ] Running command: systemctl enable ceph.target
```

```
ceph-mon ceph # ceph-deploy --overwrite-conf mds create ceph-osd3:ceph-mds3
```

...

```
[ceph-osd3][INFO ] Running command: systemctl enable ceph-mds@ceph-mds3
```

```
[ceph-osd3][INFO ] Running command: systemctl start ceph-mds@ceph-mds3
```

```
[ceph-osd3][INFO ] Running command: systemctl enable ceph.target
```

```
ceph-mon ceph # ceph -s
```

cluster:

id: 2137b4f3-ea3e-4721-bf88-98bfe6419420

health: HEALTH\_OK

services:

mon: 4 daemons, quorum ceph-mon,ceph-osd1,ceph-osd2,ceph-osd3 (age 53m)

mgr: ceph-osd1(active, since 53m), standbys: ceph-mon, ceph-osd3, ceph-osd2

**mds: 6 up:standby**

osd: 3 osds: 3 up (since 53m), 3 in (since 14h)

rgw: 4 daemons active (ceph-mon, ceph-osd1, ceph-osd2, ceph-osd3)

task status:

data:

pools: 5 pools, 138 pgs

objects: 226 objects, 1.0 MiB

usage: 3.0 GiB used, 57 GiB / 60 GiB avail

pgs: 138 active+clean

→ Creamos pool's de almacenamiento ==> cephfs-metadatos → cephfs-datos

```
ceph-mon ceph # ceph osd pool create cephfs-metadatos 10 10
```

```
pool 'cephfs-metadatos' created
```

```
ceph-mon ceph # ceph osd pool create cephfs-datos 10 10
```

```
pool 'cephfs-datos' created
```

**ceph-mon ceph # ceph osd pool stats**

pool .rgw.root id 3  
nothing is going on

pool default.rgw.control id 4  
nothing is going on

pool default.rgw.meta id 5  
nothing is going on

pool default.rgw.log id 6  
nothing is going on

pool pool-test id 9  
nothing is going on

pool cephfs-metadatos id 12  
nothing is going on

pool cephfs-datos id 13  
nothing is going on

**ceph-mon ceph # ceph osd pool ls**

.rgw.root  
default.rgw.control  
default.rgw.meta  
default.rgw.log  
pool-test  
cephfs-metadatos  
cephfs-datos

→ **Creamos** → **FileSystem|cephfs** → **cephfs-metadatos** → **cephfs-datos**

**ceph-mon ceph # ceph fs ls**

No filesystems enabled

**ceph-mon ceph # ceph fs new cephfs-01 cephfs-metadatos cephfs-datos**

new fs with metadata pool 12 and data pool 13

**ceph-mon ceph # ceph fs ls**

**name: cephfs-01, metadata pool: cephfs-metadatos, data pools: [cephfs-datos ]**

**ceph-mon ceph # ceph -s**

cluster:  
id: 2137b4f3-ea3e-4721-bf88-98bfe6419420  
health: HEALTH\_OK

services:



```

mon: 4 daemons, quorum ceph-mon,ceph-osd1,ceph-osd2,ceph-osd3 (age 50m)
mgr: ceph-mon(active, since 59m), standbys: ceph-osd1, ceph-osd3, ceph-osd2
mds: cephfs-01:1 {0=ceph-osd1=up:active} 5 up:standby
osd: 3 osds: 3 up (since 58m), 3 in (since 24h)
rgw: 4 daemons active (ceph-mon, ceph-osd1, ceph-osd2, ceph-osd3)

```

task status:

scrub status:

mds.ceph-osd1: idle

data:

pools: 7 pools, 158 pgs

objects: 248 objects, 1.0 MiB

usage: 3.1 GiB used, 57 GiB / 60 GiB avail

pgs: 158 active+clean

→ **Copiamos la clave de autenticación para poder montar el FileSystem** →  
**ceph.client.admin.keyring ==> ceph.client.key**

```
ceph-mon ceph # cat ceph.client.admin.keyring
```

```
[client.admin]
```

```
key = AQD/z6hgxsd8GRAAmq/f7TS6XtFq6lC8AC6DpA==
```

```
caps mds = "allow *"
```

```
caps mgr = "allow *"
```

```
caps mon = "allow *"
```

```
caps osd = "allow *"
```

```
ceph-mon ceph # ceph-authtool --print-key /etc/ceph/ceph.client.admin.keyring >
```

```
ceph.client.key
```

```
ceph-mon ceph # cat /etc/ceph/ceph.client.key
```

```
AQD/z6hgxsd8GRAAmq/f7TS6XtFq6lC8AC6DpA==
```

```
ceph-mon ceph # chmod 0400 /etc/ceph/ceph.client.key
```

```
ceph-mon ceph # ceph auth get client.admin
```

```
exported keyring for client.admin
```

```
[client.admin]
```

```
key = AQD/z6hgxsd8GRAAmq/f7TS6XtFq6lC8AC6DpA==
```

```
caps mds = "allow *"
```

```
caps mgr = "allow *"
```

```
caps mon = "allow *"
```

```
caps osd = "allow *"
```

```
ceph-mon ceph # mkdir /mnt/mi_cephfs-mon
```

```
ceph-mon ceph # mkdir /mnt/mi_cephfs-osd1
```

```
ceph-mon ceph # mkdir /mnt/mi_cephfs-osd2
```

```
ceph-mon ceph # mkdir /mnt/mi_cephfs-osd3
```

```
ceph-mon ceph # ssh ceph-osd1 mkdir /mnt/mi_cephfs-osd1
ceph-mon ceph # ssh ceph-osd2 mkdir /mnt/mi_cephfs-osd2
ceph-mon ceph # ssh ceph-osd3 mkdir /mnt/mi_cephfs-osd3
```

```
ceph-mon ceph # scp /etc/ceph/ceph.client.key ceph-osd1:/etc/ceph
ceph-mon ceph # scp /etc/ceph/ceph.client.key ceph-osd2:/etc/ceph
ceph-mon ceph # scp /etc/ceph/ceph.client.key ceph-osd3:/etc/ceph
```

```
ceph-mon ceph # yum install ceph-fuse
ceph-mon ceph # ssh ceph-osd1 yum install ceph-fuse -y
ceph-mon ceph # ssh ceph-osd2 yum install ceph-fuse -y
ceph-mon ceph # ssh ceph-osd3 yum install ceph-fuse -y
```

```
ceph-mon ceph # mount -t ceph ceph-mon:6789:/ /mnt/mi_cephfs-mon/ -o
name=admin,secretfile=/etc/ceph/ceph.client.key -vvv
parsing options: rw,name=admin,secretfile=/etc/ceph/ceph.client.key
```

```
ceph-mon ceph # df -hT
```

S.ficheros	Tipo	Tamaño	Usados	Disp	Uso%	Montado en
devtmpfs	devtmpfs	1,4G	0	1,4G	0%	/dev
tmpfs	tmpfs	1,4G	0	1,4G	0%	/dev/shm
tmpfs	tmpfs	1,4G	8,7M	1,4G	1%	/run
tmpfs	tmpfs	1,4G	0	1,4G	0%	/sys/fs/cgroup
/dev/mapper/centos-root	xfs	17G	2,3G	15G	14%	/
/dev/vda1	xfs	1014M	199M	816M	20%	/boot
tmpfs	tmpfs	279M	0	279M	0%	/run/user/0
192.168.10.160:6789:/	ceph	18G	0	18G	0%	/mnt/mi_cephfs-mon

→ Creamos un testigo en **ceph-mon** y **Desmontamos** para probar, y en otro nodo.

```
ceph-mon ceph # touch /mnt/mi_cephfs-mon/Desde-ceph-mon.txt
ceph-mon ceph # umount /mnt/mi_cephfs-mon
```

```
ceph-mon ceph # mount -t ceph ceph-osd1:6789:/ /mnt/mi_cephfs-osd1/ -o
name=admin,secretfile=/etc/ceph/ceph.client.key -vvv
parsing options: rw,name=admin,secretfile=/etc/ceph/ceph.client.key
```

```
ceph-mon ceph # df -hT
```

S.ficheros	Tipo	Tamaño	Usados	Disp	Uso%	Montado en
devtmpfs	devtmpfs	1,4G	0	1,4G	0%	/dev
tmpfs	tmpfs	1,4G	0	1,4G	0%	/dev/shm
tmpfs	tmpfs	1,4G	8,7M	1,4G	1%	/run
tmpfs	tmpfs	1,4G	0	1,4G	0%	/sys/fs/cgroup
/dev/mapper/centos-root	xfs	17G	2,3G	15G	14%	/
/dev/vda1	xfs	1014M	199M	816M	20%	/boot
tmpfs	tmpfs	279M	0	279M	0%	/run/user/0
192.168.10.161:6789:/	ceph	18G	0	18G	0%	/mnt/mi_cephfs-osd1

```
ceph-mon ceph # ls /mnt/mi_cephfs-mon/  
Desde-ceph-mon.txt
```

→ **Desmontamos de nuevo !**

```
ceph-mon ceph # umount /mnt/mi_cephfs-osd1
```

```
ceph-mon ceph # ceph fs status
```

```
cephfs-01 - 0 clients
```

```
=====
```

```
+-----+-----+-----+-----+-----+-----+
```

```
| Rank | State | MDS | Activity | dns | inos |
```

```
+-----+-----+-----+-----+-----+-----+
```

```
| 0 | active | ceph-osd1 | Reqs: 0/s | 11 | 14 |
```

```
+-----+-----+-----+-----+-----+-----+
```

```
+-----+-----+-----+-----+
```

```
| Pool | type | used | avail |
```

```
+-----+-----+-----+-----+
```

```
| cephfs-metadatos | metadata | 1536k | 17.9G |
```

```
| cephfs-datos | data | 0 | 17.9G |
```

```
+-----+-----+-----+-----+
```

```
+-----+
```

```
| Standby MDS |
```

```
+-----+
```

```
| ceph-osd2 |
```

```
| ceph-osd3 |
```

```
| ceph-mds1 |
```

```
| ceph-mds2 |
```

```
| ceph-mds3 |
```

```
+-----+
```

```
MDS version: ceph version 14.2.21 (5ef401921d7a88aea18ec7558f7f9374ebd8f5a6) nautilus  
(stable)
```

```
ceph-mon ceph # ceph fs get cephfs-01
```

```
Filesystem 'cephfs-01' (1)
```

```
fs_name cephfs-01
```

```
epoch 36
```

```
flags 12
```

```
created 2021-06-01 09:59:51.077794
```

```
modified 2021-06-01 09:59:52.677270
```

```
tableserver 0
```

```
root 0
```

```
session_timeout 60
```

```
session_autoclose 300
```

```
max_file_size 1099511627776
```

```
min_compat_client -1 (unspecified)
```

```

last_failure    0
last_failure_osd_epoch    0
compatcompat={},rocompat={},incompat={1=base v0.20,2=client writeable ranges,3=default file
layouts on dirs,4=dir inode in separate object,5=mds uses versioned encoding,6=dirfrag is stored in
omap,8=no anchor table,9=file layout v2,10=snaprealm v2}
max_mds        1
in             0
up             {0=264172}
failed
damaged
stopped
data_pools     [13]
metadata_pool12
inline_data    disabled
balancer
standby_count_wanted    1
[mds.ceph-osd1{0:264172} state up:active seq 643 addr
[v2:192.168.10.161:6800/845818693,v1:192.168.10.161:6801/845818693]]

```

→ **Volvemos a montar y de forma global** → `ceph-osd{1,2,3}`

```

ceph-mon ceph # mount -t ceph ceph-osd1:6789:/ /mnt/mi_cephfs-osd1/ -o
name=admin,secretfile=/etc/ceph/ceph.client.key -vvv
parsing options: rw,name=admin,secretfile=/etc/ceph/ceph.client.key

```

```

ceph-mon ceph # mount -t ceph ceph-osd2:6789:/ /mnt/mi_cephfs-osd2/ -o
name=admin,secretfile=/etc/ceph/ceph.client.key -vvv
parsing options: rw,name=admin,secretfile=/etc/ceph/ceph.client.key

```

```

ceph-mon ceph # mount -t ceph ceph-osd3:6789:/ /mnt/mi_cephfs-osd3/ -o
name=admin,secretfile=/etc/ceph/ceph.client.key -vvv
parsing options: rw,name=admin,secretfile=/etc/ceph/ceph.client.key

```

```

ceph-mon ceph # df -hT

```

```

S.ficheros      Tipo  Tamaño Usados  Disp  Uso%  Montado en
devtmpfs        devtmpfs  1,4G    0  1,4G  0% /dev
tmpfs           tmpfs    1,4G    0  1,4G  0% /dev/shm
tmpfs           tmpfs    1,4G   8,7M  1,4G  1% /run
tmpfs           tmpfs    1,4G    0  1,4G  0% /sys/fs/cgroup
/dev/mapper/centos-root xfs      17G   2,3G  15G  14% /
/dev/vda1       xfs     1014M  199M  816M  20% /boot
tmpfs           tmpfs    279M    0  279M  0% /run/user/0
192.168.10.161:6789:/ ceph    18G    0  18G  0% /mnt/mi_cephfs-osd1
192.168.10.162:6789:/ ceph    18G    0  18G  0% /mnt/mi_cephfs-osd2
192.168.10.163:6789:/ ceph    18G    0  18G  0% /mnt/mi_cephfs-osd3

```

```

ceph-mon ceph # umount /mnt/mi_cephfs-osd*

```

→ Montaje con → `ceph-fuse`

`ceph-mon ceph # mkdir /mnt/fuse`

→ Crear otra estructura LVM → `ceph-osd1` → `/dev/vdc`

Creamos un disco de reserva → `/dev/vdc`

`ceph-osd1 ~ # ceph-volume inventory`

Device Path	Size	rotates	available	Model name
<code>/dev/vdc</code>	20.00 GB	True	True	
<code>/dev/vdd</code>	20.00 GB	True	True	
<code>/dev/vda</code>	25.00 GB	True	False	
<code>/dev/vdb</code>	20.00 GB	True	False	

`ceph-osd1 ~ # lsblk /dev/vd[b-z]`

NAME	TYPE	MOUNTPOINT	MAJ:MIN	RM	SIZE	RO
<code>vdb</code>			252:16	0	20G	0 disk
<code>└─ceph--4744afc0--6372--44b5--a74e--7bd0d30d5ad7-osd--block--8acc47ed--3737--4f21--8f80--922a176b9c28</code>			253:6	0	20G	0 lvm
<code>vdc</code>			252:32	0	20G	0 disk
<code>vdd</code>			252:48	0	20G	0 disk

`ceph-osd1 ~ # ceph-volume lvm create --data /dev/vdc`

Running command: `/usr/bin/ceph-authtool --gen-print-key`

Running command: `/usr/bin/ceph --cluster ceph --name client.bootstrap-osd --keyring /var/lib/ceph/bootstrap-osd/ceph.keyring -i - osd new 59c9fb12-d0b9-4588-b5d7-423bdd7713a5`

Running command: `/usr/sbin/vgcreate --force --yes ceph-d91c5050-eb48-4d55-a103-2b6d2f5b00d3 /dev/vdc`

stdout: Physical volume `"/dev/vdc"` successfully created.

stdout: Volume group `"ceph-d91c5050-eb48-4d55-a103-2b6d2f5b00d3"` successfully created

Running command: `/usr/sbin/lvcreate --yes -l 5119 -n osd-block-59c9fb12-d0b9-4588-b5d7-423bdd7713a5 ceph-d91c5050-eb48-4d55-a103-2b6d2f5b00d3`

stdout: Logical volume `"osd-block-59c9fb12-d0b9-4588-b5d7-423bdd7713a5"` created.

...

Running command: `/usr/bin/systemctl enable ceph-volume@lvm-3-59c9fb12-d0b9-4588-b5d7-423bdd7713a5`

stderr: Created symlink from `/etc/systemd/system/multi-user.target.wants/ceph-volume@lvm-3-59c9fb12-d0b9-4588-b5d7-423bdd7713a5.service` to `/usr/lib/systemd/system/ceph-volume@.service`.

Running command: `/usr/bin/systemctl enable --runtime ceph-osd@3`

stderr: Created symlink from `/run/systemd/system/ceph-osd.target.wants/ceph-osd@3.service` to `/usr/lib/systemd/system/ceph-osd@.service`.

Running command: `/usr/bin/systemctl start ceph-osd@3`

--> `ceph-volume lvm activate successful for osd ID: 3`

--> ceph-volume lvm create successful for: /dev/vdc

### ceph-osd1 ~ # ceph-volume inventory

Device Path	Size	rotates	available	Model name
/dev/vdd	20.00 GB	True	True	
/dev/vda	25.00 GB	True	False	
/dev/vdb	20.00 GB	True	False	
/dev/vdc	20.00 GB	True	False	

### ceph-osd1 ~ # ceph osd tree

ID	CLASS	WEIGHT	TYPE	NAME	STATUS	REWEIGHT	PRI	AFF
-1		0.07794	root	default				
-3		0.03897	host	ceph-osd1				
0	hdd	0.01949	osd.0	up	1.00000	1.00000		
3	<b>hdd</b>	<b>0.01949</b>	<b>osd.3</b>	<b>up</b>	<b>1.00000</b>	<b>1.00000</b>		
-5		0.01949	host	ceph-osd2				
1	hdd	0.01949	osd.1	up	1.00000	1.00000		
-7		0.01949	host	ceph-osd3				
2	hdd	0.01949	osd.2	up	1.00000	1.00000		

→ **Sustitución de disco averiado en ceph-osd1** → /dev/vdb ==> /dev/vdc

### ceph-mon ceph # ceph -s

cluster:

id: 2137b4f3-ea3e-4721-bf88-98bfe6419420

health: HEALTH\_WARN

1 osds down

Degraded data redundancy: 80/744 objects degraded (10.753%), 20 pgs degraded, 70 pgs undersized

services:

mon: 4 daemons, quorum ceph-mon,ceph-osd1,ceph-osd2,ceph-osd3 (age 3m)

mgr: ceph-osd3(active, since 95m), standbys: ceph-mon, ceph-osd2, ceph-osd1

mds: cephfs-01:1 {0=ceph-mds2=up:active} 5 up:standby

osd: 4 osds: 3 up (since 3m), 4 in (since 23m)

rgw: 4 daemons active (ceph-mon, ceph-osd1, ceph-osd2, ceph-osd3)

task status:

scrub status:

mds.ceph-mds2: idle

data:

pools: 7 pools, 158 pgs

objects: 248 objects, 1.0 MiB

usage: 4.1 GiB used, 76 GiB / 80 GiB avail

pgs: 80/744 objects degraded (10.753%)

88 active+clean

50 active+undersized  
 20 active+undersized+degraded

**ceph-mon ceph # ceph osd tree**

ID	CLASS	WEIGHT	TYPE	NAME	STATUS	REWEIGHT	PRI-AFF
-1		0.07794	root	default			
-3		0.03897	host	ceph-osd1			
0	hdd	0.01949	osd.0		down	1.00000	1.00000
3	hdd	0.01949	osd.3		up	1.00000	1.00000
-5		0.01949	host	ceph-osd2			
1	hdd	0.01949	osd.1		up	1.00000	1.00000
-7		0.01949	host	ceph-osd3			
2	hdd	0.01949	osd.2		up	1.00000	1.00000

**ceph-mon ceph # ceph osd unset noout**

noout is unset

**ceph-mon ceph # ceph osd crush reweight osd.0 0**

reweighted item id 0 name 'osd.0' to 0 in crush map

**ceph-mon ceph # ceph osd tree**

ID	CLASS	WEIGHT	TYPE	NAME	STATUS	REWEIGHT	PRI-AFF
-1		0.05846	root	default			
-3		0.01949	host	ceph-osd1			
0	hdd	0	osd.0		down	1.00000	1.00000
3	hdd	0.01949	osd.3		up	1.00000	1.00000
-5		0.01949	host	ceph-osd2			
1	hdd	0.01949	osd.1		up	1.00000	1.00000
-7		0.01949	host	ceph-osd3			
2	hdd	0.01949	osd.2		up	1.00000	1.00000

**ceph-mon ceph # ceph osd out osd.0 0**

osd.0 is already out. osd.0 is already out.

**ceph-mon ceph # ceph osd crush remove osd.0**

removed item id 0 name 'osd.0' from crush map

**ceph-mon ceph # ceph osd tree**

ID	CLASS	WEIGHT	TYPE	NAME	STATUS	REWEIGHT	PRI-AFF
-1		0.05846	root	default			
-3		0.01949	host	ceph-osd1			
3	hdd	0.01949	osd.3		up	1.00000	1.00000
-5		0.01949	host	ceph-osd2			
1	hdd	0.01949	osd.1		up	1.00000	1.00000
-7		0.01949	host	ceph-osd3			
2	hdd	0.01949	osd.2		up	1.00000	1.00000
0		0	osd.0		down	0	1.00000

**ceph-mon ceph # ceph -s**

cluster:

id: 2137b4f3-ea3e-4721-bf88-98bfe6419420

health: HEALTH\_OK

services:

mon: 4 daemons, quorum ceph-mon,ceph-osd1,ceph-osd2,ceph-osd3 (age 14m)

mgr: ceph-osd3(active, since 105m), standbys: ceph-mon, ceph-osd2, ceph-osd1

mds: cephfs-01:1 {0=ceph-mds2=up:active} 5 up:standby

osd: 4 osds: 3 up (since 14m), 3 in (since 4m)

rgw: 4 daemons active (ceph-mon, ceph-osd1, ceph-osd2, ceph-osd3)

task status:

scrub status:

mds.ceph-mds2: idle

data:

pools: 7 pools, 158 pgs

objects: 248 objects, 1.0 MiB

usage: 3.1 GiB used, 57 GiB / 60 GiB avail

pgs: 158 active+clean

**ceph-mon ceph # ceph auth del osd.0**

updated

**ceph-mon ceph # ceph osd rm osd.0**

removed osd.0

**ceph-mon ceph # ceph osd tree**

ID	CLASS	WEIGHT	TYPE	NAME	STATUS	REWEIGHT	PRI	AFF
-1		0.05846	root	default				
-3		0.01949	host	ceph-osd1				
3	hdd	0.01949	osd.3		up	1.00000	1.00000	
-5		0.01949	host	ceph-osd2				
1	hdd	0.01949	osd.1		up	1.00000	1.00000	
-7		0.01949	host	ceph-osd3				
2	hdd	0.01949	osd.2		up	1.00000	1.00000	

**ceph-mon ceph # ceph osd df**

ID	CLASS	WEIGHT	REWEIGHT	SIZE	RAW	USE	DATA	OMAP	META	AVAIL	%USE
3	hdd	0.01949	1.00000	20 GiB	1.0 GiB	33 MiB	23 KiB	1024 MiB	19 GiB	5.16	1.00 158 up
1	hdd	0.01949	1.00000	20 GiB	1.0 GiB	33 MiB	11 KiB	1024 MiB	19 GiB	5.16	1.00 158 up
2	hdd	0.01949	1.00000	20 GiB	1.0 GiB	33 MiB	11 KiB	1024 MiB	19 GiB	5.16	1.00 158 up
				TOTAL	60 GiB	3.1 GiB	99 MiB	47 KiB	3.0 GiB	57 GiB	5.16

MIN/MAX VAR: 1.00/1.00 STDDEV: 0



**→ Configuración Actual.**

```
ceph-mon ceph # grep -v "#" /etc/ceph/ceph.conf | sed -e '/^$/d' | grep -v ';'
[global]
fsid = 2137b4f3-ea3e-4721-bf88-98bfe6419420
public_network = 192.168.10.0/24
mon_initial_members = ceph-mon
mon_host = 192.168.10.160
auth_cluster_required = cephx
auth_service_required = cephx
auth_client_required = cephx
mon_allow_pool_delete = true
mon_warn_on_pool_pg_num_not_power_of_two = false
osd pool default pg num = 1024
osd pool default pgp num = 1024
```

**BIBLIOGRAFIA:**

[https://www.server-world.info/en/note?os=CentOS\\_8&p=ceph15&f=1](https://www.server-world.info/en/note?os=CentOS_8&p=ceph15&f=1)  
<https://ralph.blog.imixs.com/2020/02/28/howto-install-ceph-on-centos-7/>  
<https://www.howtoforge.com/tutorial/how-to-build-a-ceph-cluster-on-centos-7/>  
<https://dokk.org/documentation/ceph/v14.0.1/mgr/dashboard/>  
[https://access.redhat.com/documentation/en-us/red\\_hat\\_ceph\\_storage/2/html/ceph\\_file\\_system\\_guide\\_technology\\_preview/what\\_is\\_the\\_ceph\\_file\\_system\\_cephfs](https://access.redhat.com/documentation/en-us/red_hat_ceph_storage/2/html/ceph_file_system_guide_technology_preview/what_is_the_ceph_file_system_cephfs)

**Creative Commons****Reconocimiento-NoComercial-CompartirIgual 3.1 ESPAÑA**

© 2021 by carlos briso. Usted es libre de copiar, distribuir y comunicar públicamente la obra y hacer obras derivadas bajo las condiciones siguientes:

- a) Debe reconocer y citar al autor original.
  - b) No puede utilizar esta obra para fines comerciales (incluyendo su publicación, a través de cualquier medio, por entidades con fines de lucro.
  - c) Si altera o transforma esta obra o genera una obra derivada, sólo puede distribuir la obra generada bajo una licencia idéntica a ésta. Al reutilizar o distribuir la obra, tiene que dejar bien claro los términos de la licencia de esta obra.
- Alguna de estas condiciones puede no aplicarse si se obtiene el permiso del titular de los derechos de autor. Los derechos derivados de usos legítimos u otras limitaciones no se ven afectados por lo anterior. Licencia completa en castellano.

→ La información contenida en este documento y los derivados de éste se proporcionan tal cual son y los autores no asumirán responsabilidad alguna si el usuario o lector hace mal uso de éstos.