

CEPH pg,object repair - ceph status 3

MANUALLY REPAIR OBJECT PG Repair 3

..... 3

PG 3

..... 3

object 4

..... 4

ceph -s pg 4

가 PG 5

PG repair 5

PLACEMENT GROUP STATES 5

..... 6

CEPH pg,object repair status

- ceph

MANUALLY REPAIR OBJECT PG Repare

```
health HEALTH_ERR 1 pgs inconsistent; 2 scrub errors
```

가 .

PG

```
$ sudo ceph health detail
```

```
HEALTH_ERR 1 pgs inconsistent; 2 scrub errors
pg 17.1c1 is active+clean+inconsistent, acting [21,25,30] 2 scrub errors
```

```
PG 17.1c1 OSD 21,25,30 .
```

, OSD logs .

```
$ grep -Hn 'ERR' /var/logs/ceph/ceph-osd.21.log
```

가 rotate grep zgrep .

```
log [ERR] : 17.1c1 shard 21: soid
58bcc1c1/rb.0.90213.238e1f29.00000001232d/head//17 digest 0 != known digest
3062795895
log [ERR] : 17.1c1 shard 25: soid
58bcc1c1/rb.0.90213.238e1f29.00000001232d/head//17 digest 0 != known digest
3062795895
```

digest가 3062795895 digest 0 .

object

1,2,3 가 .

1. PG: 17.1c1 (PG)
2. OSD ID (PG가 OSD)
3. rb.0.90213.238e1f29.00000001232d (Object name)

Object .

```
$ sudo find /var/lib/ceph/osd/ceph-21/current/17.1c1_head/ -name 'rb.0.90213.238e1f29.00000001232d*' -ls
```

```
671193536 4096 -rw-r-- 1 root root 4194304 Feb 14 01:05 /var/lib/ceph/osd/ceph-21/current/17.1c1_head/DIR_1/DIR_C/DIR_1/DIR_C/rb.0.90213.238e1f29.00000001232d__head_58BCC1C1__1
```

object size

MD5(rb.0.90213.238e1f29.00000001232d__head_58BCC1C1__1) .

1. PG OSD stop
2. (ceph-osd -i --flush-journal)
3. bad object .
4. OSD start

```
$ ceph pg repair 17.1c1
```

ceph -s pg

```
[root@storage004 ~] ceph -s
```

```
health HEALTH_ERR
5 pgs inconsistent
9 scrub errors
```

- active: Request
- clean: 가 Replica Replication
- inconsistent: Replica 가 가 가 . 가
scrub errors 가 .

가 PG

```
[root@storage004 ~] ceph health detail
HEALTH_ERR 5 pgs inconsistent; 9 scrub errors
pg 9.180 is active+clean+inconsistent, acting [43,36,18] # pg 9.180
pg 9.e1 is active+clean+inconsistent, acting [20,41,23] # pg 9.e1
pg 2.12 is active+clean+inconsistent, acting [24,37,15] # pg 2.12
pg 9.3e3 is active+clean+inconsistent, acting [32,26,7] # pg 9.3e3
pg 9.265 is active+clean+inconsistent, acting [24,10,3] # pg 9.265
9 scrub errors
```

PG repair

```
[root@storage004 ~] ceph pg repair 9.180
instructing pg 9.180 on osd.43 to repair

[root@storage004 ~] ceph pg repair 9.e1
instructing pg 9.e1 on osd.20 to repair

[root@storage004 ~] ceph pg repair 2.12
instructing pg 2.12 on osd.24 to repair

[root@storage004 ~] ceph pg repair 9.3e3
instructing pg 9.3e3 on osd.32 to repair

[root@storage004 ~] ceph pg repair 9.265
instructing pg 9.265 on osd.24 to repair

# repair
[root@storage004 ~] ceph health detail
HEALTH_OK
```

PLACEMENT GROUP STATES

When checking a cluster's status (e.g., running `ceph -w` or `ceph -s`), Ceph will report on the status of the placement groups. A placement group has one or more states. The optimum state for placement groups in the placement group map is `active + clean`.

- Creating - Ceph가
- Active - Ceph가
- Clean - Ceph가
- Down - Ceph가 가
- Replay - Ceph가 OSD가 가
- Splitting - Ceph가

- Scrubbing - Ceph가
- Degraded - Ceph
- Inconsistent - Ceph
 (: 가)
- Peering -
- Repair - Ceph가 . (if possible).
- Recovering - Ceph /
- Backfill - Ceph . Backfill
- Wait-backfill - Backfill
- Backfill-too-full - OSD가
- Incomplete - ceph
 가 OSD ,
 minb_size
- Stale -
 가
- Remapped - CRUSH가 OSD
- Undersized -
- Peered - min_size
 IO 가
 pg가 min_size

- <https://idchowto.com/?p=31090>
- <https://access.redhat.com/solutions/1589113>

From: <https://atli.kr/dokuwiki/> - AllThatLinux!
Permanent link: https://atli.kr/dokuwiki/doku.php/ceph_pg_object_repair_%EB%8B%A8%EA%B3%84%EB%B3%84_%EC%A0%95%EB%A6%AC_ceph_status_%EB%A9%94%EC%8B%9C%EC%A7%80_%EC%84%A4%EB%AA%85
Last update: 2021/07/19 12:12

