

知 UIS存储集群健康度告警

刘子溪 2020-08-26 发表

组网及说明

uis0715

问题描述

uis前台页面看见告警

名称	存储集群健康度告警
描述	存储集群出现故障。
级别	重要
类型	分布式存储资源告警
告警状态	已确认
告警来源	分布式存储"分布式存储_192.168.220.1"
最新告警时间	2020-07-27 21:35:29
首次告警时间	2020-07-27 21:35:29
重复次数	1
告警原因	存储集群健康度告警，可能原因：1、存储集群监控节点出现异常；2、存储集群数据盘出现异常。
恢复建议	1、重启存储集群。
维护经验	

sdf数据盘异常

The screenshot shows the '节点管理' (Node Management) section of the storage cluster monitoring interface. It displays a table of storage nodes and a detailed view of the disks.

名称	状态	数据平衡	机架	存储外网IP地址	磁盘状态	存储容量	存储利用率	CPU利用率	内存
CVK-2	正常	启用	rack0	7.7.7.3	10 10 0 0 0...	17.46TB	52.03%	3.19%	251.30GB
CVK-5	正常	启用	rack0	7.7.7.2	10 10 0 0 0...	17.46TB	52.51%	3.46%	251.30GB
CVK-1	异常	启用	rack0	7.7.7.1	10 9 1 0 0...	17.46TB	47.42%	4.74%	251.30GB

磁盘名	硬盘池	使用类型	状态	容量	磁盘利用率	类型	接口类型	位置	阵列卡	磁盘灯状	操作
sdf	diskpool_hdd	数据盘	异常	2.18TB	0.00%	HDD	SAS	8.6	LSH-0	OFF	⚙️ ✖️ 🗑️ 📄
sdb	diskpool_hdd	缓存盘	正常	446.63GB	0.00%	SSD	SATA	8.2	LSH-0	OFF	⚙️ ✖️ 🗑️ 📄
sdc	diskpool_hdd	缓存盘	正常	446.63GB	0.00%	SSD	SATA	8.3	LSH-0	OFF	⚙️ ✖️ 🗑️ 📄

过程分析

1、ceph -s 一个osd down

```
[root@CVK-1 ~]# ceph -s
cluster:
  id: 226dfc09-23c5-44c9-b30e-6ea2cee66608
  health: HEALTH_WARN
         526593/12484578 objects misplaced (4.218%)
         Degraded data redundancy: 50660/12484578 objects degraded (0.406%), 900 pgs unclean, 125 pgs degraded, 125 pgs undersized

services:
  mon: 5 daemons, quorum CVK-5,CVK-2,CVK-6,CVK-4,CVK-3
  mgr: CVK-3(active), standbys: CVK-2, CVK-5, CVK-4, CVK-6
  osd: 48 osds: 47 up, 47 in; 900 remapped pgs

data:
  pools: 2 pools, 3072 pgs
  objects: 4063K objects, 16195 GB
  usage: 54862 GB used, 50135 GB / 102 TB avail
  pgs: 50660/12484578 objects degraded (0.406%)
      526593/12484578 objects misplaced (4.218%)
      2172 active+clean
      775 active+remapped+backfilling
      125 active+undersized+degraded+remapped+backfilling

to:
  client: 512 KB/s rd, 3790 B/s wr, 129 op/s rd, 7 op/s wr
  recovery: 158 MB/s, 30 objects/s
```

2、ceph osd tree

```
[root@CVK-1 ~]# ceph osd tree
ID CLASS WEIGHT TYPE NAME STATUS REWEIGHT PRI-AFF
-16 0 0 root maintain
-15 104.71655 root diskpool_hdd
-17 104.71655 rack rack0.diskpool_hdd
-21 17.45276 host cvk-1.diskpool_hdd
 5 hdd 2.18159 osd.5 up 1.00000 1.00000
11 hdd 2.18159 osd.11 up 1.00000 1.00000
17 hdd 2.18159 osd.17 down 0 1.00000
23 hdd 2.18159 osd.23 up 1.00000 1.00000
29 hdd 2.18159 osd.29 up 1.00000 1.00000
35 hdd 2.18159 osd.35 up 1.00000 1.00000
36 hdd 2.18159 osd.36 up 1.00000 1.00000
42 hdd 2.18159 osd.42 up 1.00000 1.00000
```

3. 看下日志

```
[root@CVK-1 ~]# cd /var/log
[~][root@CVK-1 log]# vi message
Jul 28 02:56:24 CVK-1 ceph-osd.sh: 2020-07-28 02:56:24.655573 7feb6ebff700 28746 47 ERROR osd.17 6784 handle_breakdown_evt osd assert
Jul 28 02:56:24 CVK-1 kernel: sd 0:2:5:0: [sd] tag#12 BRM Debug afi stat 0x2d, data len requested/completed 0x40000/0x0
Jul 28 02:56:24 CVK-1 kernel: sd 0:2:5:0: [sd] tag#12 FAIL IO Result: hostbyte=Drive_OK, devicebyte=OK, sense=
Jul 28 02:56:24 CVK-1 kernel: sd 0:2:5:0: [sd] tag#12 [sense key : detected Medium Error [current]]
Jul 28 02:56:24 CVK-1 kernel: sd 0:2:5:0: [sd] tag#12 Medium Error: unrecovered read error
Jul 28 02:56:24 CVK-1 kernel: sd 0:2:5:0: [sd] tag#12 CDB: Read(16) 88 00 00 00 00 00 71 13 7e 00 00 00 02 00 00 00
Jul 28 02:56:24 CVK-1 kernel: print_req_error: critical medium error, dev sdf, sector 1897102848
Jul 28 02:56:24 CVK-1 kernel: device-mapper: flushcache: flushcache uncached disk IO error: io error -5 block 1896896000 R/W READ
Jul 28 02:56:24 CVK-1 kernel: device-mapper: flushcache: flushcache uncached disk IO error: io error -5 block 1896896008 R/W READ
Jul 28 02:56:24 CVK-1 kernel: device-mapper: flushcache: flushcache uncached disk IO error: io error -5 block 1896896016 R/W READ
Jul 28 02:56:24 CVK-1 kernel: device-mapper: flushcache: flushcache uncached disk IO error: io error -5 block 1896896024 R/W READ
Jul 28 02:56:24 CVK-1 kernel: device-mapper: flushcache: flushcache uncached disk IO error: io error -5 block 1896896032 R/W READ
Jul 28 02:56:24 CVK-1 kernel: device-mapper: flushcache: flushcache uncached disk IO error: io error -5 block 1896896040 R/W READ
Jul 28 02:56:24 CVK-1 kernel: device-mapper: flushcache: flushcache uncached disk IO error: io error -5 block 1896896048 R/W READ
Jul 28 02:56:24 CVK-1 kernel: device-mapper: flushcache: flushcache uncached disk IO error: io error -5 block 1896896056 R/W READ
Jul 28 02:56:24 CVK-1 kernel: device-mapper: flushcache: flushcache uncached disk IO error: io error -5 block 1896896064 R/W READ
Jul 28 02:56:24 CVK-1 kernel: device-mapper: flushcache: flushcache uncached disk IO error: io error -5 block 1896896072 R/W READ
Jul 28 02:56:24 CVK-1 kernel: device-mapper: flushcache: flushcache uncached disk IO error: io error -5 block 1896896080 R/W READ
Jul 28 02:56:24 CVK-1 kernel: device-mapper: flushcache: flushcache uncached disk IO error: io error -5 block 1896896088 R/W READ
Jul 28 02:56:24 CVK-1 kernel: device-mapper: flushcache: flushcache uncached disk IO error: io error -5 block 1896896096 R/W READ
Jul 28 02:56:24 CVK-1 kernel: device-mapper: flushcache: flushcache uncached disk IO error: io error -5 block 1896896104 R/W READ
Jul 28 02:56:24 CVK-1 kernel: device-mapper: flushcache: flushcache uncached disk IO error: io error -5 block 1896896112 R/W READ
Jul 28 02:56:24 CVK-1 kernel: device-mapper: flushcache: flushcache uncached disk IO error: io error -5 block 1896896120 R/W READ
```

发现有medium error 磁盘坏道，需要进行换盘操作

解决方法

进行换盘操作