

# 知 磁盘性能不足导致etcd集群异常进而触发pod容器重启

厉梦如 2018-09-11 发表

## 组网及说明

cloudos版本: E1138H01

## 问题描述

cloudos环境发现容器不定期重启

## 过程分析

1.分析pod容器重启原因: 发现是Etcd无法及时向从节点发送心跳会导致集群状态不稳定, 集群服务重启, 进而导致Pod重启。

下图为集群主节点etcd服务的日志, 可以看出心跳不能及时发送给从节点。

```
root@cloudos18803:~# systemctl status etcd2 -l
● etcd2.service: Etcd Service
   Loaded: loaded (/usr/lib/systemd/system/etcd2.service; enabled; vendor preset: disabled)
   Active: active (running) since Tue 2018-09-08 19:12:06 CST; 2 weeks 6 days ago
     Main PID: 3560 (etcd)
    CGroup: /system.slice/etcd2.service
            └─3560 /usr/bin/etcd --name=etcd3 --initial-advertise-peer-urls=http://10.133.188.3:2380 --listen-peer-urls=http://10.133.188.3:2380 --listen-client-urls=http://10.133.188.3:2379,http://127.0.0.1:2379 --advertise-client-urls=http://10.133.188.3:2379 --initial-cluster-token=etcd-cluster --initial-cluster=etcd1=http://10.133.188.1:2380,etcd2=http://10.133.188.2:2380,etcd3=http://10.133.188.3:2380 --initial-cluster-state=new

May 29 10:21:53 cloudos18803 etcd[3560]: failed to send out heartbeat on time (deadline exceeded for 12.388495ms)
May 29 10:21:53 cloudos18803 etcd[3560]: server is likely overloaded
May 29 10:21:58 cloudos18803 etcd[3560]: failed to send out heartbeat on time (deadline exceeded for 5.536645ms)
May 29 10:21:58 cloudos18803 etcd[3560]: server is likely overloaded
May 29 10:21:58 cloudos18803 etcd[3560]: failed to send out heartbeat on time (deadline exceeded for 5.646113ms)
May 29 10:21:58 cloudos18803 etcd[3560]: server is likely overloaded
May 29 10:22:14 cloudos18803 etcd[3560]: failed to send out heartbeat on time (deadline exceeded for 118.107314ms)
May 29 10:22:14 cloudos18803 etcd[3560]: server is likely overloaded
May 29 10:22:14 cloudos18803 etcd[3560]: failed to send out heartbeat on time (deadline exceeded for 118.26724ms)
May 29 10:22:14 cloudos18803 etcd[3560]: server is likely overloaded
```

2.分析etcd异常原因:

Etcd的wal\_fsync\_durations指标用于标识ed对磁盘操作状态, 官方给出的建议99%的延时应该少于10ms。

查看方法为:

| Web访问: ip:2379/metrics

| 命令行: curl ip:2379/metrics | grep fsync

下面三张图分别为三个节点上该指标的延时分布: 输出显示延时极少落在16ms以内的, 主要集中在32ms跟64ms之间, 此外64~128ms、128~256ms、1~2s等区间均有分布, 导致了集群状态的不稳定。

```
# HELP etcd_wal_fsync_durations_seconds The latency distributions of fsync called by wal.
# TYPE etcd_wal_fsync_durations_seconds histogram
etcd_wal_fsync_durations_seconds_bucket[le="0.001"] 0
etcd_wal_fsync_durations_seconds_bucket[le="0.002"] 0
etcd_wal_fsync_durations_seconds_bucket[le="0.004"] 0
etcd_wal_fsync_durations_seconds_bucket[le="0.008"] 0
etcd_wal_fsync_durations_seconds_bucket[le="0.016"] 106835
etcd_wal_fsync_durations_seconds_bucket[le="0.032"] 3.820172e+06
etcd_wal_fsync_durations_seconds_bucket[le="0.064"] 4.679439e+06
etcd_wal_fsync_durations_seconds_bucket[le="0.128"] 4.737618e+06
etcd_wal_fsync_durations_seconds_bucket[le="0.256"] 4.755969e+06
etcd_wal_fsync_durations_seconds_bucket[le="0.512"] 4.758262e+06
etcd_wal_fsync_durations_seconds_bucket[le="1.024"] 4.759203e+06
etcd_wal_fsync_durations_seconds_bucket[le="2.048"] 4.760925e+06
etcd_wal_fsync_durations_seconds_bucket[le="4.096"] 4.76211e+06
etcd_wal_fsync_durations_seconds_bucket[le="8.192"] 4.76211e+06
etcd_wal_fsync_durations_seconds_bucket[le="+Inf"] 4.76211e+06
etcd_wal_fsync_durations_seconds_sum 127931.77732487605
etcd_wal_fsync_durations_seconds_count 4.76211e+06
# HELP etcd_wal_fsync_durations_seconds The latency distributions of fsync called by wal.
# TYPE etcd_wal_fsync_durations_seconds histogram
etcd_wal_fsync_durations_seconds_bucket[le="0.001"] 0
etcd_wal_fsync_durations_seconds_bucket[le="0.002"] 0
etcd_wal_fsync_durations_seconds_bucket[le="0.004"] 0
etcd_wal_fsync_durations_seconds_bucket[le="0.008"] 0
etcd_wal_fsync_durations_seconds_bucket[le="0.016"] 476585
etcd_wal_fsync_durations_seconds_bucket[le="0.032"] 3.809628e+06
etcd_wal_fsync_durations_seconds_bucket[le="0.064"] 4.646586e+06
etcd_wal_fsync_durations_seconds_bucket[le="0.128"] 4.720843e+06
etcd_wal_fsync_durations_seconds_bucket[le="0.256"] 4.751611e+06
etcd_wal_fsync_durations_seconds_bucket[le="0.512"] 4.754297e+06
etcd_wal_fsync_durations_seconds_bucket[le="1.024"] 4.754823e+06
etcd_wal_fsync_durations_seconds_bucket[le="2.048"] 4.756737e+06
etcd_wal_fsync_durations_seconds_bucket[le="4.096"] 4.75899e+06
etcd_wal_fsync_durations_seconds_bucket[le="8.192"] 4.7587e+06
etcd_wal_fsync_durations_seconds_bucket[le="+Inf"] 4.7587e+06
etcd_wal_fsync_durations_seconds_sum 127036.13129084474
etcd_wal_fsync_durations_seconds_count 4.7597e+06
# HELP etcd_wal_fsync_durations_seconds The latency distributions of fsync called by wal.
# TYPE etcd_wal_fsync_durations_seconds histogram
etcd_wal_fsync_durations_seconds_bucket[le="0.001"] 0
etcd_wal_fsync_durations_seconds_bucket[le="0.002"] 0
etcd_wal_fsync_durations_seconds_bucket[le="0.004"] 0
etcd_wal_fsync_durations_seconds_bucket[le="0.008"] 0
etcd_wal_fsync_durations_seconds_bucket[le="0.016"] 189518
etcd_wal_fsync_durations_seconds_bucket[le="0.032"] 3.838254e+06
etcd_wal_fsync_durations_seconds_bucket[le="0.064"] 4.655385e+06
etcd_wal_fsync_durations_seconds_bucket[le="0.128"] 4.723128e+06
etcd_wal_fsync_durations_seconds_bucket[le="0.256"] 4.745625e+06
etcd_wal_fsync_durations_seconds_bucket[le="0.512"] 4.749422e+06
etcd_wal_fsync_durations_seconds_bucket[le="1.024"] 4.750138e+06
etcd_wal_fsync_durations_seconds_bucket[le="2.048"] 4.751918e+06
etcd_wal_fsync_durations_seconds_bucket[le="4.096"] 4.754432e+06
etcd_wal_fsync_durations_seconds_bucket[le="8.192"] 4.754442e+06
etcd_wal_fsync_durations_seconds_bucket[le="+Inf"] 4.754442e+06
etcd_wal_fsync_durations_seconds_sum 130070.59040839928
etcd_wal_fsync_durations_seconds_count 4.754442e+06
```

## 解决方法

问题定位在磁盘性能不足, 造成etcd集群心跳超时, 进而导致pod重启。所以处理方法:

1、缓解措施:

修改etcd超时时间, 降低pod重启概率, 但会增大集群间数据同步耗时增加, 影响系统可靠性灵敏度, 且随着业务量增加、系统长期运行磁盘损耗增大, pod重启概率会逐渐增加。

2、彻底解决:

更换RAID卡（缓存不小于1GB），提升磁盘性能。