

知 H3C VCFC网络存在异常可能导致控制器集群震荡业务受影响的技术公告

VCFC 刘玉娟 2020-02-25 发表

问题描述

【产品型号】

H3C VCFC控制器

【涉及版本】

强控OpenFlow组网中，E2180P16及之前版本

弱控EVPN组网中，E2503及之前版本

【问题描述】

当控制器集群发生震荡时，可能无法打开VCFC 登陆页面或打开时有明显卡顿；登陆后在如下控制器信息页面可以看到集群状态异常，在“active”和“down”两个状态之间不停切换，或者配置恢复状态为“false”。

The screenshot shows the 'Controller Information' section of the H3C management interface. Under 'Cluster Information', it displays the cluster name as 'VCFC集群', IP as '255.255.255.255', and 'Cluster Mode' as 'Multi-Region'. Below this, the 'Controller Information' table lists three controllers with their IP, role, region, priority, and status. The 'Cluster Status' column shows alternating 'active' and 'down' states for each controller.

IP	控制器名称	集群角色	Region	Region...	优先级	集群网卡	主OpenFl...	总OpenFl...	集群状态	配置恢复...	Region连通性
99.1.1.133	133	Leader	777777	eth0	0	0	active	true	...
99.1.1.132	132	Leader	region1	Subordin...	888888	eth0	3	6	active	true	green
99.1.1.131	131	Leader*	region1	Master	999999	eth0	3	6	active	true	green

在系统日志中设置高级搜索条件，发现短时间内频繁有控制器断连重连的日志：

The screenshot shows the 'Advanced Search' dialog box over a log viewer. The search criteria are set to 'region' in the 'Subject' field and 'The connection' in the 'Content' field. The main log viewer shows several entries related to controller connections.

级别	日期/时间	IP	业务名称	主题	内容
INFO	2020-02-17 17:58:25	99.1.1.132	CON_REGION	region	The connection between the current controller and controller 99.
INFO	2020-02-17 17:58:20	99.1.1.131	CON_REGION	region	The connection between the current controller and controller 99.
INFO	2020-02-17 17:58:18	99.1.1.132	CON_REGION	region	The connection between the current controller and controller 99.
INFO	2020-02-17 17:58:18	99.1.1.131	CON_REGION	region	The connection between the current controller and controller 99.

[2020-02-18 14:24:48:714] INFO region REGION_CONTROLLER_STATUS_UP
The connection between the current controller and controller 10.47.156.2 in the region is normal.
[2020-02-18 14:24:21:854] SERIOUS region REGION_CONTROLLER_STATUS_DOWN
The connection between the current controller and controller 10.47.156.4 in the region is abnormal.
[2020-02-18 14:24:33:919] INFO region REGION_CONTROLLER_STATUS_UP
The connection between the current controller and controller 10.47.156.4 in the region is normal.
[2020-02-18 14:24:57:800] INFO region REGION_CONTROLLER_STATUS_UP
The connection between the current controller and controller 10.47.156.4 in the region is normal.
[2020-02-18 14:24:57:801] SERIOUS region REGION_CONTROLLER_STATUS_DOWN
The connection between the current controller and controller 10.47.156.4 in the region is abnormal.
[2020-02-18 14:24:57:803] INFO region REGION_CONTROLLER_STATUS_UP
The connection between the current controller and controller 10.47.156.4 in the region is normal.
[2020-02-18 14:27:55:732] SERIOUS region REGION_CONTROLLER_STATUS_DOWN
The connection between the current controller and controller 10.47.156.3 in the region is abnormal.
[2020-02-18 14:27:12:364] SERIOUS region REGION_CONTROLLER_STATUS_DOWN
The connection between the current controller and controller 10.47.156.1 in the region is abnormal.
由网络不稳定引起的集群震荡是随机的，但持续且频繁的集群震荡将对控制器的稳定运行产生不良影响。强控主机Overlay的组网下可能导致现网业务出现大量虚机丢包或不通的严重故障；强控网络Overlay和弱控组网下可能出现物理网元状态为“inactive”的现象，影响新增业务的下发。

原因分析

虚机发送ARP报文上送控制器处理，控制器使用计数器记录收到的ARP Packet-in消息，每次获取一定数量的消息在各软件模块之间同步。计数器的计数值由正值反转为负值后，当网络存在异常如ARP攻

击或虚机异常迁移，短时间内大量ARP上送控制器时，由于代码考虑不周，控制器后台每次获取的待同步消息都会包含已经同步处理过的消息，造成大量消息在软件模块间同步，内存占用过高，引发集群震荡。

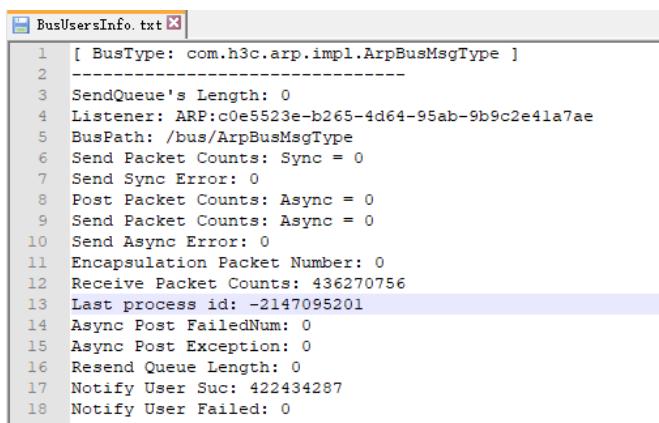
规避措施/解决方案

【规避措施】

如果局点使用的控制器为本技术公告涉及的版本，在升级控制器版本彻底解决之前，建议各代表处工程师按照每星期一次的频率主动进行检查。

检查方法：

1、导出控制器最近一天的诊断日志，解压并查看以下文件内容。\\BusKeystore\\BusUserInfo.txt.zip，如下图所示：



```
1 [ BusType: com.h3c.arp.impl.ArpBusMsgType ]
2 -----
3 SendQueue's Length: 0
4 Listener: ARP:c0e5523e-b265-4d64-95ab-9b9c2e41a7ae
5 BusPath: /bus/ArpBusMsgType
6 Send Packet Counts: Sync = 0
7 Send Sync Error: 0
8 Post Packet Counts: Async = 0
9 Send Packet Counts: Async = 0
10 Send Async Error: 0
11 Encapsulation Packet Number: 0
12 Receive Packet Counts: 436270756
13 Last process id: -2147095201
14 Async Post FailedNum: 0
15 Async Post Exception: 0
16 Resend Queue Length: 0
17 Notify User Suc: 422434287
18 Notify User Failed: 0
```

2、查找所有Last process id字段，如下图所示；若该字段为除-1以外的负数代表已经反转，若已达10亿量级，说明该值有反转为负的风险。ARP计数反转为负并不会直接引起控制器故障，但计数反转后网络中出现短时间内大量ARP上送控制器的异常情况，则可能会引发控制器集群震荡。

恢复方法：

如果Last process id字段为除-1以外的负数或已达10亿量级，需要联系解决方案支持部L3，尽快重启控制器整集群，重启前请提交网络变更电子流，并提前申请操作窗口。

【解决方案】

- 1、强控OpenFlow组网中，请升级到E2180P17（含）之后版本解决。
- 2、弱控VPN组网中，请升级到E2504（含）之后版本解决。