EIGRP路由震荡排错案例

目录

硬件平台

路由器及多层交换机

软件版本

所有

案例简介

在一些客户的实际网络中，存在一些由于策略配置而导致的在某些情况的下的路由震荡，路由环路。这些一般是由于一些特殊的需求配置已经不规则的拓扑导致。 通常情况下我们仔细分析选路规则，路由策略规则就可以解决问题。此类问题有明显的特征（路由环路/震荡），解决此类问题主要根据所学路由协议的基础，一步一步分析便可解决。 以下是一个真实案例，为保护客户资料，路由器输出信息是在实验环境中抓取。

信息及描述：

1. 客户有两台路由器互联通过两条链路互联，F0/0是百兆链路，F1/0是十兆链路，网络协议是使用的是EIGRP协议
2. 云网络（外部路由器）向下发送一条192.168.1.0的路由，在Router B上做了管理距离（AD）调整，要求去往192.168.1.0的流量必须优先选择ROUTER A

故障描述：

当网络云（外部路由器）与ROUTER A之间的链路断开后，发现两台路由器关于某个EIGRP路由192.168.1.0/24无法正常收敛，路由有震荡。

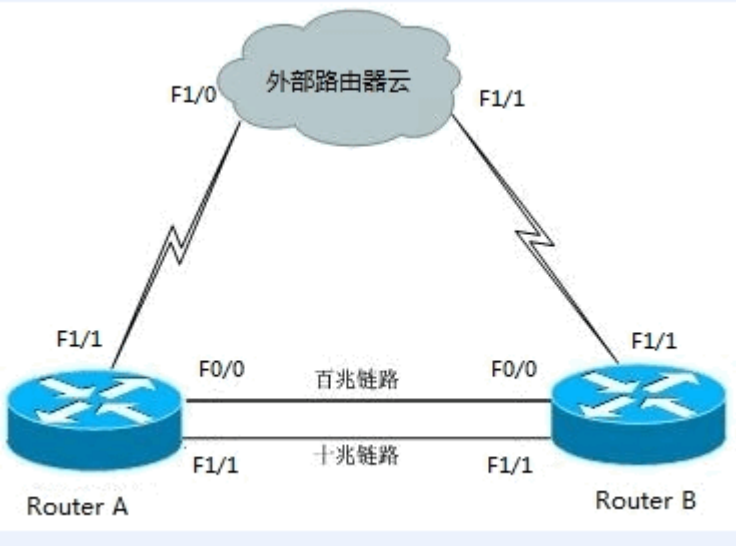

地址规划：

路由器接口	IP 地址	路由器接口	IP 地址
Router A F0/0	12.1.1.1/24	Router B F0/0	12.1.1.2/24
Router A F1/0	21.1.1.1/24	Router B F1/0	21.1.1.2/24
Router A F1/1	13.1.1.1/24	路由器云 F1/0	13.1.1.3/24
Router B F1/1	23.1.1.2/24	路由器云 F1/1	23.1.1.3/24

故障诊断步骤

首先检查故障现象，当Router A的上行链路down掉以后，发现Router A去往192.168.1.0/24的路由出现路由翻动，next hop 始终在12.1.1.2 (F0/0)和 21.1.1.2 (F1/0)之间切换

```
Router_A#show ip route 192.168.1.0
Routing entry for 192.168.1.0/24
  Known via "eigrp 100", distance 90, metric 158720, type internal
  Redistributing via eigrp 100
  Last update from 12.1.1.2 on FastEthernet0/0, 00:00:00 ago
  Routing Descriptor Blocks:
  * 12.1.1.2, from 12.1.1.2, 00:00:00 ago, via FastEthernet0/0
      Route metric is 158720, traffic share count is 1
      Total delay is 5200 microseconds, minimum bandwidth is 100000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 2
Router_A#show ip route 192.168.1.0
% Network not in table
Router_A#show ip route 192.168.1.0
Routing entry for 192.168.1.0/24
  Known via "eigrp 100", distance 90, metric 389120, type internal
  Redistributing via eigrp 100
  Last update from 21.1.1.2 on FastEthernet1/0, 00:00:00 ago
  Routing Descriptor Blocks:
  * 21.1.1.2, from 21.1.1.2, 00:00:00 ago, via FastEthernet1/0
      Route metric is 389120, traffic share count is 1
      Total delay is 5200 microseconds, minimum bandwidth is 10000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 2
```

检查Router B的路由表发现Router B 关于192.168.1.0/24的next hop 也始终在21.1.1.1 和23.1.1.3之间切换

```
Router_B#show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, + - replicated route

Gateway of last resort is not set

      12.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        12.1.1.0/24 is directly connected, FastEthernet0/0
L        12.1.1.2/32 is directly connected, FastEthernet0/0
      13.0.0.0/24 is subnetted, 1 subnets
D        13.1.1.0 [90/30720] via 23.1.1.3, 00:10:44, FastEthernet1/1
      21.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        21.1.1.0/24 is directly connected, FastEthernet1/0
L        21.1.1.2/32 is directly connected, FastEthernet1/0
      23.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        23.1.1.0/24 is directly connected, FastEthernet1/1
L        23.1.1.2/32 is directly connected, FastEthernet1/1
D     192.168.1.0/24 [90/156160] via 23.1.1.3, 00:00:00, FastEthernet1/1
Router_B#show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, + - replicated route

Gateway of last resort is not set

      12.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        12.1.1.0/24 is directly connected, FastEthernet0/0
L        12.1.1.2/32 is directly connected, FastEthernet0/0
      13.0.0.0/24 is subnetted, 1 subnets
D        13.1.1.0 [90/30720] via 23.1.1.3, 00:11:27, FastEthernet1/1
      21.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        21.1.1.0/24 is directly connected, FastEthernet1/0
L        21.1.1.2/32 is directly connected, FastEthernet1/0
      23.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        23.1.1.0/24 is directly connected, FastEthernet1/1
L        23.1.1.2/32 is directly connected, FastEthernet1/1
D     192.168.1.0/24 [89/391680] via 21.1.1.1, 00:00:00, FastEthernet1/0
```

此时一定是出现了路由优选的环路即Router A 与 Router B之间在不同的时间相互优选了对方为最佳路由，导致路由震荡，此时我们细分析一下路由的选路原则即可发现问题：

当路由条目一致时，优先选择管理距离比较低的路由并添加进入路由表，此时在Router_B上进行了AD的修改（通过Router A的两个链路学到的这条路由的AD均为89，默认为90），所以正常时刻Router_B会有优选Router A作为next-hop，但是Router_A有两个不同的链路，所以需要继续比较，当AD一样时会优先选择Metric值最小的Next-hop作为最佳路由并且添加到路由表当中，所以正常时刻（即Router A上行链路没有down的时刻 ），Router B会优选Router A的F0/0 作为最优路由

```
Router_B#show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, + - replicated route

Gateway of last resort is not set

      12.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        12.1.1.0/24 is directly connected, FastEthernet0/0
L        12.1.1.2/32 is directly connected, FastEthernet0/0
      13.0.0.0/24 is subnetted, 1 subnets
D        13.1.1.0 [90/30720] via 23.1.1.3, 00:00:12, FastEthernet1/1
                  [90/30720] via 12.1.1.1, 00:00:12, FastEthernet0/0
      21.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        21.1.1.0/24 is directly connected, FastEthernet1/0
L        21.1.1.2/32 is directly connected, FastEthernet1/0
      23.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        23.1.1.0/24 is directly connected, FastEthernet1/1
L        23.1.1.2/32 is directly connected, FastEthernet1/1
D     192.168.1.0/24 [89/158720] via 12.1.1.1, 00:00:10, FastEthernet0/0
```

检查Router_B的EIGRP 拓扑表

```
Router_B#show ip eigrp topology
EIGRP-IPv4 Topology Table for AS(100)/ID(23.1.1.2)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 12.1.1.0/24, 1 successors, FD is 28160
        via Connected, FastEthernet0/0
P 13.1.1.0/24, 2 successors, FD is 30720
        via 12.1.1.1 (30720/28160), FastEthernet0/0
        via 23.1.1.3 (30720/28160), FastEthernet1/1
        via 21.1.1.1 (261120/28160), FastEthernet1/0
P 21.1.1.0/24, 1 successors, FD is 258560
        via Connected, FastEthernet1/0
P 23.1.1.0/24, 1 successors, FD is 28160
        via Connected, FastEthernet1/1
P 192.168.1.0/24, 1 successors, FD is 158720
        via 12.1.1.1 (158720/156160), FastEthernet0/0, serno 19409
        via 21.1.1.1 (389120/156160), FastEthernet1/0
        via 23.1.1.3 (156160/128256), FastEthernet1/1
```

12.1.1.1 成为去往192.168.1.0/24的后继路由器（Successor），21.1.1.1 成为去往192.168.1.0/24的可行后继路由器（Feasible Successor）

当Router_A的上行链路down 掉以后，Router_A无可行后继路由器，会把192.168.1.0 /24置为不可达，(Metric 4294967295)，发送给所有的邻居

```
Router_A#show ip eigrp ev
*Oct  8 20:20:11.215: %SYS-5-CONFIG_I: Configured from console by console
Event information for AS 100:
1    20:20:13.599 State change: Local origin Successor Origin
2    20:20:13.599 Metric set: 192.168.1.0/24 4294967295
3    20:20:13.599 Active net/peers: 192.168.1.0/24 131072
4    20:20:13.599 FC not sat Dmin/met: 4294967295 389120
5    20:20:13.599 Find FS: 192.168.1.0/24 158720
6    20:20:13.599 Rcv query met/succ met: 4294967295 4294967295
7    20:20:13.599 Rcv query dest/nh: 192.168.1.0/24 21.1.1.2
8    20:20:13.595 Change queue emptied, entries: 1
9    20:20:13.595 Metric set: 192.168.1.0/24 389120
10   20:20:13.595 Update reason, delay: new if 133120
11   20:20:13.595 Update sent, RD: 192.168.1.0/24 158720
12   20:20:13.595 Update reason, delay: metric chg 133120
13   20:20:13.595 Update sent, RD: 192.168.1.0/24 158720
14   20:20:13.595 Route install: 192.168.1.0/24 21.1.1.2
15   20:20:13.591 Send reply: 192.168.1.0/24 12.1.1.2
16   20:20:13.591 FC sat rdbmet/succmet: 389120 156160
17   20:20:13.591 FC sat nh/ndbmet: 21.1.1.2 158720
18   20:20:13.591 Find FS: 192.168.1.0/24 158720
19   20:20:13.591 Rcv query met/succ met: 4294967295 4294967295
20   20:20:13.591 Rcv query dest/nh: 192.168.1.0/24 12.1.1.2
21   20:20:13.551 Change queue emptied, entries: 1
22   20:20:13.551 Metric set: 192.168.1.0/24 158720
```

Router_B收到此不可达路由后会撤销掉此路由，并且优选从路由器云学习此路由，并且把去往外部路由器的接口作为next-hop （ 23.1.1.3 ），metric 156160 ，AD 90

```
Router_B#show ip eigrp topology
IP-EIGRP Topology Table for AS(100)/ID(23.1.1.2)

Codes: P - Passive, A - Active, U - Update, Q - Query, R -
       r - reply Status, s - sia Status

P 12.1.1.0/24, 1 successors, FD is 28160
        via Connected, FastEthernet0/0
P 13.1.1.0/24, 1 successors, FD is 30720
        via 23.1.1.3 (30720/28160), FastEthernet1/1
P 21.1.1.0/24, 1 successors, FD is 258560
        via Connected, FastEthernet1/0
P 23.1.1.0/24, 1 successors, FD is 28160
        via Connected, FastEthernet1/1
P 192.168.1.0/24, 1 successors, FD is 156160
        via 23.1.1.3 (156160/128256), FastEthernet1/1
```

此时路由器B收敛完成向Router_A宣告此路由,由于Router A 已经失去了Successor ，并且没有Feasible Successor，所以Router_A 会选择Router_B 作为新的Successor，同样的选举原则，Router_A会优选Fa0/0 （ 百兆链路 ）作为最佳路由,此时A的metric 为156160+256*10 （ delay ）=158720

```
rcvupdate: 192.168.1.0/24 via 12.1.1.2 metric 158720/156160
DUAL: Find FS for dest 192.168.1.0/24. FD is 4294967295, RD is 4294967295 found
DUAL: RT installed 192.168.1.0/24 via 12.1.1.2
DUAL: Send update about 192.168.1.0/24.  Reason: metric chg
DUAL: Send update about 192.168.1.0/24.  Reason: new if
DUAL: Removing dest 192.168.1.0/24, nexthop 21.1.1.2
DUAL: dest(192.168.1.0/24) not active
DUAL: rcvupdate: 192.168.1.0/24 via 21.1.1.2 metric 389120/156160
DUAL: Find FS for dest 192.168.1.0/24. FD is 158720, RD is 158720
DUAL:     12.1.1.2 metric 158720/156160
DUAL:     21.1.1.2 metric 389120/156160 found Dmin is 158720
```

```
Router_A#show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static rou
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

     21.0.0.0/24 is subnetted, 1 subnets
C       21.1.1.0 is directly connected, FastEthernet1/0
     23.0.0.0/24 is subnetted, 1 subnets
D       23.1.1.0 [90/30720] via 12.1.1.2, 12:02:38, FastEthernet0/0
     12.0.0.0/24 is subnetted, 1 subnets
C       12.1.1.0 is directly connected, FastEthernet0/0
D    192.168.1.0/24 [90/158720] via 12.1.1.2, 00:00:00, FastEthernet0/0
```

同时由于Router_A学到了新的路由，需要向其所有邻居(除了Successor)通告此路由，向Successor宣告一条毒性逆转路由（Metirc 为4294967295），所以Router_A 又通过Fa 1/0接口发送此路由更新，Router_B从Fa1/0收到此路由更新，会优选从此接口学到的路由(由于之前做过AD调整，从Router_A 学到的路由均为AD 89，优选于从路由器云学到的路由)，所以Router_B撤销掉之前的最佳路由（通过23.1.1.3），优选21.1.1.2 （F1/0）作为最佳路由，此时metric 更大，为10M链路的带宽已经RouterA 和Router B之间的Delay之和 ，metric 为391680

```
Router_B#show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-
       ia - IS-IS inter area, * - candidate default, U - per-user static ro
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

     21.0.0.0/24 is subnetted, 1 subnets
C       21.1.1.0 is directly connected, FastEthernet1/0
     23.0.0.0/24 is subnetted, 1 subnets
C       23.1.1.0 is directly connected, FastEthernet1/1
     12.0.0.0/24 is subnetted, 1 subnets
C       12.1.1.0 is directly connected, FastEthernet0/0
D    192.168.1.0/24 [89/391680] via 21.1.1.1, 00:00:00, FastEthernet1/0
```

由于B有了新的Successor ，所以向除Successor 之外所有邻居通告此路由(F0/0)，向Successor 宣告一条毒性逆转路由（Metirc 为4294967295），同样的道理Router_A会从F0/0 收到一个新的update 带有新的metric (391680)，此Metirc 要比原有Successor 的Metirc 158720要大，但是仍然是同一个successor (12.1.1.2)

```
Router_A#show ip eigrp topology
IP-EIGRP Topology Table for AS(100)/ID(21.1.1.1)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 12.1.1.0/24, 1 successors, FD is 28160
        via Connected, FastEthernet0/0
P 13.1.1.0/24, 1 successors, FD is 33280
        via 12.1.1.2 (33280/30720), FastEthernet0/0
        via 21.1.1.2 (263680/30720), FastEthernet1/0
P 21.1.1.0/24, 1 successors, FD is 258560
        via Connected, FastEthernet1/0
P 23.1.1.0/24, 1 successors, FD is 30720
        via 12.1.1.2 (30720/28160), FastEthernet0/0
        via 21.1.1.2 (261120/28160), FastEthernet1/0
P 192.168.1.0/24, 1 successors, FD is 158720, U
        via 12.1.1.2 (158720/156160), FastEthernet0/0
      ▌ via 21.1.1.2 (389120/156160), FastEthernet1/0
```

由于此Metric比原有的Feasible Successor（21.1.1.2）的metric（391620）还要大，所有R1忽略这条update，优选21.1.1.2 作为新的Successor，并且向此接口发送毒性逆转路由，

Router B 上debug eigrp fsm

```
DUAL: rcvquery: 192.168.1.0/24 via 12.1.1.1 metric 4294967295/4294967295, RD is 4294967295
DUAL: send REPLY(r1/n1) about 192.168.1.0/24 to 12.1.1.1
DUAL: rcvquery: 192.168.1.0/24 via 21.1.1.1 metric 4294967295/4294967295, RD is 4294967295
DUAL: send REPLY(r1/n1) about 192.168.1.0/24 to 21.1.1.1
DUAL: rcvreply: 192.168.1.0/24 via 21.1.1.1 metric 4294967295/4294967295
```

Router B收到此毒性逆转路由，删掉21.1.1.2 作为next-hop，并且重新选择23.1.1.3作为最优路由，产生选路环路，Router A的next hop 始终在12.1.1.2 和21.1.1.2 之间切换，Router B 的next hop 始终在23.1.1.2 和 21.1.1.2 之间切换，出现路由震荡

```
260  20:51:12.283 Rcv reply met/succ met: 4294967295 4294967295
261  20:51:12.283 Rcv reply dest/nh: 192.168.1.0/24 21.1.1.1
262  20:51:12.223 Send reply: 192.168.1.0/24 12.1.1.1
263  20:51:12.223 Rcv query met/succ met: 4294967295 4294967295
264  20:51:12.223 Rcv query dest/nh: 192.168.1.0/24 12.1.1.1
265  20:51:12.187 RDB delete: 192.168.1.0/24 12.1.1.1
266  20:51:12.187 Clr handle num/bits: 2 0x2
267  20:51:12.187 Clr handle dest/cnt: 192.168.1.0/24 1
268  20:51:12.187 Rcv reply met/succ met: 4294967295 4294967295
269  20:51:12.187 Rcv reply dest/nh: 192.168.1.0/24 12.1.1.1
270  20:51:12.171 Send reply: 192.168.1.0/24 21.1.1.1
271  20:51:12.171 Rcv query met/succ met: 4294967295 4294967295
272  20:51:12.171 Rcv query dest/nh: 192.168.1.0/24 21.1.1.1
273  20:51:12.155 Clr handle num/bits: 0 0x6
274  20:51:12.155 Clr handle dest/cnt: 192.168.1.0/24 2
275  20:51:12.155 Rcv reply met/succ met: 156160 128256
276  20:51:12.155 Rcv reply dest/nh: 192.168.1.0/24 23.1.1.3
277  20:51:12.091 Metric set: 192.168.1.0/24 4294967295
278  20:51:12.091 Active net/peers: 192.168.1.0/24 196608
```

DUAL Rule 1:
Whenever a Router chooses a new successor, it informs all its other neighbors about the new reported distance.

DUAL Rule 2:

Every time a Router selects a successor, it sends a poison update to its successor (a poison reverse).

DUAL Rule 3:
A poison update is sent to all neighbors on the interface through which the successor is reachable unless split-horizon is turned
综上所述，因为拓扑中存在不对称链路百兆链路和十兆链路，并且做了特殊配置，两个链路上都做了AD优先级的配置，此问题的解决办法有很多，修改AD以及修改metric都可避免此类问题发生，不再赘述。

经验总结

对于路由协议相关命令的输出，逐跳查找，show ip eigrp event可以记录协议内部计算的信息， 结合metric可以计算出所经过的路径，根据原理分析便可解决问题。

相关命令

```
show ip route
show ip route x.x.x.x
show ip eigrp topology
show ip eigrp {AS} topology x.x.x.x
show ip eigrp events
debug eigrp packet
debug eigrp fsm
```

相关错误信息

无

其他相关文档

EIGRP排错流程
http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094613.shtml