

# ASR9000 ping丢包troubleshooting

## 目录

[硬件平台](#)  
[软件版本](#)  
[案例介绍](#)  
[问题分析思路](#)  
[问题总结](#)  
[经验总结](#)  
[相关命令](#)

## 硬件平台

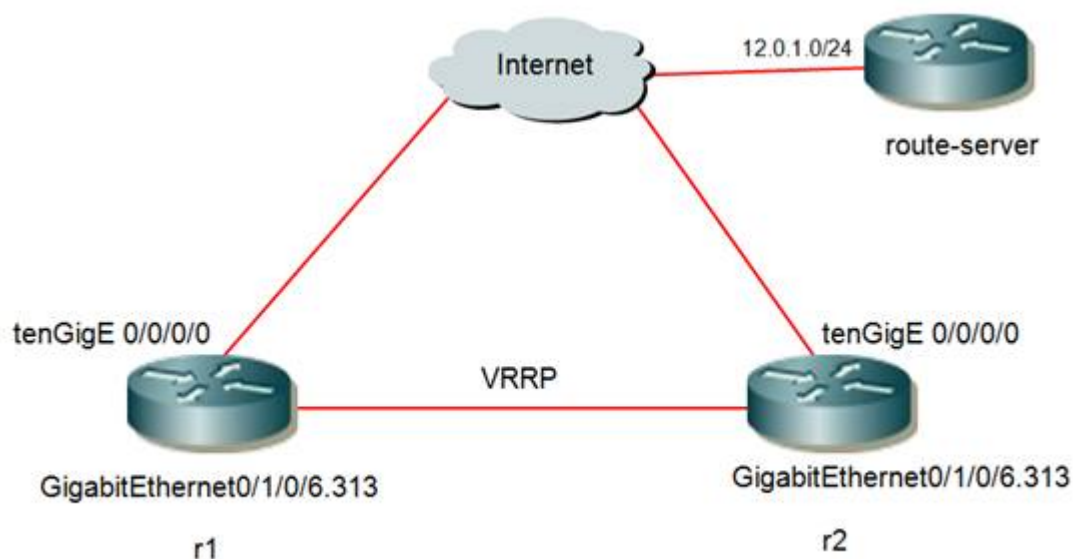
ASR9000

## 软件版本

4.2.0

## 案例介绍

拓扑示例:




问题,客户从 外网internet ping  
3 个 vrrp subnet的地址时候始终只能通一个IP地址:  
vrrp virtual IP :2.2.2.129

其他不能ping通的IP地址

active physical rp address: 2.2.2.130  
backup physical rp address: 2.2.2.131

截取部分TOPO图说明到2.2.2.131 不通的问题：



```
vrrp virtual IP :2.2.2.129
```

```
RP/0/RSP0/CPU0:r1#show run router vrrp interface gigabitEthernet 0/1/0/6.313
```

```
Tue Mar 27 11:35:27.676 Beijing
```

```
router vrrp
```

```
interface GigabitEthernet0/1/0/6.313
```

```
address-family ipv4
```

```
vrrp 113
```

```
priority 120
```

```
preempt delay 10
```

```
address 2.2.2.129
```

```
!
```

```
!
```

```
!
```

```
!
```

```
RP/0/RSP0/CPU0:r2#
```

```
RP/0/RSP0/CPU0:r2#show run router vrrp interface gigabitEthernet 0/1/0/6.313
```

```
Tue Mar 27 11:35:27.676 Beijing
```

```
router vrrp
```

```
interface GigabitEthernet0/1/0/6.313
```

```
address-family ipv4
```

```
vrrp 113
```

```
preempt delay 10
```

```
address 2.2.2.129
```

```
!
```

```
!
```

```
!
```

```
!
```

## 问题分析思路

### A. 数据包丢在什么地方

1.1.1.48(R1) 其中一个上行接口tenGigE 0/0/0/0 :ipv4 address 3.3.3.66 255.255.255.252

```
1.1.1.48: vrrp active
```

```
RP/0/RSP0/CPU0:r1#
RP/0/RSP0/CPU0:r1#show run interface gigabitEthernet
0/1/0/6.313
Tue Mar 27 09:52:10.186 Beijing
interface GigabitEthernet0/1/0/6.313
service-policy input default
ipv4 address 2.2.2.130 255.255.255.240 route-tag 500
ipv4 verify unicast source reachable-via any
encapsulation dot1q 313
!
```

```
RP/0/RSP0/CPU0:r1#
```

```
RP/0/RSP0/CPU0:r1#show vrrp interface gigabitEthernet 0/1/0/6.313 detail | utility egrep
Master
Tue Mar 27 09:51:00.516 Beijing
State is Master
Mar 23 03:04:57.960 Beijing Backup -> Master Master down timer expired
Master router is local
Master Down Timer 3.531 (3 x 1 + 136/256)
```

### 1.1.1.49(R2) 其中一个上行接口TenGigE 0/0/0/0 :ipv4 address 3.3.3.74 255.255.255.252

```
1.1.1.49: vrrp backup
```

```
RP/0/RSP0/CPU0:r2#show run int gigabitEthernet 0/1/0/6.313
Tue Mar 27 09:42:58.874 UTC
interface GigabitEthernet0/1/0/6.313
service-policy input default
ipv4 address 2.2.2.131 255.255.255.240 route-tag 500
ipv4 verify unicast source reachable-via any=====>
encapsulation dot1q 313
!
```

```
RP/0/RSP0/CPU0:r2#show vrrp interface gigabitEthernet 0/1/0/6.313 detail | utility egrep
Master
Tue Mar 27 09:51:41.125 UTC
Master router is 2.2.2.130, priority 120
Master Down Timer 3.609 (3 x 1 + 156/256)
```

### 测试Server的源地址网段：

```
route-server>show ip route c
12.0.0.0/8 is variably subnetted, 2509 subnets, 11 masks
C 12.0.1.0/24 is directly connected, GigabitEthernet0/1
route-server>
```

```
1.1.1.49: R2 12.0.1.0 :
```

```
RP/0/RSP0/CPU0:r2#show route ipv4 12.0.1.0
Tue Mar 27 09:56:01.428 UTCRouting entry for 0.0.0.0/0
=====> Known via "ospf 100", distance 110,
metric 1, candidate default path Tag 100, type extern 2 Installed Mar 22 15:51:04.265 for 4d18h
Routing Descriptor Blocks 3.3.3.73, from 2.2.2.8, via TenGigE0/0/0/0 Route metric is 1 3.3.3.77,
from 2.2.2.9, via TenGigE0/1/0/0 Route metric is 1 No advertising protos. RP/0/RSP0/CPU0:r2#
route-server>trace 2.2.2.131 Type escape sequence to abort. Tracing the route to 2.2.2.131 1
gateway.cbbtier3.att.net (12.0.1.202) [AS 7018] 4 msec 0 msec 4 msec 2 n54ny401me3-
cbbtier3.ip.att.net (12.89.5.13) [AS 7018] 8 msec 16 msec 16 msec 3 cr1.n54ny.ip.att.net
(12.123.2.6) [MPLS: Label 16092 Exp 1] 80 msec 80 msec 76 msec 4 cr2.cgcil.ip.att.net
(12.122.1.2) [MPLS: Labels 23252/16494 Exp 1] 80 msec 80 msec 84 msec 5 cr1.cgcil.ip.att.net
```

```
(12.122.2.53) [MPLS: Labels 23524/16494 Exp 1] 80 msec 80 msec 76 msec 6 cr2.dvmco.ip.att.net
(12.122.31.85) [MPLS: Labels 23794/16494 Exp 1] 80 msec 80 msec 80 msec 7 cr1.slkut.ip.att.net
(12.122.30.25) [MPLS: Labels 16216/16494 Exp 1] 80 msec 80 msec 80 msec 8 cr2.la2ca.ip.att.net
(12.122.30.30) [MPLS: Labels 0/16494 Exp 1] 80 msec 84 msec 80 msec 9 cr84.la2ca.ip.att.net
(12.123.30.249) [MPLS: Labels 0/16333 Exp 1] 76 msec 80 msec 80 msec 10 gar2.lsrca.ip.att.net
(12.122.129.49) 80 msec 80 msec 80 msec 11 12.118.130.86 [AS 7018] 388 msec 388 msec 384 msec 12
219.158.96.221 [AS 4837] 380 msec 396 msec 392 msec 13 219.158.96.229 [AS 4837] 384 msec 376
msec 376 msec 14 219.158.10.38 [AS 4837] 368 msec 372 msec 372 msec 15 120.84.0.50 [AS 17816]
380 msec 388 msec 376 msec 16 3.3.3.66 [AS 17622] 384 msec 380 msec 404 msec=====> [
(1.1.1.48)R1, (1.1.1.48) R1 gigabitEthernet 0/1/0/6.313 (1.1.1.49) R2 gigabitEthernet
0/1/0/6.313 ] 17 * * * 18 * * * 19 * * * 20 * * * 21 * * * 22 * * * 23 * * * 24 * * * 25 * * *
26 * * * 27 * * * 28 * * * 29 * * * 30 * * *
```

**检查: r2 的gigabitEthernet 0/1/0/6.313**

```
RP/0/RSP0/CPU0:r2#show cef ipv4 drops location 0/1/CPU0 | inc RPF Tue Mar 27 10:06:02.511 UTC
RPF drops packets : 38410840 RPF suppressed drops packets : 0
RP/0/RSP0/CPU0:r2# RP/0/RSP0/CPU0:r2#show cef ipv4 drops location 0/1/CPU0 | inc RPF Tue Mar
27 10:06:09.591 UTC RPF drops packets : 38412257 RPF suppressed drops packets
: 0 RP/0/RSP0/CPU0:r2#
```

**B.关于到2.2.2.130 不通的问题:**

```
route-server>trace 2.2.2.130
Type escape sequence to abort.
Tracing the route to 2.2.2.130
 1 gateway.cbbtier3.att.net (12.0.1.202) [AS 7018] 0 msec 0 msec 0 msec
 2 n54ny401me3-cbbtier3.ip.att.net (12.89.5.13) [AS 7018] 8 msec 16 msec 16 msec
 3 cr1.n54ny.ip.att.net (12.123.2.6) [MPLS: Label 16092 Exp 1] 80 msec 108 msec 84 msec
 4 cr2.cgcil.ip.att.net (12.122.1.2) [MPLS: Labels 23256/16494 Exp 1] 80 msec 80 msec 80 msec
 5 cr1.cgcil.ip.att.net (12.122.2.53) [MPLS: Labels 21629/16494 Exp 1] 80 msec 80 msec 80
msec
 6 cr2.dvmco.ip.att.net (12.122.31.85) [MPLS: Labels 21370/16494 Exp 1] 84 msec 80 msec 80
msec
 7 cr1.slkut.ip.att.net (12.122.30.25) [MPLS: Labels 20076/16494 Exp 1] 80 msec 76 msec 80
msec
 8 cr2.la2ca.ip.att.net (12.122.30.30) [MPLS: Labels 0/16494 Exp 1] 80 msec 80 msec 80 msec
 9 cr84.la2ca.ip.att.net (12.123.30.249) [MPLS: Labels 0/16333 Exp 1] 76 msec 72 msec 84 msec
10 gar2.lsrca.ip.att.net (12.122.129.49) 80 msec 76 msec 80 msec
11 12.118.130.86 [AS 7018] 316 msec 316 msec 320 msec
12 219.158.97.9 [AS 4837] 304 msec 296 msec 304 msec
13 219.158.11.153 [AS 4837] 296 msec 284 msec 280 msec
14 219.158.19.82 [AS 4837] 284 msec * 288 msec
15 120.82.0.150 [AS 17816] 288 msec 292 msec 292 msec
16 3.3.3.74 [AS 17622] 304 msec 304 msec 300 msec=====>
[
(1.1.1.49) R2, (1.1.1.49)R2gigabitEthernet 0/1/0/6.313
(1.1.1.48)R1gigabitEthernet 0/1/0/6.313
]
17 * * *
18 * * *
19 * * *
20 * * *
21 * * *
22 * * *
23 * * *
24 * * *
25 * * *
26 * * *
```

```
27 * * *
28 * * *
29 * * *
30 * * *
route-server>
```

和问题A是对称的

### C. 至于为什么到虚拟VRRP地址2.2.2.129 可以ping通

是因为数据包没有绕gigabitEthernet 0/1/0/6.313

```
route-server>trace 2.2.2.129
|
Type escape sequence to abort.
Tracing the route to 2.2.2.129
|
 1 gateway.cbbtier3.att.net (12.0.1.202) [AS 7018] 4 msec 0 msec 4 msec
 2 n54ny401me3-cbbtier3.ip.att.net (12.89.5.13) [AS 7018] 4 msec 0 msec 0 msec
 3 cr1.n54ny.ip.att.net (12.123.2.6) [MPLS: Label 16092 Exp 1] 72 msec 68 msec 76 msec
 4 cr2.cgcil.ip.att.net (12.122.1.2) [MPLS: Labels 23256/16494 Exp 1] 72 msec 72 msec 72 msec
 5 cr1.cgcil.ip.att.net (12.122.2.53) [MPLS: Labels 21629/16494 Exp 1] 76 msec 72 msec 72
msec
 6 cr2.dvmco.ip.att.net (12.122.31.85) [MPLS: Labels 21370/16494 Exp 1] 72 msec 72 msec 72
msec
 7 cr1.slkut.ip.att.net (12.122.30.25) [MPLS: Labels 20076/16494 Exp 1] 76 msec 76 msec 72
msec
 8 cr2.la2ca.ip.att.net (12.122.30.30) [MPLS: Labels 0/16494 Exp 1] 72 msec 72 msec 72 msec
 9 cr84.la2ca.ip.att.net (12.123.30.249) [MPLS: Labels 0/16333 Exp 1] 72 msec 72 msec 72 msec
10 gar2.lsrca.ip.att.net (12.122.129.49) 68 msec 72 msec 72 msec
11 12.118.130.86 [AS 7018] 268 msec 268 msec 272 msec
12 219.158.96.245 [AS 4837] 268 msec 272 msec 276 msec
13 219.158.3.121 [AS 4837] 252 msec 256 msec 256 msec
14 219.158.19.86 [AS 4837] 260 msec 256 msec 256 msec
15 120.84.0.34 [AS 17816] 336 msec 344 msec 336 msec
16 * *
 3.3.3.66 [AS 17622] 284 msec=====>
route-server>
```

## 问题总结

出现上面的原因是因为上联某台路由器到

58.248.19.128/28

是负载均衡的, 路由器选择哪条路径就由 CEF HASH的结果得出:

HASH 的因子 包括 (源地址 + 目的地址+.....)

```
12.0.1.x,2.2.2.130
12.0.1.x,2.2.2.131
```

这两对虽然SOURCE 一样 ,但是destination不同

HASH 到不同link上了

## 经验总结

loose mode urpf 的规则:

1. loose mode 只查路由表中有无匹配，不查进入接口.
2. 但是如果source 所属是本地直连,loose mode 也要检查进入接口.
3. 默认是不会用default route 作urpf的 的检查依据的,需要allow-default 开启

## 相关命令

```
show ip route  
show cef ipv4 drops location x/x/x  
traceroute x.x.x.x
```