某医科大学客户SR6608路由器直连偶尔丢包问题

丢包 **傅昆** 2014-09-24 发表

一　组网：

二　问题描述：

某医科大学客户反馈，我司SR6608(版本为Release 3303P10)作为出口，内网SR6608与juniperMX960直连，近期网络出现不定期无规律性大量丢包。

为排查问题，在SR6608与MX960之间添加一台S5800交换机，并在S5800上做流统计。T1/0/27口匹配源地址10.0.0.90 ，目的地址10.0.0.91的ICMP包，T1/0/25口匹配源地址10.0.0.91目的地址10.0.0.90的ICMP包。

故障发生时时，在SR6608(10.0.0.91)上PING MX960(10.0.0.90) 300个包，通过S5800上流统计发现SR66的ping包发到MX960上，并且MX960回复了SR66，但在SR66上显示有15%的丢包。

三　过程分析：

通过分析SR6608的诊断信息和日志文件，发现如下异常信息：

Ten-GigabitEthernet5/1/0 current state: UP
Line protocol current state: UP
Description: to_MX960.bak
The Maximum Transmit Unit is 1500
Internet Address is 10.0.0.91/29 Primary
IP Packet Frame Type: PKTFMT_ETHNT_2, Hardware Address: 7425-8a37-9107
IPv6 Packet Frame Type: PKTFMT_ETHNT_2, Hardware Address: 7425-8a37-9107
Media type is optical fiber, loopback not set, promiscuous mode not set
10G, Full, link type is force link
Output flow-control is disabled, input flow-control is disabled
XFP Transceiver Info:
Vendor name: H3C Port hardware type: 10G_BASE_SR_XFP
Ordering Name:
Output queue : (Urgent queuing : Size/Length/Discards) 0/100/0
Output queue : (Protocol queuing : Size/Length/Discards) 0/500/0
Output queue : (FIFO queuing : Size/Length/Discards) 0/1024/0
Ethernet port mode: LAN
Last clearing of counters: 14:15:32 Mon 09/22/2014
  Last 300 seconds input rate 31365648.00 bytes/sec, 250925184 bits/sec, 327737.37 packets/sec
  Last 300 seconds output rate 35101172.00 bytes/sec, 280809376 bits/sec, 35228.91 packets/sec
  Input: 319793474 packets, 57076296928 bytes, 42418703 no buffers
    2 broadcasts, 255 multicasts, 0 pauses
    0 errors, 0 runts, 0 giants
    0 crc, 0 align errors, 0 overruns
    0 dribbles, 0 drops
  Output:74979687 packets, 75239231406 bytes
    0 broadcasts, 166 multicasts, 0 pauses
    0 errors, 0 underruns, 0 collisions
    0 deferred, 0 lost carriers

以上黑体字信息表明，SR6608很可能收到了异常报文，导致某CPU线程异常繁忙，甚至超出了其处理能力，设备上曾有接口流量超带宽的情况发生：

more logfile/logfile10.log | i peak
%@131665%Sep 16 20:08:48:453 2014 AYD-SR6608 DMON/7/MSG: -Slot=3; GigabitEthernet3/2/1 output rate peak : 113173750 Bps
%@132790%Sep 17 16:10:17:627 2014 AYD-SR6608 DMON/7/MSG: -Slot=3; GigabitEthernet3/2/1 output rate peak : 122721240 Bps
%@135665%Sep 19 18:00:45:878 2014 AYD-SR6608 DMON/7/MSG: -Slot=3; GigabitEthernet3/2/1 output rate peak : 122731708 Bps
%@135667%Sep 19 18:01:45:877 2014 AYD-SR6608 DMON/7/MSG: -Slot=3; GigabitEthernet3/2/1 output rate peak : 122735479 Bps
%@135668%Sep 19 18:02:15:876 2014 AYD-SR6608 DMON/7/MSG: -Slot=3; GigabitEthernet3/2/1 output rate peak : 122743324 Bps
%@135669%Sep 19 18:02:45:876 2014 AYD-SR6608 DMON/7/MSG: -Slot=3; GigabitEthernet3/2/1 output rate peak : 122744331 Bps
%@135670%Sep 19 18:04:45:875 2014 AYD-SR6608 DMON/7/MSG: -Slot=3; GigabitEthernet3/2/1

output rate peak : 122746834 Bps
%@135671%Sep 19 18:05:15:874 2014 AYD-SR6608 DMON/7/MSG: -Slot=3; GigabitEthernet3/2/1
output rate peak : 122749913 Bps
%@135851%Sep 19 22:49:15:666 2014 AYD-SR6608 DMON/7/MSG: -Slot=3; GigabitEthernet3/2/3
output rate peak : 103651020 Bps
%@135853%Sep 19 22:49:45:666 2014 AYD-SR6608 DMON/7/MSG: -Slot=3; GigabitEthernet3/2/3
output rate peak : 120973568 Bps

正常情况下，该接口入方向每秒十几万个包：
[AYD-SR6608-hidecmd]dis counters rate in int

Interface	Total(pkts/sec)	Broadcast(pkts/sec)	Multicast(pkts/sec)
GE3/1/1	1358	--	--
GE3/1/2	18481	--	--
GE3/2/1	6858	--	--
GE3/2/3	114479	--	--
XGE5/1/0	119519	--	--

但是出问题的时候，该接口有几十万个包，而且都是小包：
[AYD-SR6608-hidecmd]dis cou ra in int

Interface	Total(pkts/sec)	Broadcast(pkts/sec)	Multicast(pkts/sec)
GE3/1/1	133	--	--
GE3/1/2	1571	--	--
GE3/2/1	1489	--	--
GE3/2/3	11103	--	--
XGE5/1/0	666727	--	--

[AYD-SR6608-hidecmd]display interface ten 5/1/0
Ten-GigabitEthernet5/1/0 current state: UP
Line protocol current state: UP
Description: to_MX960.bak
The Maximum Transmit Unit is 1500
Internet Address is 10.0.0.91/29 Primary
IP Packet Frame Type: PKTFMT_ETHNT_2, Hardware Address: 7425-8a37-9107
IPv6 Packet Frame Type: PKTFMT_ETHNT_2, Hardware Address: 7425-8a37-9107
Media type is optical fiber, loopback not set, promiscuous mode not set
10G, Full, link type is force link
Output flow-control is disabled, input flow-control is disabled
XFP Transceiver Info:
Vendor name: H3C   Port hardware type: 10G_BASE_SR_XFP
Ordering Name:
Output queue : (Urgent queuing : Size/Length/Discards) 0/100/0
Output queue : (Protocol queuing : Size/Length/Discards) 0/500/0
Output queue : (FIFO queuing : Size/Length/Discards) 0/1024/0
Ethernet port mode: LAN
Last clearing of counters: 19:59:01 Mon 09/22/2014
  Last 5 seconds input rate 52079560.00 bytes/sec, 416636480 bits/sec, 654111.
18 packets/sec
  Last 5 seconds output rate 4357302.00 bytes/sec, 34858416 bits/sec, 11164.20
packets/sec
  Input: 611440361 packets, 105768981789 bytes, 52462056 no buffers
      22 broadcasts, 487 multicasts, 0 pauses
      0 errors, 0 runts, 0 giants
      0 crc, 0 align errors, 0 overruns
      0 dribbles, 0 drops
  Output:258973728 packets, 264481015964 bytes
      0 broadcasts, 318 multicasts, 0 pauses
      0 errors, 0 underruns, 0 collisions
      0 deferred, 0 lost carriers

为了进一步确认攻击流量的类型，在该接口配置了net stream和firewall（相关配置详见相关配置指导手册）
在出现问题的时候，收集了如下信息：
_hidecmd
dsplay session statistics slot 5

display ip net ca slot 5
display interface ten 5/1/0
display cout rate in int
tshow mainboard slot 5 vcpu all cpu

说明：上述命令连续执行5遍，其中dis ip net ca slot 5显示内容较多，至少显示如下内容

```
[AYD-SR6608-hidecmd]dis ip net ca slot 5
IP netstream cache information:
  Stream active timeout (in seconds)  : 1800
  Stream inactive timeout (in seconds): 30
  Stream max entry number             : 2600000
    IP active stream entry number     : 2283142
  MPLS active stream entry number     : 0
    L2 active stream entry number     : 0
  IPL2 active stream entry number     : 0
    IP stream entry been counted      : 6416180
  MPLS stream entry been counted      : 0
    L2 stream entry been counted      : 0
  IPL2 stream entry been counted      : 0
  Last statistics reset time          : Never

IP packet size distribution (151524745 total packets):
 1-32   64   96  128  160  192  224  256  288  320  352  384  416  448  480
 .002 .580 .189 .035 .009 .003 .002 .002 .002 .001 .001 .001 .001 .001 .001

  512  544  576 1024 1536 2048 2560 3072 3584 4096 4608 >4608
 .001 .001 .006 .009 .144 .000 .000 .000 .000 .000 .000 .000

  Protocol         Total   Packets   Stream  Packets Active(sec) Idle(sec)
                 Streams      /Sec     /Sec  /stream   /stream    /stream
-------------------------------------------------------------------------------
  IP-other            21         0       0       8        25        153
  UDP-other       1915615        51       2      18        19        169
  TCP-other        549296        12       0      16        11        102
  ICMP             19514         0       0      16        56        132
  GRE                  6         0       0     394        57        254
  UDP-Echo          2038         0       0       1         0        182
  UDP-Time             1         0       0      15        71        164
  UDP-DNS        2613162        28       3       7        39        179
  UDP-NTP           1283         0       0       1         4        146
  UDP-NetBios         97         0       0       6        11        117
  UDP-SNMP           207         0       0       3        10        144
  UDP-XDMCP            2         0       0      14        48         68
  UDP-Biff             2         0       0       1         0        187
  UDP-Who              1         0       0       1         0        336
  UDP-SysLog          11         0       0       1         0        202
  UDP-Talk             1         0       0       1         0        183
  TCP-Echo            54         0       0       1         0         55
  TCP-Daytime         20         0       0       3        19         52
  TCP-FTP            581         0       0       3         7        102
  TCP-Telnet         129         0       0       2         8        143
  TCP-SMTP            90         0       0       9         9         43
  TCP-Time             7         0       0       2         3         93
  TCP-WHOIS            1         0       0       1         0        362
  TCP-DNS            673         0       0       7         8        115
  TCP-HTTP       1313309        42       1      22        12         43
  TCP-Hostname         1         0       0       1         0         34
  TCP-POP3            58         0       0      35         7         54

Type DstIP(Port)            SrcIP(Port)            Pro ToS If(Direct)  Pkts
     DstMAC(VLAN)           SrcMAC(VLAN)
     TopLblType(IP/MASK)    Lbl-Exp-S-List
-------------------------------------------------------------------------------
IP   119.60.206.229(4784)   10.23.37.203(14445)    17  0   XGE5/1/0(I) 2
```

```
IP   116.9.78.230(3784)     10.3.5.222(8091)       17  0   XGE5/1/0(I) 1
IP   222.210.108.43(62687)  10.24.68.242(9073)     17  0   XGE5/1/0(I) 2
IP   208.109.255.29(53)     33.68.186.228(1297)    17  0   XGE5/1/0(I) 14
IP   208.109.255.29(53)     101.246.216.79(3863)   17  0   XGE5/1/0(I) 14
IP   216.69.185.29(53)      157.228.87.92(3680)    17  0   XGE5/1/0(I) 15
IP   208.109.255.29(53)     139.120.216.39(3078)   17  0   XGE5/1/0(I) 13
IP   216.69.185.29(53)      47.189.206.1(3535)     17  0   XGE5/1/0(I) 14
IP   208.109.255.29(53)     153.152.251.64(3105)   17  0   XGE5/1/0(I) 14
IP   119.75.220.50(80)      10.25.44.250(62422)    6   0   XGE5/1/0(I) 6
IP   216.69.185.29(53)      125.107.120.15(1176)   17  0   XGE5/1/0(I) 17
IP   208.109.255.29(53)     174.21.174.235(3826)   17  0   XGE5/1/0(I) 14
IP   208.109.255.29(53)     63.94.59.191(2038)     17  0   XGE5/1/0(I) 15
  ---- More ----
```

（最后，执行display ip netstream ca slot 5，可以尽量多显示一些）

根据以上命令搜集的信息分析，故障发生的时候，G5/1/0接口上瞬间有大量的DNS攻击报文（详见黑体字数值）：

```
UDP     Session establishment rate:    113499/s            每秒钟有十几万以上的udp报文
UDP-DNS       183678834      1104      246       4        7         21    netstream统计到的这种报文最多

[AYD-SR6608-hidecmd]display  session  statistic  slot 5
Current session(s):2000008
        Current     TCP session(s): 60480
              Half-Open: 9054         Half-Close: 8264
        Current     UDP session(s): 1939073
        Current    ICMP session(s): 455
        Current   RAWIP session(s): 0

Current relation table(s): 0

Session establishment rate:    115137/s
        TCP     Session establishment rate:      1630/s
        UDP     Session establishment rate:    113499/s
        ICMP    Session establishment rate:        7/s
        RAWIP   Session establishment rate:        1/s

Received     TCP:          14557145582 packet(s)        11378169511957 byte(s)
Received     UDP:          20298838172 packet(s)        11435916580045 byte(s)
Received    ICMP:            185821464 packet(s)          16773402598 byte(s)
Received   RAWIP:            152582367 packet(s)          51844314116 byte(s)
Dropped      TCP:                    0 packet(s)                  0 byte(s)
Dropped      UDP:                    0 packet(s)                  0 byte(s)
Dropped     ICMP:                    0 packet(s)                  0 byte(s)
Dropped    RAWIP:                     0 packet(s)                   0 byte(s)

[AYD-SR6608-hidecmd]display  ip  netstream cache  slot 5
IP netstream cache information:
  Stream active timeout (in seconds)  : 1800
  Stream inactive timeout (in seconds): 30
  Stream max entry number             : 2600000
    IP active stream entry number     : 2592834
  MPLS active stream entry number     : 0
    L2 active stream entry number     : 0
  IPL2 active stream entry number     : 0
    IP stream entry been counted      : 296131864
  MPLS stream entry been counted      : 0
    L2 stream entry been counted      : 0
  IPL2 stream entry been counted      : 0
  Last statistics reset time          : Never

IP packet size distribution (3400984685 total packets):
 1-32   64   96  128  160  192  224  256  288  320  352  384  416  448  480
 .002 .559 .164 .032 .009 .003 .002 .002 .002 .001 .001 .001 .001 .001 .001
```

512 544 576 1024 1536 2048 2560 3072 3584 4096 4608 >4608
.001 .001 .008 .010 .191 .000 .000 .000 .000 .000 .000 .000

Protocol	Total Streams	Packets /Sec	Stream /Sec	Packets /stream	Active(sec) /stream	Idle(sec) /stream
IP-other	952	0	0	8	20	67
UDP-other	61305624	1745	82	21	12	0
TCP-other	19257644	551	25	21	9	41
ICMP	561462	24	0	32	61	56
OSPF	69	0	0	96	715	17
GRE	478	8	0	13242	416	36
UDP-Echo	74384	0	0	1	0	61
UDP-Time	18	0	0	2	5	132
UDP-NameSev	13	0	0	2	7	30
UDP-TACACS	11	0	0	2	6	68
UDP-DNS	184484753	1107	247	4	7	21
UDP-TFTP	10	0	0	1	5	44
UDP-SunRPC	8	0	0	1	2	30
UDP-NTP	46970	0	0	1	2	57
UDP-NetBios	4827	0	0	12	11	58
UDP-SNMP	9519	0	0	2	5	52
UDP-XDMCP	12	0	0	3	9	59
UDP-DNSIX	4	0	0	3	18	30
UDP-MobileIP	6	0	0	1	0	84
UDP-Biff	63	0	0	2	19	49
UDP-Who	6	0	0	1	0	164
UDP-SysLog	115	0	0	1	1	114
UDP-Talk	19	0	0	4	1	67
UDP-RIP	11	0	0	1	0	103
TCP-Echo	1481	0	0	1	0	103
TCP-Daytime	1222	0	0	2	7	30
TCP-CHARgen	3	0	0	2	6	30
TCP-FTP-data	58	0	0	72	19	37
TCP-FTP	14172	0	0	3	8	47
TCP-Telnet	6198	0	0	3	9	50
TCP-SMTP	1514	0	0	27	3	19
TCP-Time	461	0	0	2	2	26
TCP-WHOIS	30	0	0	2	8	41
TCP-DNS	41310	0	0	13	6	58
TCP-Gopher	151	0	0	4	1	15
TCP-Finger	159	0	0	9	12	31
TCP-HTTP	30315030	1113	40	27	8	25
TCP-Hostname	40	0	0	4	6	53
TCP-POP3	2866	0	0	13	3	24
TCP-Ident	5	0	0	3	9	30
TCP-NNTP	38	0	0	2	4	50
TCP-BGP	15	0	0	2	3	30
TCP-IRC	8	0	0	3	9	71
TCP-RshExec	81	0	0	453	205	44
TCP-RCMD	15	0	0	2	13	47
TCP-Talk	14	0	0	6	1	30
TCP-KShell	15	0	0	3	7	47

后来，Juniper工程师到现场配合问题分析，发现Juniper设备上部分端口有大量丢包情况，抓包后发现大量伪造IP的DNS攻击报文。初步定位为：内网部分服务器中中病毒导致网络终端情况发生。Juniper工程师已在其设备上进行了部分数据过滤。

四　解决方法：

已经明确了是DNS攻击导致，规避方法是删除SR6608路由器DNS相关的配置；根本的解决方法是，查出攻击源，在安全设备上阻止相关的攻击行为。