MPLS L2VPN OVER GRE OVER IPSEC

ipoe IPsec MPLS L2VPN 肖倩艳 2015-07-22 发表

客户网络有总部A和分支B分别通过MSR5660路由器接入互联网，想通过mpls l2vpn将两个网络打通，实现两个局域网互访，且需要加密流量。

客户组网拓扑图大致如下：
此需求将使用MPLS L2VPN（ldp pw）之 GRE over IPSEC实现，通过GRE over IPSEC公网隧道来承载PW。

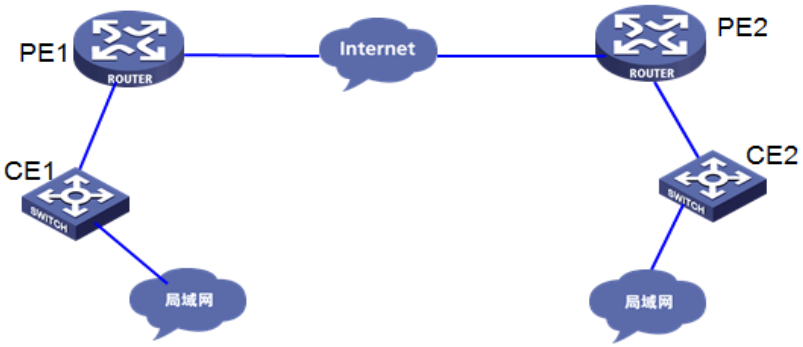

PE1配置

```
#
 sysname pe1
#
mpls lsr-id 3.3.3.3
#
mpls ldp
#
 l2vpn enable
#
interface LoopBack0
 description gre
 ip address 1.1.1.1 255.255.255.255
#
interface LoopBack1
 description ldp
 ip address 3.3.3.3 255.255.255.255
#
interface GigabitEthernet0/0
 port link-mode route
 combo enable copper
 ip address 200.1.1.2 255.255.255.252
 ipsec apply policy 1
#
interface GigabitEthernet0/1
 port link-mode route
 combo enable copper
 description to-ce1
#
interface GigabitEthernet0/1.110
 vlan-type dot1q vid 110
#
interface Tunnel10 mode gre
 ip address 5.5.5.1 255.255.255.252
 mpls enable
 mpls ldp enable
 source LoopBack0
```

```
 destination 2.2.2.2
#
xconnect-group vpn2
 connection ldp
  ac interface GigabitEthernet0/1.110
  peer 4.4.4.4 pw-id 801001111
#
 ip route-static 0.0.0.0 0 200.1.1.1
 ip route-static 4.4.4.4 32 Tunnel10 ////到ldp peer 走tunnel 10口
#
acl advanced 3002
 rule 0 permit ip source 1.1.1.1 0 destination 2.2.2.2 0 ///封装gre的源和目的触发建立ipsec
#
ipsec transform-set cdgac
 esp encryption-algorithm 3des-cbc
 esp authentication-algorithm sha1
#
ipsec policy 1 2 isakmp
 transform-set cdgac
 security acl 3002
 local-address 200.1.1.2
 remote-address 201.1.1.2
 ike-profile cdgac
#
ike profile cdgac
 keychain cdgac
 match remote identity address 201.1.1.2 255.255.255.252
 proposal 2
#
ike proposal 2
 encryption-algorithm aes-cbc-128
 dh group2
#
ike keychain cdgac
 pre-shared-key address 201.1.1.2 255.255.255.252 key cipher $c$3$XUQhTUr370G91QQqpi2T88F
DJcPtvg==
#
PE2配置
#
 sysname pe2
#
 mpls lsr-id 4.4.4.4
#
mpls ldp
#
 l2vpn enable
#
interface LoopBack0
 description GRE
 ip address 2.2.2.2 255.255.255.255
#
interface LoopBack1
 description LDP
 ip address 4.4.4.4 255.255.255.255
#
interface GigabitEthernet0/0
 port link-mode route
 combo enable copper
 ip address 201.1.1.2 255.255.255.252
 ipsec apply policy cdgac
#
interface GigabitEthernet0/1
 port link-mode route
 combo enable copper
```

```
 description to-ce2

#
interface GigabitEthernet0/1.110
 vlan-type dot1q vid 110
#
interface Tunnel10 mode gre
 ip address 5.5.5.2 255.255.255.252
 mpls enable
 mpls ldp enable
 source loopback0
 destination 1.1.1.1
#
xconnect-group vpn2
 connection ldp
  ac interface GigabitEthernet0/1.110
  peer 3.3.3.3 pw-id 801001111
#
 ip route-static 0.0.0.0 0 201.1.1.1
 ip route-static 3.3.3.3 32 Tunnel10
#
acl advanced 3002
 rule 0 permit ip source 2.2.2.2 0 destination 1.1.1.1 0
#
ipsec transform-set cdgac
 esp encryption-algorithm 3des-cbc
 esp authentication-algorithm sha1
#
ipsec policy cdgac 2 isakmp
 transform-set cdgac
 security acl 3002
 local-address 201.1.1.2
 remote-address 200.1.1.2
 ike-profile cdgac
#
ike profile cdgac
 keychain cdgac
 match remote identity address 200.1.1.2 255.255.255.252
 proposal 2
#
ike proposal 2
 encryption-algorithm aes-cbc-128
 dh group2
#
ike keychain cdgac
 pre-shared-key address 200.1.1.2 255.255.255.252 key cipher $c$3$uVIpwExz145rpaEPkx8RrzB0q
Nwktg==
#
Ce1配置
#
 sysname ce1
#
vlan 110
#
interface Vlan-interface110
 ip address 10.1.1.1 255.255.255.0
#
interface GigabitEthernet1/0/1
 port link-mode bridge
 port link-type trunk
 port trunk permit vlan 1 110
 combo enable fiber
#
Ce2配置
```

```
#
 sysname ce2
#
vlan 110
#
interface Vlan-interface110
 ip address 10.1.1.2 255.255.255.0
#
interface GigabitEthernet1/0/1
 port link-mode bridge
 port link-type trunk
 port trunk permit vlan 1 110
 combo enable fiber
#
```

测试结果

pe1侧gre触发ipsec建立成功，ldp peer地址流量走gre隧道，l2vpn pw 状态up

```
dis ike sa
   Connection-ID   Remote               Flag        DOI
-------------------------------------------------------------------------
   1               201.1.1.2            RD          IPsec
Flags:
RD--READY RL--REPLACED FD-FADING
dis ipsec sa
-------------------------------
Interface: GigabitEthernet0/0
-------------------------------


  ----------------------------
  IPsec policy: 1
  Sequence number: 2
  Mode: ISAKMP
  ----------------------------
    Tunnel id: 0
    Encapsulation mode: tunnel
    Perfect forward secrecy:
    Path MTU: 1443
    Tunnel:
        local  address: 200.1.1.2
        remote address: 201.1.1.2
    Flow:
        sour addr: 1.1.1.1/255.255.255.255  port: 0  protocol: ip
        dest addr: 2.2.2.2/255.255.255.255  port: 0  protocol: ip

    [Inbound ESP SAs]
      SPI: 2495663367 (0x94c0cd07)
      Connection ID: 4294967296
      Transform set:  ESP-ENCRYPT-3DES-CBC ESP-AUTH-SHA1
      SA duration (kilobytes/sec): 1843200/3600
      SA remaining duration (kilobytes/sec): 1843137/1966
      Max received sequence-number: 709
      Anti-replay check enable: Y
      Anti-replay window size: 64
      UDP encapsulation used for NAT traversal: N
      Status: Active

    [Outbound ESP SAs]
      SPI: 2673009478 (0x9f52e346)
      Connection ID: 4294967297
      Transform set:  ESP-ENCRYPT-3DES-CBC ESP-AUTH-SHA1
      SA duration (kilobytes/sec): 1843200/3600
      SA remaining duration (kilobytes/sec): 1843137/1966
      Max sent sequence-number: 711
      UDP encapsulation used for NAT traversal: N
      Status: Active
```

```
dis ip int b
*down: administratively down
(s): spoofing  (l): loopback
Interface              Physical Protocol IP Address      Description
GE0/0                  up       up       200.1.1.2       --
GE0/1                  up       up       --              to-ce1
GE0/1.110                up       up       --              --
GE0/2                  down     down     192.168.3.1     --
GE5/0                  down     down     --              --
GE5/1                  down     down     --              --
GE6/0                  down     down     --              --
GE6/1                  down     down     --              --
Loop0                  up       up(s)    1.1.1.1         gre
Loop1                  up       up(s)    3.3.3.3         ldp
Ser1/0                 down     down     --              --
Ser2/0                 down     down     --              --
Ser3/0                 down     down     --              --
Ser4/0                 down     down     --              --
Tun10                  up       up       5.5.5.1         --
ping -a 5.5.5.1 5.5.5.2
Ping 5.5.5.2 (5.5.5.2) from 5.5.5.1: 56 data bytes, press CTRL_C to break
56 bytes from 5.5.5.2: icmp_seq=0 ttl=255 time=7.244 ms
56 bytes from 5.5.5.2: icmp_seq=1 ttl=255 time=2.576 ms
56 bytes from 5.5.5.2: icmp_seq=2 ttl=255 time=2.429 ms
56 bytes from 5.5.5.2: icmp_seq=3 ttl=255 time=2.397 ms
56 bytes from 5.5.5.2: icmp_seq=4 ttl=255 time=2.826 ms
--- Ping statistics for 5.5.5.2 ---
5 packets transmitted, 5 packets received, 0.0% packet loss
round-trip min/avg/max/std-dev = 2.397/3.494/7.244/1.881 ms
%Jul 22 08:59:53:871 2015 pe1 PING/6/PING_STATISTICS: Ping statistics for 5.5.5.2: 5 packets tra
nsmitted, 5 packets received, 0.0% packet loss, round-trip
min/avg/max/std-dev = 2.397/3.494/7.244/1.881 ms.
dis l2vpn pw
Flags: M - main, B - backup, H - hub link, S - spoke link, N - no split horizon
Total number of PWs: 1
1 up, 0 blocked, 0 down, 0 defect, 0 idle, 0 duplicate

Xconnect-group Name: vpn2
Peer           PW ID/Rmt Site    In/Out Label    Proto   Flag  Link ID  State
4.4.4.4        801001111         917631/917631   LDP     M     1        Up
Pe2侧测试结果同pe1
dis ike sa
   Connection-ID   Remote               Flag        DOI
-------------------------------------------------------------------------
   1               200.1.1.2            RD          IPsec
Flags:
RD--READY RL--REPLACED FD-FADING
dis ipsec sa
-------------------------------
Interface: GigabitEthernet0/0
-------------------------------
 -------------------------------
 IPsec policy: cdgac
 Sequence number: 2
 Mode: ISAKMP
 -------------------------------
   Tunnel id: 0
   Encapsulation mode: tunnel
   Perfect forward secrecy:
   Path MTU: 1443
   Tunnel:
      local  address: 201.1.1.2
      remote address: 200.1.1.2
   Flow:
```

```
    sour addr: 2.2.2.2/255.255.255.255  port: 0  protocol: ip
    dest addr: 1.1.1.1/255.255.255.255  port: 0  protocol: ip

  [Inbound ESP SAs]
    SPI: 2673009478 (0x9f52e346)
    Connection ID: 4294967296
    Transform set:  ESP-ENCRYPT-3DES-CBC ESP-AUTH-SHA1
    SA duration (kilobytes/sec): 1843200/3600
    SA remaining duration (kilobytes/sec): 1843136/1896
    Max received sequence-number: 735
    Anti-replay check enable: Y
    Anti-replay window size: 64
    UDP encapsulation used for NAT traversal: N
    Status: Active

  [Outbound ESP SAs]
    SPI: 2495663367 (0x94c0cd07)
    Connection ID: 4294967297
    Transform set:  ESP-ENCRYPT-3DES-CBC ESP-AUTH-SHA1
    SA duration (kilobytes/sec): 1843200/3600
    SA remaining duration (kilobytes/sec): 1843135/1896
    Max sent sequence-number: 733
    UDP encapsulation used for NAT traversal: N
    Status: Active
  dis ip int b
*down: administratively down
(s): spoofing  (l): loopback
Interface              Physical Protocol IP Address      Description
GE0/0                  up       up       201.1.1.2       --
GE0/1                  up       up       --              --
GE0/1.110                up       up       --              --
GE0/2                  down     down     192.168.2.1     --
GE5/0                  down     down     --              --
GE5/1                  down     down     --              --
GE6/0                  down     down     --              --
GE6/1                  down     down     --              --
Loop0                  up       up(s)    2.2.2.2         GRE
Loop1                  up       up(s)    4.4.4.4         LDP
Ser1/0                 down     down     --              --
Ser2/0                 down     down     --              --
Ser3/0                 down     down     --              --
Ser4/0                 down     down     --              --
Tun10                  up       up       5.5.5.2         --
ping -a 5.5.5.2 5.5.5.1
Ping 5.5.5.1 (5.5.5.1) from 5.5.5.2: 56 data bytes, press CTRL_C to break
56 bytes from 5.5.5.1: icmp_seq=0 ttl=255 time=5.598 ms
56 bytes from 5.5.5.1: icmp_seq=1 ttl=255 time=3.794 ms
56 bytes from 5.5.5.1: icmp_seq=2 ttl=255 time=3.066 ms
56 bytes from 5.5.5.1: icmp_seq=3 ttl=255 time=2.787 ms
56 bytes from 5.5.5.1: icmp_seq=4 ttl=255 time=3.242 ms
--- Ping statistics for 5.5.5.1 ---
5 packets transmitted, 5 packets received, 0.0% packet loss
round-trip min/avg/max/std-dev = 2.787/3.697/5.598/1.006 ms
%Jul 22 08:59:24:816 2015 pe2 PING/6/PING_STATISTICS: Ping statistics for 5.5.5.1: 5 packets tra
nsmitted, 5 packets received, 0.0% packet loss, round-trip min/avg/max/std-dev =
2.787/3.697/5.598/1.006 ms.
dis l2vpn pw
Flags: M - main, B - backup, H - hub link, S - spoke link, N - no split horizon
Total number of PWs: 1
1 up, 0 blocked, 0 down, 0 defect, 0 idle, 0 duplicate

Xconnect-group Name: vpn2
Peer           PW ID/Rmt Site    In/Out Label    Proto   Flag  Link ID  State
3.3.3.3        801001111         917631/917631   LDP     M     1        Up
```

最终需求，两个ce网络可达
[ce1]ping -a 10.1.1.1 10.1.1.2
Ping 10.1.1.2 (10.1.1.2) from 10.1.1.1: 56 data bytes, press CTRL_C to break
56 bytes from 10.1.1.2: icmp_seq=0 ttl=255 time=12.646 ms
56 bytes from 10.1.1.2: icmp_seq=1 ttl=255 time=7.242 ms
56 bytes from 10.1.1.2: icmp_seq=2 ttl=255 time=7.387 ms
56 bytes from 10.1.1.2: icmp_seq=3 ttl=255 time=6.654 ms
56 bytes from 10.1.1.2: icmp_seq=4 ttl=255 time=6.986 ms

--- Ping statistics for 10.1.1.2 ---
5 packets transmitted, 5 packets received, 0.0% packet loss
round-trip min/avg/max/std-dev = 6.654/8.183/12.646/2.245 ms
[ce1]%Jul 22 08:55:50:880 2015 ce1 PING/6/PING_STATISTICS: Ping statistics for 10.1.1.2: 5 packe
ts transmitted, 5 packets received, 0.0% packet loss, round-trip
min/avg/max/std-dev = 6.654/8.183/12.646/2.245 ms.

ping -a 10.1.1.2 10.1.1.1
Ping 10.1.1.1 (10.1.1.1) from 10.1.1.2: 56 data bytes, press CTRL_C to break
56 bytes from 10.1.1.1: icmp_seq=0 ttl=255 time=95.774 ms
56 bytes from 10.1.1.1: icmp_seq=1 ttl=255 time=14.967 ms
56 bytes from 10.1.1.1: icmp_seq=2 ttl=255 time=11.184 ms
56 bytes from 10.1.1.1: icmp_seq=3 ttl=255 time=12.997 ms
56 bytes from 10.1.1.1: icmp_seq=4 ttl=255 time=7.828 ms

--- Ping statistics for 10.1.1.1 ---
5 packets transmitted, 5 packets received, 0.0% packet loss
round-trip min/avg/max/std-dev = 7.828/28.550/95.774/33.694 ms
%Jul 22 08:52:24:489 2015 H3C PING/6/PING_STATISTICS: Ping statistics for 10.1.1.1: 5 packets tr
ansmitted
, 5 packets received, 0.0% packet loss, round-trip min/avg/max/std-dev = 7.828/28.550/95.774/33.694
ms.

第一：PE设备的mpls lsr-id 一定要配置，否则l2vpn pw 状态up不起来
第二：LDP PEER流量需要写路由指向tunnel口
第三：ac interface GigabitEthernet0/1.110 ////如果ac链路变更了，peer命令需要重新下发
　　peer 4.4.4.4 pw-id 801001111///即此条命令需要重新配置