




TANet 100G

主節點備援

iBGP 台大區網



```
RP/0/RP0/CPU0:TANet-NTU-ASR9912-01#sh ip ro 0.0.0.0
Tue Nov  8 14:59:10.266 CST

Routing entry for 0.0.0.0/0
  Known via "bgp 1659", distance 200, metric 0, candidate default path, type internal
  Installed Oct 25 15:03:05.452 for 1w6d
  Routing Descriptor Blocks
  192.192.61.82, from 192.192.61.82
    route metric is 0
  No advertising protos.
```

Default Route 往臺北主節點

iBGP

台大區網

* BGP local-preference

- * 越高優先權越大，影響 Outgoing 路徑，僅在此 AS 有效。

```
RP/0/RP0/CPU0:TANet-NTU-ASR9912-01#sh bgp neighbor 192.192.61.82 routes
  Network          Next Hop          Metric LocPrf Weight Path
*>i0.0.0.0/0       192.192.61.82    400          0 i
Processed 1 prefixes, 1 paths

RP/0/RP0/CPU0:TANet-NTU-ASR9912-01#sh bgp neighbor 192.192.61.86 routes
  Network          Next Hop          Metric LocPrf Weight Path
* i0.0.0.0/0       192.192.61.86    200          0 i
Processed 1 prefixes, 1 paths
```

BGP Best Path

```
RP/0/RPO/CPU0:TANet-NTU-ASR9912-01#sh ip bgp 0.0.0.0
Thu Nov 17 08:55:56.318 CST
BGP routing table entry for 0.0.0.0/0
Versions:
  Process          bRIB/RIB   SendTblVer
  Speaker          1020798    1020798
Last Modified: Apr 14 10:42:25.944 for 30w6d
Paths: (4 available, best #1)
  Advertised to peers (in unique update groups):
    140.112.0.69
  Path #1: Received by speaker 1
  Advertised to peers (in unique update groups):
    140.112.0.69
  Local
    192.192.61.82 from 192.192.61.82 (192.192.60.21)
    Origin IGP, localpref 400, valid, internal, best, group-best
    Received Path ID 0, Local Path ID 1, version 1020798
  Path #2: Received by speaker 1
  Not advertised to any peer
  Local, (received-only)
    192.192.61.82 from 192.192.61.82 (192.192.60.21)
    Origin IGP, localpref 100, valid, internal
    Received Path ID 0, Local Path ID 0, version 0
  Path #3: Received by speaker 1
  Not advertised to any peer
  Local
    192.192.61.86 from 192.192.61.86 (192.192.60.31)
    Origin IGP, localpref 200, valid, internal
    Received Path ID 0, Local Path ID 0, version 0
  Path #4: Received by speaker 1
  Not advertised to any peer
  Local, (received-only)
    192.192.61.86 from 192.192.61.86 (192.192.60.31)
    Origin IGP, localpref 100, valid, internal
    Received Path ID 0, Local Path ID 0, version 0
```

臺北主節點 best path

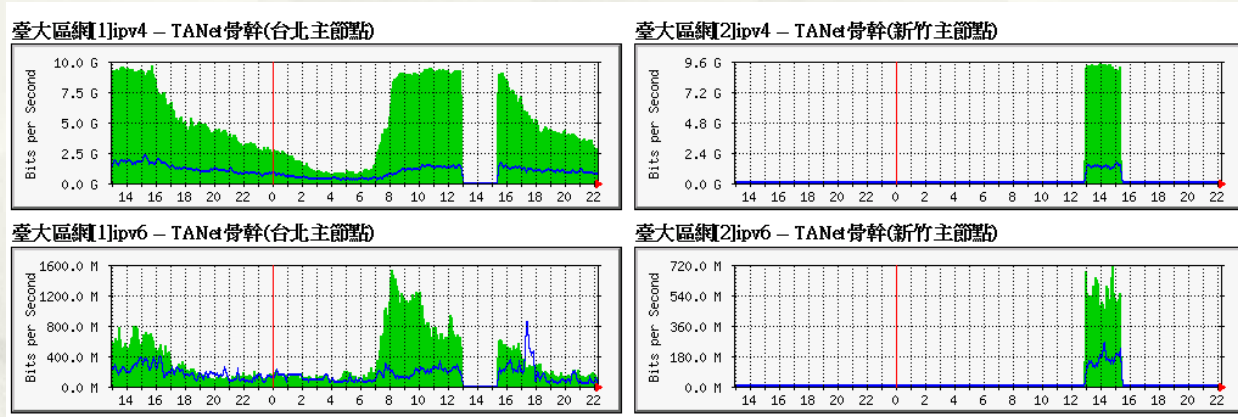
臺北主節點

新竹主節點

新竹主節點


iBGP 台大區網 備援實例

* 2017/05/10




* TAnet NOC 公告

臺灣學術網路 - 網路維運中心 **TAnet**
Taiwan Academic Network - Network Operate Center

 最新消息


[解除]5/10 12:53亞太A4及A7光纜斷線修護



2022/04、05 市網 DDoS 攻擊事件

市網 DDoS 攻擊事件


- * 2022/04/01 14:00~14:30
- * 2022/04/28 15:00~15:05
- * 2022/05/09 09:30 10:00
- * 2022/05/10 09:50 14:00
- * 2022/05/13 14:20 17:00



市網 DDoS 事件 攻擊來源分析


* Protocol: UDP

Bar: Protocol In Packets History



* Source Port: DNS, NTP 放大攻擊

Bar: Source Port In Packets



市網 DDoS 事件 攻擊目標分析


* 攻擊目的 IP

Bar: Dest_IP Packets History

用戶單位網段：203.72.61.0/24

Chinese Name

臺北市立永春高級中學



* 每次不盡相同

市網 DDoS 事件

BGP Session 中斷

- * 區網 ASR 100G主節點 BGP: **發生斷線**
 - * LC/0/0/CPU0:Apr 28 14:55:40.950 : bfd_agent[129]: %L2-BFD-6-SESSION_STATE_DOWN : BFD session to neighbor 192.192.61.82 on interface Bundle-Ether2020.1 has gone down. Reason: Control timer expired
- * 區網 ASR 其他 ISP BGP : **維持正常**
- * 原因分析
 - * **主節點 BGP 啟用 BFD: Bidirectional Forwarding Detection**
 - * **sub-second failure detection** between two adajecnt routers.


```
router bgp 1659
  bfd minimum-interval 200 (每秒發送 5次)
  bfd multiplier 5
  neighbor-group TANet1659_v4_Primary
  remote-as 1659
  bfd fast-detect
  neighbor-group TANet1659_v6_Primary
  remote-as 1659
  bfd fast-detect
```

市網遭受 DDoS 攻擊 緩解方法

* 北區 ASOC 流量清洗

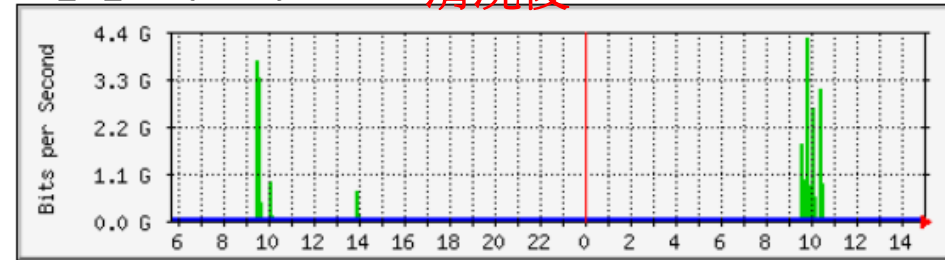
ASR_to_DP(dirty)

清洗前



DP_to_ASR(clean)

清洗後



- * ASR 至 Gigamon 分流器頻寬: 40G 擴增至 80G
- * 暫時移除 ASR 與主節點 BGP 之 BFD 設定

BFD Config

- * 僅有 OSPF 有設定 BFD

```
router ospf TANet-MPLS-IGP
```

```
address-family ipv4 unicast
```

```
area 0
```

```
interface HundredGigE0/0/0/0.3
```

```
bfd fast-detect
```

- * BGP 維持預設值

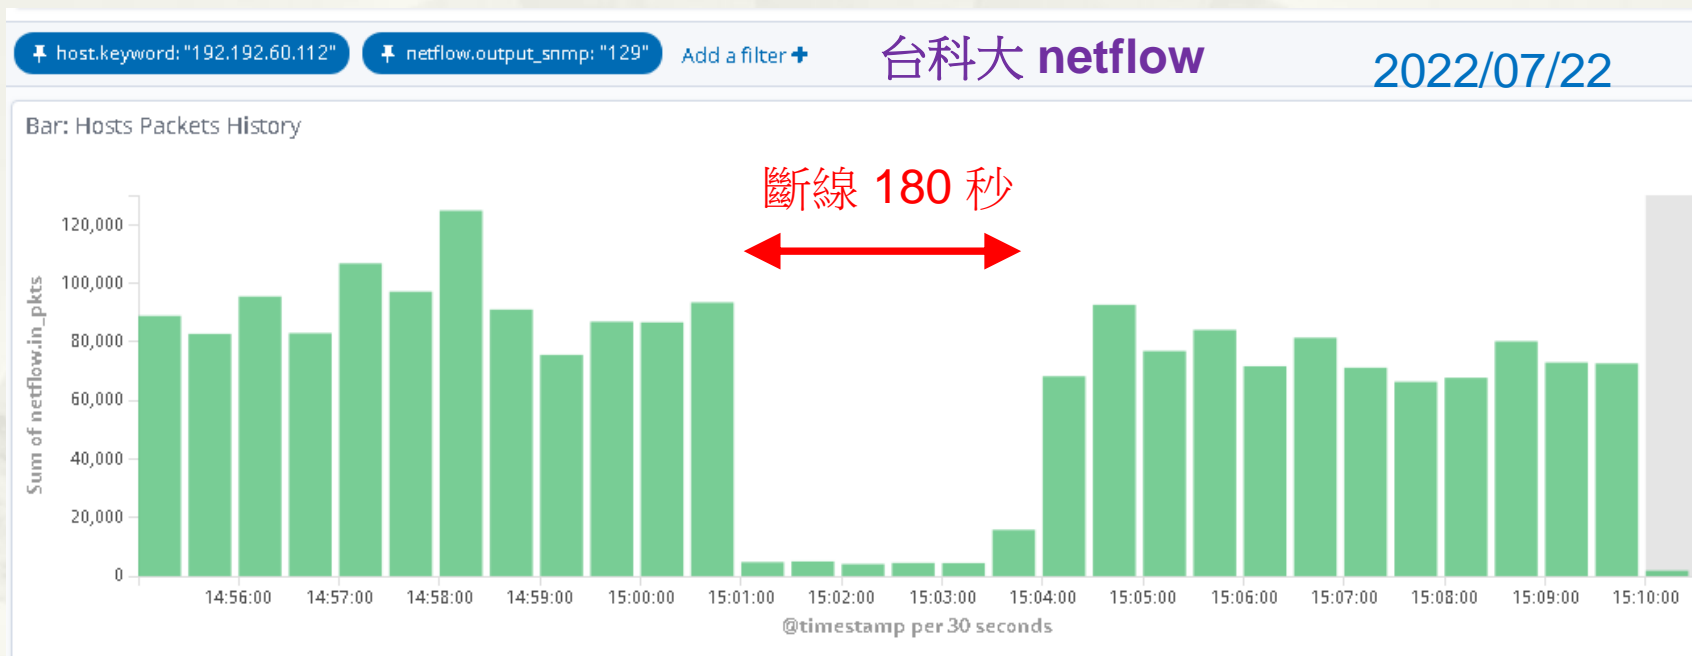
- * Periodic TCP keepalives 60 sec , Hold time 180 sec

移除 **BGP BFD** 偵測後 臺北主節點斷線

2022/07/22 15:00

臺北主節點斷線

- * TAnet 100G 臺北主節點切換至新竹主節點，BGP 斷線約 180 秒




- * Why 180 秒?

- * BGP Default Periodic TCP keepalives 60 sec、Hold time 180 sec


2022/11/18 16:37~18:17

臺北主節點斷線


臺大區網[1]ipv4 -- TANet骨幹(台北主節點)




臺大區網[2]ipv4 -- TANet骨幹(新竹主節點)



臺大區網[1]ipv6 -- TANet骨幹(台北主節點)



臺大區網[2]ipv6 -- TANet骨幹(新竹主節點)



呂思瑩 <chiyinlu@mail.moe.gov.tw>

davisyou@ntu.edu.tw; 'hshu'; 'scpei'; '張睿華 Lisa' ▾

RE: [External] RE: 臺北主節點斷線

您已於 2022/11/25 下午 05:25 回覆此訊息。



[External] (障礙解除通報)RE: [External] DFA4光纖(台北主節點-台灣大學)有光衰現象，台北主節點ROADM有告警產生 (157 KB) Outlook 項目

子興老師好

11/19 為誤植，應為 11/18

如同剛剛電話中討論的，


11/19 傍晚 4 點~7 點間的 TP1RC 網路服務主節點從 TP 切換成 HC，

主要原因是 DFA4 光纖光衰問題，以至於 11/19 傍晚 4 點多系統自動切換連結至新竹主節點。

16:37~18:17 100G 實體介面

共計 12次 up/down

- * LC/0/0/CPU0:Nov 18 16:37:44.481 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0, changed state to Down
- * LC/0/0/CPU0:Nov 18 17:01:55.892 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0, changed state to Up
- * LC/0/0/CPU0:Nov 18 17:07:43.085 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0, changed state to Down
- * LC/0/0/CPU0:Nov 18 17:09:52.073 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0, changed state to Up
- * LC/0/0/CPU0:Nov 18 17:09:53.876 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0, changed state to Down
- * LC/0/0/CPU0:Nov 18 17:13:47.830 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0, changed state to Up
- * LC/0/0/CPU0:Nov 18 17:13:56.990 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0, changed state to Down
- * LC/0/0/CPU0:Nov 18 17:13:57.067 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0, changed state to Up
- * LC/0/0/CPU0:Nov 18 17:13:58.007 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0, changed state to Down
- * LC/0/0/CPU0:Nov 18 17:13:58.062 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0, changed state to Up
- * LC/0/0/CPU0:Nov 18 17:42:53.906 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0, changed state to Down
- * LC/0/0/CPU0:Nov 18 17:48:13.254 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0, changed state to Up
- * LC/0/0/CPU0:Nov 18 17:48:13.583 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0, changed state to Down
- * LC/0/0/CPU0:Nov 18 17:48:13.638 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0, changed state to Up
- * LC/0/0/CPU0:Nov 18 17:53:33.286 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0, changed state to Down
- * LC/0/0/CPU0:Nov 18 18:12:23.903 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0, changed state to Up
- * LC/0/0/CPU0:Nov 18 18:12:24.093 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0, changed state to Down
- * LC/0/0/CPU0:Nov 18 18:12:24.193 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0, changed state to Up
- * LC/0/0/CPU0:Nov 18 18:12:30.484 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0, changed state to Down
- * LC/0/0/CPU0:Nov 18 18:12:30.562 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0, changed state to Up
- * LC/0/0/CPU0:Nov 18 18:12:30.569 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0, changed state to Down
- * LC/0/0/CPU0:Nov 18 18:12:30.629 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0, changed state to Up
- * LC/0/0/CPU0:Nov 18 18:12:31.539 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0, changed state to Down
- * LC/0/0/CPU0:Nov 18 18:12:31.615 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0, changed state to Up

The background features a traditional Chinese landscape painting, likely a 'shanshui' style, rendered in a light, monochromatic tone. The painting is presented as if on a fan, with the central part of the image being the widest and tapering towards the top and bottom. The scene depicts a mountainous landscape with a river or path winding through it, surrounded by trees and small structures. The overall aesthetic is serene and classical.

實體線路中斷 (第一次)

ASR Log

* 實體線路中斷(第一次):

- * LC/0/0/CPU0:Nov 18 16:37:44.481 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0, changed state to Down
- * LC/0/0/CPU0:Nov 18 16:37:44.481 : ifmgr[218]: %PKT_INFRA-LINEPROTO-5-UPDOWN : Line protocol on Interface HundredGigE0/0/0/0, changed state to Down
- * LC/0/0/CPU0:Nov 18 16:37:44.481 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0.3, changed state to Down
- * LC/0/0/CPU0:Nov 18 16:37:44.481 : ifmgr[218]: %PKT_INFRA-LINEPROTO-5-UPDOWN : Line protocol on Interface HundredGigE0/0/0/0.3, changed state to Down
- * LC/0/0/CPU0:Nov 18 16:37:44.481 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0.2, changed state to Down
- * LC/0/0/CPU0:Nov 18 16:37:44.481 : ifmgr[218]: %PKT_INFRA-LINEPROTO-5-UPDOWN : Line protocol on Interface HundredGigE0/0/0/0.2, changed state to Down
- * LC/0/0/CPU0:Nov 18 16:37:44.481 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0.1, changed state to Down
- * LC/0/0/CPU0:Nov 18 16:37:44.482 : ifmgr[218]: %PKT_INFRA-LINEPROTO-5-UPDOWN : Line protocol on Interface HundredGigE0/0/0/0.1, changed state to Down

ASR Log

- * **OSFP Neighbor Down(BFD): 200ms 之内**
 - * LC/0/0/CPU0:Nov 18 16:37:44.484 : bfd_agent[129]: %L2-BFD-6-ADJACENCY_DELETE : Adjacency to neighbor 172.16.1.82 on interface HundredGigE0/0/0/0.3 was deleted
 - * RP/0/RP0/CPU0:Nov 18 16:37:44.487 : ospf[1022]: %ROUTING-OSPF-5-ADJCHG : Process TANet-MPLS-IGP, Nbr 172.16.0.21 on HundredGigE0/0/0/0.3 in area 0 from FULL to DOWN, Neighbor Down: interface down or detached, vrf default vrfid 0x600000000
 - * LC/0/0/CPU0:Nov 18 16:37:44.492 : bfd_agent[129]: %L2-BFD-6-SESSION_REMOVED : BFD session to neighbor 172.16.1.82 on interface HundredGigE0/0/0/0.3 has been removed
 - * LC/0/0/CPU0:Nov 18 16:37:45.548 : vic[389]: %PLATFORM-VIC-4-SIGNAL : Interface HundredGigE0/0/0/0, Detected Signal failure

ASR Log

* BGP ipv6 TCP keepalives 60 sec:


- * RP/0/RP0/CPU0:Nov 18 16:38:54.997 : tcp[469]: %IP-TCP_NSR-5-DISABLED : 2001:288:0:1659:192:192:61:80:179 <-> 2001:288:0:1659:192:192:61:81:54639:: NSR disabled for TCP connection because Retransmission threshold exceeded
- * RP/0/RP0/CPU0:Nov 18 16:38:54.997 : bgp[1059]: %ROUTING-BGP-3-NBR_NSR_DISABLED : NSR disabled on neighbor 2001:288:0:1659:192:192:61:81 due to 'ip-tcp' detected the 'warning' condition 'NSR is down because the retransmission threshold exceeded (probably because downstream RP is not healthy)'
- * RP/0/RP1/CPU0:Nov 18 16:38:54.997 : bgp[1059]: %ROUTING-BGP-5-NBR_NSR_DISABLED_STANDBY : NSR disabled on neighbor 2001:288:0:1659:192:192:61:81 on standby RP due to Peer closing down the session (VRF: default)

* BGP ipv4 TCP keepalives 60 sec:

- * RP/0/RP0/CPU0:Nov 18 16:40:01.852 : tcp[469]: %IP-TCP_NSR-5-DISABLED : 192.192.61.81:179 <-> 192.192.61.82:45213:: NSR disabled for TCP connection because Retransmission threshold exceeded
- * RP/0/RP0/CPU0:Nov 18 16:40:01.852 : bgp[1059]: %ROUTING-BGP-3-NBR_NSR_DISABLED : NSR disabled on neighbor 192.192.61.82 due to 'ip-tcp' detected the 'warning' condition 'NSR is down because the retransmission threshold exceeded (probably because downstream RP is not healthy)'
- * RP/0/RP1/CPU0:Nov 18 16:40:01.851 : bgp[1059]: %ROUTING-BGP-5-NBR_NSR_DISABLED_STANDBY : NSR disabled on neighbor 192.192.61.82 on standby RP due to Peer closing down the session (VRF: default)

ASR Log

- * BGP ipv4/ipv6 neighbor down: 實體線路中斷後三分鐘(Hold time 180 sec)
 - * RP/0/RP0/CPU0:Nov 18 16:40:29.867 : bgp[1059]:
%ROUTING-BGP-5-ADJCHANGE_DETAIL : neighbor
2001:288:0:1659:192:192:61:81 Down - Peer closing
down the session (VRF: default; AFI/SAFI: 2/1) (AS:
1659)
 - * RP/0/RP0/CPU0:Nov 18 16:40:51.627 : bgp[1059]:
%ROUTING-BGP-5-ADJCHANGE_DETAIL : neighbor
192.192.61.82 Down - Peer closing down the session
(VRF: default; AFI/SAFI: 1/1) (AS: 1659)

The background features a traditional Chinese landscape painting, likely a 'shanshui' style, rendered in a light, monochromatic tone. The scene is depicted on a fan shape, with a central mountain peak and various smaller peaks, trees, and a winding path or river. The overall aesthetic is serene and classical.

實體線路恢復 (最後一次)

ASR Log

* 實體線路恢復(最後一次):

- * LC/0/0/CPU0:Nov 18 18:12:31.615 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0, changed state to Up
- * LC/0/0/CPU0:Nov 18 18:12:31.616 : ifmgr[218]: %PKT_INFRA-LINEPROTO-5-UPDOWN : Line protocol on Interface HundredGigE0/0/0/0, changed state to Up
- * LC/0/0/CPU0:Nov 18 18:12:31.617 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0.3, changed state to Up
- * LC/0/0/CPU0:Nov 18 18:12:31.619 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0.2, changed state to Up
- * LC/0/0/CPU0:Nov 18 18:12:31.619 : ifmgr[218]: %PKT_INFRA-LINK-3-UPDOWN : Interface HundredGigE0/0/0/0.1, changed state to Up
- * LC/0/0/CPU0:Nov 18 18:12:31.621 : ifmgr[218]: %PKT_INFRA-LINEPROTO-5-UPDOWN : Line protocol on Interface HundredGigE0/0/0/0.3, changed state to Up
- * LC/0/0/CPU0:Nov 18 18:12:31.624 : ifmgr[218]: %PKT_INFRA-LINEPROTO-5-UPDOWN : Line protocol on Interface HundredGigE0/0/0/0.2, changed state to Up
- * LC/0/0/CPU0:Nov 18 18:12:31.625 : ifmgr[218]: %PKT_INFRA-LINEPROTO-5-UPDOWN : Line protocol on Interface HundredGigE0/0/0/0.1, changed state to Up

ASR Log

* OSFP Neighbor Up(BFD): 200ms 之內

- * RP/0/RP0/CPU0:Nov 18 18:12:39.801 : mpls_ldp[1211]: %ROUTING-LDP-5-NSR_SYNC_START : Initial synchronization started for 1 peers
- * RP/0/RP0/CPU0:Nov 18 18:12:41.161 : ospf[1022]: %ROUTING-OSPF-5-ADJCHG : Process TANet-MPLS-IGP, Nbr 172.16.0.21 on HundredGigE0/0/0/0.3 in area 0 from LOADING to FULL, Loading Done, vrf default vrfid 0x60000000
- * LC/0/0/CPU0:Nov 18 18:12:41.170 : bfd_agent[129]: %L2-BFD-6-SESSION_DAMPENING_ON : Session to neighbor 172.16.1.82 on interface HundredGigE0/0/0/0.3 entered Dampened state (initial: 2000 ms,secondary: 5000 ms,maximum: 120000 ms).
- * LC/0/0/CPU0:Nov 18 18:12:43.607 : bfd_agent[129]: %L2-BFD-6-SESSION_DAMPENING_OFF : Session to neighbor 172.16.1.82 on interface HundredGigE0/0/0/0.3 moved out of Dampened state.
- * LC/0/0/CPU0:Nov 18 18:12:43.612 : bfd_agent[129]: %L2-BFD-6-SESSION_STATE_UP : BFD session to neighbor 172.16.1.82 on interface HundredGigE0/0/0/0.3 is up

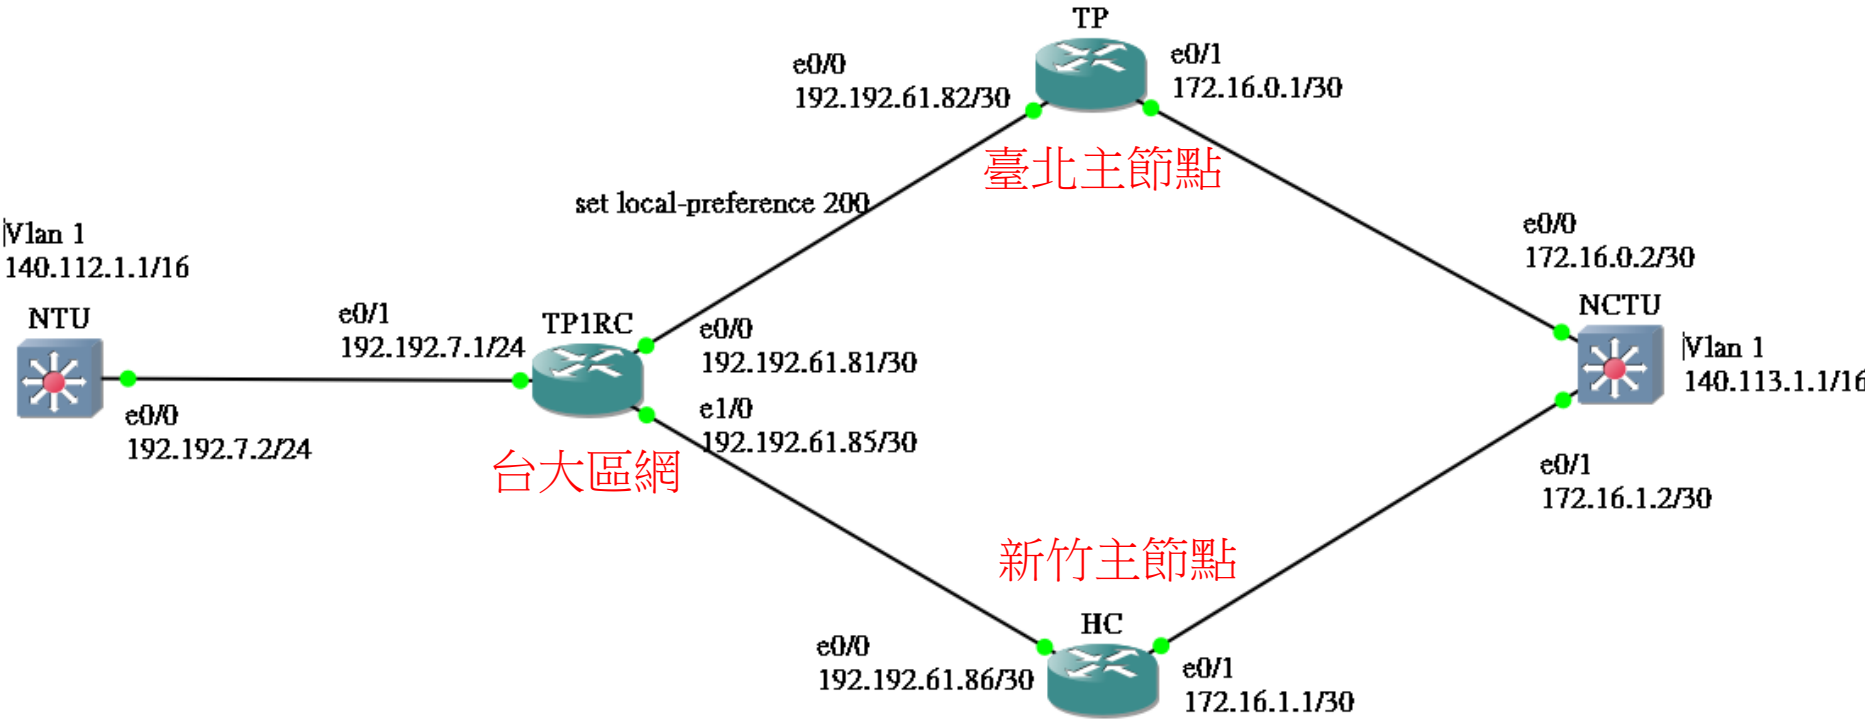
* BGP ipv4/ipv6 neighbor up: 實體線路恢復後三分鐘(Hold time 180 sec)

- * RP/0/RP0/CPU0:Nov 18 18:16:37.417 : bgp[1059]: %ROUTING-BGP-5-ADJCHANGE_DETAIL : neighbor 192.192.61.82 Up (VRF: default; AFI/SAFI: 1/1) (AS: 1659)
- * RP/0/RP0/CPU0:Nov 18 18:17:01.282 : bgp[1059]: %ROUTING-BGP-5-ADJCHANGE_DETAIL : neighbor 2001:288:0:1659:192:192:61:81 Up (VRF: default; AFI/SAFI: 2/1) (AS: 1659)

解決方法

- * 方案1:縮短 BGP Periodic TCP Timer
- * 方案2:拉長 BFD 時間
- * 方案3:備援線路不設定 BFD

GNS3 模擬器




BGP Periodic TCP Timer

- * BGP Periodic TCP Timer(Default)
 - * keepalives 60 sec
 - * Hold time 180 sec
 - * When a BGP connection negotiate the holdtimer between two BGP peers started, the smaller of the two holdtimers will be chosen.
- * If link down, it will immediate see neighbour down.

方案1.縮短 BGP Periodic TCP Timer

- * Change Default Timer
 - * keepalives 10 sec 、 Hold time 30 sec
 - * TP(config)# router bgp 1659
 - * TP(config-router)# neighbor 192.192.61.81 timers 10 30
- * 將 TP 介面 e0/0 上 IP 設定移除:
 - * NTU# ping 140.113.1.1 repeat 100000




斷線約 30秒


為何區網 100G 介面 Down BGP 不會立即切換?

原因: 區網與主節點間接介面為
Layer2 Only


台大區網 Gigamon 分流架構 80Gbps Inline Network to Inline Tool (外對內)



方案3.備援線路不需設定 BFD



主線路	備援線路	應變措施
正常	正常	N/A
異常	正常	需快速切換至備援
正常	異常	毋須切換
異常	異常	快速切換也無用



簡報完畢
謝謝