Agilent N2X

N2X Core Routing - BGP-4 MPLS VPN scenario with integrated traffic

Application Note

Agilent Technologies

Introduction

Test Objective

The objective of this application note is to demonstrate the power of N2X (in particular the BGP-4 changes that were made in the 6.10 release) in simulating a high scaled VPN network from the edge to the core.

In the first part of the application note, we will use N2X to set up the control plane, using:

- OSPF to simulate the core,
- BGP-4 to simulate the PE routers and CE routes, and
- LDP to set up the LSPs used for forwarding the VPN traffic.

We will then create IPv4 traffic from the edge to the core, and MPLS labeled IPv4 traffic from the core to the edge.

The second part of the application note will introduce the concept of the "super VRF", which can be using to dramatically improve the scalability of core to edge traffic creation.

The third (and final) part of the application note will introduce the powerful CreateL3BgpMplsVpnTraffic QuickTool, and how it can be used to set up all of the traffic in a highly scaled VPN scenario in merely a few clicks of the mouse.

Target User

This test is most relevant to POC (Proof of Concept) users, or other users who need to quickly build up a large BGP-4 MPLS VPN topology through the GUI without resorting to using the TCL scripting API. It aims to show how easy it is to use N2X for all BGP-4 MPLS VPN testing and how to use aggregated objects to represent groups of peers and routes. Also included is a full high-scale traffic integration, to show how users can easily simulate a real world VPN network scenario.

Key Features

New aggregated BGP-4 pool objects:

- Peer pools
- Route profiles (IPv4, VPNv4)
- VPN VRF pools (including "Super VRF")
- CreateL3BgpMplsVpnTraffic QuickTool

Test Set-up

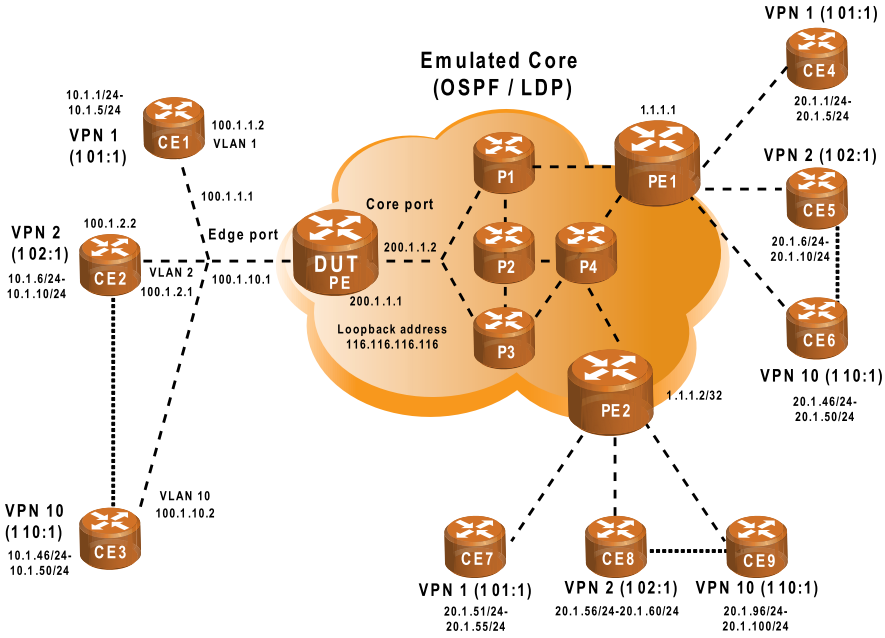

DUT Configuration:

See **Appendix A** at the end of the document.

Equipment Required

N2X Equipment

- 2 x 10/100/1000 Ethernet ports per group (1 core port, 1 edge port)

N2X Software

- N2X Packets and Protocols – 6.11 System Release
- CreateL3BgpMplsVpnTraffic QuickTool

Slides will be presented, and available on the web at a later date

Device Under Test (DUT)

Cisco GSR-12008 router (revision 53.50 or later)
IOS (tm) GS Software (GSR-P-M), Version 12.0(30)S1, RELEASE SOFTWARE (fc1)

Instructions

Part #1: Emulate scaled BGP-4 MPLS VPN topology with traffic

In the first part of the application note we will show how quickly you can simulate a high-scale BGP-4 VPN scenario using N2X. We will build the configuration incrementally from the link layer, to the control plane and routing protocols and finally the bi-directional traffic.

Select ports

► Step 1

Click "Ports" from the main window top tool bar. In Port Selection dialog, select 2 test ports. The first selected port will represent the edge port, and the second selected port will represent the core port (please refer to the topology diagram).

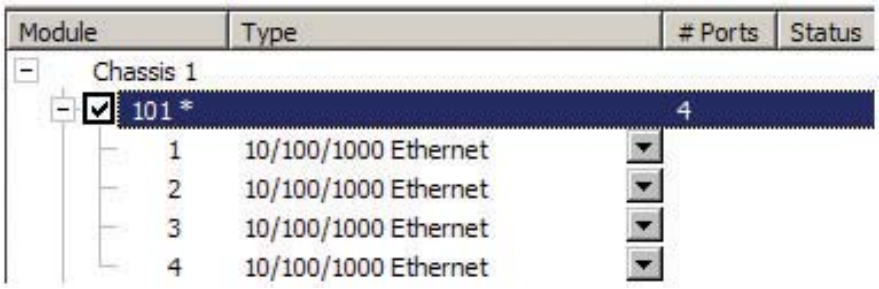

Configure physical layer

► Step 2

Click "Physical Layer" on the left hand Setup pane of the main application. In the Physical configuration dialog, select each port individually and click "Configure". Change the Media Type to "SFP" if using optical fibre, or "RJ45" if using CAT-5 copper wire.

Note: For SFP, Step 3 and 4 must be performed. For RJ45, skip step 3 and 4.

► Step 3

Click "Turn All Lasers Off".

► Step 4

Click "Turn All Lasers On".

► Step 5

Close the Physical configuration dialog.

Configure link layer

► Step 6

Click "Link Layer" on the left hand Setup pane of the main application. In the Link configuration dialog, select the "Ethernet" tab and then click "LAN/VLAN Addresses". Select the Tester row on the edge port and click "Add" to bring up the address pools dialog.

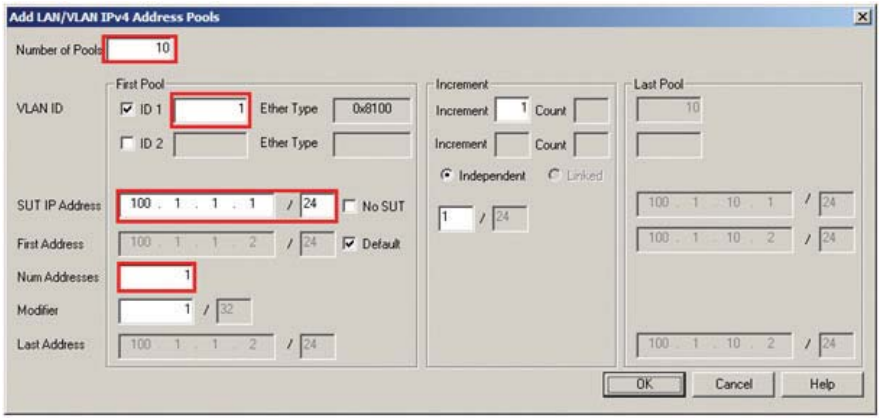

► Step 7

Modify the following values in the address pool dialog, to simulate the 10 VLANs on the edge of the network:

- Number of Pools = 10
- VLAN ID 1 = 1
- SUT IP Address = 100.1.1.1
- Num Addresses = 1

► Step 8

Remove the native Ethernet address pool on the edge port (i.e. the default one without the VLAN ID that was there before you added the 10 VLAN address pools).

► Step 9

Select the native tester address pool on the core port and click "Edit" to bring up the address pools dialog.

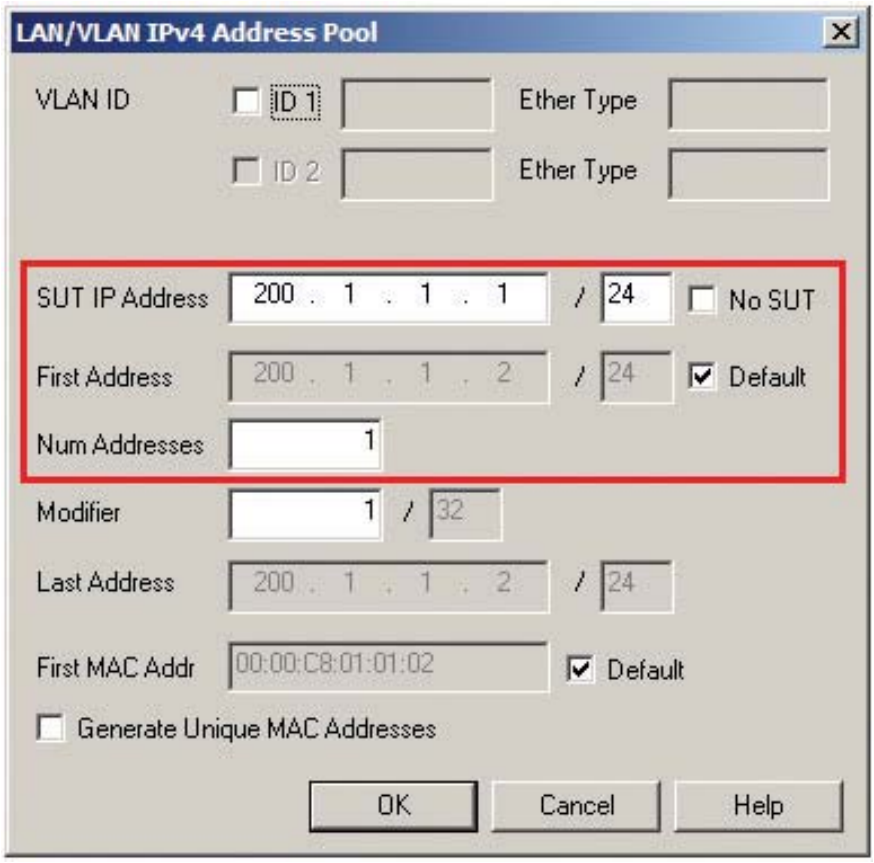

► Step 10

Modify the following values in the address pool dialog, to simulate the link on the core of the network:

- SUT IP Address = 200.1.1.1
- First Address = 200.1.1.2
- Num Addresses = 1

► Step 11

Resolve ARP and disable NDP as shown in the screenshots below. Ensure that the link layer button of the main application is not yellow or red, which indicates a problem with the configuration.

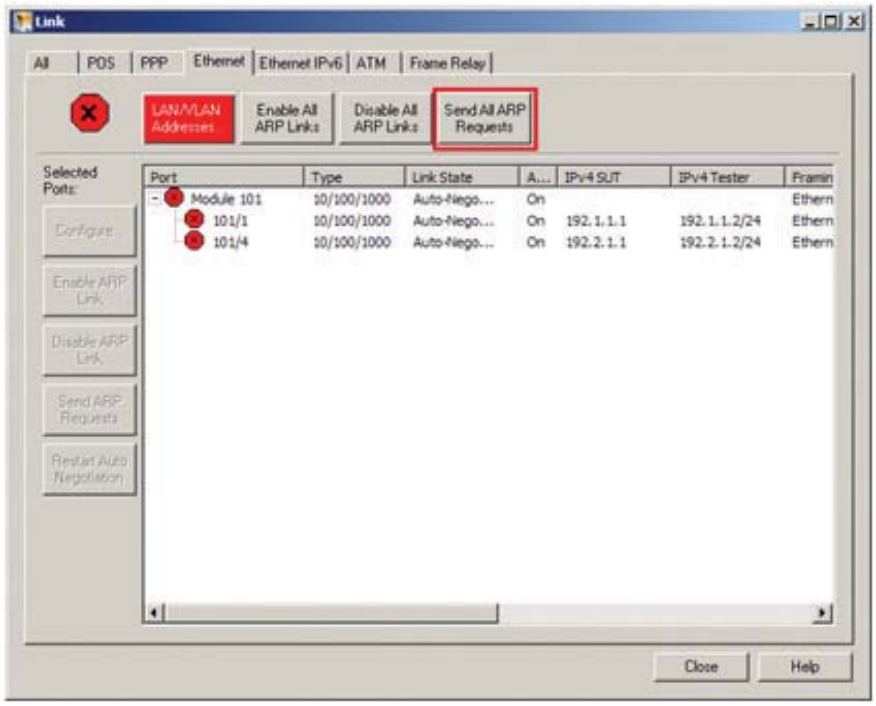

Edge – Advertise CEs on VLANs with BGP-4

► Step 12

Click "Emulation" on the left hand Setup pane of the main application. Select the edge port, and click "New" on the toolbar of the Emulation pane to bring up the New Emulations dialog.

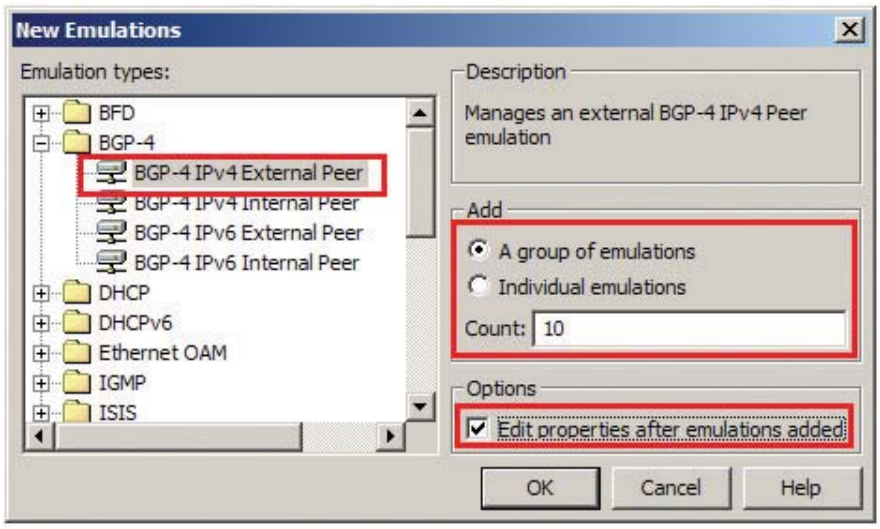

► Step 13

Modify the following values in the dialog to add a BGP-4 peer pool of size 10 representing the CEs on the edge of the network:

- Emulation type = BGP-4 IPv4 External Peer
- Add = A group of emulations
- Count = 10
- Select "Edit properties after emulations added"

► Step 14

Click "OK" to add the peer pool. The dialog below will be displayed.

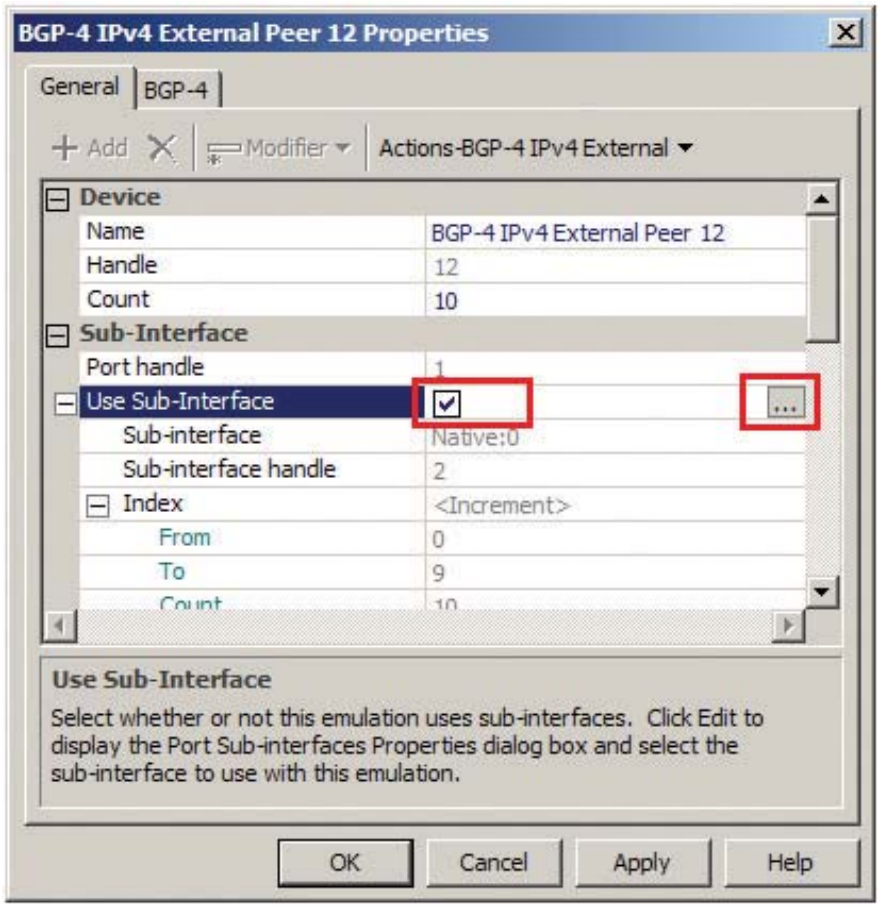

► Step 15

Select "Use Sub-Interface" to add the BGP-4 peer pool to the previously created VLANs. Click "..." to open the sub-interfaces selection dialog.

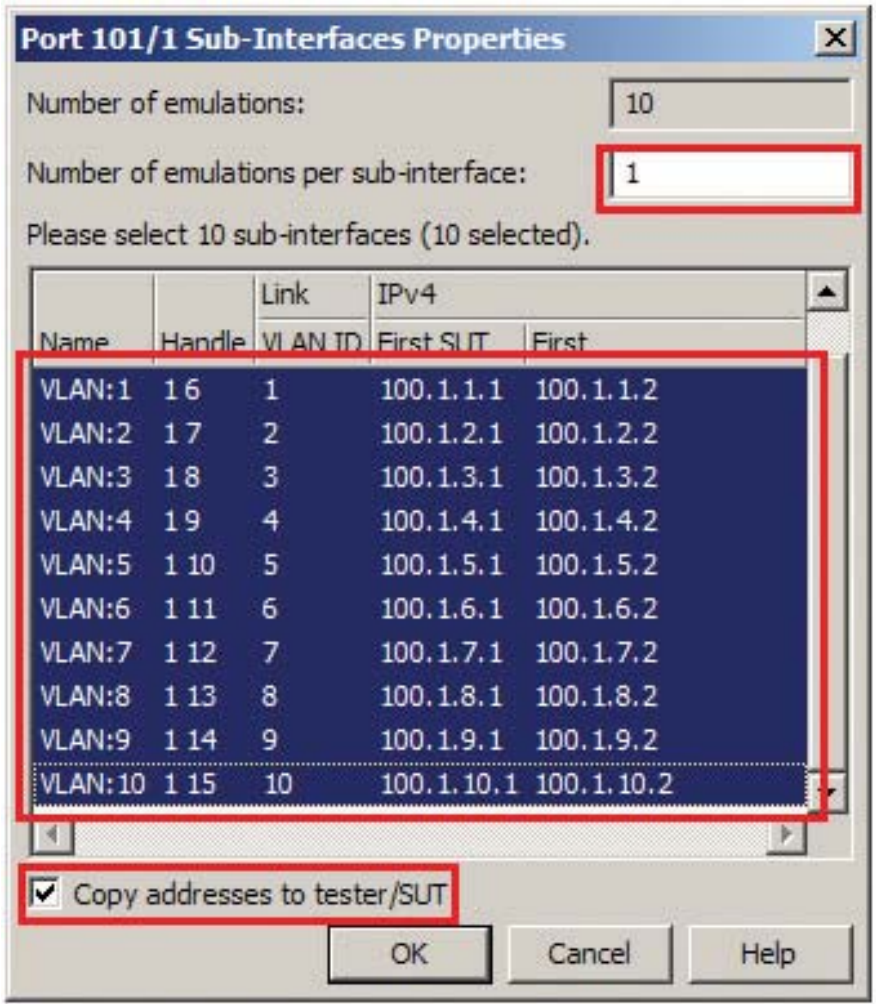

► Step 16

Multi-select all the VLANs in the list, and modify the following values in the dialog to add the BGP-4 peer pool over the VLANs:

- Number of emulations per sub-interface = 1
- Select "Copy addresses to tester/SUT"

► Step 17

Change to the "BGP-4" tab in the peer pool properties dialog.

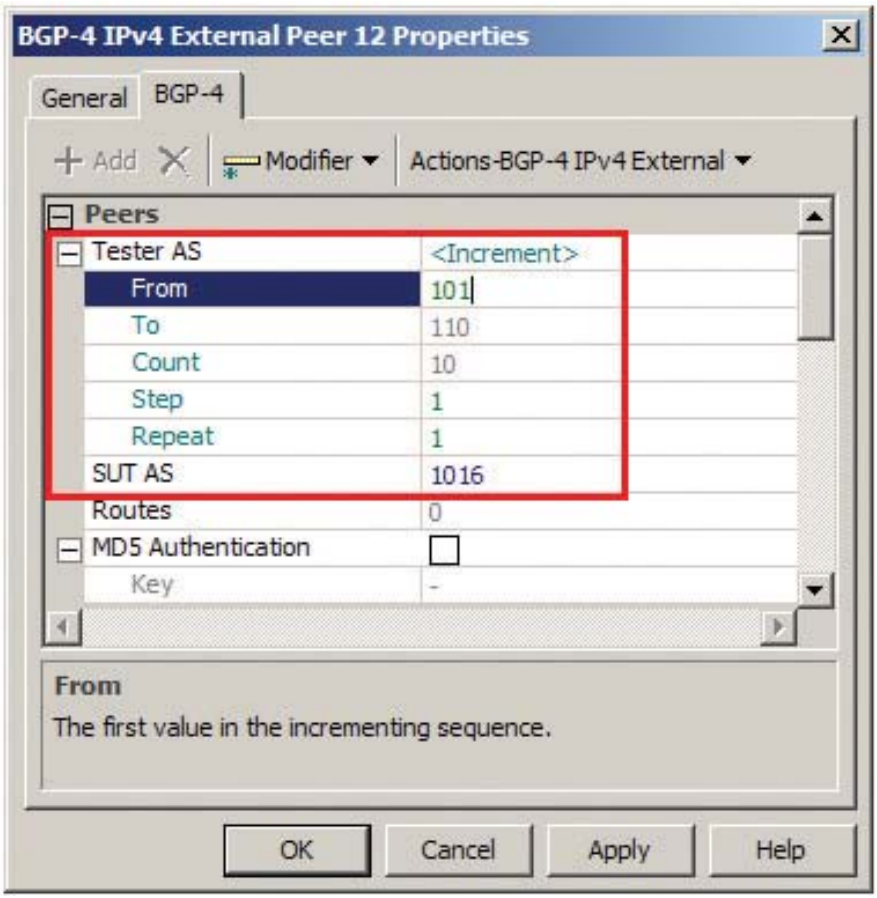

► Step 18

Modify the following values to set the AS numbers:

- Tester AS range = 101-110
- SUT AS = 1016

► Step 19

Click "OK" to apply the changes to the BGP-4 peer pool.

Edge – Advertise IPv4 routes behind CEs

► Step 20

Select the peer pool, and click the arrow next to "New" on the toolbar of the Emulation pane. Select "BGP-4 IPv4 Route Profile".

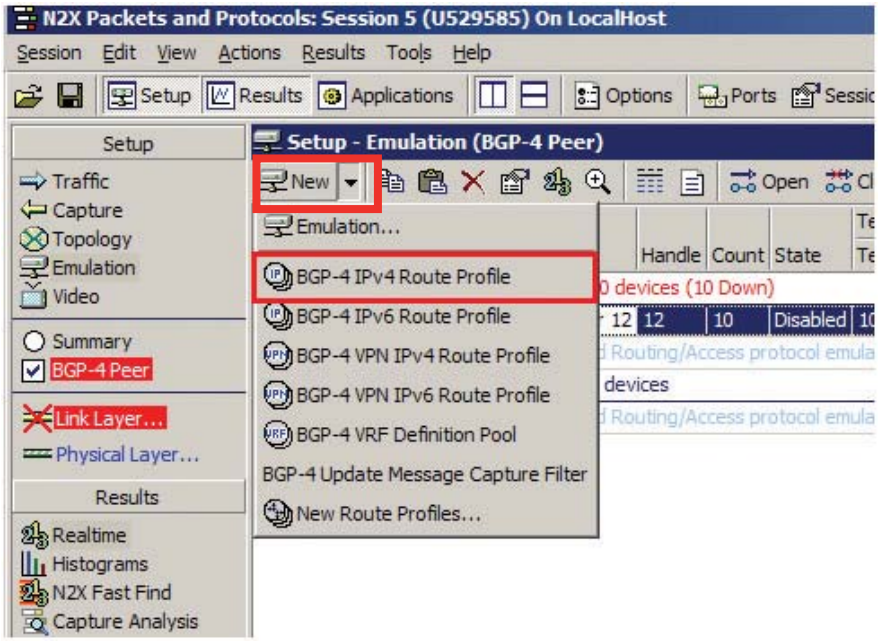

► Step 21

Double-click on the created route profile to edit it.

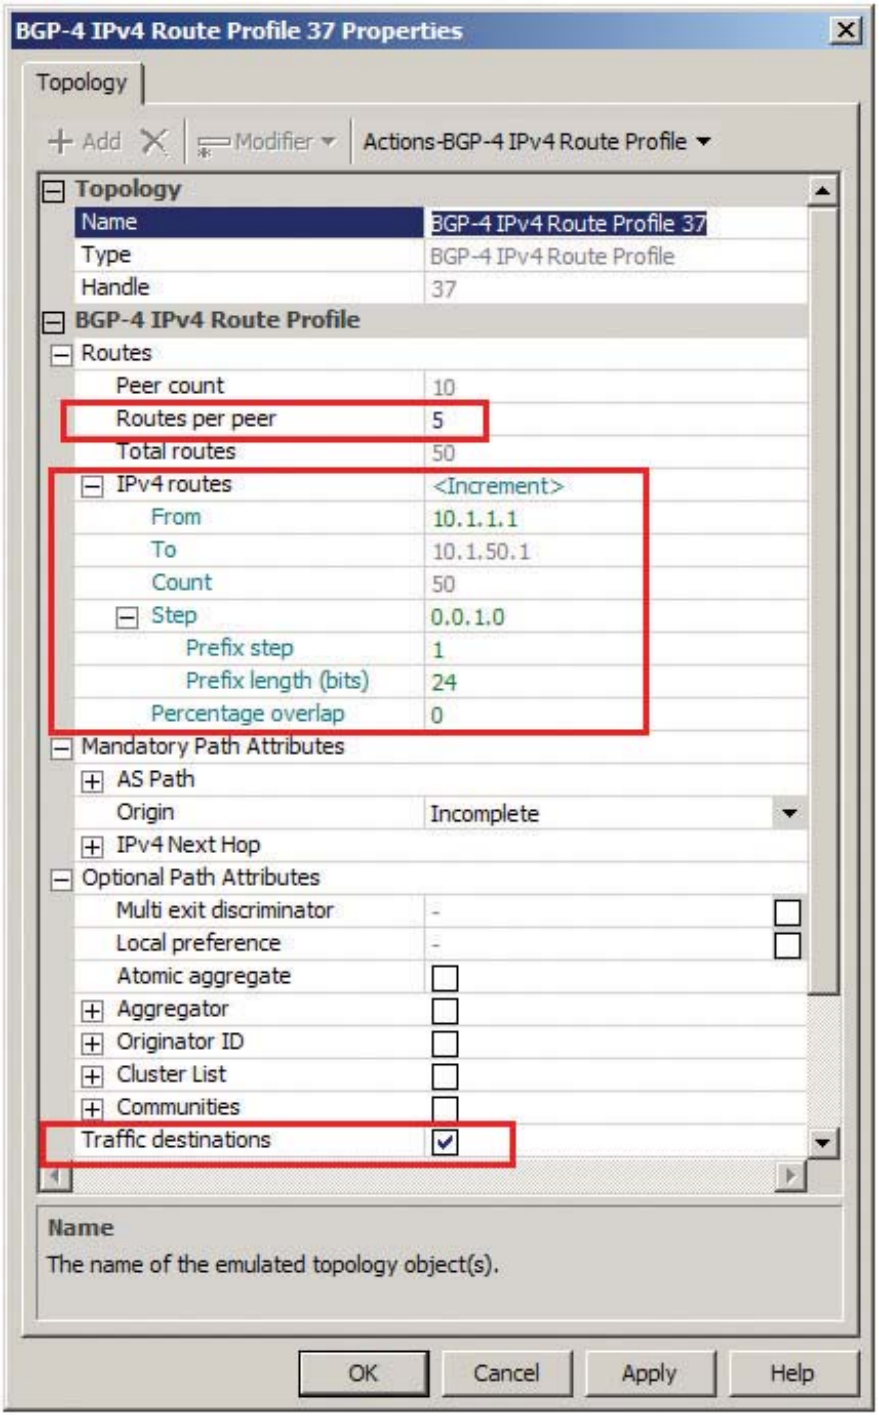

► Step 22

Modify the following values in the route profile to simulate routes behind each CE:

- Routes per peer = 5
- IPv4 routes
  - From = 10.1.1.1
  - Prefix step = 1
  - Prefix length = 24
- Select "Traffic destinations"

Core – Advertise PEs using OSPF

► Step 23

Click "Emulation" on the left hand Setup pane of the main application. Select the core port, and click "New" on the toolbar of the Emulation pane. From the New Emulations dialog, add a single OSPFv2 Router.

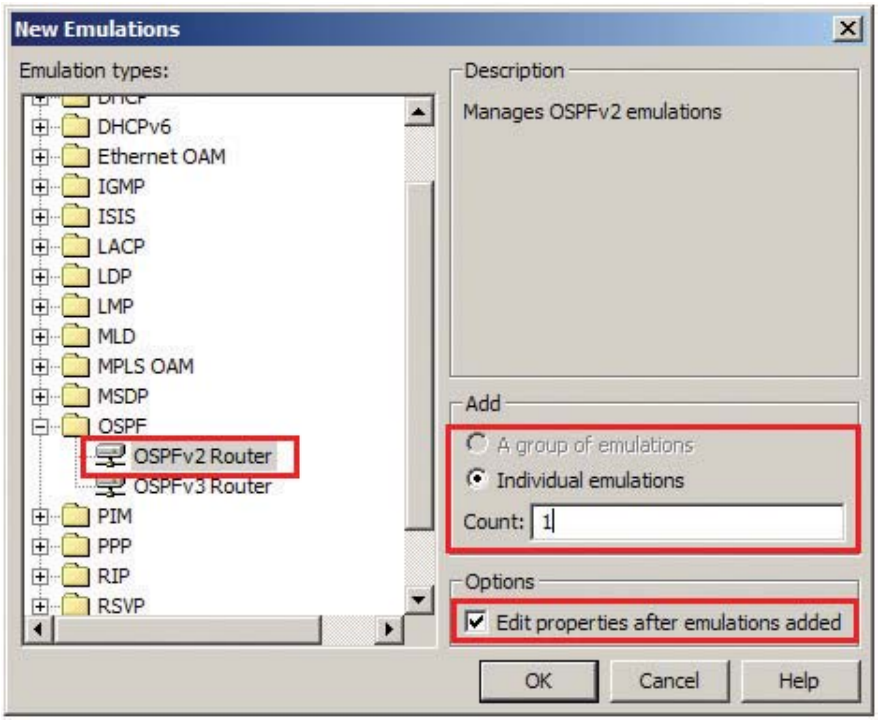

► Step 24

Modify the following values in the dialog to add an OSPFv2 router to advertise the PEs in the core:

- Emulation type = OSPFv2 Router
- Add = Individual emulations
- Count = 1
- Select "Edit properties after emulations added"

► Step 25

Click "OK" to add the OSPFv2 router. The dialog below will be displayed.

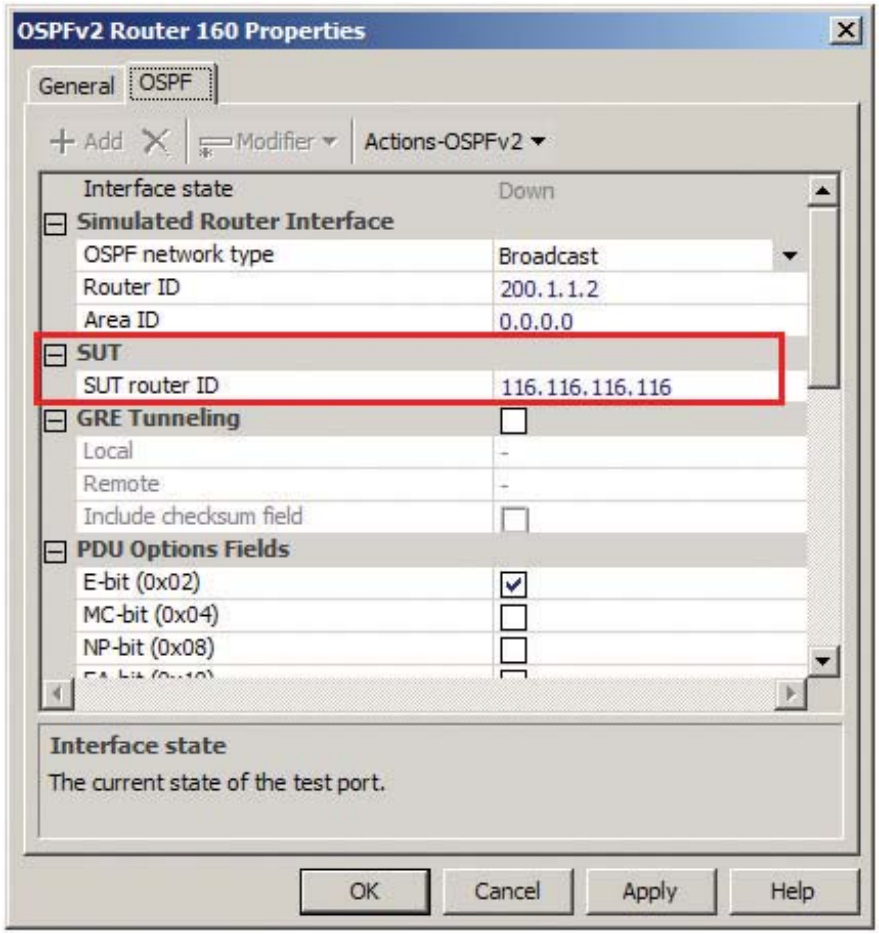

► Step 26

Select the "OSPF" tab, and modify the following values to make the OSPFv2 router peer with the DUT router:

- SUT router ID = 116.116.116.116 (loopback address of DUT)

► Step 27

Click "OK" to apply the changes to the OSPFv2 router.

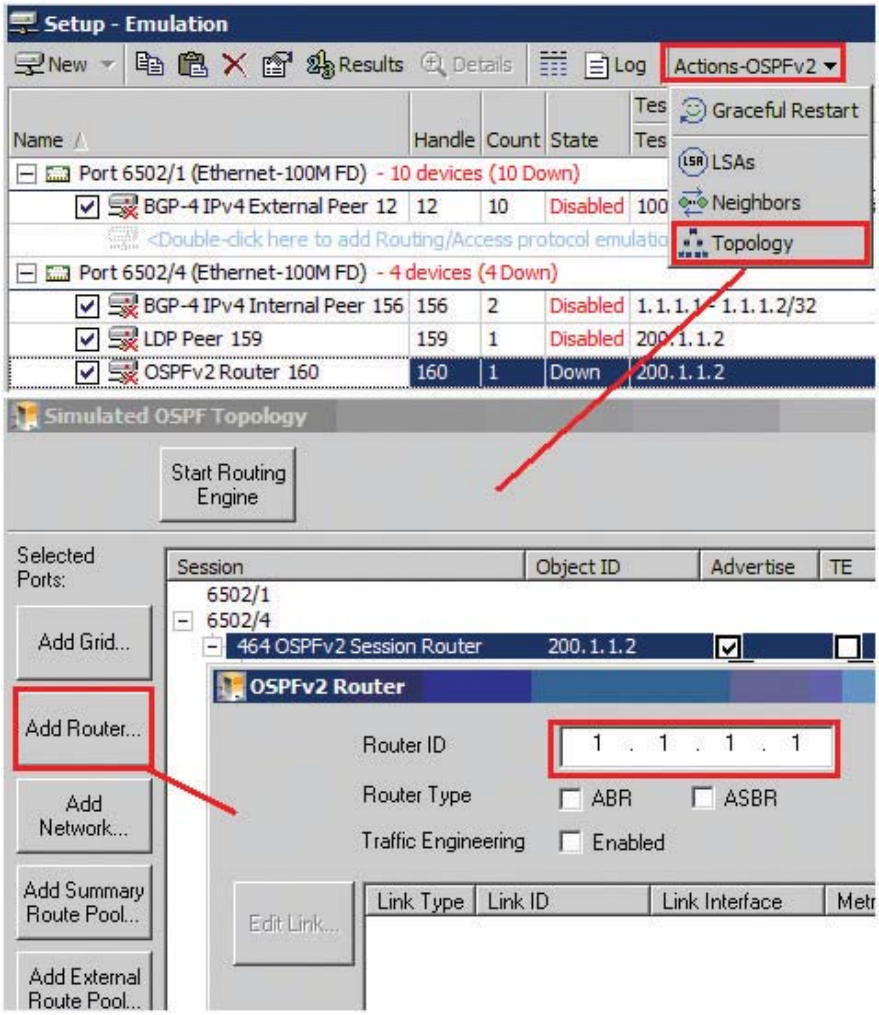

The following steps will add the OSPFv2 representing PE1 and PE2 behind the session router.

► Step 28

Click "Topology" from the Actions-OSPFv2 menu.

► Step 29

Select the "OSPFv2 Session Router", and click "Add Router".

► Step 30

Set the Router ID to 1.1.1.1 which is the loopback address of simulated PE1 (see topology diagram).

► Step 31

Click "OK" to add the router.

► Step 32

Repeat the last 3 steps to simulate PE2, using Router ID 1.1.1.2 this time.

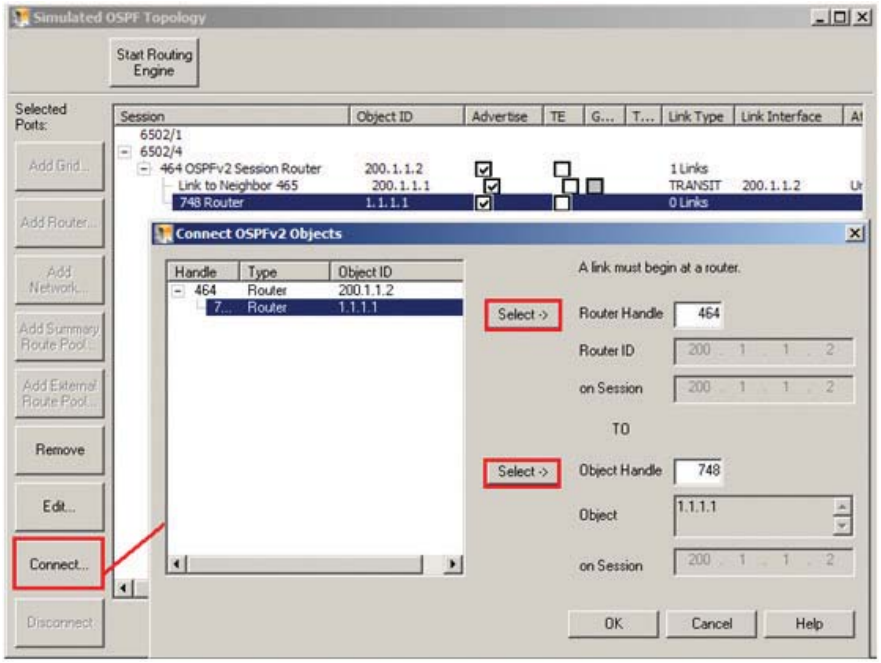

The following steps will connect PE1 and PE2 to the session router.

► Step 33

Click "Connect" which brings up the Connect OSPFv2 Objects dialog.

► Step 34

Select the session router (200.1.1.2) on the left, then click the top "Select" button.

▶ Step 35

Select the simulated PE (1.1.1.1) on the left, then click the bottom "Select" button.

▶ Step 36

Click "OK" to connect the routers.

▶ Step 37

Repeat the last 3 steps for the second PE router (1.1.1.2).

Core – Set up I-BGP sessions between PE1, PE2 and the DUT PE

▶ Step 38

Click "Emulation" on the left hand Setup pane of the main application. Select the core port, and click "New" on the toolbar of the Emulation pane to bring up the New Emulations dialog.

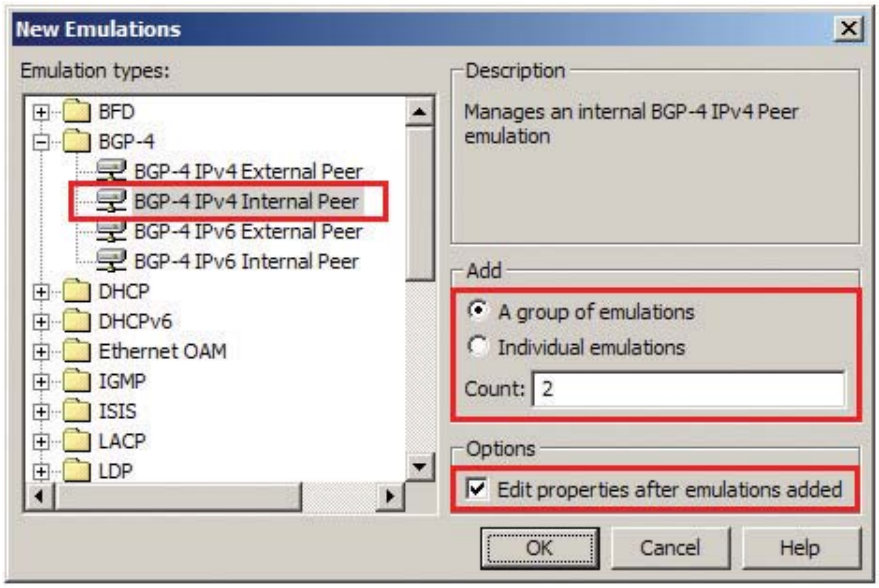

▶ Step 39

Modify the following values in the dialog to add an I-BGP peer pool of size 2 to the core port:

- Emulation type = BGP-4 IPv4 Internal Peer
- Add = A group of emulations
- Count = 2
- Select "Edit properties after emulations added"

► Step 40

Click "OK" to add the peer pool. The dialog below will be displayed.

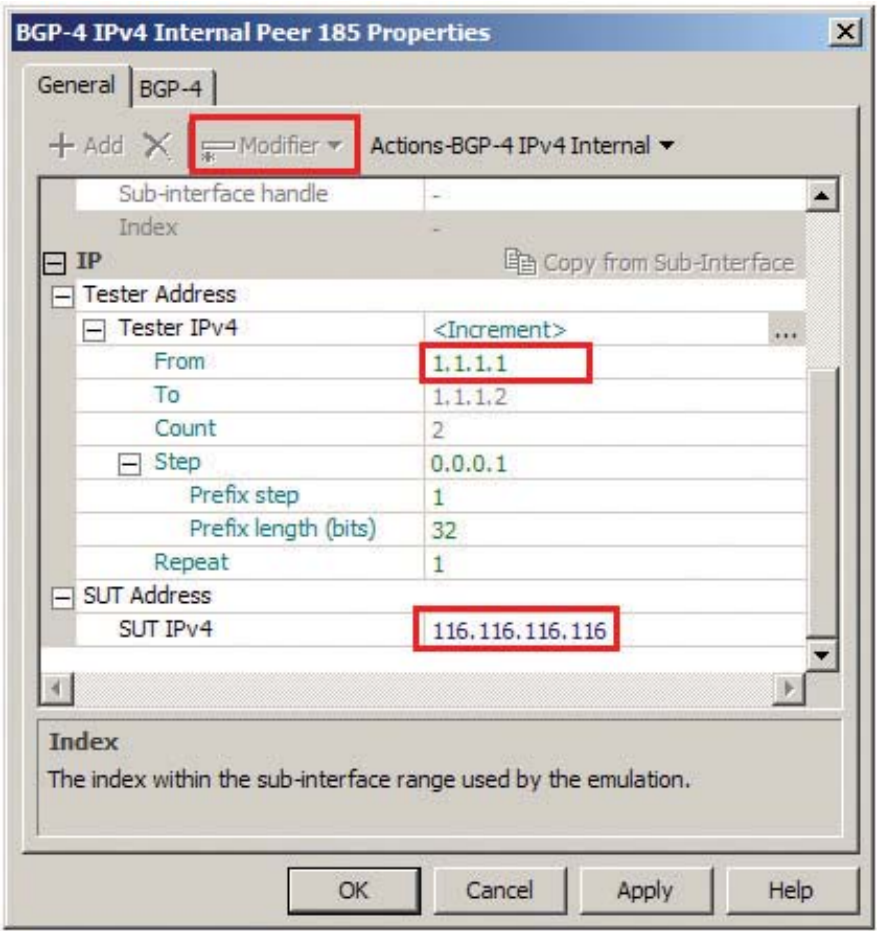

► Step 41

Modify the following values to set the tester and SUT IP addresses of the peer pool to match the addresses simulated through OSPF:

- Tester IPv4 = 1.1.1.1 (address of PE1 simulated through OSPF)
- Ensure that To value for Tester IPv4 shows IP address for PE2 (1.1.1.2)
- Change the Modifier of the SUT IPv4 field to "None", and set the SUT IP address to the loopback address of the DUT (116.116.116.116)

► Step 42

Change to the "BGP-4" tab of the configuration dialog.

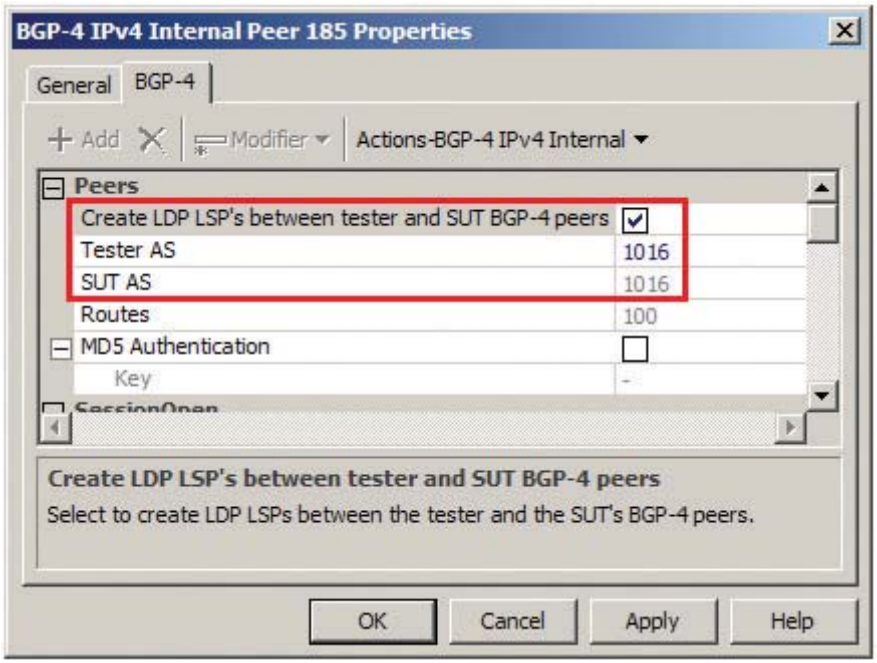

► Step 43

Modify the following values to set the AS numbers and enable auto LDP LSP creation:

- Tester AS = 1016
- SUT AS = 1016 (will change to automatically match tester AS)
- Select "Create LDP LSP's between tester and SUT BGP-4 peers"

Core – Advertise CEs on core side with VPNv4 route profile

► Step 44

Select the peer pool, and click the arrow next to "New" on the toolbar of the Emulation pane. Select "BGP-4 VPN IPv4 Route Profile".

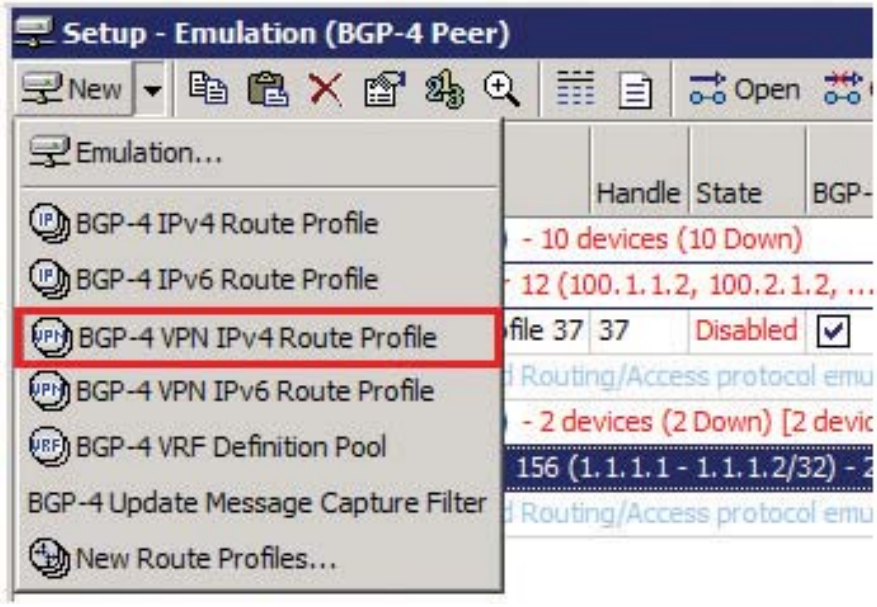

► Step 45

Double-click on the created route profile to edit it.

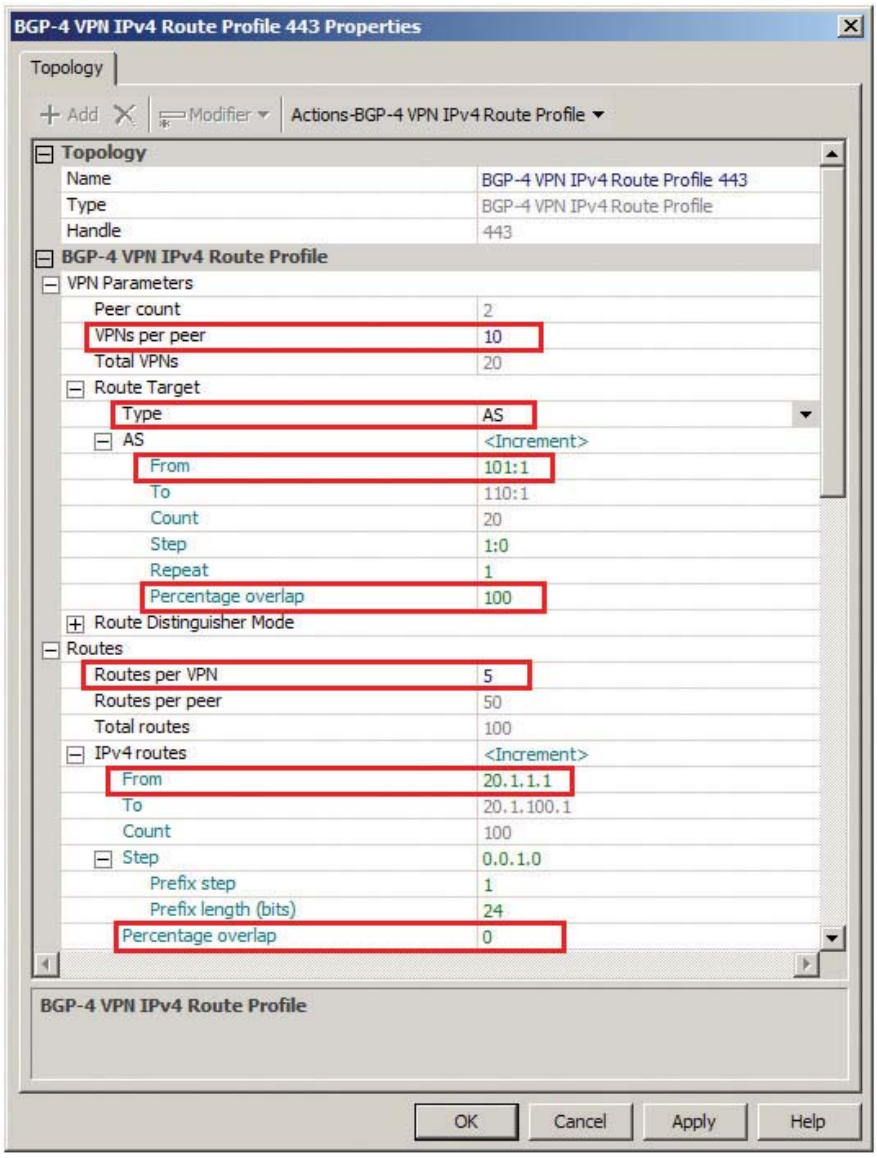

► Step 46

Modify the following values in the route profile to simulate VPN labeled routes behind simulated PE1 and PE2:

- VPNs per peer = 10
- Route Target
  - Type = AS
  - From = 101:1
  - Percentage overlap = 100 (simulating access to same VPNs from each simulated PE)
- Routes per VPN = 5
- IPv4 routes
  - From = 20.1.1.1
  - Percentage overlap = 0 (unique routes per VPN in this particular application note, but because it is a VPN, routes could overlap)

► Step 47

Click "OK" to add the VPNv4 route profile.

Core – Add VRF pool to store incoming VPN routes

► Step 48

Right-click on the VPNv4 route profile, and click "Create Matching VRF Pool". This is a short-cut which creates a VRF pool with the Import Route Target range matching the Export Route Target range of the VPNv4 route profile.

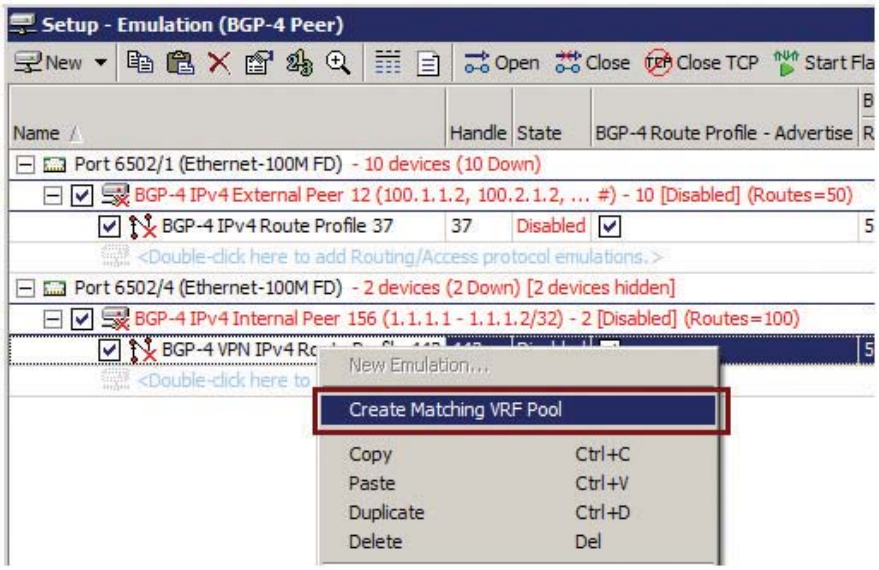

► Step 49

Edit the VPN VRF pool.

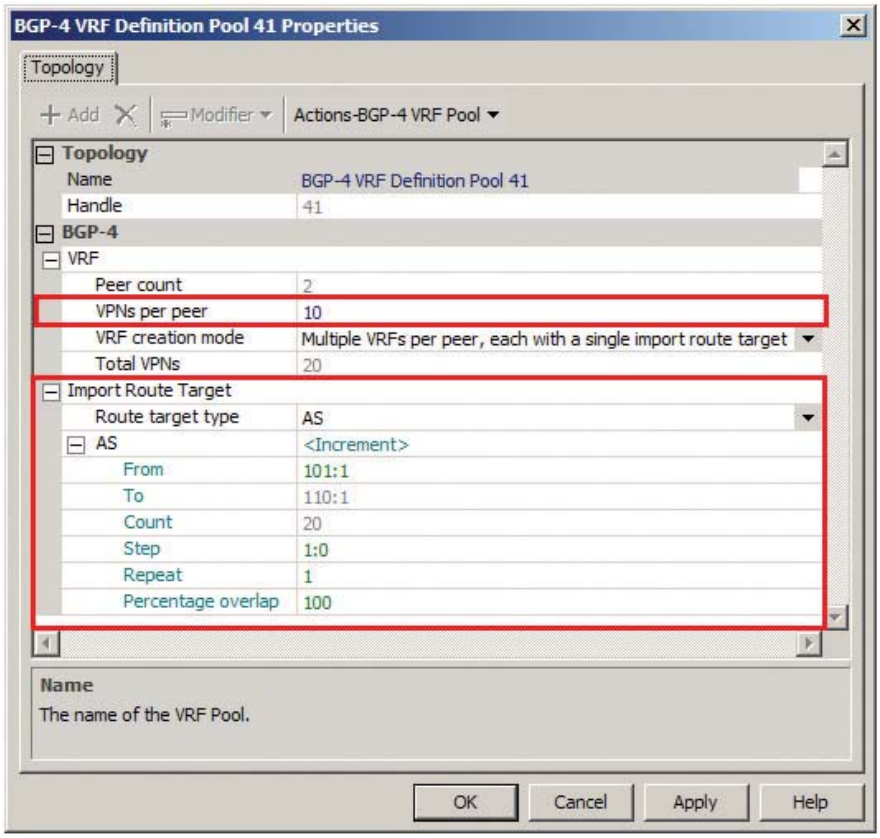

► Step 50

Observe that the highlighted values have been automatically inherited from the VPNv4 route profile. Click "OK" or "Cancel" to close the VPN VRF pool properties dialog.

Bring up the control plane

▶ Step 51

Click the "Summary" radio button on the left of the screen (above BGP-4, OSPF and LDP) to show all protocol sessions. Start the Routing Engine by clicking the "Routing" button on the top toolbar of the main application.

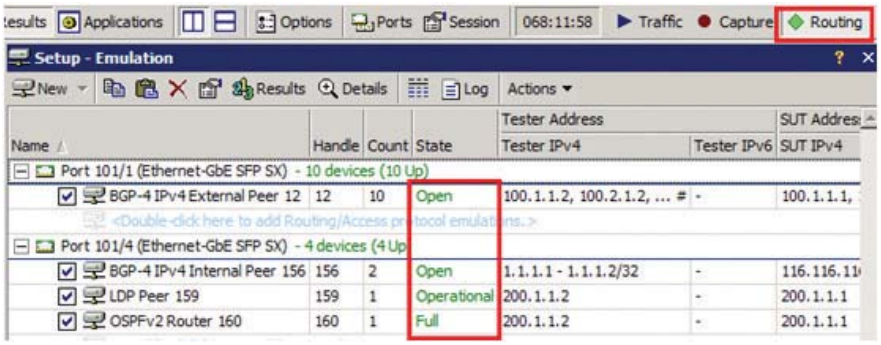

▶ Step 52

Wait for the routing and MPLS protocol sessions to converge and reach their final state:

- OSPF will reach the "Full" state
- LDP will reach the "Operational" state
- BGP-4 will reach the "Open" state (if this doesn't occur, manually open the BGP-4 peers from the GUI)

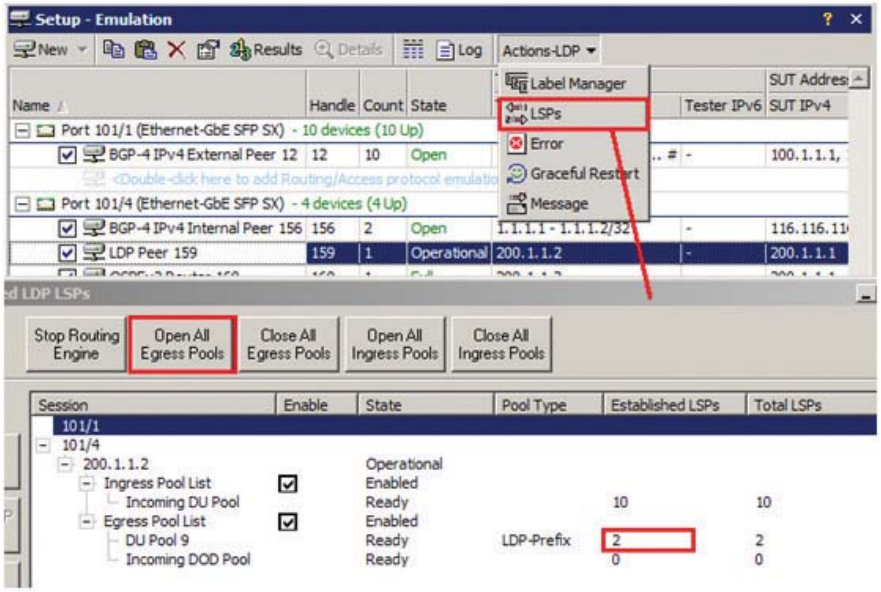

The following steps will ensure that the LDP LSPs are open.

► Step 53

Select the LDP peer, and then select "LSPs" from the Actions-LDP menu on the toolbar. You should see two LSPs in the Established LSPs column (the incoming DU pool in the Ingress Pool List may show more LSPs, as the DUT may create additional LSPs for other reachable destinations).

► Step 54

If the LSPs haven't established, click "Open All Egress Pools" and they will open.

Configure traffic from edge to core (the difficult way)

► Step 55

Add an IPv4 traffic mesh by clicking on "Traffic" on the Setup pane on the top-left of the main application, clicking the arrow next to "New" on the Traffic toolbar, and selecting "IPv4 mesh".

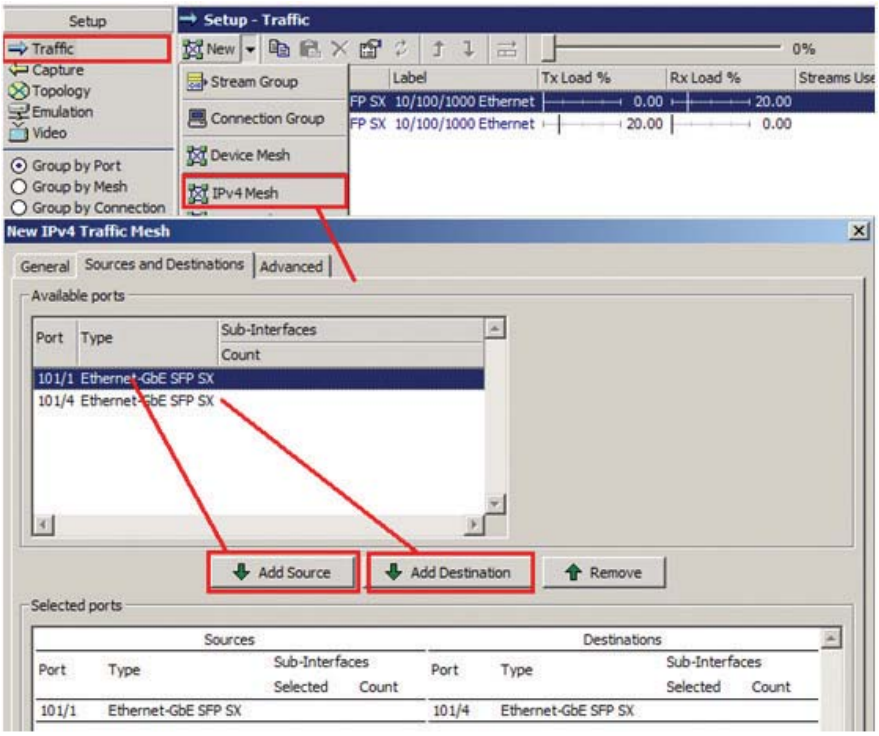

► Step 56

Change to the "Sources and Destinations" tab, and configure the mesh by doing the following:

- Select the edge port and click "Add Source".
- Select the core port and click "Add Destination".
- Click "OK" to create the IPv4 traffic mesh.

NOTE: Remember how you selected Traffic Destinations when you added the VPNv4 route profile before? This was done to make the mesh configuration really simple here – the mesh will automatically transmit to all routes which have "Traffic Destinations" selected on the destination port.

► Step 57

You will notice in the bottom-left of the screen that the mesh has automatically created 20 stream groups. These represent each VLAN on the edge side transmitting to the set of routes (that belong to the same VPN as the VLAN) behind each PE on the core side. There are 10 VLANs on the left, and they are each transmitting to the routes behind each CE on the edge which belong to the same VPN as them. Since there are 2 CEs on the edge belong to each VPN (1 behind PE1, and 1 behind PE2), this results in 20 stream groups.

Unfortunately, there isn't 100% traffic integration in this scenario, and some manual editing of the stream groups is required. The IPv4 mesh uses the "default" VLAN ID and L3 Source address from the link layer of the source port, which is the first address pool that has been added. This means that while the destination addresses are correct, the correct VLANs aren't transmitting to their corresponding VPNs so the router will not forward the traffic! You will need to now manually edit the created stream groups so that the correct VLANs are transmitting to the correct VPNs.

Name	Packet	VLAN IDs	L3 Source	L3 Destination	Streams	Connections	Length
TrafficMesh 65 (10.00%)							
Port 101/1 (10.00% of TX line rate)							
AGT_CONSTANT_PROFILE15 (148809.5 Fps)							
TrafficMesh 65/1	IPv4/Ethernet	1	10.1.1.1	20.1.1.1-20.1.5.1	1		L2: 64
TrafficMesh 65/2	IPv4/Ethernet	2	10.1.6.1	20.1.6.1-20.1.10.1	1		L2: 64
TrafficMesh 65/3	IPv4/Ethernet	3	10.1.11.1	20.1.11.1-20.1.15.1	1		L2: 64
TrafficMesh 65/4	IPv4/Ethernet	4	10.1.16.1	20.1.16.1-20.1.20.1	1		L2: 64
TrafficMesh 65/5	IPv4/Ethernet	5	10.1.21.1	20.1.21.1-20.1.25.1	1		L2: 64
TrafficMesh 65/6	IPv4/Ethernet	6	10.1.26.1	20.1.26.1-20.1.30.1	1		L2: 64
TrafficMesh 65/7	IPv4/Ethernet	7	10.1.31.1	20.1.31.1-20.1.35.1	1		L2: 64
TrafficMesh 65/8	IPv4/Ethernet	8	10.1.36.1	20.1.36.1-20.1.40.1	1		L2: 64
TrafficMesh 65/9	IPv4/Ethernet	9	10.1.41.1	20.1.41.1-20.1.45.1	1		L2: 64
TrafficMesh 65/10	IPv4/Ethernet	10	10.1.46.1	20.1.46.1-20.1.50.1	1		L2: 64
TrafficMesh 65/11	IPv4/Ethernet	1	10.1.1.1	20.1.51.1-20.1.55.1	1		L2: 64
TrafficMesh 65/12	IPv4/Ethernet	2	10.1.6.1	20.1.56.1-20.1.60.1	1		L2: 64
TrafficMesh 65/13	IPv4/Ethernet	3	10.1.11.1	20.1.61.1-20.1.65.1	1		L2: 64
TrafficMesh 65/14	IPv4/Ethernet	4	10.1.16.1	20.1.66.1-20.1.70.1	1		L2: 64
TrafficMesh 65/15	IPv4/Ethernet	5	10.1.21.1	20.1.71.1-20.1.75.1	1		L2: 64
TrafficMesh 65/16	IPv4/Ethernet	6	10.1.26.1	20.1.76.1-20.1.80.1	1		L2: 64
TrafficMesh 65/17	IPv4/Ethernet	7	10.1.31.1	20.1.81.1-20.1.85.1	1		L2: 64
TrafficMesh 65/18	IPv4/Ethernet	8	10.1.36.1	20.1.86.1-20.1.90.1	1		L2: 64
TrafficMesh 65/19	IPv4/Ethernet	9	10.1.41.1	20.1.91.1-20.1.95.1	1		L2: 64
TrafficMesh 65/20	IPv4/Ethernet	10	10.1.46.1	20.1.96.1-20.1.100.1	1		L2: 64

► Step 58

Change the following values:

- Change "VLAN IDs" and "L3 Source addresses" so that traffic is transmitted from correct VLAN/source IP address to correct destination IP addresses (e.g. VLAN ID 1 is on VPN1, so need to select "VLAN ID and source IP address" from VPN1 to transmit to destination routes which are also in VPN1 on the core side). This can be done via inline editing in the front panel (refer to the topology diagram to ensure the correct VLANs are transmitting to the correct VPN).
- Repeat this for all stream groups (or as many as you want, as it can get tedious).

Configure traffic from core to edge

► Step 59

Add a BGP-4 MPLS VPN traffic mesh by clicking on "Traffic" on the Setup pane on the top-left of the main application, clicking the arrow next to "New" on the Traffic toolbar, and selecting "BGP-4 MPLS VPN Mesh".

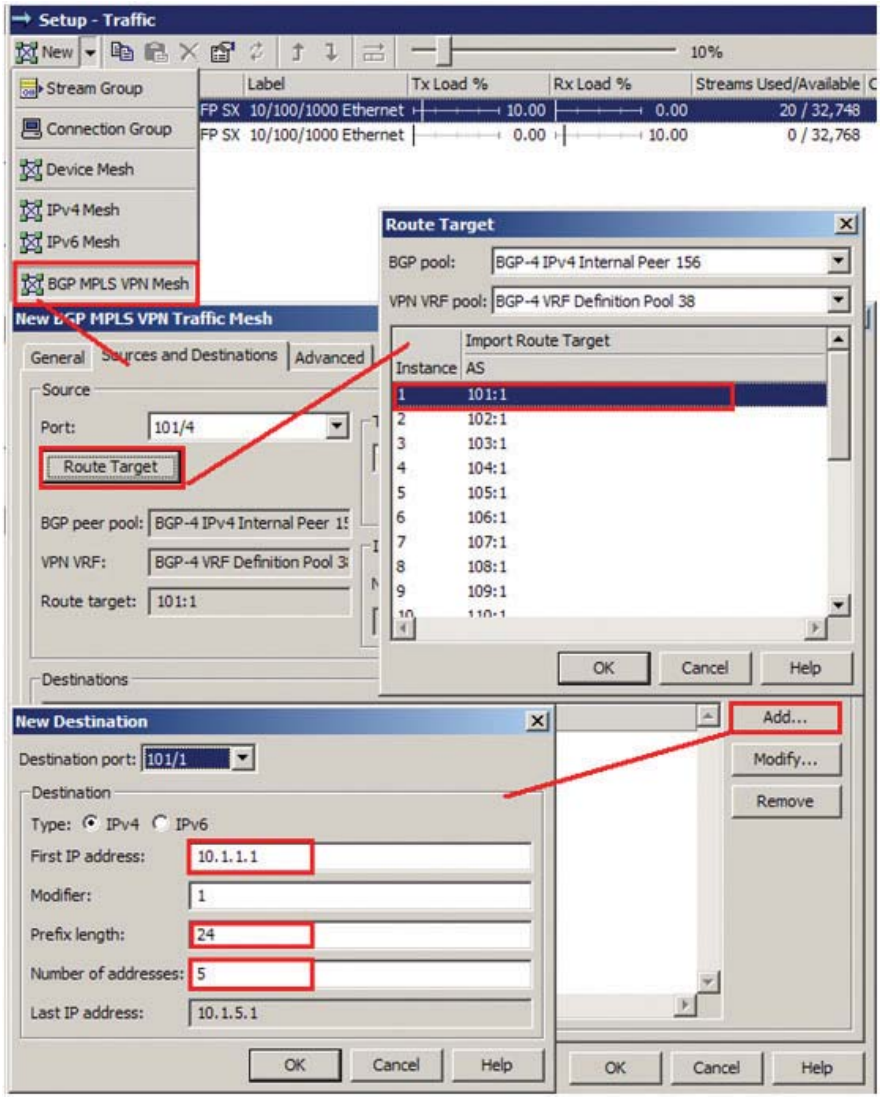

► Step 60

Change to the "Sources and Destinations" tab. Click "Route Target" and select the route target on the first VPN (101:1) as the source. Click "OK" to add the source.

► Step 61

Click "Add" in the "Destinations" tab to add a new destination. Configure the parameters to the following values:

- First IP address = 10.1.1.1 (i.e. first address in VLAN on edge port which belongs to the same VPN)
- Prefix length = 24
- Number of addresses = 5 (as we had configured 5 routes per peer on edge port)

► Step 62

Click "OK" to add the destination, and click "OK" again in the main mesh configuration dialog to add the BGP-4 MPLS VPN mesh.

► Step 63

Repeat the previous 4 steps for other 9 VPNs, ensuring that for each mesh, you select the correct route target for the VPN as the source, and the correct destination IP address range on the VLAN that belongs to the same VPN (or do as many as you want, as it can get tedious).

► Step 64

Observe that we have created a single stream group per VPN, which is not very scalable as we may run out of stream group resources if we have a lot of VPNs. We will see how the super VRF concept can help improve scalability in this area in the next section of the application note.

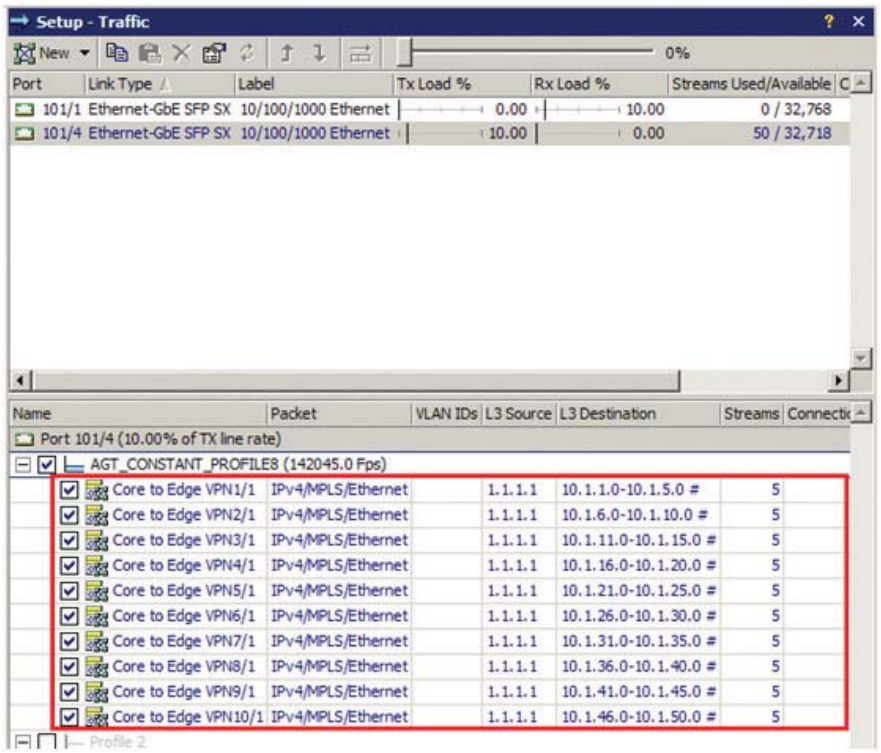

Start traffic and analyse results

► Step 65

Start the traffic by clicking the "Traffic" button on the top toolbar of the main application.

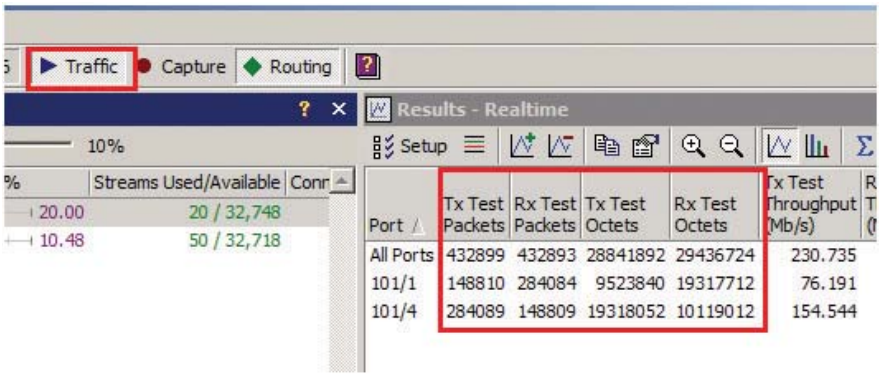

► Step 66

Observe in the results pane that 100% of traffic that is being transmitted from the edge port is being received on the core port, and vice-versa.

► Step 67

Change to the Capture view by clicking on "Capture" on the Setup pane on the top-left of the main application.

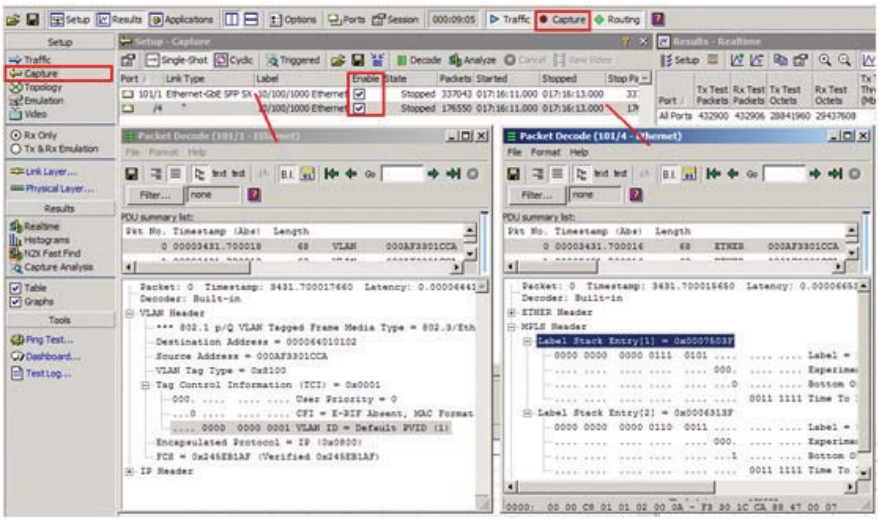

► Step 68

Select the ports to capture by selecting the checkboxes in the Enable column.

► Step 69

Start capture by clicking the "Capture" button (in between Traffic and Routing). The capture buffer will fill up pretty quickly, and capture will stop automatically.

► Step 70

Double-click on the core port in the Capture view to view the packets captured there. Observe that the capture buffer contains MPLS labeled traffic with a 2 label stack (inner VPN label, and outer LSP tunnel label).

► Step 71

Double-click on the edge port in the Capture view to view the packets captured there. Observe that the capture buffer contains standard IP packets with the correct VLAN IDs.

Part #2: Using super VRF to improve core to edge traffic scalability

In this part of the application note we will show how a super VRF can be used to improve the scalability of the core to edge traffic. Note that this part of the application note builds on part #1, and requires that to be completed prior to commencing this application note.

Change the VPN VRF pool to 'Super VRF' mode

► Step 72

Stop Traffic, and remove the existing core to edge BGP-4 MPLS VPN traffic meshes.

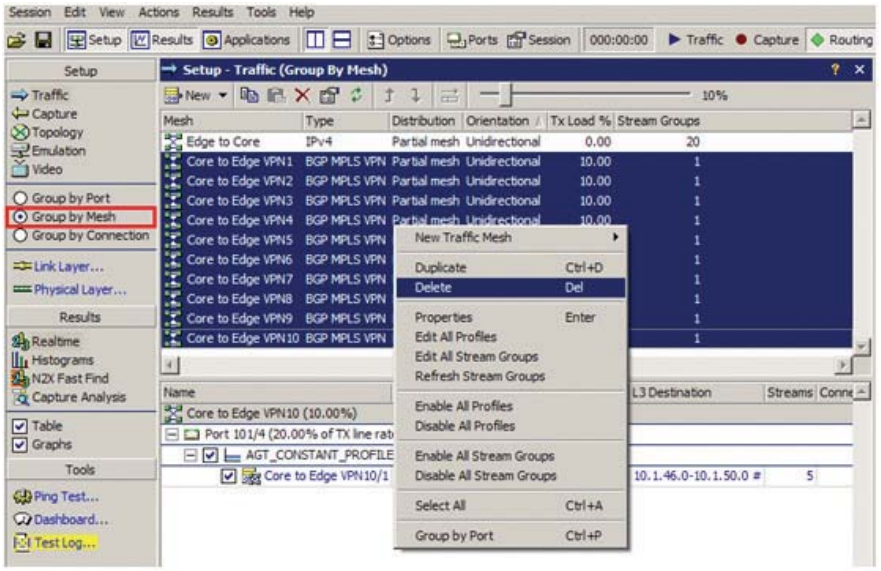

► Step 73

Disable the VPN VRF pool and double-click on it to edit it.

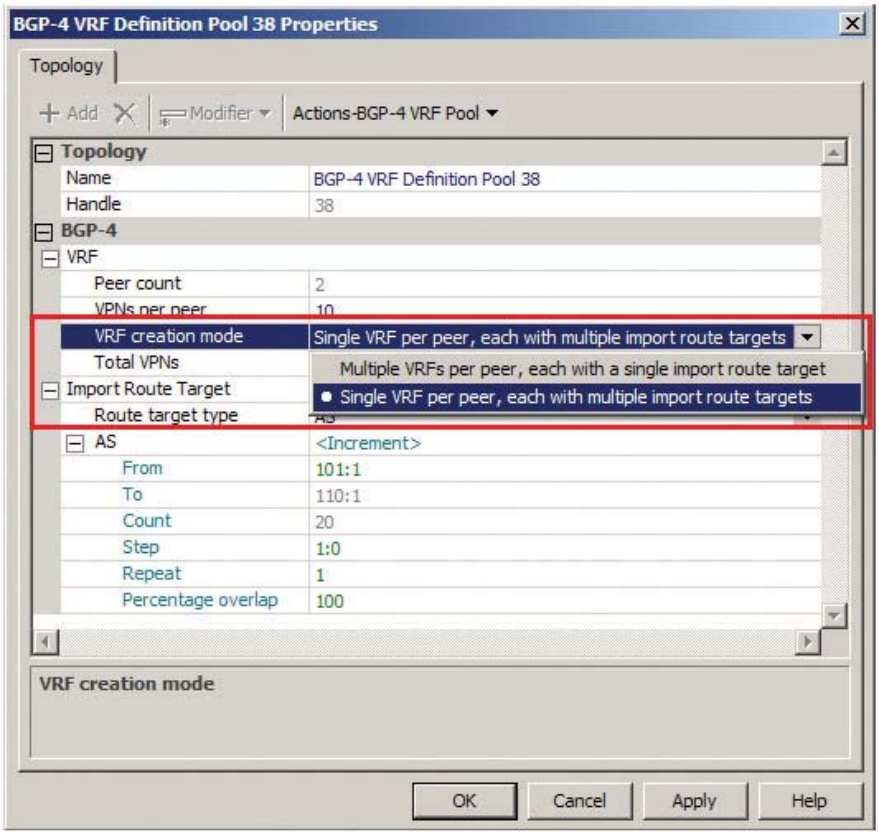

► Step 74

Put the VPN VRF pool into "Super VRF" mode by changing the VRF creation mode to "Single VRF per peer". Instead of there being one VRF per VPN per PE, a super VRF allows us to create a single VRF for all VPNs attached to that PE (with an import route target range).

► Step 75

Click "OK" to apply the changes to the VRF pool. Re-enable the VPN VRF pool again.

View the VPN VRF table.

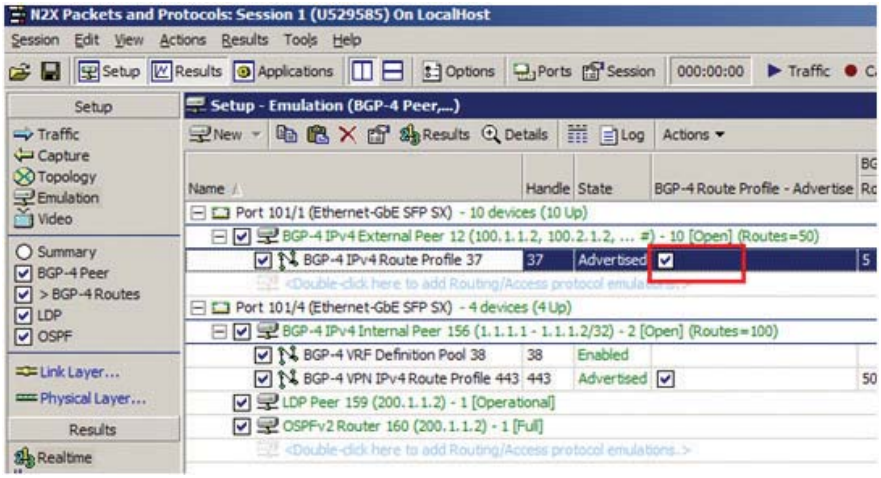

► Step 76

Re-advertise edge-side routes to re-populate the super VRF table on the core side. You can do this by un-checking the Advertise checkbox to withdraw the routes, and checking it again to re-advertise the routes.

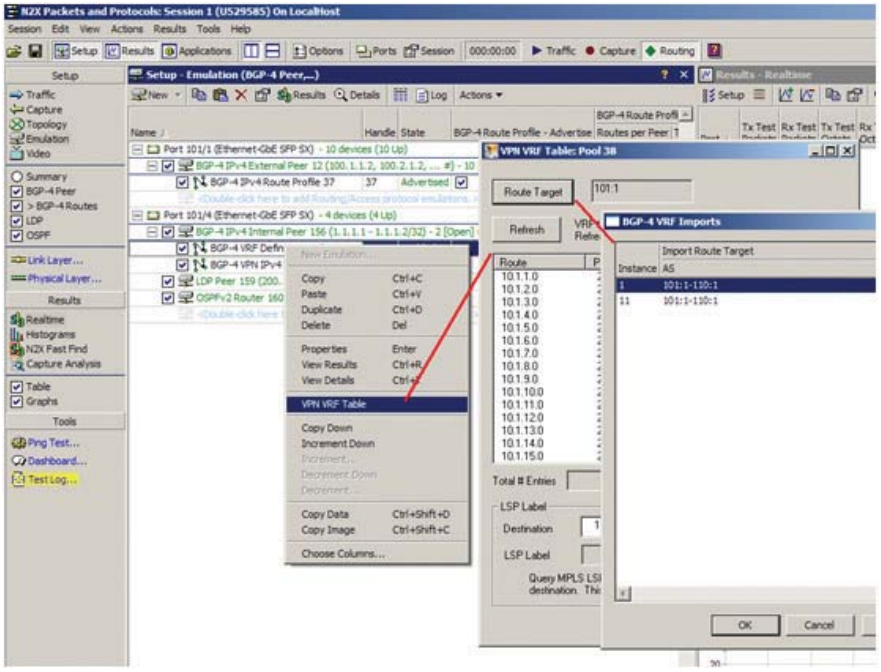

The following steps will view the VPN VRF table to ensure that it is correctly populated with routes from the edge.

► Step 77

Right-click on the VRF pool, and select "VPN VRF Table".

► Step 78

Click "Route Target" in the VPN VRF Table dialog, and select the VRF range representing the PE whose VRF table you want to see.

► Step 79

Click "OK" and the VRF table for that PE will be displayed.

Re-configure core to edge traffic using super VRF

► Step 80

Re-configure the traffic from the core to the edge using super VRF as a source. Add a BGP-4 MPLS VPN traffic mesh by clicking on "Traffic" on the Setup pane on the top-left of the main application, clicking the arrow next to "New" on the Traffic toolbar, and selecting "BGP-4 MPLS VPN Mesh".

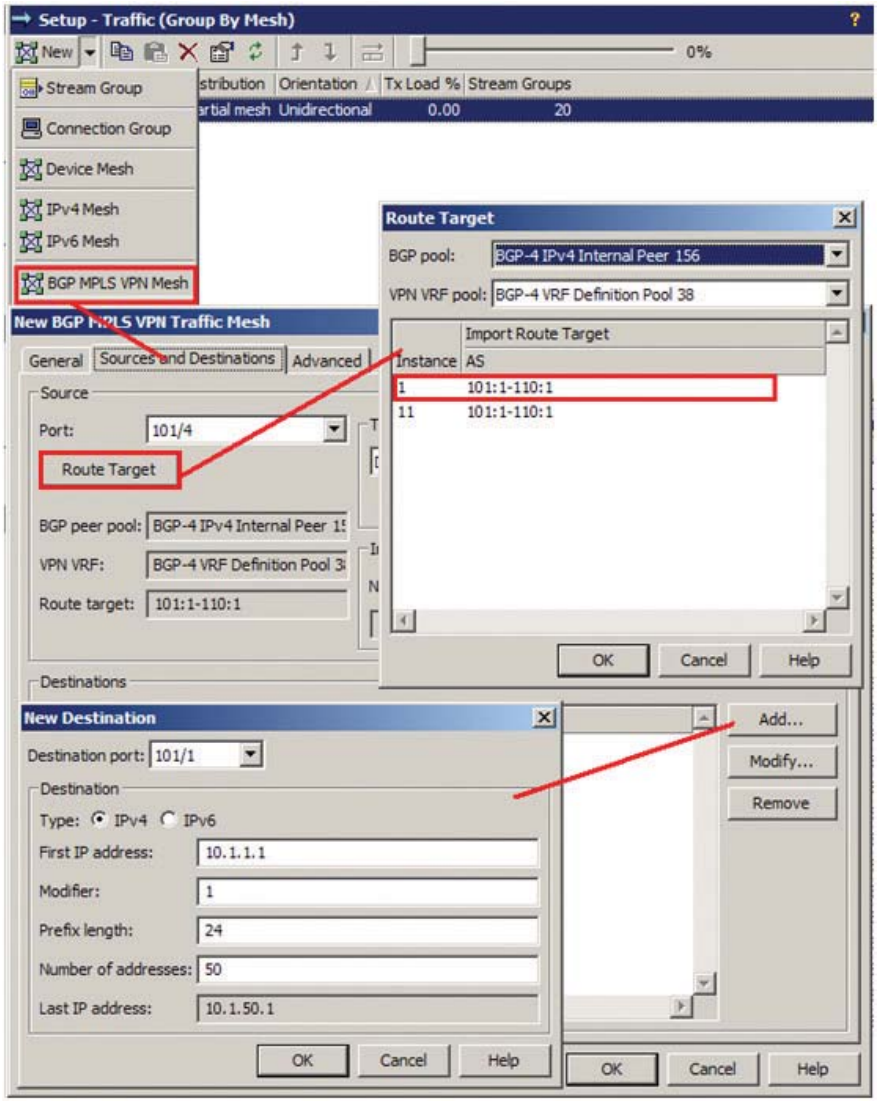

► Step 81

Change to the "Sources and Destinations" tab. Click "Route Target" and select the super VRF route target range on PE1 (101:1-X10:1) as the source. Click "OK" to add the source.

► Step 82

Click "Add" in the "Destinations" tab to add a new destination. Configure parameters to the following values:

- First IP address = 10.1.1.1 (i.e. first address in VLAN on edge port which belongs to the same VPN)
- Prefix length = 24
- Number of addresses = 50 (as we are transmitting to all destination VPNs using the same super VRF)

► Step 83

Click "OK" to add the destination, and click "OK" again in the main mesh configuration dialog to add the BGP-4 MPLS VPN mesh.

Normally we would create another traffic mesh to simulate traffic from PE2, but we will skip this step now to save time.

► Step 84

Observe the stream groups that have been created. Note that by using a super VRF on the core port, we have significantly saved stream group resources as we only need to add a single BGP-4 MPLS VPN traffic mesh (and thus stream group) per PE, not a single BGP-4 MPLS VPN traffic mesh (and thus stream group) per VPN. This means that the number of stream groups will always be proportional to the number of PEs, not VPNs, and we can scale up our scenario to 1000's of VPNs without much of a performance impact on the tester.

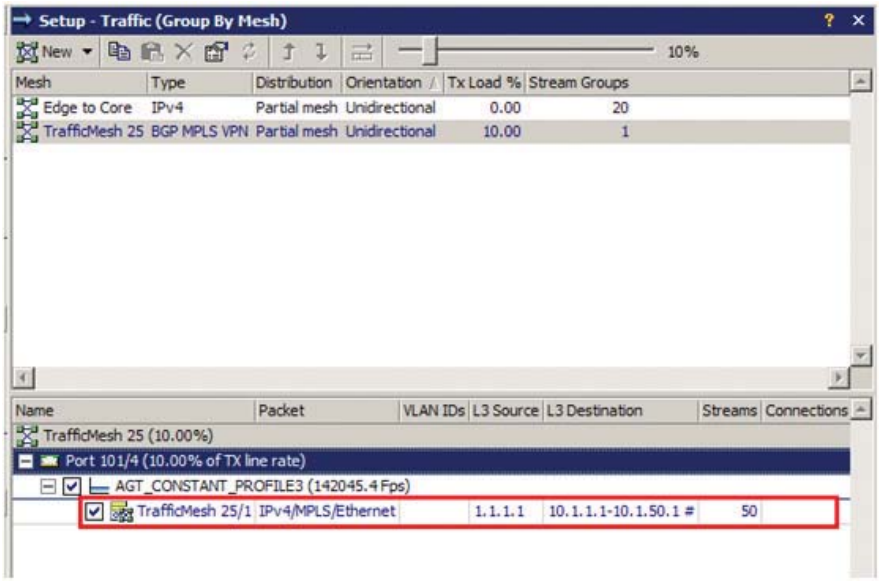

Start traffic and analyse results

► Step 85

Start the traffic by clicking the "Traffic" button on the top toolbar of the main application.

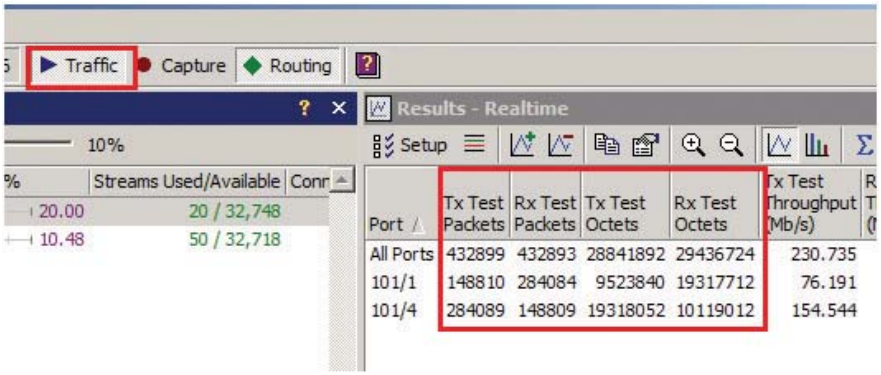

► Step 86

Observe in the Results pane that 100% of traffic that is being transmitted from the edge port is being received on the core port, and vice-versa.

► Step 87

Change to the Capture view by clicking on "Capture" on the Setup pane on the top-left of the main application.

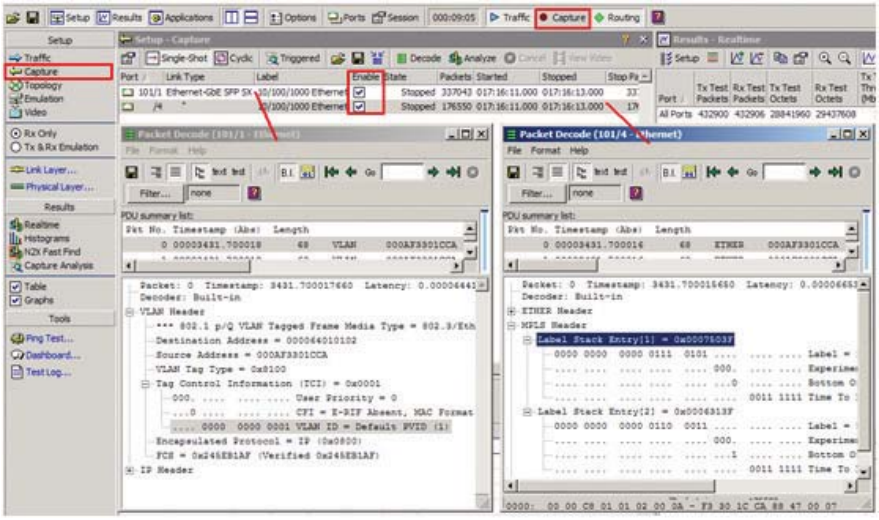

► Step 88

Select the ports to capture by selecting the checkboxes in the Enable column.

► Step 89

Start capture by clicking the "Capture" button (in between Traffic and Routing). The capture buffer will fill up pretty quickly, and capture will stop automatically.

► Step 90

Double-click on the core port in the Capture view to view the packets captured there. Observe that the capture buffer contains MPLS labeled traffic with a 2 label stack (inner VPN label, and outer LSP tunnel label).

► Step 91

Double-click on the edge port in the Capture view to view the packets captured there. Observe that the capture buffer contains standard IP packets with the correct VLAN IDs.

Part #3: Using L3MPLS VPN Traffic configuration QuickTool

In this part of the application note we will show the L3MPLS VPN Traffic configuration QuickTool can be used to completely automate the setup of core to edge and edge to core traffic in the previous application notes. Note that this part of the application note builds on part #2, and requires that to be completed prior to commencing this application note.

Clean up previously created traffic

► Step 92

Stop Traffic, and remove all traffic stream groups and meshes previously created in parts 1 and 2 of this application note.

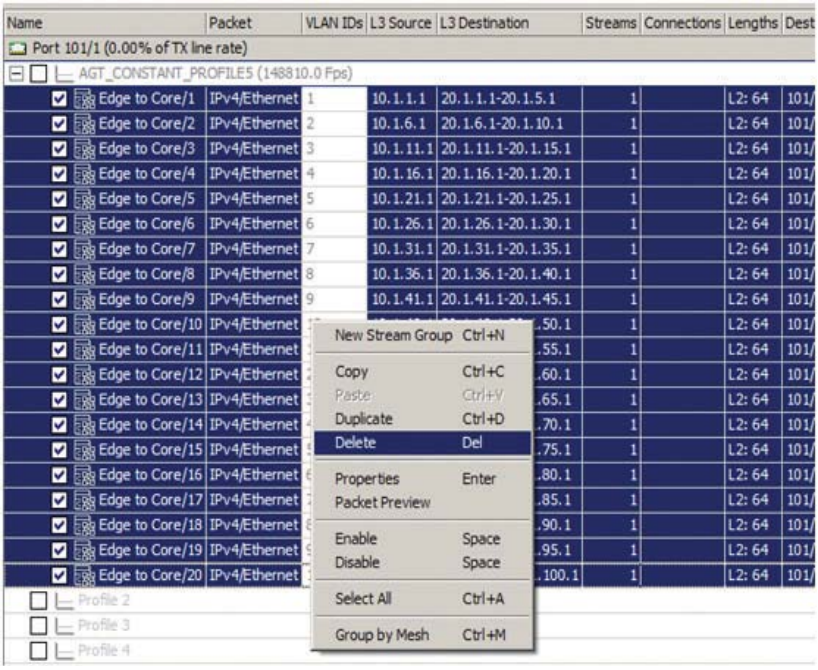

Launch and configure the CreateL3BgpMplsVpnTraffic QuickTool

► Step 93

Launch the CreateL3BgpMplsVpnTraffic QuickTool. You will see the following screen.

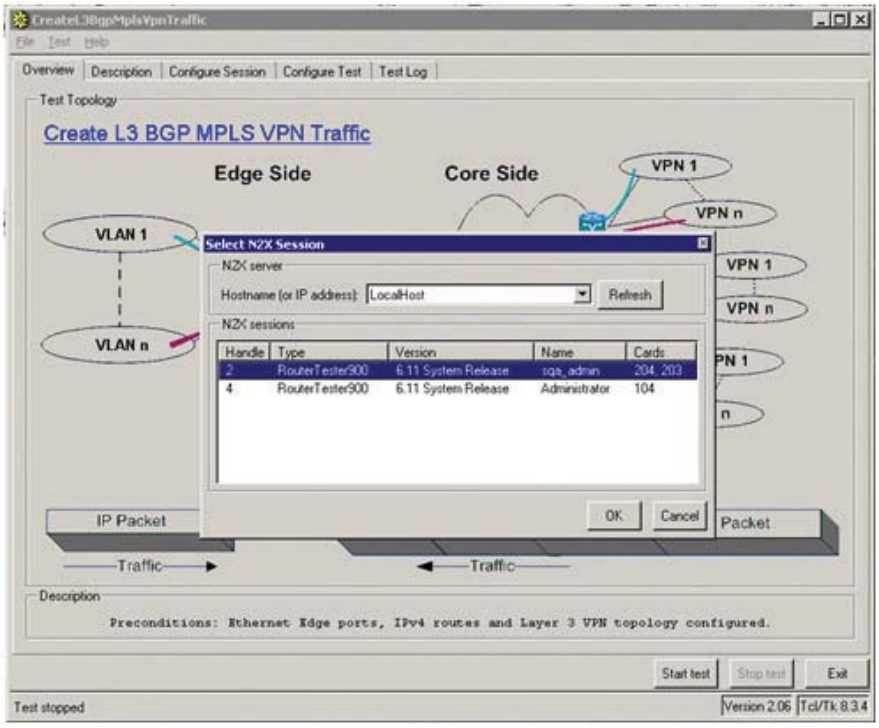

► Step 94

Attach the QuickTool to your N2X session using the Select N2X Session dialog. The Allocate Test Ports dialog will now come up.

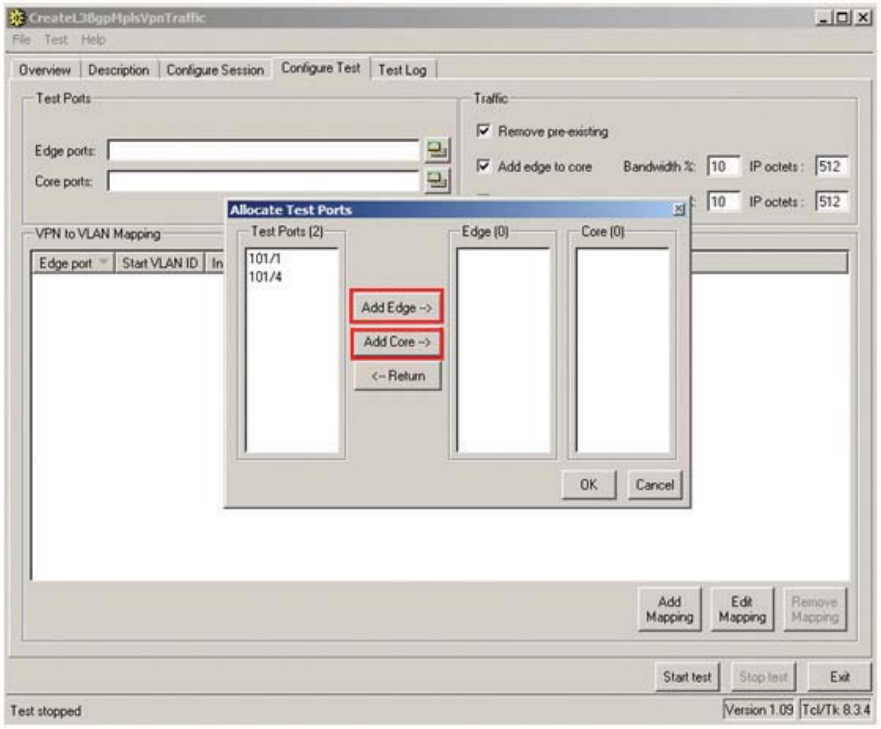

► Step 95

Select the edge port in the Test Ports list and click "Add Edge". Select the core port in the Test Ports list and click "Add Core". Click "OK" to complete allocation of test ports.

► Step 96

Change to the "Configure Test" tab.

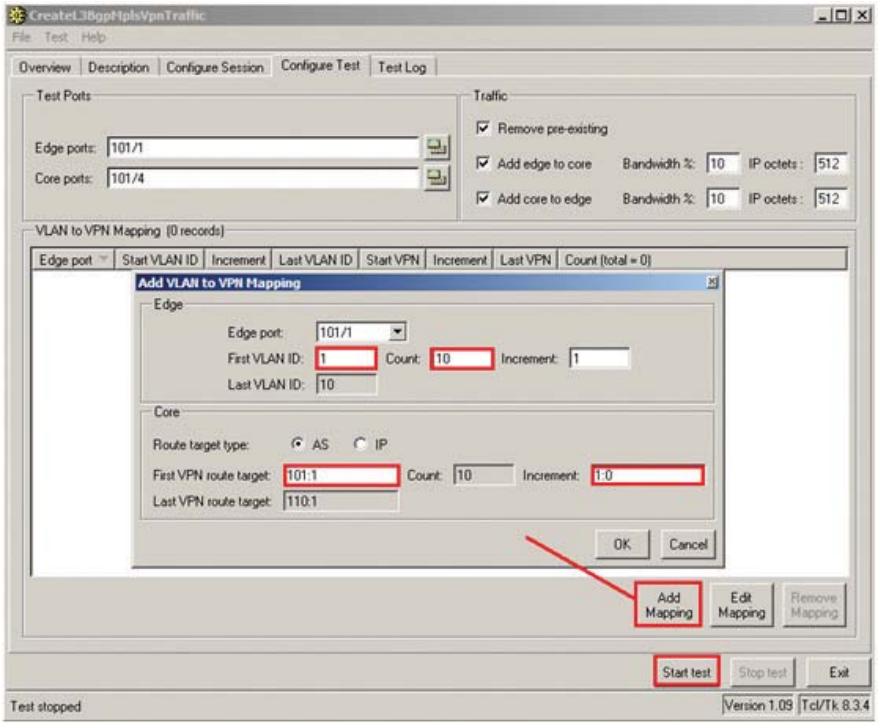

► Step 97

Click "Add Mapping" to bring up the Add VLAN to VPN Mapping dialog. Configure the mappings like this:

- Edge mapping (VLAN)
  - Edge port = <Your edge port>
  - First VLAN ID = 1
  - Count = 10
  - Increment = 1
- Core mapping (VPN)
  - Route target type = AS
  - First VPN route target = 101:1
  - Count = 10
  - Increment = 1:0

► Step 98

Click "Start test" and observe traffic meshes/stream groups that have been created in N2X GUI. This has significantly simplified the traffic configuration, especially the traffic from the edge to the core (i.e. all the correct VLAN IDs have been inserted into the traffic to match the destination IP addresses which are on the same VPN).

Port 101/1 (10.00% of TX line rate)					
101/1 Edge_to_Core CONSTANT_PROFILE6 (10.00% of TX line rate)					
101/1->101/4_VLAN_1->101:1 SG 80	IPv4/Ethernet	1	100.1.1.2	20.1.1.1-20.1.5.1	
101/1->101/4_VLAN_1->101:1 SG 81	IPv4/Ethernet	1	100.1.1.2	20.1.51.1-20.1.55.1	
101/1->101/4_VLAN_2->102:1 SG 82	IPv4/Ethernet	2	100.2.1.2	20.1.6.1-20.1.10.1	
101/1->101/4_VLAN_2->102:1 SG 83	IPv4/Ethernet	2	100.2.1.2	20.1.56.1-20.1.60.1	
101/1->101/4_VLAN_3->103:1 SG 84	IPv4/Ethernet	3	100.3.1.2	20.1.11.1-20.1.15.1	
101/1->101/4_VLAN_3->103:1 SG 85	IPv4/Ethernet	3	100.3.1.2	20.1.61.1-20.1.65.1	
101/1->101/4_VLAN_4->104:1 SG 86	IPv4/Ethernet	4	100.4.1.2	20.1.16.1-20.1.20.1	
101/1->101/4_VLAN_4->104:1 SG 87	IPv4/Ethernet	4	100.4.1.2	20.1.66.1-20.1.70.1	
101/1->101/4_VLAN_5->105:1 SG 88	IPv4/Ethernet	5	100.5.1.2	20.1.21.1-20.1.25.1	
101/1->101/4_VLAN_5->105:1 SG 89	IPv4/Ethernet	5	100.5.1.2	20.1.71.1-20.1.75.1	
101/1->101/4_VLAN_6->106:1 SG 90	IPv4/Ethernet	6	100.6.1.2	20.1.26.1-20.1.30.1	
101/1->101/4_VLAN_6->106:1 SG 91	IPv4/Ethernet	6	100.6.1.2	20.1.76.1-20.1.80.1	
101/1->101/4_VLAN_7->107:1 SG 92	IPv4/Ethernet	7	100.7.1.2	20.1.31.1-20.1.35.1	
101/1->101/4_VLAN_7->107:1 SG 93	IPv4/Ethernet	7	100.7.1.2	20.1.81.1-20.1.85.1	
101/1->101/4_VLAN_8->108:1 SG 94	IPv4/Ethernet	8	100.8.1.2	20.1.36.1-20.1.40.1	
101/1->101/4_VLAN_8->108:1 SG 95	IPv4/Ethernet	8	100.8.1.2	20.1.86.1-20.1.90.1	
101/1->101/4_VLAN_9->109:1 SG 96	IPv4/Ethernet	9	100.9.1.2	20.1.41.1-20.1.45.1	
101/1->101/4_VLAN_9->109:1 SG 97	IPv4/Ethernet	9	100.9.1.2	20.1.91.1-20.1.95.1	
101/1->101/4_VLAN_10->110:1 SG 98	IPv4/Ethernet	10	100.10.1.2	20.1.46.1-20.1.50.1	
101/1->101/4_VLAN_10->110:1 SG 99	IPv4/Ethernet	10	100.10.1.2	20.1.96.1-20.1.100.1	

► Step 99

Experiment with other parts of the QuickTool. Some things to try:

- Remove pre-existing traffic
- Add only some VLAN to VPN mappings the first time the test is run, then incrementally add other mappings.

Appendix A – Router configuration

NB: For the interests of paper conservation, this only includes the configuration for one group (i.e. a single edge port, and a single core port). Additional configuration settings will need to be added for the remaining groups.

```
!
version 12.0
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname GSR-12008
!
boot-start-marker
boot system flash slot1:gsr-p-mz.120-30.S1.bin
boot bootldr bootflash:gsr-boot-mz.120-30.S1.bin
boot-end-marker
!
redundancy
 mode rpr-plus
logging console emergencies
enable password gsr12000
!
username all
monitor event-trace rlc all enable
!
ip vrf v101
 rd 101:1
 route-target export 101:1
 route-target import 101:1
 bgp next-hop Loopback0
!
ip vrf v102
 rd 102:1
 route-target export 102:1
 route-target import 102:1
 bgp next-hop Loopback0
!
ip vrf v103
 rd 103:1
 route-target export 103:1
 route-target import 103:1
 bgp next-hop Loopback0
!
ip vrf v104
 rd 104:1
 route-target export 104:1
 route-target import 104:1
 bgp next-hop Loopback0
!
ip vrf v105
 rd 105:1
 route-target export 105:1
 route-target import 105:1
 bgp next-hop Loopback0
!
ip vrf v106
 rd 106:1
 route-target export 106:1
 route-target import 106:1
 bgp next-hop Loopback0
!
ip vrf v107
 rd 107:1
 route-target export 107:1
 route-target import 107:1
 bgp next-hop Loopback0
!
ip vrf v108
 rd 108:1
 route-target export 108:1
 route-target import 108:1
 bgp next-hop Loopback0
!
ip vrf v109
 rd 109:1
 route-target export 109:1
 route-target import 109:1
 bgp next-hop Loopback0
!
ip vrf v110
 rd 110:1
 route-target export 110:1
 route-target import 110:1
 bgp next-hop Loopback0
!
!
!
ip subnet-zero
ip cef table hardware resource-failure action punt
ip multicast-routing distributed
frame-relay switching
clns routing
mpls label protocol ldp
ipv6 unicast-routing
ipv6 multicast-routing
!
interface Loopback0
 ip address 116.116.116.116 255.255.255.255
 no ip route-cache
!
interface Tunnel1
 ip unnumbered Loopback0
 tunnel destination 1.1.1.1
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng priority 0 0
 tunnel mpls traffic-eng bandwidth  100
 tunnel mpls traffic-eng path-option 1 dynamic
!
interface Tunnel2
 ip unnumbered Loopback0
 tunnel destination 1.1.1.2
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng priority 0 0
 tunnel mpls traffic-eng bandwidth  100
 tunnel mpls traffic-eng path-option 1 dynamic
!
interface FastEthernet0/0
 no ip address
!
interface FastEthernet0/0.1
 encapsulation dot1Q 1
 ip vrf forwarding v101
 ip address 100.1.1.1 255.255.255.0
 no cdp enable
!
interface FastEthernet0/0.2
 encapsulation dot1Q 2
 ip vrf forwarding v102
 ip address 100.1.2.1 255.255.255.0
 no cdp enable
!
interface FastEthernet0/0.3
 encapsulation dot1Q 3
 ip vrf forwarding v103
 ip address 100.1.3.1 255.255.255.0
 no cdp enable
!
interface FastEthernet0/0.4
 encapsulation dot1Q 4
 ip vrf forwarding v104
 ip address 100.1.4.1 255.255.255.0
 no cdp enable
!
interface FastEthernet0/0.5
 encapsulation dot1Q 5
 ip vrf forwarding v105
 ip address 100.1.5.1 255.255.255.0
 no cdp enable
!
interface FastEthernet0/0.6
 encapsulation dot1Q 6
 ip vrf forwarding v106
 ip address 100.1.6.1 255.255.255.0
 no cdp enable
!
interface FastEthernet0/0.7
 encapsulation dot1Q 7
 ip vrf forwarding v107
 ip address 100.1.7.1 255.255.255.0
 no cdp enable
!
interface FastEthernet0/0.8
 encapsulation dot1Q 8
 ip vrf forwarding v108
 ip address 100.1.8.1 255.255.255.0
 no cdp enable
!
interface FastEthernet0/0.9
 encapsulation dot1Q 9
 ip vrf forwarding v109
 ip address 100.1.9.1 255.255.255.0
 no cdp enable
!
interface FastEthernet0/0.10
 encapsulation dot1Q 10
```

```
 ip vrf forwarding v110
 ip address 100.1.10.1 255.255.255.0
 no cdp enable
!
interface FastEthernet0/1
 ip address 200.1.1.1 255.255.255.0
 ip directed-broadcast
 negotiation auto
 mpls label protocol ldp
 mpls traffic-eng tunnels
 tag-switching ip
 no cdp enable
!
interface Ethernet0
 ip address 146.223.197.15 255.255.248.0
 ip access-group 198 in
 ip access-group 199 out
 no ip directed-broadcast
 no ip proxy-arp
 ip route-cache cef
 no cdp enable
!
!
autonomous-system 1016
!
router ospf 1000
 router-id 116.116.116.116
 log-adjacency-changes
 passive-interface Loopback0
 network 200.1.1.0 0.0.0.255 area 0
!
router bgp 1016
 bgp router-id 116.116.116.116
 bgp log-neighbor-changes
 neighbor 1.1.1.1 remote-as 1016
 neighbor 1.1.1.1 update-source Loopback0
 neighbor 1.1.1.2 remote-as 1016
 neighbor 1.1.1.2 update-source Loopback0
!
 address-family ipv4
 redistribute connected
 redistribute static
 neighbor 1.1.1.1 activate
 neighbor 1.1.1.2 activate
no auto-summary
 no synchronization
 exit-address-family
 !
 address-family vpnv4
 neighbor 1.1.1.1 activate
 neighbor 1.1.1.1 send-community extended
 neighbor 1.1.1.2 activate
 neighbor 1.1.1.2 send-community extended
exit-address-family
 !
 address-family ipv4 vrf v101
 redistribute connected
 redistribute static
 neighbor 100.1.1.2 remote-as 101
 neighbor 100.1.1.2 activate
 no auto-summary
 no synchronization
 exit-address-family
 !
 address-family ipv4 vrf v102
 redistribute connected
 redistribute static
 neighbor 100.1.2.2 remote-as 102
 neighbor 100.1.2.2 activate
 no auto-summary
 no synchronization
 exit-address-family
 !
 address-family ipv4 vrf v103
 redistribute connected
 redistribute static
 neighbor 100.1.3.2 remote-as 103
 neighbor 100.1.3.2 activate
 no auto-summary
 no synchronization
 exit-address-family
 !
 address-family ipv4 vrf v104
 redistribute connected
 redistribute static
 neighbor 100.1.4.2 remote-as 104
 neighbor 100.1.4.2 activate
 no auto-summary
 no synchronization
 exit-address-family
 !
 address-family ipv4 vrf v105
 redistribute connected
 redistribute static
 neighbor 100.1.5.2 remote-as 105
 neighbor 100.1.5.2 activate
 no auto-summary
 no synchronization
 exit-address-family
 !
 address-family ipv4 vrf v106
 redistribute connected
 redistribute static
 neighbor 100.1.6.2 remote-as 106
 neighbor 100.1.6.2 activate
 no auto-summary
 no synchronization
 exit-address-family
 !
 address-family ipv4 vrf v107
 redistribute connected
 redistribute static
 neighbor 100.1.7.2 remote-as 107
 neighbor 100.1.7.2 activate
 no auto-summary
 no synchronization
 exit-address-family
 !
 address-family ipv4 vrf v108
 redistribute connected
 redistribute static
 neighbor 100.1.8.2 remote-as 108
 neighbor 100.1.8.2 activate
 no auto-summary
 no synchronization
 exit-address-family
 !
 address-family ipv4 vrf v109
 redistribute connected
 redistribute static
 neighbor 100.1.9.2 remote-as 109
 neighbor 100.1.9.2 activate
 no auto-summary
 no synchronization
 exit-address-family
 !
 address-family ipv4 vrf v110
 redistribute connected
 redistribute static
 neighbor 100.1.10.2 remote-as 110
 neighbor 100.1.10.2 activate
 no auto-summary
 no synchronization
 exit-address-family
!
ip classless
ip route 146.223.72.0 255.255.248.0 146.223.197.1
!
access-list 198 permit tcp any any eq telnet
access-list 198 permit udp any any
access-list 198 deny   ip any any
access-list 199 deny   ip any any
snmp-server engineID local
00000009020000D0FF65C400
snmp-server enable traps sonet
!
control-plane
!
banner login ^CCConsult users (run "show
users") before modifying config^C
banner motd ^CCisco GSR-12008^C
!
line con 0
 exec-timeout 0 0
 logging synchronous
 no history
line aux 0
 no history
line vty 0 4
 exec-timeout 60 0
 password letmein
 logging synchronous
 login
!
no cns aaa enable
end
```

www.agilent.com/find/n2x

Agilent N2X

Agilent's N2X multi-service tester combines leading-edge services with carrier grade infrastructure testing and emulation. The N2X solution set allows network equipment manufacturers and service providers to more comprehensively test new services end-to-end, resulting in higher quality of service and lower network operating costs.

Warranty and Support

Hardware Warranty

All N2X hardware is warranted against defects in materials and workmanship for a period of 1 year from the date of shipment.

Software Warranty

All N2X software is warranted for a period of 90 days. The applications are warranted to execute and install properly from the media provided.
This warranty only covers physical defects in the media, whereby the media is replaced at no charge during the warranty period.

Software Updates

With the purchase of any new system controller Agilent will provide 1 year of complimentary software updates. At the end of the first year you can enroll into the Software and Suport Agreement (SSA) contract for continuing software product enhancements.

Support

Technical support is available throughout the support life of the product. Support is available to verify that the equipment works properly, to help with product operation, and to provide basic measurement assistance for the use of the specified capabilities, at no extra cost, upon request.

Ordering Information

To order and configure the test system consult your local Agilent field engineer.

Sales, Service and Support

United States:
Agilent Technologies
Test and Measurement Call Center
P.O. Box 4026
Englewood, CO 80155-4026
1-800-452-4844

Canada:
Agilent Technologies Canada Inc.
2660 Matheson Blvd. E
Mississauga, Ontario
L4W 5M2
1-877-894-4414

Europe:
Agilent Technologies
European Marketing Organisation
P.O. Box 999
1180 AZ Amstelveen
The Netherlands
(31 20) 547-2323

United Kingdom
07004 666666

Japan:
Agilent Technologies Japan Ltd.
Measurement Assistance Center
9-1, Takakura-Cho, Hachioji-Shi,
Tokyo 192-8510, Japan
Tel: (81) 426-56-7832
Fax: (81) 426-56-7840

Latin America:
Agilent Technologies
Latin American Region Headquarters
5200 Blue Lagoon Drive, Suite #950
Miami, Florida 33126
U.S.A.
Tel: (305) 269-7500
Fax: (305) 267-4286

Asia Pacific:
Agilent Technologies
19/F, Cityplaza One, 1111 King's Road,
Taikoo Shing, Hong Kong, SAR
Tel: (852) 3197-7777
Fax: (852) 2506-9233

Australia/New Zealand:
Agilent Technologies Australia Pty Ltd
347 Burwood Highway
Forest Hill, Victoria 3131
Tel: 1-800-629-485 (Australia)
Fax: (61-3) 9272-0749
Tel: 0-800-738-378 (New Zealand)
Fax: (64-4) 802-6881

Printed in USA November 06, 2008
5989-9761EN

Agilent Technologies