S-38.3192
ITGuru Exercise (3: Building the MPLS BGP VPN)

Spring 2006

Original version: Johanna Nieminen and Timo Viipuri (2005)
Modified: Timo-Pekka Heikkinen, Juha Järvinen and Yavor Ivanov (2006)

Task Description

In ITGuru MPLS VPN creation is highly automated. The VPN Creation Wizard configures automatically e.g. BGP peers and VRF tables, and even the assignment of IP-addresses and autonomous system numbers is handled by the Wizard. This means that you should have a very clear view of how the network and the VPNs should be built before starting the wizard. Once the VPN has been created, it requires a lot of manual work to clear the configurations from the routers and build the VPN again if something went wrong in the first time.

When creating a VPN with the Wizard, you first have to decide what sites you want to connect. You must select one LSR router for each site (in the MPLS core network) that will be designated as a Provider Edge (PE) router. These PE routers will be automatically configured as BGP peers and a VRF configuration will be created in each of them.

Before you start configuring the wizard, do the following:

- If you have some separate server objects in the access networks, remove them. In the exercise we will utilize the server object available in the LAN object itself. If you have used an Ethernet switch to connect the separate server and the LAN, remove the switch as well.

- Add another IP gateway to Berlin and connect it with the LAN object and with the PE. This gateway is used as an endpoint for the IP VPN tunnel (this is not realistic and it’s only done because ITGuru has its limitations). You should now have two different *ethernet4_slip8_gtwy*-gateways in Berlin: one for the MPLS VPN and another for the IP VPN. Be careful when assigning new IP addresses for the interfaces: it’s best to first select the newly created links and then use AutoAssign (**Protocols->IP->Addressing->Auto Assign IP addresses**).

- Configure OSPF for the new links between Berlin LAN, Berlin tunnel-GW and Berlin PE.

- Name the devices in the access networks consistently. In this way it is much easier to read the configuration reports, for instance the IP-addresses assigned for the devices.

Step-by-step configuration guide

Step 1

Remove all previous LSPs and Traffic Mapping Configurations (if you made them in the previous exercise). The wizard will automatically create a full mesh of LSPs between the VPN sites.

Step 2

Create backup links between the PEs (LERs) and Ps (LSRs). Don't forget that 3 hops are required between any 2 subnets. If needed add new nodes. Assign IP addresses ONLY to the newly created links and enable OSPF on the interfaces.

NOTE: As stated before, inside the Berlin subnet you have to put an additional gateway (*ethernet4_slip8_gtwy*) which will be used as a gateway for the IP VPN. Connect it to the LAN and the PE and on both interfaces enable OSPF. Check if everything is working (ping).

Step 3

Applied ONLY to the sites Lisbon, Berlin, Rome

1. Disable OSPF between the CEs and PEs. Note that in Berlin subnet OSPF should only be disabled between the MPLS-VPN gateway (CE) and PE (not between the tunnel GW and PE).

2. Assign AS numbers to all autonomous systems.

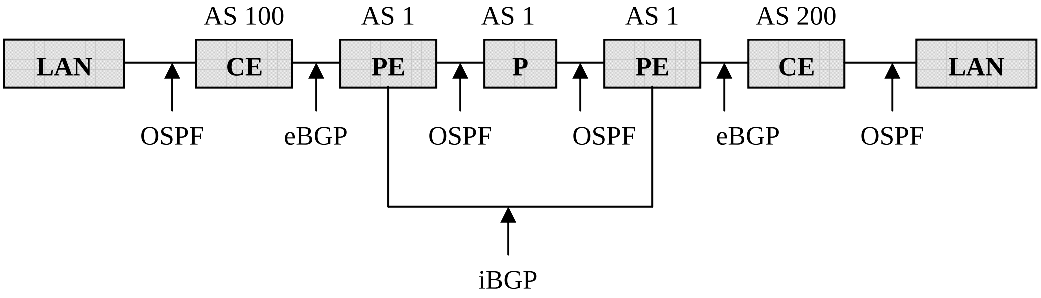

Assign AS numbers. A group of routers with a common set of administrative policies is called an Autonomous System (AS). Each AS is identified by a number, which you can specify on the individual routers comprising the AS. You can group the routers in your network into different ASs, each running a different interior gateway protocol (such as RIP, OSPF, or IGRP) within the AS and BGP between the border routers. AS numbers can be any integer value between 1 and 65,535. If the Autonomous System Number attribute has the default value of Auto Assigned, a random value between 1 and 65,535 is assigned to each router. For proper BGP operation, you should manually assign each BGP speaker an AS number using the Protocols > BGP > Configure AS Number operation. [See *BGP Model User Guide*]

Step 4

MPLS/BGP VPN

The approach relies on taking customer IP datagrams from a given site, looking up the destination IP address of the datagram in a forwarding table, then sending that datagram to its destination across the provider's network using an LSP.

Use the MPLS VPN wizard (**Protocols->MPLS->Deploy MPLS VPNs**…) to create the MPLS/BGP VPNs. Select layer 3 VPN. If you still have routers with AS number set to auto-assigned, here you can give a value.

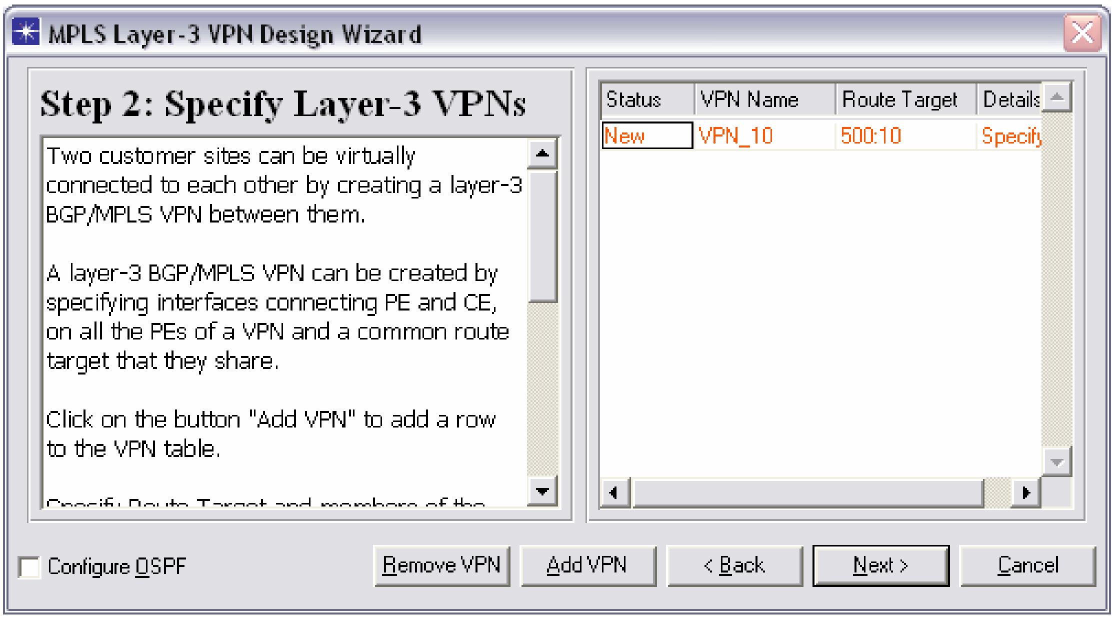

MPLS Layer-3 VPN Design Wizard

Step 2: Specify Layer-3 VPNs

Two customer sites can be virtually connected to each other by creating a layer-3 BGP/MPLS VPN between them.

A layer-3 BGP/MPLS VPN can be created by specifying interfaces connecting PE and CE, on all the PEs of a VPN and a common route target that they share.

Click on the button "Add VPN" to add a row to the VPN table.

Specify Route Target and members of the

Status	VPN Name	Route Target	Details
New	VPN_10	500:10	Specify

☐ Configure OSPF

Add a VPN and uncheck the Configure OSPF box as shown on the picture. Before continuing you have to set the VPN details in the relevant column. Create a full mesh VPN between the PEs. For the interface attribute you have to select the one that points to the CE. Open the BGP parameters on each PE. You should see the following parameters:

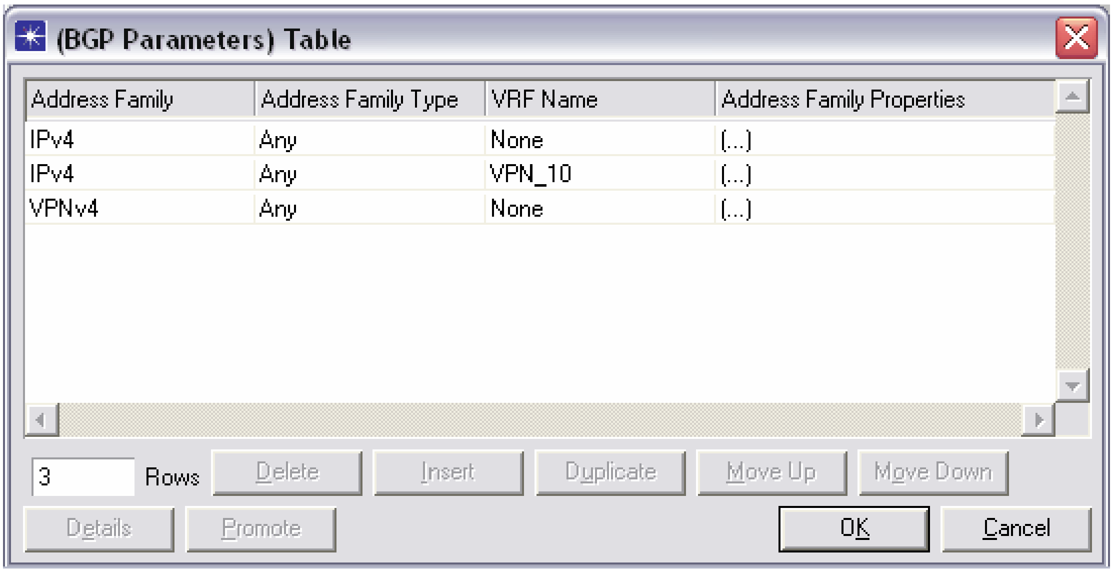

(BGP Parameters) Table

Address Family	Address Family Type	VRF Name	Address Family Properties
IPv4	Any	None	(...)
IPv4	Any	VPN_10	(...)
VPNv4	Any	None	(...)

3 Rows

BGP can be used to distribute routing information about different address families, such as IPv4 and VPNv4. The first line specifies the IPv4 addresses and how they should be treated. Here should be defined information for the iBGP peers. Second line is specific for the created VPN and defines the eBGP neighbors. The last line is related to the VPN-IPv4 addresses (When an ingress PE router receives an IPv4 route from a device within a VPN, it converts it into a VPN-IPv4 route by adding the route distinguisher prefix to the route).

Check the Neighbor information for all address families and be sure that it contains the correct peers (in *IPv4* and *VPNv4*: iBGP peers ONLY, and in the IPv4 parameters responsible for the VPN: eBGP peers ONLY).

IPv4 and VPNv4 - iBGP peers

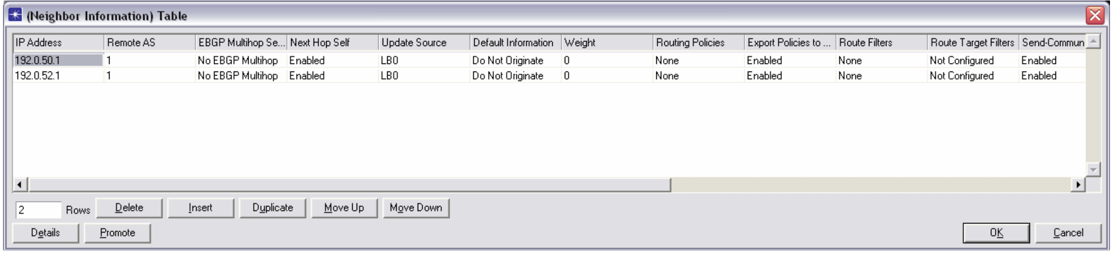

(Neighbor Information) Table

IP Address	Remote AS	EBGP Multihop Se...	Next Hop Self	Update Source	Default Information	Weight	Routing Policies	Export Policies to ...	Route Filters	Route Target Filters	Send-Commun
192.0.50.1	1	No EBGP Multihop	Enabled	LB0	Do Not Originate	0	None	Enabled	None	Not Configured	Enabled
192.0.52.1	1	No EBGP Multihop	Enabled	LB0	Do Not Originate	0	None	Enabled	None	Not Configured	Enabled

VPN – eBGP peers

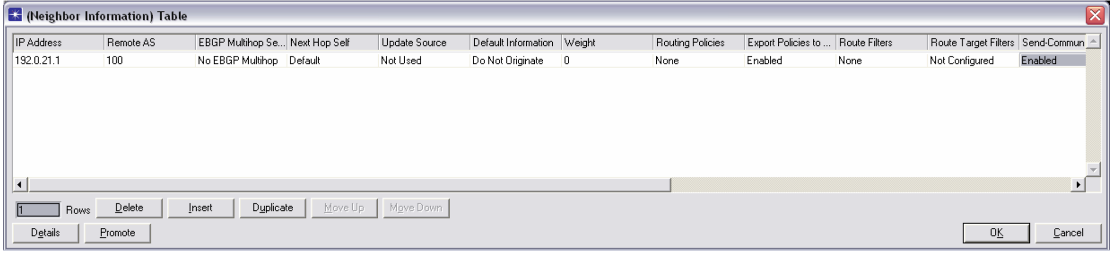

(Neighbor Information) Table

IP Address	Remote AS	EBGP Multihop Se...	Next Hop Self	Update Source	Default Information	Weight	Routing Policies	Export Policies to ...	Route Filters	Route Target Filters	Send-Commun
192.0.21.1	100	No EBGP Multihop	Default	Not Used	Do Not Originate	0	None	Enabled	None	Not Configured	Enabled

Neighbor Information. Configures the router's BGP neighbors. Unlike most other routing protocols, BGP does not automatically discover its neighbors. Instead, neighbors are manually configured. To configure two routers as neighbors, you should configure each as the neighbor of the other. When BGP is run between routers of the same autonomous system, it is termed an Internal BGP (IBGP). If the participating routers belong to different ASs, it is an External BGP (EBGP). When a BGP speaker receives an UPDATE message from an IBGP peer, the receiving BGP speaker does not forward the routing information in the UPDATE message to other IBGP peers. Because of this, each BGP speaker needs BGP connections to all other BGP speakers in its own AS to maintain a consistent view of the routing information within the AS. [See *BGP Model User Guide*]

NOTE: If you need to change any records in the Neighbor Information table, be sure to specify the correct interface (for iBGP – Loopback, for eBGP – the physical interface). Define update source: Loopback for IBGP and for EBGP physical interface ("not used"). Enable also the Send-community attribute! Set BGP start time (use 160 seconds).

<u>Step 5</u>

1. Configure eBGP between each CE and PE (use the correct AS numbers. NOTE: each site should have its own AS number, different from the other sites). Use the BGP wizard to create eBGP peers (**Protocols->BGP->Configure EBGP Peers**).

2. Redistribute the routing information accordingly. NOTE: For Berlin_CE (the VPN_Gateway) you have to enable <u>ONLY</u> BGP redistribution for directly connected peers. Do not enable redistribution on the PEs!

3. Select relevant statistics and test the setup by using ping from Lisbon to Berlin.

Step 6

Configure the IP VPN tunnels (hint: there's an object in the *internet_toolbox* palette…):

1. Include IP VPN Config.

2. Configure the tunnel parameters

3. Test the configuration

Again, select relevant statistics and test the setup by using ping from Paris to Berlin and from Athens to Berlin.

Testing the VPN

It is recommended that you test the VPNs at this point (if you haven't already) with simple IP Traffic Demands (or with anything that you find appropriate, for example ping or VoIP). First, select two site gateways belonging to the VPN and create a bidirectional IP traffic flow between these points ('**Traffic->Create IP Traffic Flows**'). Record the statistics that you want from the network by right clicking the mouse and selecting '**Choose Individual DES Statistics**'. Choose at least the throughput statistics from '**Path Statistics->LSP**' and from '**Link Statistics->point-to-point**'. In this way you can examine whether the traffic flows through the LSPs. To see if any traffic goes through a VPN select '**Global Statistics->VPN**' or '**Node Statistics->IP VPN Tunnel'**. If it seems that there is no traffic on the LSPs or VPNs, check if there is traffic on the links between the end points. It might be that conventional IP routing is used instead of the MPLS LSPs if you have not done the configuration properly.

If the connection between the sites gateways seems to be working, try another test. Create an IP Traffic Demand between the LANs of two cities belonging to the VPN and check whether the traffic flows through all links that it should. Check also the VRF tables created by the PE routers. You can export the VRF table information by right clicking the mouse and selecting '**Edit Attributes->Reports->VRF Table->Export at End of Simulation**'. You may also export other tables, such as IP Forwarding Table and OSPF Routing Table. Make sure that you have all the necessary destinations in the IP Forwarding Table. If something seems to be wrong, check the IP and BGP routing tables (click on a router and set **Reports->BGP Routing Table** to Export at End of Simulation).

Hints

You can check quickly the routing domains and protocol configuration by choosing '**View->Visualize Protocol Configuration->IP Routing Domains**'. The appearing Routing Domain Legend shows the meaning of the symbols. In this view you are able to see the routing domains, points where static routes are used and points where redistribution is used. In order to visualize BGP peers, choose '**View->Visualize Protocol Configuration->BGP peers**'. The BGP peers should now be connected with green, dotted lines. When trying to figure out what IP address belongs to which network node use the network browser **'View->Show Network Browser'**.

- You must NOT touch the IP-addresses that the Wizard has assigned for the devices in the core network. For instance, if you select '**Clear IP addresses on all interfaces**' and after this use AutoAssign, this will clear and reassign ALL IP-addresses in the network, including the addresses of the core devices. The Wizard has configured BPG peers, VRF tables etc., and thus the IP-addresses of these peers MUST NOT be changed; otherwise you have to manually reconfigure the devices.

- After AutoAssign you can create a configuration report (**Protocols->IP->Configuration Reports->Select All)** to check that there are no overlapping IP-addresses in the network.

- If you need to assign IP-addresses to some interfaces (for example, if you create new routers in the access network), you may use AutoAssign, as long as you do not erase the existing addresses. Check from the configuration reports or directly from the router configurations that IP-addresses were really created in the right interfaces.

- If you have problems when testing the network (e.g. no route to Berlin_LAN before at certain time) try adjusting the start times of traffic demands or the routing protocols. You will experience problems when trying to ping from Athens to Berlin before the routing protocols have converged.

Exercise Sessions

The exercise session for this task will be arranged on Thursday, 9th February at 14 o'clock in computer class Maari-A.

Handout Requirements

The exercise should be returned before the beginning of the next exercise session (23rd February, 14 o'clock). Send the exercise package as an e-mail attachment to Yavor Ivanov (yivanov@netlab.hut.fi) provided that the size of the attachment is reasonable. You have to pack the files with the command:

```
tar -cf - -C ~ op_models | gzip > 3_building_vpn.tar.gz
```

Remember to include only the relevant files (we do not want all the backups). Try to unpack the files and then run the project on some other computer (or with another group member's account) to make sure that the packaging is done properly.