



Praktikum RouterLab SS 2008

Work Sheet 4: Policy-based Routing, Border Gateway Protocol (BGP)

Table 1: Assignment of devices to groups

Group	Ham-Cloud	Muc-Cloud
Router	ham-rj1, ham-rj2	muc-rj1, muc-rj2
Switches	ham-sc1	muc-sc1
Loadgens	loadgen102, loadgen104	loadgen106, loadgen107
IP range	10.1.0.0/16	10.2.0.0/16
VLAN IDs	Each team can use up to 9 VLAN IDs, where last 2 digits reflect team ID <i>Example: Team 09 has VLANs 109, 209, . . . , 909</i>	

Question 1: (50 Points) *Basic BGP Configuration*

In this question, we will establish one BGP session between all pairs of router.

- (a) Enable IP connectivity between adjacent routers in a similar manner as in the previous work sheet! For this purpose, please have a look at Figure 1. Draw a topology map that shows your IP address assignment. Your configuration must satisfy the following conditions:
 1. All pairs of routers that are connected with a link in Figure 1(b) must be able to ping each other.
 2. Assign IP addresses in a way such that there is no (!) overlap in IP address range. The IP address ranges used for interfaces on different links in Figure 1(b) must be disjoint.
 3. Install and configure Quagga on your *loadgens*. Activate the `bgpd` daemon in `/etc/quagga/daemons`, as it will be needed for the BGP setup later. For more information on how to configure Quagga, have a look at Question 1b) of the previous work sheet.
 4. Configure the IP address `192.168.1.1/24` for the loopback interface `lo` on *loadgen102* (Muc-Cloud: *loadgen106*). Apart from this, exclusively use IP addresses from the range `10.1.0.0/16` (Muc-Cloud: `10.2.0.0/16`).
 5. Configure separate tagged VLANs for each pair of routers, i.e. for each “arc” in Figure 1(a). Thus, you will have to configure 6 VLANs. Have a look at Question 1d) of the previous work sheet to find out how to configure tagged VLAN and “subinterfaces” on the *loadgens*.
- (b) Read Section 5.1.2 of the BGP RFC 4271. In your opinion: Is BGP a pure distance-vector or a pure link-state routing protocol? Explain in 2 or 3 sentences!
- (c) According to the topology of Figure 1(b), each router is assigned to an AS. Altogether, we have four ASs (65001 – 65004), where each AS consists of a single router. All ASs are “fully meshed”. After enabling IP connectivity, establish one BGP session for each pair of routers! For this purpose, have a look at the Cisco and Juniper BGP manuals that are provided on the course website. Focus on the minimum configurations to establish a BGP session between two routers!

Some hints:

- *Juniper*: The AS number is specified within `routing-options`, while BGP neighbors are specified within `protocols bgp`. We recommend to add one `group` for each BGP neighbor which will allow you to configure separate policies for each BGP neighbor later on.

- *Quagga*: Configuring Quagga is similar to Cisco. The BGP configuration can be done after using the command `router bgp <AS number>`. To specify neighbors use `neighbor <ip> remote-as <remote asn>`. To inject a prefix `192.168.1.0/24` into BGP, use the command `network 192.168.1.0/24!`
- (d) Which AS paths are learned by the four routers to reach network `192.168.1.0/24`? Which AS paths are selected as best? Explain why!
- To display the BGP routing tables we recommend the command `show route protocol bgp terse` for Juniper routers and `show ip bgp` on Quagga (Cisco). Of course, you can also use the standard commands to display the routing tables.

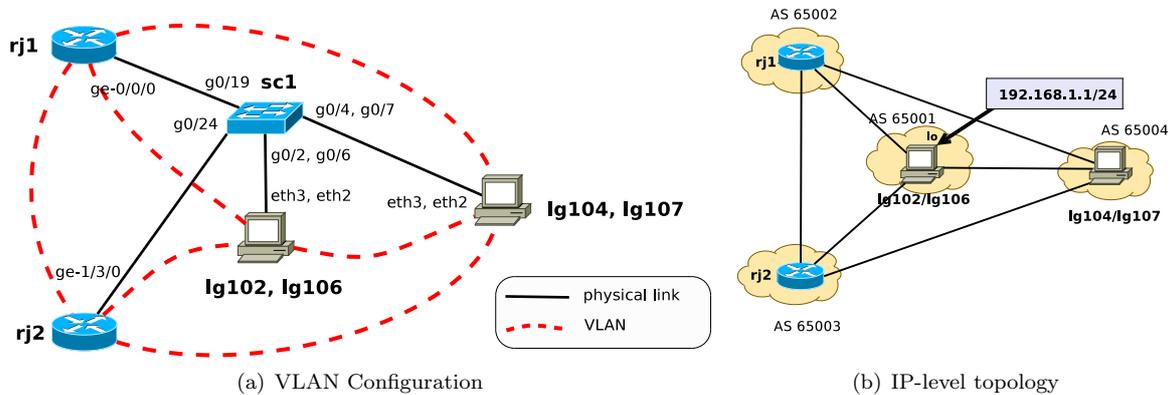


Figure 1: Topology – Basic Configuration.

Question 2: (50 Points) *BGP Routing Policies, Routing Instabilities*

Frequently, it is not desired that packets are sent along the shortest AS path – fewest number of AS hops – towards a certain destination. Rather economic aspects may need to be considered. The following part illustrates how policies are configured and how they are used. In addition, it will show that inconsistent policy configurations may cause routing instabilities.

- (a) We assume that AS 65004 is a customer of AS 65002 and AS 65003. In general this means that AS 65004 has to pay for traffic that uses the link to AS 65002 and AS 65003. The two providers AS 65002 and AS 65003 want to maximize their profit by sending as much traffic as possible over AS 65001 (prefix `192.168.1.0/24`) over AS 65004.
- Use *local-preference* to achieve that AS 65002 and AS 65003 send as much traffic as possible over AS 65001. For this purpose, you may use either a local-pref value of 50 or 200.¹
- What can you see in the BGP routing tables of AS 65002 and AS 65003? Can you figure out what the default local-preference value for BGP routes is?
- Question 2d) provides some hints on how to set local-preference values. Note that contrary to 2d), you should not configure local-preference values for a specific AS path.
- (b) AS 65004 notices that its providers AS 65002 and AS 65003 send all their traffic to `192.168.1.0/24` over its AS. Therefore, it wants to prevent some routes from being announced to its providers AS 65002 and AS 65003.
- Configure AS 65004 such that routes/AS paths originated by AS 65001 are not announced to AS 65002 and AS 65003. Which routes are now used by AS 65002 and AS 65003? Submit the BGP routing tables!
- Use the commands `neighbor filter-list` and `ip as-path access-list`! More information can be found in the Cisco BGP documentation. The code snippets of Question 2d) may also be helpful.
- (c) Remove all policies from Question 2a) and 2b)!
- (d) Configure policies such that AS 65002, 65003 and 65004 all prefer the AS path of length 2 in the counter-clockwise direction! This requires the following policies:

- AS 65002 applies *local-pref* 200 to routes with AS path "65003 65001"

¹You may need the commands `clear ip bgp * soft out` on Quagga or `clear bgp neighbor` on Juniper

- AS 65003 applies *local-pref* 200 to routes with AS path "65004 65001"
- AS 65004 applies *local-pref* 200 to routes with AS path "65002 65001"

Make sure that local-preference values are exclusively applied for the paths "65003 65001", "65004 65001" and "65002 65001"!

The following example illustrates how to configure a routing policy on Juniper routers that prefers all routes (*local-pref* 500) that use the AS path 1111 2222 and that are learned from neighbor 1.2.3.4.

- **Juniper:**

```
policy-options {
  policy-statement <policy-name> {
    from {
      protocol bgp;
      as-path <reg-expr-name>;
    }
    then {
      local-preference 500;
    }
  }
  as-path <reg-expr-name> "1111 2222";
}
```

Finally, to apply the policy *policy-name*, use an *import* statement within the appropriate group (group for neighbor 1.2.3.4) in the BGP configuration.

- **Cisco/Quagga:**

```
router bgp <AS number>
  neighbor 1.2.3.4 route-map <route-map-name> in
  !
  ip as-path access-list <access-list-name> permit ^1111_2222$
  !
  route-map <route-map-name> permit <line number>
  match as-path <access-list-name>
  set local-preference 500
```

More information can be found in the Cisco and Juniper documentation!

- (e) Activate debug output on rj2 by configuring the following:

```
protocols {
  bgp {
    traceoptions {
      file <file-name> size 5m;
      flag packets detail;
    }
  }
}
```

This will store debug output for BGP in the file *<file-name>*. To observe the latest debug messages on the console, use the command *monitor start <file-name>*. Wait 3 to 4 minutes and then use *monitor stop* to stop the debug output. Copy the output shown in the console to a text file and submit the file!

Before you finish change to the underlying UNIX shell on the Juniper router and delete the file *<file-name>* in the directory */var/log*.

- (f) Explain what you see in the debug output! Why is the routing not stable? Explain briefly (4-5 sentences) why this BGP configuration will not produce stable results! You may also have a look at Section 3.1 of the Sigcomm'99 paper ("An Analysis of BGP Convergence Properties") that is provided on the web site.
- (g) What do you think? Is it likely that such a routing instability occurs in the real Internet? Why or why not? (1-2 sentences)

For submission details please check the FAQ:

http://www.net.t-labs.tu-berlin.de/teaching/ss08/RL_labcourse/

Due Date: Friday, June 6th, 2008, 08:00 am