

Praktikum RouterLab SS 2009 Work Sheet 5: Border Gateway Protocol (BGP)

The purpose of this work sheet is to introduce the Border Gateway Protocol (BGP), which is the dynamic routing protocol used for inter-domain routing.

Important Notes:

All configuration files and dumps of commands must be provided!

Table 1: Assignment of devices to groups

Group	Ham-Cloud	Muc-Cloud
Router	ham-rc1, ham-rj1, ham-rj2	muc-rc1, muc-rj1, muc-rj2
Switches	ham-sc1, ham-sc2	muc-sc1, muc-sc2
Loadgens	loadgen10[2/3/4]-ham	loadgen10[2/3/4]-muc
IP range	10.1.0.0/16	10.2.0.0/16
VLAN IDs	$x00 - x99$ where x is your team number	

Question 1: (30 Points) *BGP Basics*

Read Section 5 of the BGP RFC 4271, then answer the following:

- List the categories of attributes of a path and explain briefly each one.
- Which are the mandatory attributes?
- Is BGP a pure distance-vector or a pure link-state routing protocol? Explain in 2 or 3 sentences.
- When should the attribute `ATOMIC_AGGREGATE` be included in an announcement?

Read Section 9.1.2.2 of the BGP RFC 4271, then have a look at the following dump of a BGP RIB:

```
BGP table version is 0, local router ID is 203.119.0.116
Status codes: s suppressed, d damped, h history, * valid, > best, i -internal,
r RIB-failure, S Stale, R Removed
Origin codes: i -IGP, e -EGP, ? -incomplete
```

```
Network          Next Hop Metric LocPrf Weight Path
* > 3.0.0.0       193.0.4.28          0 12654 34225 80 e
* 3.0.0.0        193.0.4.28          0 701 703 80 i
* 3.0.0.0        193.0.4.28          0 3456 8765 8690 80 e
```

- Which of the three routes is the best one (ignore attributes that are not listed)?
- Which rule of section 9.1.2.2 of RFC 4271 makes the tie-break?
- Explain briefly (4 - 6 sentences) what is a *Route-Reflector* and why it is used.

Question 2: (40 Points) *Basic BGP Configuration*

In this question, we will develop a simple BGP configuration in order to establish one BGP session between all pairs of router.

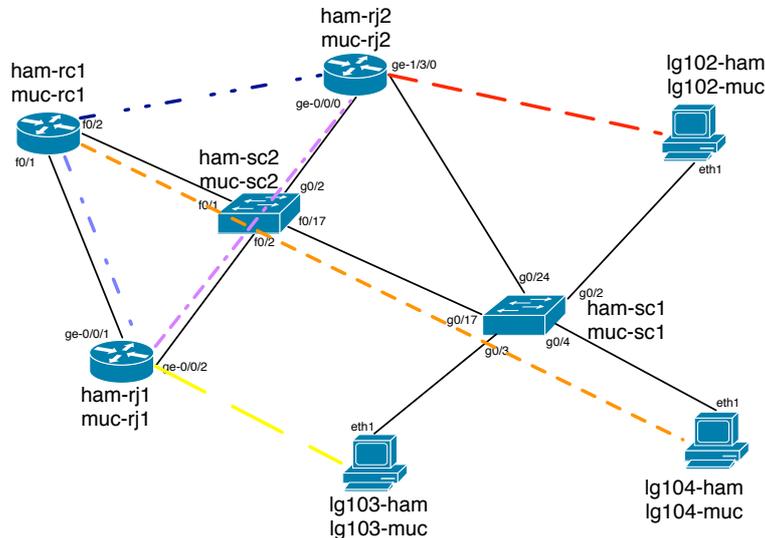


Figure 1: Topology – Basic Configuration.

Enable IP connectivity between adjacent routers/loadgens in a similar manner as in the previous work sheet. For this purpose, please have a look at Figure 1, where each colored¹ dashed line indicates a different logical subnetwork. Establish IP connectivity between the (two) devices that are inside the same logical subnetwork. Draw a topology map that shows your IP address assignment. Your configuration must satisfy the following conditions:

- When necessary configure separate VLANs.
- 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 virtual links in Figure 1 must be disjoint.
- All pairs of routers that are connected with a link in Figure 1 must be able to ping each other.

Note that routing between subnetworks is not yet enabled.

- Provide a topology map and all the details of your configuration including the indication of which kind of VLAN (trunk/access) is used on each link.
- For each loadgen of your cloud: traceroute every other configured interface (the ones *you* have given IP addresses) of the routers and loadgens of your cloud (You have to submit also the dumps).
- Take the trace of *loadgen102-xxx* from the previous point. Which interfaces of your cloud are reachable? Which are not? Why?

According to the topology of Figure 2, each router-loadgen pair is assigned to an AS. Altogether, we have three ASs (65003/65006/65009), and they are “fully meshed”.

Configure every router to establish one BGP session with every other router in order to provide full connectivity to all subnetworks. 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, but do not forget to look at how to inject IGP routes into BGP.

¹Please check the online version of this worksheet.

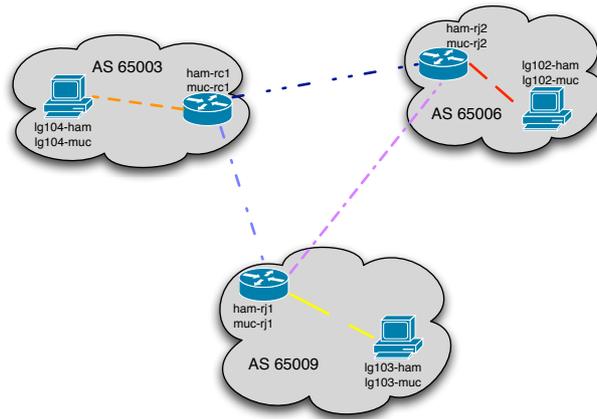


Figure 2: Topology – BGP Configuration.

- d) For each loadgen of your cloud: traceroute every other configured interface (the ones *you* have given IP addresses) of the routers and loadgens of your cloud (you have to submit also the dumps).

Look at the BGP routing table of `xxx-rc1` (dump to be submitted).

- e) How many routes are available toward each loadgen?
 f) Which one is used? and why?
 g) You should observe also some routes marked as *RIB-Failure*. Can you explain what does it mean and explain why this is happening?

Question 3: (30 Points) BGP Routing Policies

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.

We assume that AS 65003 is a customer of AS 65006 and AS 65009. In general this means that AS 65003 has to pay for traffic that uses the link to AS 65006 and AS 65009. The two providers AS 65006 and AS 65009 want to maximize their profit by sending as much traffic as possible over AS 65003. For example, traffic destined for AS 65006 and AS 65009, should be sent over AS 65003. Use *local-preference* attribute to achieve that AS 65006 and AS 65009 send as much traffic as possible over AS 65003.

- a) What can you see in the BGP routing tables of all the three routers (you must provide a dump of all of them)?
 b) Can you figure out what the default local-preference value for BGP routes is?

For each loadgen of your cloud: traceroute every other configured interface (the ones *you* have given IP addresses) of the routers and loadgens of your cloud (you have to submit also the dumps).

- c) Which routes have change compared to the previous set of traceroute?

AS 65003 notices that its providers AS 65006 and AS 65009 send all their traffic over it. Therefore, it wants to prevent some routes from being announced to its providers AS 65003 and AS 65006. Configure AS 65003 such that routes originated by AS 65003 and AS 65009 are not announced. **Hint:** Use the commands `neighbor filter-list` and `ip as-path access-list`. More information can be found in the Cisco BGP documentation.

- d) Which routes are now used by AS 65006 and AS 65009? Submit the BGP routing tables of all routers.

For each loadgen of your cloud: traceroute every other loadgen (only loadgens!) of your cloud (you have to submit also the dumps).

e) Are the routes the same as for Question 2? Explain.

For submission details please check the FAQ:

http://www.net.t-labs.tu-berlin.de/teaching/ss09/RL_labcourse/

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