KTH CSC

BGP-Intro-lab

Juniper version

Group Nr	
Name1	
Name2	
Name3	
Name4	
Date	
Grade	
Instructor's Signature	

Table of Contents

1 Goals

The lab is an introduction to BGP and the interaction between IGP and BGP.

The lab uses IGP to distribute internal routes within an AS, EBGP to announce own routes to other routing domains and IBGP to distribute external routes within an AS. A primary goal of the lab is that you understand the roles of each protocol.

Before you begin this lab, please consult documents [1], [2] and [3]. You should also have read the lecture notes and RFC pages about BGP. Extra material from vendors published on the web are also useful.

2 Introduction

Connect from your workstation via telnet to your router. You can use terminal<1/2>.netlab.csc.kth.se to access the terminal servers.

You need to use telnet to a specific port number to get attached to your router. Ask the teacher if you do not know the port numbers or the password to the router.

3 Network setup

You have been assigned the 10.X.0.0/20 network block. You have a total of 4 costumers, and each one wants 254 addresses. See the next page for the network layout. Use the netmap topology [3] to configure your routers.

Your internal network is configured as a square and you have three external connections (called peerings). One peer is your Tier1 provider, your connection to the global internet. The other two peers are neigbor groups.

Your uplink provider uses the 192.71.24.32/27 network for the peering. The uplink uses 192.71.24.62, while group A uses 192.71.24.34, group B uses 192.71.24.38, etc. For your neighbours you will have to come to an agreement on which network to use for the peering. This net is often called DMZ or demilitarized Zone. Create an address map for your network and configure the interfaces. Remember to assign the loopback interfaces of the routers addresses to use as router-id and make sure the first address of each customer network is configured (while the rest is a

default route). The customer networks should be configured on RTX3.

Make sure you communicate the customer networks to your peers so you can simplify testing later.

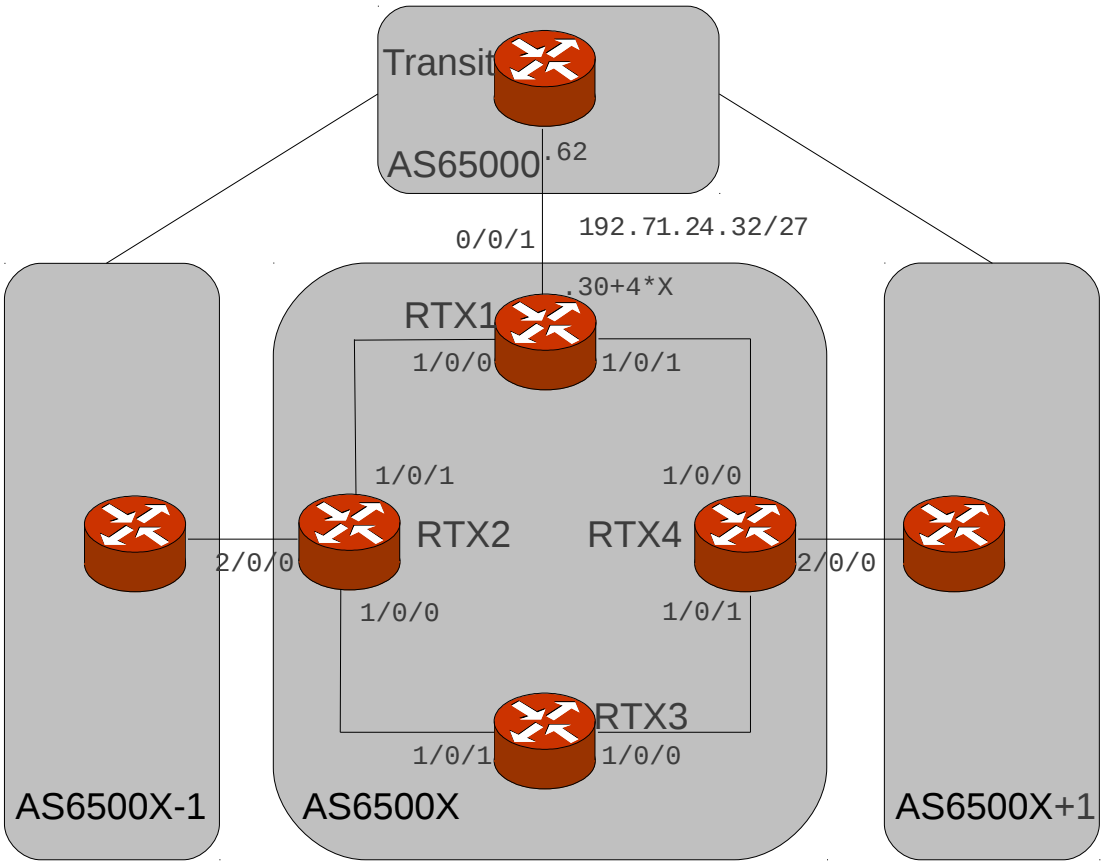

4 Setup IGP

The next step is to start routing inside your AS. You are free to use any of the routing protocols we have previously deployed (RIP, OSPF or ISIS). Make sure you are not leaking any routing information to the neighbouring AS. Also verify connectivity throughout your network. For obvious reasons you will not have any connectivity in neighbouring networks at this point.

5 Setup EBGP

You are now ready to start EBGP and begin peering with your neighbours, both neighbor groups and transit. Set a router-id and turn on BGP on your border routers and bring up the peering to neighbouring ASes. Do not peer with any internal servers. This means that you setup three peerings in total. Note that RTX3 is an internal router and should NOT run EBGP.

You should use default next-hop (not next-hop self).

Study one of the BGP peerings.

- What state is your session in?
- What is the session keep-alive value?

- What hold time value has been negotiated with your neighbors?
- Can you tell who initiated the TCP session?

Now create a policy to spread information about your networks to your neighbours. That is, export all IGP routes to BGP in all border routers.

- Verify that the BGP route table looks ok. Which command is best for this?
- Can you ping/traceroute between all addresses? Why/Why not?
  Some examples:
  - Between customers?
  - Between RTX4 and RT(X+1)2?
  - Between RTX4 and RT(X+2)2?
  - Between RTX1 and RT(X+2)1?

Explain why some destinations are reachable and some not.

Milestone 1, report to a lab assistant.

---

6 IBGP

Now configure your routers with IBGP. You do this to distribute external routes and to distribute origin routes to border routers. You do this by establishing IBGP full mesh between all routers. Use the local address (router-id) for IBGP peering.

- Verify that you get a complete routing table in your internal router and that you can ping/traceroute all destinations.
- Verify with traceroute/ping record-route that the path between customers go as expected. (Hint: via transit or peering?)

Milestone 2, report to a lab assistant.

---

7 Prefix filtering

Now create prefix filters. This filter should contain the addresses of the other groups. Make sure you only accept updates containing these addresses from your neighbouring AS. This will prevent other groups from stealing traffic from your AS by announcing a shorter AS path than the real target. It will also stop you from using them as a transit AS (as you do not accept routes to AS65000 through them).

Protect yourself from being used as a transit by filtering the routes you announce in similar fashion.

Milestone 3, report to a lab assistant.

---

8 Aggregate your routes

Aggregate your routes so that you only distribute a minimal set of routes to your neighbours.

Milestone 4, report to a lab assistant.

---

9 The RIPE database

For this kind of filtering the RIPE database (http://www.ripe.net/db/index.html) can be helpful. It can also be accessed using the whois tool.

- Check who owns the lab network (192.71.24.0/24)
- Which AS is the lab network in?
- What other nets belong to the same AS as the lab net?
- Would you expect AS1653 to announce the lab network?
- Why/Why not?
- What would the AS path look like in that case?
- What is the relationship between KTHLAN and SUNET?
- What is the relationship between NORDUNET and SUNET?

Milestone 5, report to a lab assistant.

---

10 References

[1] KTH CSC Router lab Introduction - Juniper version
[2] KTH CSC Router lab Reference - Juniper version
[3] KTH CSC Router Netmap Topology