Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Networked Systems Group

Spring 2017 Prof. L. Vanbever/ T. Bühler, R. Birkner, T. Holterbach, R. Meier

Communication Networks

Group project 1: Build your own Internet

Deadline: May 7 2017 at 11.59pm

1 Introduction

In this assignment, you will learn how to build and operate a layer-3 network and how different networks, managed by different organizations, interconnect with each other.

More particularly, you will first learn how to set-up a valid forwarding state within an Autonomous System (Fig 1(b)). To do that, you will use an intra-domain routing protocol: OSPF. Then, you will learn how to set-up a valid forwarding state between different ASes so that an end-host in one AS (e.g. ETH) can communicate with an end host in another AS (e.g. Google) (Fig 1(c)). For this task you will use the only inter-domain routing protocol available today: BGP. We will build a mini Internet using virtual routers running the Quagga software routing suite [1]).

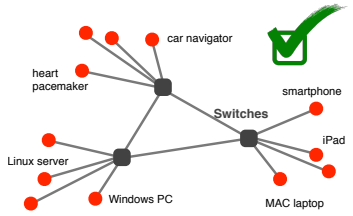

(a) A local network is composed of end-hosts and switches. For instance, all the hosts connected in a building can form a local network, and each floor is connected to another floor via a switch.

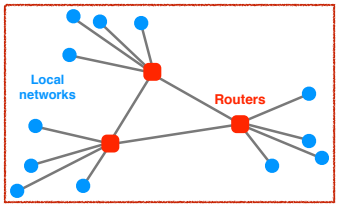

(b) A network is composed of sub-networks and IP Routers. For instance, a large company can have several buildings. Inside each buildings there is a local network, and routers interconnect those local networks, such that two people in two different buildings can communicate.

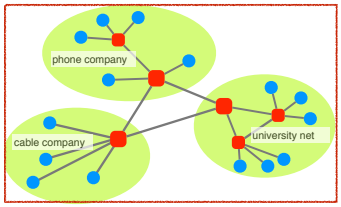

(c) IP routers also interconnect networks of different organizations. For instance, Orange and Deutsche Telekom, two major European ISPs, are interconnected with IP routers. This is the Internet!

Figure 1: The Internet is a network of networks. In each case, different protocols and concepts are used to build and update a forwarding state that ensures connectivity.

Traditional IP routers cannot be configured through OpenFlow, the protocol you used to configure the L2 switches in the first practical assignment. Instead, you need to configure them through a Common Line Interface (CLI). Each IP router vendor (e.g. Cisco or Juniper) or software routing suite (e.g. Quagga) has its own CLI. Fortunately, those CLIs are similar, and if you know how to configure a router using the Quagga CLI, you can easily configure a Cisco or a Juniper router as well. In a separate document[1], we give you a crash course on how to configure a Quagga router. The rest of this document is organized as follow. We first describe the setup you

[1] available at `http://comm-net.ethz.ch/assignment2/quagga_tuto.pdf`

will have to use (Section 2). Section 3 explains how you can access and configure your Quagga routers and Section 4 presents verification tools. We then list the questions in Section 5. Finally, we give general information (including submission instructions) in Section 6.

2 *Your* mini Internet

Each group will manage its own AS, just like a normal network operator would. The assignment will be divided in two parts. You will first have to enable end-to-end connectivity inside your own AS, using OSPF. Then, you will have to establish BGP sessions with your neighboring ASes and configure BGP policies such that you respect the business relationships (e.g. customer/provider, peer/peer). At the end, any host located in any AS should be able to reach (i.e., ping) any other host in any other AS. The AS paths used between two hosts must follow the business relationships defined between ASes.

Each group has its own AS The number of your AS is your group number. If you are group 28, you will have the AS 28. Routers and hosts composing your AS are already running in your VM. The topology of your network follows the swiss geography: your routers are located in swiss cities like Zürich, Lausanne or Olten. Figure 2 shows how your network looks like. Each group has a /8 prefix that it can use. If you are group X, then the prefix X.0.0.0/8 is yours! For example group 22 has the prefix 22.0.0.0/8. If you assign the IP addresses to your routers, you will have to use a different subnet (all belonging to your /8 prefix) between each pair of routers. The subnets to use are indicated in Figure 2. For example, between WINT and STGA, you must use the subnet X.0.4.0/24. The interface in WINT that is connected to STGA must have the IP address X.0.4.1 and the interface in STGA that is connected to WINT must have the IP address X.0.4.2 (with X your group number).

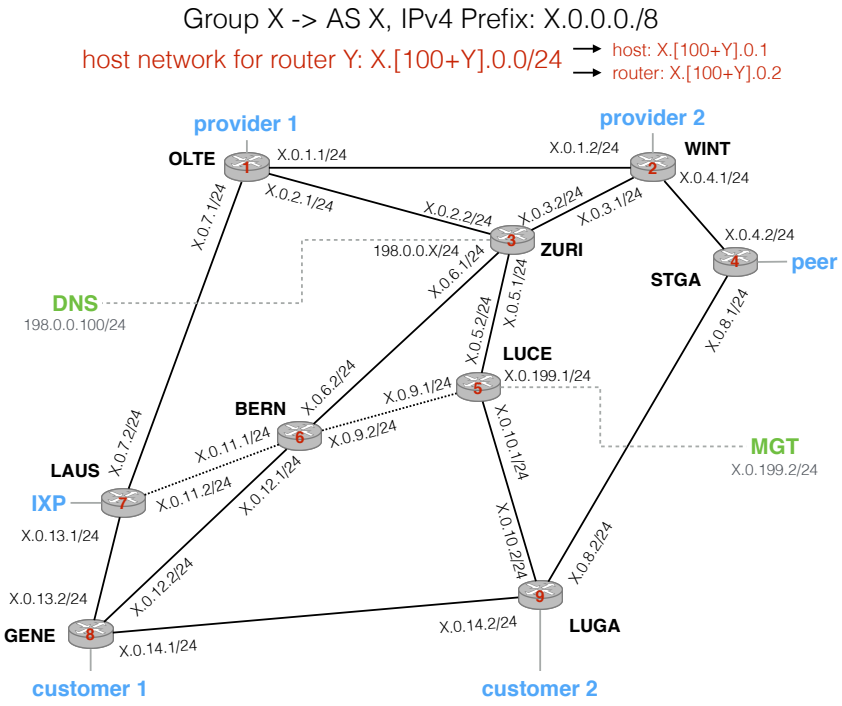

Figure 2: Each group will have to manage an entire AS. Your AS is composed of 9 routers. A /8 prefix has been assigned to each group. You can use it to configure your local networks. One host is also connected to each router. The subnets you must use for each of your local networks are indicated on each interface.

One host is also connected to each router. Between each router and its host, you will have to use the subnet X.[100+Y].0.0/24, where X is your group number, and Y is the ID of the router (IDs are shown on each router, for example the ID of LUGA is 9). Then, the host will have the

IP address X.[100+Y].0.1 and the interface of the router that is connected to this host will have the IP address X.[100+Y].0.2. As an example, you are group 15 and you want to configure the host connected to LUCE. This host will have the IP address 15.105.0.1/24 and the interface of the LUCE route connected to this host will have the IP address 15.105.0.2/24.

Each link between two routers can support up to 100Gb/s, except between LAUS-BERN and BERN-LUCE, where old 1Gb/s links are still deployed.

Putting your networks together: Building a mini Internet Some of your routers have external connections to your neighboring ASes. OLTE and WINT are connected to two providers (or peers if you are a Tier1 AS). GENE and LUGA are connected to two customers. STGA is connected to a peer, and LAUS is connected to an IXP.

In order to build the mini-Internet, you will have to configure eBGP sessions on these external links. Figure 3 shows how the mini-Internet you will build looks like.

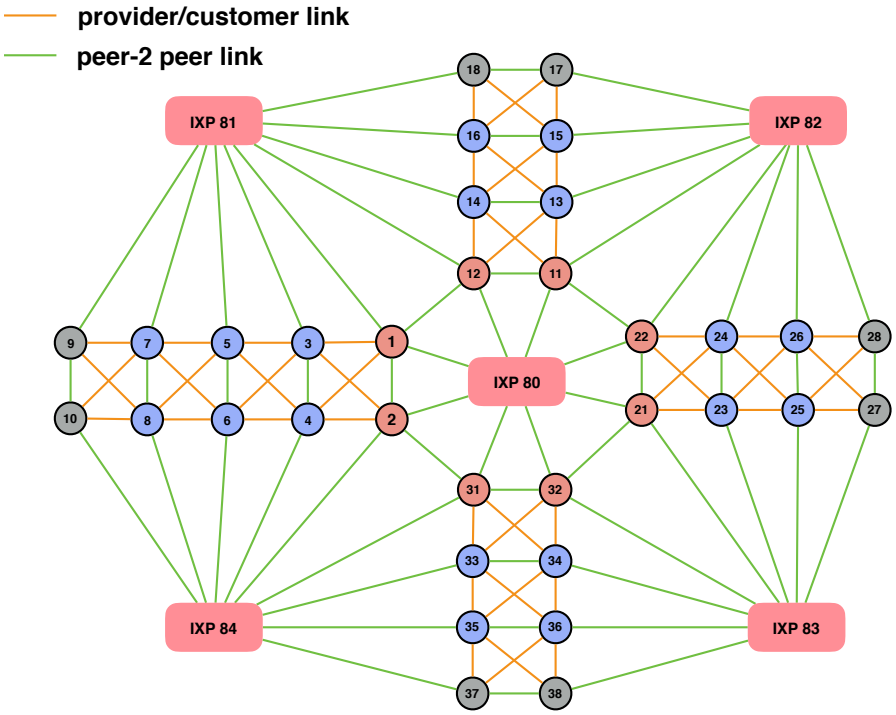

Figure 3: This figure shows the AS-level topology of our mini-Internet. Eight groups will be Tier1 ASes, eight other groups will be stub ASes, the other groups will have peers, customers and providers. Stub ASes are managed by the Communication Networks TA team.

Group 1, 2, 11, 12, 21, 22, 31, 32 are Tier1 ASes, meaning their neighboring ASes are either peers or customers. Group 9, 10, 17, 18, 27, 28, 37, 38 are stub ASes, meaning their neighboring ASes are either peers or providers but they have no customers. We (the TA team) will take care of them. All the other groups will have peers, customers and providers. For example, group 5 has two providers (3 and 4), two peers (6 and the IXP 81) and two customers (7 and 8). The file `as_connections`, available at `http://comm-net.ethz.ch/assignment2/as_connections`, shows all the connections between ASes. For each connection, it shows its type (peer, customer or provider), which router is connected to the neighboring AS, and what subnet should be used between the two ASes. Table 1 is an example of what the file `as_connections` tells you if you are group 6 (AS 6).

Based on Table 1 we can see that AS 6 is not a Tier1. It has two peers (AS5 and IXP84), two customers (AS7 and AS8) and two providers (AS3 and AS4). AS6 is connected to AS7 via its router GENE. Between AS6 and AS7, the subnet 179.24.32.0/24 must be used. AS6 is connected to AS3 via its router OLTE, and the subnet 179.24.25.0/24 must be used. You will have to talk with your neighboring ASes to decide who takes which IP address in the subnet between you and your neighboring AS. AS7 is also connected to the IXP84 via its router LAUS. In this case,

NotTier1	Prov1	6	3	OLTE	179.24.25.0/24
NotTier1	Prov2	6	4	WINT	179.24.28.0/24
NotTier1	Cust1	6	7	GENE	179.24.32.0/24
NotTier1	Cust2	6	8	LUGA	179.24.33.0/24
NotTier1	Peer1	6	5	STGA	179.24.31.0/24
NotTier1	IXPOut	6	84	LAUS	180.84.0.6/24

Table 1: An example of what you can find in the file `as_connections`

you must configure the IP address 180.84.0.6/24 on the interface of LAUS connected to the IXP. If you are connected to a Stub AS, the IP address (and not just the subnet) you must configure is also specified in this file.

There are five IXPs within our mini Internet. The IXP number is actually its AS number. For example, IXP82 has the AS number 82. The IP address of the IXP Route Server is 180.X.0.X with X the IXP number. For example the Route Server of IXP83 has he IP address 180.83.0.83. One IXP is connected to all the Tier1 ASes, allowing them to be connected in a full-mesh fashion. Each other IXP is connected to eight (or nine) ASes. For example, IXP81 is connected to AS 1, 3, 5, 7, 9, 12, 14, 16, 18. This enables these ASes to peer between them (as long as they respect the BGP customer/provider policies), instead of using (and paying!) their providers. The following example illustrates the benefit of being connected to an IXP: AS7 can send traffic to AS18 via the IXP81, instead of paying AS 5 to send the traffic via the path 7-5-3-1-12-14-16-18 if IXP81 is not used.

3 Configure your network

In this section, we show you how you can access your Quagga routers to configure them.

Each group has its own Virtual Machine (VM) As for the first practical assignment, we have prepared VMs for you. The routers composing your AS are already running in your VM and we have setup our server so that each AS is connected to its neighboring ASes only.

To access your VM, please follow the same procedure than for the first assignment, *i.e.* use the login `root`, and the port $2000 + X$, X being your group number. For example for group 11, here is how you can access the VM with ssh.

```
> ssh -p 2011 root@samichlaus.ethz.ch
```

We have sent you your new password by email. For some groups, we changed the group number to match the topology depicted in Figure 3. Please verify your group number on the Google sheet[2]. If you cannot connect to your VM, please report it in the slack channel `#exercises`. If you want to simplify the access to your VM, please use SSH key authentication[3], but do not change your password. If you want to download an entire directory (e.g. the `configs` directory) from your VM to the current directory of your own machine, you can use scp. Y is $2000 + X$ and X your group number:

```
> scp -r -P Y root@samichlaus.ethz.ch:~/path_to_the_directory .
```

If you use `tmux` [2], a terminal multiplexer, please do not use the session *mininext*. The virtual network is running in this session. If you need to reboot your VM, please contact Thomas Holterbach (`@thomas` on Slack).

How to access your routers and your hosts When your are in your VM, you can use the script `go_to.sh` to move to a router or a host. For example, with the following command, you

[2] https://docs.google.com/spreadsheets/d/1Q7ci5n3pkndnNrhCkBDfRUxgizR0xOtgpGov1fEow28/edit#gid=0
[3] Ask us or Google it if you want to know more about this.

will access the router WINT.

```
> ./go_to.sh WINT
```

WINT is a router, to access the CLI of WINT, just use the following command.

```
WINT> vtysh
```

Once you are in the CLI of a router, you can see its interface with the command `show run`. In the case of the WINT router, the interface connected to OLTE is named `olte`, the interface connected to ZURI is named `zuri`, the interface connected to STGA is named `stga`, the interface connected to the host is named `host` and the interface which is connected to another AS is named `ebgp_peer`. Use `exit` to leave the CLI, and another `exit` to leave the router.

From your VM, you can also go to a host. For example, if you want to go to the host which is connected to LUGA, just use the following command.

```
> ./go_to.sh LUGA-host
```

When you are in the host, you can use `ifconfig` to see the interface which is connected to the LUGA router. In this case, the name of the interface is `luga`.

The Quagga configuration files You can save your current Quagga configuration into a configuration file with the following command.

```
router# write file
```

You will have to do that for each router. The configuration files are named `Quagga.conf` and are located in the `config` directory (one for each router). If you want to download this entire directory from your VM to the current directory of your own machine, you can use the scp command from above.

Important: When you want to configure your router, do not edit the config files directly. Instead, always use the Quagga CLI. We recommend you to regularly save your configuration (the directory `configs`) into your own machine, as your configuration will be reset should we have to reboot your VM. In case of a reboot, you can quickly put back your configurations in the routers by copy/pasting your configurations into the Quagga CLI.

4 Verify your configuration

As any network operator, you must verify that your configuration does what you want, and debug it in case something goes wrong. We offer you several tools that you can use to verify your configuration (tools that network operators do use in practice). In this section, we show you how to use these tools.

The management VM We have setup a management VM which will enable you to launch `pings` or `traceroutes` from any AS (and not necessarily your AS), towards any destination in the mini-Internet. This will help you to know the paths used *towards* your network. The management VM is connected to each AS via the interface `mgt` of the router LUCE. The IP address of this interface is pre-configured and follows the convention X.0.199.1/24 (see Fig. 2), with X your group number. For example if you are group 15, your pre-configured IP address on the interface `mgt` at LUCE will be 15.0.199.1/24. The X.0.199.1/24 subnet must be reachable from anywhere in your network. You must therefore add it in your OSPF configuration. To access the management VM, use the following command.

```
> ssh -p 2099 students@samichlaus.ethz.ch
```

The password will be available in the `#exercices` Slack channel. Inside the VM, there is one virtual interface connected to each group via the `mgt` interface at LUCE. Those virtual interfaces are named gX, with X the group number they are connected to. For example, the virtual interface connected to group 11 is named `g11`. You can see all of them with `ifconfig`. With this management VM, you are able to launch pings and traceroutes from any group to any other group. To launch pings and traceroutes, you will use `nping` [3]. Nping is a very powerful networking tool as it gives you the ability to craft almost any field of the packets. In your case,

when you will launch a ping or a traceroute, you will have to set the source IP address, the destination IP address, the destination MAC address, the interface to use in the management VM, as well as a few additional parameters to customize your measurements. When you launch a measurement, you must set the destination MAC address to 02:02:02:02:03:X with X the AS number from where you want to launch your measurements. We strongly recommend you to look at the nping manual page using `man nping` to understand all the possible options. For example, it you want to launch a ping, from AS10, towards the host connected to OLTE in AS28, you will have to use the following command.

```
> nping --dest-mac 02:02:02:02:03:10 --interface g10 --source-ip 10.0.199.2 --dest-ip
28.101.0.1 -v0
```

The source IP is the IP of the interface `g10`, and the destination MAC address is the MAC address of the interface `mgt` of the router LUCE in AS10. To launch traceroutes, just add the parameter `--traceroute` (or `--tr`). For example, the following command launches a traceroute from AS38 towards the host connected to LAUS in AS37.

```
> nping --dest-mac 02:02:02:02:03:38 --source-ip 38.0.199.2 --dest-ip 37.107.0.1
--interface g38 -v0 --tr
```

If you want, you can also set the initial TTL value by yourself with the option `--ttl`.

The DNS service We have setup a DNS server that everyone can use. The DNS server is located in a VM connected to the interface `dns` of the router ZURI. The IP address on that interface is pre-configured, you do not need to modify it. Also, each host is pre-configured to use the DNS server. Only hosts use the DNS server, routers do not. Of course, hosts can only use the DNS server if they can reach the network 198.0.0.0/24 (where the DNS server is located). As such, do not forget to include this network in your OSPF configuration. This DNS service will help you to decode the traceroute output. Instead of a list of IP addresses, you will end up with the name of the corresponding routers. For example, 19.0.1.2 will be translated into WINT-olte.group19, because this is the IP address configured on the interface `olte` of the router WINT in the AS19.

The looking glass We have setup a looking glass service. In practice, looking glasses are servers remotely accessible which display the routing information of an IP router. For example, SWITCH, the swiss educational network, gives a public access to its looking glass[4]. This is useful to see how your BGP advertisements look like from a remote point of view. For this assignment, we make publicly available on our website (under the `looking_glass` directory) one file per group and per router showing the result of a `show ip bgp`. The files for group X are in the directory GX, there exists one file for each router. For example, if you want to get the result of a `show ip bgp` at BERN for group 23, download the file `http://comm-net.ethz.ch/assignment2/looking_glass/G23/BERN`[5]. These files are updated every minute.

The connectivity matrix We also provide you with a connectivity matrix, which shows you whether two ASes can ping between each other. Before you setup the eBGP sessions, everything will be red, meaning you can't communicate with another group. As soon as you will setup the eBGP sessions with your neighbors, the matrix will turn green for some pairs of ASes. At the end of this assignment, we hope to see this matrix completely green! The matrix is available at `http://comm-net.ethz.ch/assignment2/matrix/matrix.html` and is updated every minute.

5 Questions

The assignment will be split in three parts with a class-wide "Internet Hackathon"[6] being the middle one. The first part (easier) must be finished **before** the Internet Hackathon. It mainly involves setting up your OSPF and iBGP configuration. The second part will be done **during**

[4] https://www.switch.ch/network/tools/lookingglass/

[5] The same file ending with `.txt` is also available so that you can open it with your browser.

[6] April 12, 2017 at 6pm

the Hackathon. It involves bringing your eBGP sessions up with your neighboring ASes and advertising your prefixes to your neighbors and the IXPs. The third and last part will be done **after** the Hackathon. It involves implementing your BGP policies according to the business relationships that you have with your neighbors. In addition to these three parts, we also ask you to answer to two questions related to BGP security. The last one is a bonus question. To help you, we give you a crash course on how to configure Quagga routers in a separate document available at `http://comm-net.ethz.ch/assignment2/quagga_tuto.pdf`

For each question, you must explain in your report what you configured in your routers. You must also provide the results of the `pings` and `traceroutes` that you will have to do to confirm that your configuration works well. To submit your results, send us your final router configurations (the full directory `configs`), in a zip or tar.gz archive as well as your report (pdf).

5.1 Before the Hackathon (3.5 points)

Question 1.1 (1 point)

Configure OSPF to enable an end-to-end connectivity between all the hosts inside your AS. Before configuring OSPF, you will have to configure all the IP addresses in each interface of your routers and hosts. In each host, you will also have to configure a default gateway. For example, if you are group 1 and you want to configure the IP address and the default gateway of the host connected to ZURI, you will have to use the following commands.

```
> sudo ifconfig zuri 1.103.0.1/24
> sudo route add default gw 1.103.0.2
```

Be sure that each host can ping its directly connected router. Then, you can start configuring OSPF.

Test the host-to-host connectivity with `ping`, and look at the paths used with `traceroute`. Make sure you can reach the DNS server and the management VM from any host in your network. From now on, always prefer to launch `traceroute` from the hosts because they can use the DNS service (routers do not). If one host can't access the DNS server because the OSPF configuration is not ready yet, run `traceroute` with the option **`-n`** so that it does not try to translate each IP address found on the path.

Question 1.2 (1 point)

You expect lots of traffic flowing between your customers (connected to GENE and LUGA) and your providers[7] (connected to OLTE and WINT). To avoid congestions, you must (i) never use the 1Gb/s links except if you have no other option and (ii) configure your routers to do load-balancing. To do load-balancing, you can configure OSPF weights. You must configure two disjoint paths with the same cost between OLTE and GENE, OLTE and LUGA, WINT and GENE, WINT and LUGA. In other words, between each customer/provider pair, two disjoint paths must be used. Additional non-disjoint paths are allowed as long as there are at least two disjoint paths between each pair. Verify your configuration with `traceroute`.

Note: GENE-BERN-ZURI-OLTE and GENE-LUGA-LUCE-ZURI-WINT-OLTE are **not** disjoint because they both cross ZURI.

Question 1.3 (1 point)

A delay-sensitive application is continuously sending some traffic between the host connected to LAUS, the host connected to LUCE and the host connected to BERN. To make your client happy, you decide to leverage your 1Gb/s links and to forward the traffic between these hosts on the path with the shortest physical distance, which is LAUS-BERN-LUCE. Do traffic engineering

[7] For Tier1 ASes, we are referring to the peers connected to the IXP80, as Tier1 ASes do not have providers.

such that the traffic between LAUS-host and BERN-host, LAUS-host and LUCE-host as well as BERN-host and LUCE-host always uses the 1Gb/s links. Any other traffic must **not** use the 1Gb/s links. Also, the load-balancing you configured in Question 1.2 must still work.

Verify your configuration with `traceroute`. For example, the traffic between GENE-host and LUCE-host must **not** cross any 1Gb/s link, while the traffic between LUCE-host and LAUS-host must traverse the two 1Gb/s links. In your report, briefly explain and justify your configuration.

Hint: for this question, you can either configure appropriate OSPF weights or use static routes instead.

Question 1.4 (0.5 point)

Configure internal BGP sessions (iBGP) between all pairs of routers (full-mesh). Verify that each of your router does have an iBGP session with all the other routers with the command `sh ip bgp summary`.

To answer this question, you will have to use the `update-source` command when you configure the internal BGP sessions. We explain why and how to configure it in our quagga tutorial.

5.2 During the Hackathon (1 point)

Question 2.1 (0.5 point)

Configure the external BGP sessions (eBGP) with your neighboring ASes (including the IXPs). You must negotiate with your neighboring ASes and agree on which IP addresses should be used by you and your peer. The information about where and with who you are supposed to have an eBGP session is available at `http://comm-net.ethz.ch/pdfs/assignment2/as_connections`. Once the eBGP sessions are up, advertise your prefix to your peers. You must only advertise the /8 that has been assigned to you. Unfortunately, if you `redistribute ospf` routes into BGP, you will advertise all the /24 prefixes to your peers. Initially, your routers won't let you advertise your /8 prefix, as a router does not advertise an unreachable prefix (and this is the case with your routers, as they only know how to reach a few /24 prefixes belonging to your /8). To force your routers to advertise your /8, you must configure a static route to this prefix, and set the next hop to null. At the mean time, your peers should advertise to you their /8 prefix, as well as all the /8 prefixes they have learned (since there are no BGP policies yet). Verify that you do indeed receive these advertisements with `show ip route`, and that these advertisements are correctly propagated through your iBGP sessions. Test your connectivity from your hosts towards hosts in other ASes using `ping` and `traceroute`.

To answer this question, you will have to use the `next-hop-self` command when you configure the external BGP sessions. We explain why and how to configure it in our quagga tutorial.
Reminder: the IP address of the IXP Route Server is 180.X.0.X with X the IXP number.

Question 2.2 (0.5 points)

By default, we have configured the IXPs to not relay your BGP advertisements to their other peers. To announce a prefix to another peer via an IXP, you must specify it using a BGP community value. IXPs are configured to relay a BGP advertisement to a peer X if the advertisement has a community value equal to N:X with N the IXP number. For example, if you are AS2 and you want to advertise a prefix to AS21 via the IXP80, you must add the community value 80:21 in your BGP advertisements.

Use the community values to send BGP advertisements to the peers connected to you through an IXP. Confirm that these peers do receive your BGP advertisements by looking at the looking glass.

5.3 After the Hackathon (5 points)

Question 3.1 (2 points)

Configure your local-preference as well as the exportation rules to implement the customer/provider and peer/peer business relationships with your neighbors [4]. The connections you have through your IXP must be considered as peer-to-peer connections.

Hint: For the exportation rules, we advise you to use BGP communities to keep track of where the routes have been learned. Verify with `traceroutes` that the paths used do respect the business relationships.

To test that your BGP policies work correctly, you can use the management VM to launch `traceroute` from another AS. For example, you can launch a `traceroute` from one of your customers towards one of your peers, and verify that your AS forwards the packet directly to your peer and not to your provider. The looking glass will also help to verify your BGP advertisements. In your report, explain us your configuration, *e.g.*, what BGP communities you used (if any) and what local-preference you configured.

Question 3.2 (1 points)

The AS topology (Fig.3) shows four main regions (top, bottom, left and right). Configure your BGP policies such that you can leverage your connection with your IXP at LAUS. You do want to peer through this IXP with ASes that are located in another region. However, for business reasons, you do **not** want to peer through this IXP with ASes that are located in the same region. To not peer through the IXP with ASes in the same region, you must (i) not advertise them any prefixes and (ii) deny any advertisements coming from them. Explain your configuration in the report.

Verify that your configuration works properly by looking at the looking glass and launching `traceroutes` from your AS, or from another AS towards your AS with the management VM. To check whether you properly configured (ii), we have configured the stubs ASes to advertise their prefix to all the ASes connected to their IXP. If your configuration is correct, you should not see their advertisement coming from the IXP when you do a `sh ip bgp`.

Question 3.3 (1 points)

In this question, the goal will be to configure BGP policies in order to influence the **inbound** traffic destined to your **own** prefix.

For groups who are Tier1: IXP80 is actually a Software-Defined Internet Exchange Point (SDX)[5] which gives you a very fined-grained control over your traffic that is traversing it. To use the services offered by the SDX (IXP80), you want the inbound traffic coming from another Tier1 and destined to your own prefix to traverse IXP80.

For the other groups: Configure BGP policies such that the inbound traffic destined to your own prefix uses the provider connected to OLTE in priority.

You can test your configuration by launching `traceroutes` using the management VM. For example, you can launch a `traceroute` from a group which is reachable only via your Tier1 peers (if you are a Tier1) or your providers (if you are not a Tier1) towards an IP belonging to your prefix, and see if the traceroute does use the correct peer or provider. The looking glass can also be very useful. In your report, briefly explain us your configuration.

Question 3.4 (1 points)

In this question, the goal will be to configure BGP policies in order to influence the **outbound** traffic.

For groups who are Tier1: You have a special agreement with your two Tier1 peers which force you to send the traffic destined to their own prefix towards the direct peering link that you have with them (instead of using IXP80 or another IXP). For the rest of the traffic flowing between you and your Tier1 peers, always use the SDX (IXP80). For example, if you are AS1, the traffic destined to AS13 may cross AS12. If that is the case, make sure this traffic traverses IXP80. For the traffic destined to the prefix owned by AS12, AS1 must not use IXP80 but the direct peering link instead.

For the other groups: Force your routers to send a part of your traffic to one of your providers and another part of your traffic to the other provider.

As AS-path length can change with the time, the way you split your traffic between the two peers (or providers) should not be based on the AS-path length.

Launch `traceroutes` from your AS towards different destinations and verify that your configuration works well. In your report, briefly explain us your configuration.

5.4 What about security? (1 point). Answer in your report.

Question 4.1 (0.5 point)

How would you modify your BGP configuration to intercept traffic which is destined to someone else and normally not traversing your AS? For example, if you are AS28, how would you intercept some of the traffic flowing between AS3 and AS4?

Bonus question 4.2 (0.5 point)

For research purposes, you don't want your traffic to traverse AS35. How would you modify your BGP configuration to do that? Katz-Basset et al. [6] and Anwar et al. [7] have already done that for their research, feel free to get some inspiration from their papers.

6 General Information

6.1 If you have questions: use Slack or visit one of the exercise sessions

Use the Slack channel available at `comm-net17.slack.com`. Please do not ask questions in the `#general` channel, but use the `#exercises` channel instead. You can also send questions on Slack directly to Thomas Holterbach (`@thomas_holterbach`), Tobias Bühler (`@buehlert`), Maximilian Schütte (`@mschuette`), Alexander Dietmüller (`@adietmue`), Rüdiger Birkner (`@rbirkner`) or Roland Meier (`@roland`).

Starting from this week (April, 3 2017), you can also use the two exercises sessions (on Tuesday and Thursday) to work on the group project or to ask specific questions. We will reduce the number of theoretical questions during the group project.

6.2 Submit your work by e-mail

Send your final report and configuration by email to Laurent Vanbever (`lvanbever@ethz.ch`), Thomas Holterbach (`thomahol@ethz.ch`) and Tobias Bühler (`buehlert@ethz.ch`). Make sure that your email includes a zip or tar.gz archive containing a PDF report as well as all your configuration files (the directory `configs`). Just as a reminder, for each question, you must explain your configuration, and show screenshots of your `pings` and `traceroutes` (you can also just copy/past). Please make sure that your PDF report includes your group number as well as the name of the members composing your group. The maximum length for your PDF report is 3 to 4 A4 pages but can be more if you have a lot of screenshots.

Important: the subject of your email must follow this format: **[comm_net] groupX project 1**, where X is your group number.

6.3 Grading

This assignment will be graded and counts as 10% towards your final Communication Networks grade. There are a maximum of 10 points (plus half a bonus point). Each group member will receive the same grade: $\min\{1 + \frac{\sum pts}{2}, 6\}$

6.4 Academic integrity

We adopt a strict zero tolerance policy when it comes to cheating. Cheating will immediately translate to the group failing the assignment and being reported to ETH administration. In particular, you can only do your assignment with the other members of your group. Do not look at other groups' configuration and do not copy configurations from anywhere. It is OK to discuss things or find help online but you must do the work by yourself.

Your configuration and report may be checked with automated tools so as to discover plagiarism. Again, **do not copy-and-paste** code, text, etc.

References

[1] Quagga Routing Suite. [Online]. Available: http://www.nongnu.org/quagga/

[2] Tmux, a terminal multiplexer. [Online]. Available: https://tmux.github.io

[3] Nmap. [Online]. Available: https://nmap.org

[4] L. Gao and J. Rexford, "Stable internet routing without global coordination," SIGMETRICS Perform. Eval. Rev., vol. 28, no. 1, pp. 307–317, Jun. 2000.

[5] A. Gupta, L. Vanbever, M. Shahbaz, S. P. Donovan, B. Schlinker, N. Feamster, J. Rexford, S. Shenker, R. Clark, and E. Katz-Bassett, "Sdx: A software defined internet exchange," SIGCOMM Comput. Commun. Rev., vol. 44, no. 4, pp. 551–562, Aug. 2014. [Online]. Available: http://doi.acm.org/10.1145/2740070.2626300

[6] E. Katz-Bassett, C. Scott, D. R. Choffnes, I. Cunha, V. Valancius, N. Feamster, H. V. Madhyastha, T. Anderson, and A. Krishnamurthy, "Lifeguard: Practical repair of persistent route failures," in Proceedings of the ACM SIGCOMM 2012 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication, ser. SIGCOMM '12. New York, NY, USA: ACM, 2012, pp. 395–406.

[7] R. Anwar, H. Niaz, D. Choffnes, I. Cunha, P. Gill, and E. Katz-Bassett, "Investigating interdomain routing policies in the wild," in Proceedings of the 2015 ACM Conference on Internet Measurement Conference, ser. IMC '15. New York, NY, USA: ACM, 2015, pp. 71–77.