Computer Networks
Jacobs University Bremen
Dr. Jürgen Schönwälder, Johannes Schauer

Course: 320301
Date: 2013-10-11
Deadline: 2013-10-18

Problem Sheet #3

Download and install cbgp and its dependency libgds. Either go to http://c-bgp.sourceforge.net/ and http://libgds.sourceforge.net/ to manually download and compile both source packages or go to http://mister-muffin.de/cbgp/ to download Debian packages for cbgp, libgds and libcsim for your architecture (i386 and amd64 are provided). Should you not have a Debian based distribution or otherwise want to compile from source, you might want to have a look at the patches located in '‘./debian/patches‘ and the arguments with which ‘./configure‘ is called in ‘./debian/rules‘. Both can be found in the ‘*.debian.tar.gz‘ files at http://mister-muffin.de/cbgp/. The patches and configure options are needed to compile cbgp and libgds on current systems. The tutorial at http://c-bgp.sourceforge.net/tutorial.php to explains how to build a topology using cbgp. You will find example scripts in ‘/usr/share/doc/cbgp/examples‘. To understand them, refer to the comments as well as the documentation at http://c-bgp.sourceforge.net/doc/html/index.html. The cbgp executable can either run in interactive mode (no arguments) or read commands that can otherwise be given in interactive mode from a script specified with ‘-c‘ or via stdin. All commands that are used in the scripts can be used in interactive mode and vice versa.

Problem 3.1: *bgp simulation of filter rules* (2+2 = 4 points)

Extract the script “topology1.cli”, which is appended to this problem sheet. It implements the following topology:

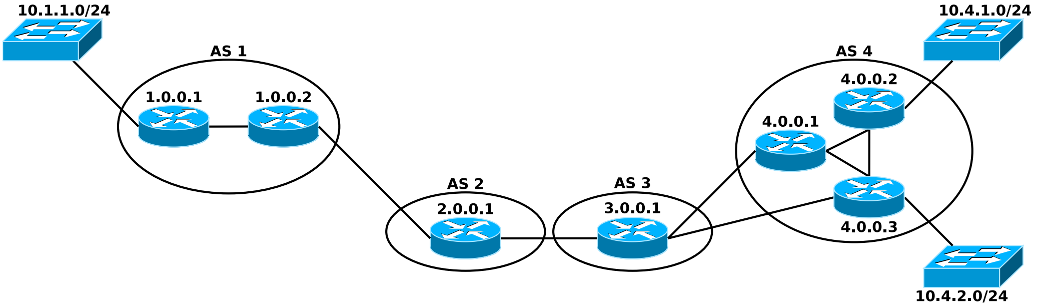

At the end of the script, traceroute is run to show how 10.1.1.0 is reached from 4.0.0.2 and how 10.4.1.0 is reached from 1.0.0.1.

a) Modify the script mentioned above by entering the following lines before “run sim”:

```
bgp router 4.0.0.1 peer 3.0.0.1
        filter in
                add-rule
                        match any
                        action "local-pref 50"
                        exit
                exit
        exit

bgp router 4.0.0.3 peer 3.0.0.1
        filter in
                add-rule
                        match any
                        action "local-pref 100"
                        exit
                exit
        exit
```

What changes and why? Does it work the same if it would be an ”out” instead of an ”in” filter? If not, why?

b) Modify the script mentioned above by entering the following lines before “run sim”:

```
bgp router 4.0.0.1 peer 3.0.0.1
        filter out
                add-rule
                        match any
                        action "metric 2"
                        exit
                exit
        exit

bgp router 4.0.0.3 peer 3.0.0.1
        filter out
                add-rule
                        match any
                        action "metric 1"
                        exit
                exit
        exit
```

What changes and why? Does it work the same if it would be an ”in” instead of an ”out” filter? If not, why?

Solution:

a) Packets from 4.0.0.2 to 10.1.1.0 will now use 4.0.0.3 instead of 4.0.0.1 to go into AS3. 3.0.0.1 advertises its route to 4.0.0.1 and to 4.0.0.3 but with the higher local-pref, 4.0.0.3 is now chosen as the exit point to AS3 from AS4.

The local-pref must be given as an ”in” filter because it influences how arriving routing advertisements coming from the outside are interpreted by the current AS.

b) Packets from 1.0.0.1 to 10.4.1.0 will now use 4.0.0.3 instead of 4.0.0.1 to go into AS4. 4.0.0.3 and 4.0.0.1 advertise their routes to AS3 but AS3 will choose the route with the lower MED.

The metric must be given as an ”out” filter because it influences the advertisement that the current AS sends to the outside to neighbors.

Problem 3.2: *bgp simulation of filter rules* (3+3 = 6 points)

a) Extend the script “topology1.cli” such that it represents the following topology.

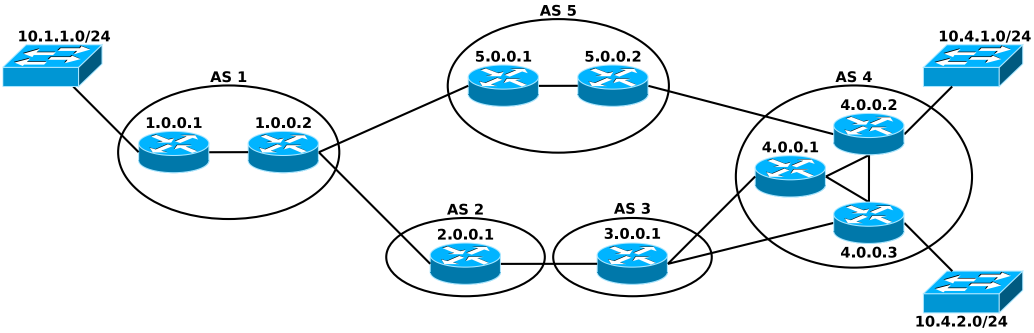

Make sure, that packets from AS1 to AS4 and vice versa now go through AS5.

b) We start with the script for the extended topology. As the owner of AS5 you decide that you do not like to route traffic from AS4 to AS1 but you have a contract with the owner of AS1 to route traffic to AS4 as usual. Implement a policy in AS5 that forbids to route traffic from AS4 to anywhere outside AS5 but still allows AS5 to communicate with AS4 and vice versa. Essentially, implement that AS5 is not a transit AS for traffic from AS4. Do not touch the routing policies of any other AS than AS5. Do not use ”metric” or ”local-pref”.

Solution:

a) The solution is in the script file “topology2.cli”.

b) The following filter achieves the goal:

```
bgp router 5.0.0.2 peer 4.0.0.2
      filter in
              add-rule
                      match "prefix in 5/8"
                      action accept
                      exit
              add-rule
                      match any
                      action deny
                      exit
              exit
      exit
```

Script topology1.cli

```
############
# TOPOLOGY #
############

# set up the topology for AS 1
net add node 1.0.0.1
net add node 1.0.0.2
net add link 1.0.0.1 1.0.0.2

# set up the topology for AS 2
net add node 2.0.0.1

# set up the topology for AS 3
net add node 3.0.0.1

# set up the topology for AS 4
net add node 4.0.0.1
net add node 4.0.0.2
net add node 4.0.0.3
net add link 4.0.0.1 4.0.0.2
net add link 4.0.0.1 4.0.0.3
net add link 4.0.0.2 4.0.0.3

# set up inter domain links
net add link 1.0.0.2 2.0.0.1
net add link 2.0.0.1 3.0.0.1
net add link 3.0.0.1 4.0.0.1
net add link 3.0.0.1 4.0.0.3

#######
# IGP #
#######

# IGP configuration for AS 1

net add domain 1 igp
net node 1.0.0.1 domain 1
net node 1.0.0.2 domain 1
net link 1.0.0.1 1.0.0.2 igp-weight --bidir 10
net domain 1 compute

# IGP configuration for AS 4

net add domain 4 igp
net node 4.0.0.1 domain 4
net node 4.0.0.2 domain 4
net node 4.0.0.3 domain 4
net link 4.0.0.1 4.0.0.2 igp-weight --bidir 10
net link 4.0.0.1 4.0.0.3 igp-weight --bidir 10
net link 4.0.0.2 4.0.0.3 igp-weight --bidir 10
net domain 4 compute

##########
# ROUTES #
##########

# static routes for AS 1

net node 1.0.0.2 route add --oif=2.0.0.1 2.0.0.1/32 0
```

```
# static routes for AS 2

net node 2.0.0.1 route add --oif=1.0.0.2 1.0.0.2/32 0
net node 2.0.0.1 route add --oif=3.0.0.1 3.0.0.1/32 0

# static routes for AS 3

net node 3.0.0.1 route add --oif=2.0.0.1 2.0.0.1/32 0
net node 3.0.0.1 route add --oif=4.0.0.1 4.0.0.1/32 0
net node 3.0.0.1 route add --oif=4.0.0.3 4.0.0.3/32 0

# static routes for AS 4

net node 4.0.0.1 route add --oif=3.0.0.1 3.0.0.1/32 0
net node 4.0.0.3 route add --oif=3.0.0.1 3.0.0.1/32 0


#######
# BGP #
#######

# full-mesh of iBGP sessions in AS 1

bgp add router 1 1.0.0.1
bgp add router 1 1.0.0.2
bgp domain 1 full-mesh

# full-mesh of iBGP sessions in AS 2

bgp add router 2 2.0.0.1
bgp domain 2 full-mesh

# full-mesh of iBGP sessions in AS 3

bgp add router 3 3.0.0.1
bgp domain 3 full-mesh

# full-mesh of iBGP sessions in AS 1

bgp add router 4 4.0.0.1
bgp add router 4 4.0.0.2
bgp add router 4 4.0.0.3
bgp domain 4 full-mesh

# Setup eBGP sessions for AS 1
bgp router 1.0.0.2
        add peer 2 2.0.0.1
        peer 2.0.0.1
                next-hop-self
                up
                exit
        exit

# Setup eBGP sessions for AS 2
bgp router 2.0.0.1
        add peer 1 1.0.0.2
        peer 1.0.0.2
                next-hop-self
                up
                exit
        add peer 3 3.0.0.1
        peer 3.0.0.1
                next-hop-self
                up
                exit
        exit

# Setup eBGP sessions for AS 3
bgp router 3.0.0.1
        add peer 2 2.0.0.1
        peer 2.0.0.1
                next-hop-self
                up
                exit
        add peer 4 4.0.0.1
        peer 4.0.0.1
                next-hop-self
                up
                exit
        add peer 4 4.0.0.3
        peer 4.0.0.3
                next-hop-self
                up
                exit
        exit

# Setup eBGP sessions for AS 4
bgp router 4.0.0.1
        add peer 3 3.0.0.1
        peer 3.0.0.1
                next-hop-self
                up
                exit
        exit

bgp router 4.0.0.3
        add peer 3 3.0.0.1
        peer 3.0.0.1
                next-hop-self
                up
                exit
        exit

# Originate own prefixes for AS 1
bgp router 1.0.0.1 add network 1/8
bgp router 1.0.0.2 add network 1/8

# Originate own prefixes for AS 2
bgp router 2.0.0.1 add network 2/8

# Originate own prefixes for AS 3
bgp router 3.0.0.1 add network 3/8

# Originate own prefixes for AS 4
bgp router 4.0.0.1 add network 4/8
bgp router 4.0.0.2 add network 4/8
bgp router 4.0.0.3 add network 4/8

# Add other networks to AS 1
bgp router 1.0.0.1 add network 10.1.1.0/24

# Add other networks to AS 4
bgp router 4.0.0.2 add network 10.4.1.0/24
bgp router 4.0.0.3 add network 10.4.2.0/24

##############
```

```
# SIMULATION #
##############

sim run

print "executing: net node 1.0.0.1 traceroute 10.4.1.0\n"

net node 1.0.0.1 traceroute 10.4.1.0

print "executing: net node 4.0.0.2 traceroute 10.1.1.0\n"

net node 4.0.0.2 traceroute 10.1.1.0
```

Script topology2.cli

```
############
# TOPOLOGY #
############

# set up the topology for AS 1
net add node 1.0.0.1
net add node 1.0.0.2
net add link 1.0.0.1 1.0.0.2

# set up the topology for AS 2
net add node 2.0.0.1

# set up the topology for AS 3
net add node 3.0.0.1

# set up the topology for AS 4
net add node 4.0.0.1
net add node 4.0.0.2
net add node 4.0.0.3
net add link 4.0.0.1 4.0.0.2
net add link 4.0.0.1 4.0.0.3
net add link 4.0.0.2 4.0.0.3

# set up the topology for AS 5
net add node 5.0.0.1
net add node 5.0.0.2
net add link 5.0.0.1 5.0.0.2

# set up inter domain links
net add link 1.0.0.2 2.0.0.1
net add link 1.0.0.2 5.0.0.1
net add link 2.0.0.1 3.0.0.1
net add link 3.0.0.1 4.0.0.1
net add link 3.0.0.1 4.0.0.3
net add link 5.0.0.2 4.0.0.2

#######
# IGP #
#######

# IGP configuration for AS 1

net add domain 1 igp
net node 1.0.0.1 domain 1
net node 1.0.0.2 domain 1
net link 1.0.0.1 1.0.0.2 igp-weight --bidir 10
net domain 1 compute

# IGP configuration for AS 4

net add domain 4 igp
net node 4.0.0.1 domain 4
net node 4.0.0.2 domain 4
net node 4.0.0.3 domain 4
net link 4.0.0.1 4.0.0.2 igp-weight --bidir 10
net link 4.0.0.1 4.0.0.3 igp-weight --bidir 10
net link 4.0.0.2 4.0.0.3 igp-weight --bidir 10
net domain 4 compute

# IGP configuration for AS 5
```

```
net add domain 5 igp
net node 5.0.0.1 domain 5
net node 5.0.0.2 domain 5
net link 5.0.0.1 5.0.0.2 igp-weight --bidir 10
net domain 5 compute

##########
# ROUTES #
##########

# static routes for AS 1

net node 1.0.0.2 route add --oif=2.0.0.1 2.0.0.1/32 0
net node 1.0.0.2 route add --oif=5.0.0.1 5.0.0.1/32 0

# static routes for AS 2

net node 2.0.0.1 route add --oif=1.0.0.2 1.0.0.2/32 0
net node 2.0.0.1 route add --oif=3.0.0.1 3.0.0.1/32 0

# static routes for AS 3

net node 3.0.0.1 route add --oif=2.0.0.1 2.0.0.1/32 0
net node 3.0.0.1 route add --oif=4.0.0.1 4.0.0.1/32 0
net node 3.0.0.1 route add --oif=4.0.0.3 4.0.0.3/32 0

# static routes for AS 4

net node 4.0.0.1 route add --oif=3.0.0.1 3.0.0.1/32 0
net node 4.0.0.3 route add --oif=3.0.0.1 3.0.0.1/32 0
net node 4.0.0.2 route add --oif=5.0.0.2 5.0.0.2/32 0

# static routes for AS 5
net node 5.0.0.1 route add --oif=1.0.0.2 1.0.0.2/32 0
net node 5.0.0.2 route add --oif=4.0.0.2 4.0.0.2/32 0

#######
# BGP #
#######

# full-mesh of iBGP sessions in AS 1

bgp add router 1 1.0.0.1
bgp add router 1 1.0.0.2
bgp domain 1 full-mesh

# full-mesh of iBGP sessions in AS 2

bgp add router 2 2.0.0.1
bgp domain 2 full-mesh

# full-mesh of iBGP sessions in AS 3

bgp add router 3 3.0.0.1
bgp domain 3 full-mesh

# full-mesh of iBGP sessions in AS 1

bgp add router 4 4.0.0.1
bgp add router 4 4.0.0.2
bgp add router 4 4.0.0.3
bgp domain 4 full-mesh
```

```
# full-mesh of iBGP sessions in AS 1

bgp add router 5 5.0.0.1
bgp add router 5 5.0.0.2
bgp domain 5 full-mesh

# Setup eBGP sessions for AS 1
bgp router 1.0.0.2
        add peer 2 2.0.0.1
        peer 2.0.0.1
                next-hop-self
                up
                exit
        add peer 5 5.0.0.1
        peer 5.0.0.1
                next-hop-self
                up
                exit
        exit

# Setup eBGP sessions for AS 2
bgp router 2.0.0.1
        add peer 1 1.0.0.2
        peer 1.0.0.2
                next-hop-self
                up
                exit
        add peer 3 3.0.0.1
        peer 3.0.0.1
                next-hop-self
                up
                exit
        exit

# Setup eBGP sessions for AS 3
bgp router 3.0.0.1
        add peer 2 2.0.0.1
        peer 2.0.0.1
                next-hop-self
                up
                exit
        add peer 4 4.0.0.1
        peer 4.0.0.1
                next-hop-self
                up
                exit
        add peer 4 4.0.0.3
        peer 4.0.0.3
                next-hop-self
                up
                exit
        exit

# Setup eBGP sessions for AS 4
bgp router 4.0.0.1
        add peer 3 3.0.0.1
        peer 3.0.0.1
                next-hop-self
                up
                exit
        exit
```

```
bgp router 4.0.0.3
        add peer 3 3.0.0.1
        peer 3.0.0.1
                next-hop-self
                up
                exit
        exit

bgp router 4.0.0.2
        add peer 5 5.0.0.2
        peer 5.0.0.2
                next-hop-self
                up
                exit
        exit

# Setup eBGP sessions for AS 5

bgp router 5.0.0.1
        add peer 1 1.0.0.2
        peer 1.0.0.2
                next-hop-self
                up
                exit
        exit

bgp router 5.0.0.2
        add peer 4 4.0.0.2
        peer 4.0.0.2
                next-hop-self
                up
                exit
        exit


# Originate own prefixes for AS 1
bgp router 1.0.0.1 add network 1/8
bgp router 1.0.0.2 add network 1/8

# Originate own prefixes for AS 2
bgp router 2.0.0.1 add network 2/8

# Originate own prefixes for AS 3
bgp router 3.0.0.1 add network 3/8

# Originate own prefixes for AS 4
bgp router 4.0.0.1 add network 4/8
bgp router 4.0.0.2 add network 4/8
bgp router 4.0.0.3 add network 4/8

# Originate own prefixes for AS 5
bgp router 5.0.0.1 add network 5/8
bgp router 5.0.0.2 add network 5/8

# Add other networks to AS 1
bgp router 1.0.0.1 add network 10.1.1.0/24

# Add other networks to AS 4
bgp router 4.0.0.2 add network 10.4.1.0/24
bgp router 4.0.0.3 add network 10.4.2.0/24

##############
# SIMULATION #
```

```
##############

sim run

print "executing: net node 1.0.0.1 traceroute 10.4.1.0\n"

net node 1.0.0.1 traceroute 10.4.1.0

print "executing: net node 4.0.0.2 traceroute 10.1.1.0\n"

net node 4.0.0.2 traceroute 10.1.1.0
```