

## 8th Assignment: Network Protocols and Architectures, WS 14/15

**Question 1:** (5 + 5 + 5 + 5 + 5 = 25 points) *Routing / The border gateway protocol*

- Why are policies in the area of intra-AS routing protocols (e. g., OSPF) of little importance?
- Why do policies on the other hand play a more important role within Inter-AS traffic compared to global traffic optimization?
- Why doesn't it make sense to optimize routes for minimum (packet) delay in intra-AS routing protocols?
- How does BGP—as a member of the path-vector protocol family—bypass the problem of routing loops?
- The de-facto standard for inter-AS routing is BGP. Why is it difficult to enforce alternatives to BGP or switch completely to a different protocol that is incompatible with BGP?

**Question 2:** (5 + 5 + 10 + 5 = 25 points) *Properties of BGP*




Figure 1: Intra-AS setup

Consider the topology shown in Figure 1. All edges have a weight of 2. An intra-AS routing protocol should be used.

- Identify the cost-optimal route from A to B.
- Suppose that the edge weight between Q2 and Q3 increases to 6. Given this case, state the cost optimal route from A to B. Which path will a packet actually take between A and B? Explain why.

As shown in Figure 2, the network is now split into two autonomous systems (AS). BGP is used as routing protocol between those autonomous systems.

- The edge weight between Q2 and Q3 is still 6. How will a packet from A to B be routed? Explain why.
- Supposing that the connection between Q2 and Q3 drops out, how are packets routed from A to B in this case? Explain why.

**Question 3:** (5 + 20 + 10 + 5 + 10 = 50 points) *BGP experiment*

We are going to explore BGP in a more practical manner by visualizing the route to the Universidad de La Habana, Cuba. First we need the IP address of the target host:




Figure 2: BGP setup

```
$ host www.uh.cu
www.uh.cu has address 200.55.139.216
```

Next, we examine the route of a host (in this example is the source a machine in the network of TU Berlin) to the target by using `traceroute`<sup>1</sup>. As you can see, we get the DNS names and IP addresses of the intermediate routers, but actually we would be interested in the AS numbers. More coming next.

Connect to `route-server.as3257.net` by using `telnet`. This server provides you with an emulation of a Cisco IOS shell with the possibility to explore BGP routes from there to any IP address in the world. Enter the following command at the prompt:

```
route-server.as3257.net>show ip bgp 200.55.139.216
```

- (a) To which AS does the `www.uh.cu` host belong to?
- (b) Which ASes are hit on the BGP route to `www.uh.cu`? State the AS number and the corresponding name of the AS. Copy the result of `show ip bgp` to your solution.
  - Hint: The line that starts with 3257 represents the BGP route. AS 3257 represents the Tiscali backbone network. To receive more detailed information about AS numbers, please visit <http://www.arin.net/> (US), or <http://www.ripe.net/whois> (Europe). The syntax to query the RIPE database to get information about the AS number needs a capital “AS” as prefix.
- (c) Visualize the result of `traceroute` in a drawing with each AS as a dashed ellipses, routers as small circles and links as lines. Include the IP addresses and AS numbers and the location information that you may guess from the routers name.

Hints:

- `route-server.as3257.net > traceroute 200.55.139.216`
- Depending on the state of the networks along the route, `traceroute` may not reach the destination network.
- The AS numbers are given in the `traceroute` output. In general, the IP to AS number mapping can be obtained by invoking `show ip bgp` for each IP address or by using a web service<sup>2</sup>.
- “ham”, “ams”, ... represent city names. “ham” is equivalent to Hamburg. Hint: Abbreviations for city names in DNS names of routers are often chosen according to airport codes<sup>3</sup>.
- Instead of using the complete IP address scheme, you can also use a prefix, e. g., 188.1.0.0/16 and label the router with the remaining part of the IP address, e. g., 33.81, 144.221, and 145.137.
- You can use geolocation databases like <http://www.iplocation.net/> in addition. (Those information might not be always accurate!)

<sup>1</sup><http://en.wikipedia.org/wiki/Traceroute>

<sup>2</sup><http://asn.cymru.com/cgi-bin/whois.cgi>

<sup>3</sup>[http://en.wikipedia.org/wiki/List\\_of\\_airports\\_by\\_IATA\\_code](http://en.wikipedia.org/wiki/List_of_airports_by_IATA_code)

- (d) Compare the route established from route-server.as3257.net to www.uh.cu with the route below, originating at TU Berlin towards www.uh.cu. Which part of the paths are identical? Where do they differ?

```

$ traceroute 200.55.139.216
traceroute to 200.55.139.216 (200.55.139.216), 30 hops max, 60 byte packets
 1 130.149.220.126 (130.149.220.126) 0.214 ms 0.183 ms 0.167 ms
 2 ta-inet.gate.tu-berlin.de (130.149.235.193) 0.654 ms 1.041 ms 1.035 ms
 3 ma-ta.gate.tu-berlin.de (130.149.126.113) 0.997 ms 1.006 ms 1.001 ms
 4 en-ma.gate.tu-berlin.de (130.149.126.69) 0.978 ms 0.971 ms 0.965 ms
 5 cr-tub1-te0-0-0-15.x-win.dfn.de (188.1.235.117) 1.342 ms 1.344 ms 1.335 ms
 6 be4193.rcr11.b015814-1.ham01.atlas.cogentco.com (149.6.142.101) 5.624 ms 6.016 ms 6.199 ms
 7 be2460.ccr42.ham01.atlas.cogentco.com (154.54.38.241) 5.539 ms
   be2198.ccr41.ham01.atlas.cogentco.com (154.54.39.5) 5.657 ms 5.659 ms
 8 be2187.ccr42.ams03.atlas.cogentco.com (154.54.74.125) 14.679 ms 15.931 ms 15.010 ms
 9 be2488.ccr42.lon13.atlas.cogentco.com (154.54.39.110) 93.473 ms 93.455 ms
   be2194.ccr41.lon13.atlas.cogentco.com (130.117.50.241) 96.224 ms
10 be2493.ccr22.bos01.atlas.cogentco.com (154.54.42.97) 98.613 ms
   be2387.ccr22.bos01.atlas.cogentco.com (154.54.44.165) 94.299 ms
   be2317.ccr41.jfk02.atlas.cogentco.com (154.54.30.185) 93.972 ms
11 be2060.ccr21.jfk05.atlas.cogentco.com (154.54.31.10) 93.354 ms
   be2096.ccr42.jfk02.atlas.cogentco.com (154.54.30.41) 95.299 ms
   be2317.ccr41.jfk02.atlas.cogentco.com (154.54.30.185) 97.516 ms
12 be2151.ccr42.jfk02.atlas.cogentco.com (154.54.40.74) 102.148 ms
   be2107.ccr42.jfk02.atlas.cogentco.com (154.54.3.93) 97.177 ms
   be2061.ccr21.jfk05.atlas.cogentco.com (154.54.3.70) 96.959 ms
13 tata.jfk05.atlas.cogentco.com (154.54.12.18) 97.795 ms 95.285 ms
   63.243.128.69 (63.243.128.69) 117.425 ms
14 tata.jfk05.atlas.cogentco.com (154.54.12.18) 101.494 ms 101.627 ms
   if-5-5.tcore1.NYY-New-York.as6453.net (216.6.90.5) 116.697 ms
15 63.243.128.69 (63.243.128.69) 127.038 ms
   if-11-2.tcore2.NYY-New-York.as6453.net (216.6.99.1) 117.026 ms 117.014 ms
16 if-11-2.tcore2.NYY-New-York.as6453.net (216.6.99.1) 121.630 ms 117.774 ms
   if-12-6.tcore1.CT8-Chicago.as6453.net (216.6.99.46) 115.610 ms
17 if-12-6.tcore1.CT8-Chicago.as6453.net (216.6.99.46) 119.976 ms 119.996 ms
   if-22-2.tcore2.CT8-Chicago.as6453.net (64.86.79.1) 114.217 ms
18 if-12-6.tcore1.CT8-Chicago.as6453.net (216.6.99.46) 123.198 ms 123.221 ms 120.307 ms
19 66.198.96.34 (66.198.96.34) 279.080 ms
   if-3-2.tcore1.W6C-Montreal.as6453.net (66.198.96.45) 120.255 ms
   66.198.96.34 (66.198.96.34) 282.755 ms
20 200.0.16.85 (200.0.16.85) 259.453 ms 260.798 ms
   66.198.96.34 (66.198.96.34) 287.252 ms
21 200.0.16.113 (200.0.16.113) 265.926 ms
   66.198.96.34 (66.198.96.34) 291.867 ms
   200.0.16.113 (200.0.16.113) 265.555 ms
22 200.0.16.85 (200.0.16.85) 267.474 ms
   200.0.16.141 (200.0.16.141) 265.580 ms
   200.0.16.113 (200.0.16.113) 265.877 ms
23 200.0.16.113 (200.0.16.113) 267.458 ms 268.718 ms *
24 * * *
25 * * *
26 * * *
27 * * *
28 * * *
29 * * cisco2800.uh.cu (200.55.139.221) 259.991 ms
30 * internet.uh.cu (200.55.139.216) 259.601 ms 262.057 ms

```

- (e) Comment on the routes to www.uh.cu regarding the geographical distance. Speculate about the reasons for the *unnecessarily* long routes.

**Due Date: Wednesday, December, 17th 2014 only until 14:00 h s. t.**

- As PDF files (no MS Office or OpenOffice files): Uploaded via ISIS (<https://www.isis.tu-berlin.de/2.0/course/view.php?id=2560>)
- Put your name, StudentID number (Matrikelnummer) **and** the name of your tutor on your solution.