Technische Universität München
Informatik VIII
Netzarchitekturen und Netzdienste
Prof. Dr.-Ing. Georg Carle
Christian Grothoff, PhD
Dr. Nils Kammenhuber

Masterkurs Rechnernetze / Master Lecture on Computer Networks (IN2097) — Tutorial

Class Assignment No. 3, WS 2009/2010

Abgabedatum / To be handed in by: Monday, 2010-01-18, 14h

Exercise 7 — Routing fundamentals

Goal: Basic understanding of principles and concepts in the context of routing.

Answer each of the following questions; *briefly* explain your answer.

a) A packet arrives at a router. Which of the following tasks does the router do with the packet, and which ones not: Route it? Forward it? Switch it?

b) Can a router have multiple routing tables?

c) Can a router have multiple forwarding tables?

d) Does a router have an IP address?

e) Do normal end hosts run routing protocols?

f) Recall that Inter-AS routing is not symmetric, due to hot-potato routing. Can we assume that, in contrast, routing *within* an autonomous system featuring OSPF routing always uses symmetrical paths?

g) Draw a simple network topology where the split horizon / poisoned reverse technique cannot prevent the count-to-infinity effect.

Exercise 8 — IP address prefixes

Goal: Understanding IP address prefixes

Can you aggregate the following IPv4 address ranges or IPv4 prefixes into an IPv4 prefix (of the type a.b.c.d/p)? If yes, which one? If no, why not?

a) 10.0.0.0 to 10.255.255.255

b) 10.0.0.0 to 10.0.0.3

c) 10.0.0.0 to 10.0.0.5

d) 192.168.0.0/24 and 192.168.1.0/24

e) 172.16.0.0/17 and 172.16.128.0/17

f) 172.16.0.0/17 and 172.16.127.0/17

g) 192.168.1.0/24 and 192.168.2.0/24

h) 10.0.0.0/16 and 10.1.0.0/17

i) 10.0.0.0/16 and 10.1.0.0/17 and 10.1.128.0/17

Exercise 9 — IP routing/forwarding

Goal: Seeing that not all paths through a network are feasible with IP routing, even when manually configuring the network

In the following IP network, the routing is manually configured such that any traffic from host x to host z is being forwarded exclusively along the bold arrows via router A, C, and E:

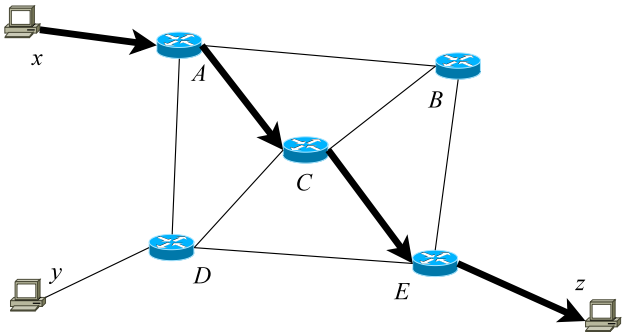

Consider IP traffic flowing from y to z. Which of the following paths can be manually configured for traffic $y \rightsquigarrow z$, and which one cannot — under the constraint that we must not alter the path for traffic from x to z? If a path is not feasible, explain why.

a) $y \to D \to E \to z$
b) $y \to D \to C \to A \to C \to E \to z$
c) $y \to D \to C \to A \to B \to E \to z$
d) $y \to D \to C \to B \to E \to z$
e) $y \to D \to A \to C \to E \to z$
f) $y \to D \to A \to C \to B \to E \to z$

Exercise 10 — (In-)Visibility of the network topology

Goal: Understanding the implications of business relationships on BGP policy routing; understanding why it is infeasible to see the entire topology of the Internet

Let us assume that the topology of the entire Internet looked like this graph:

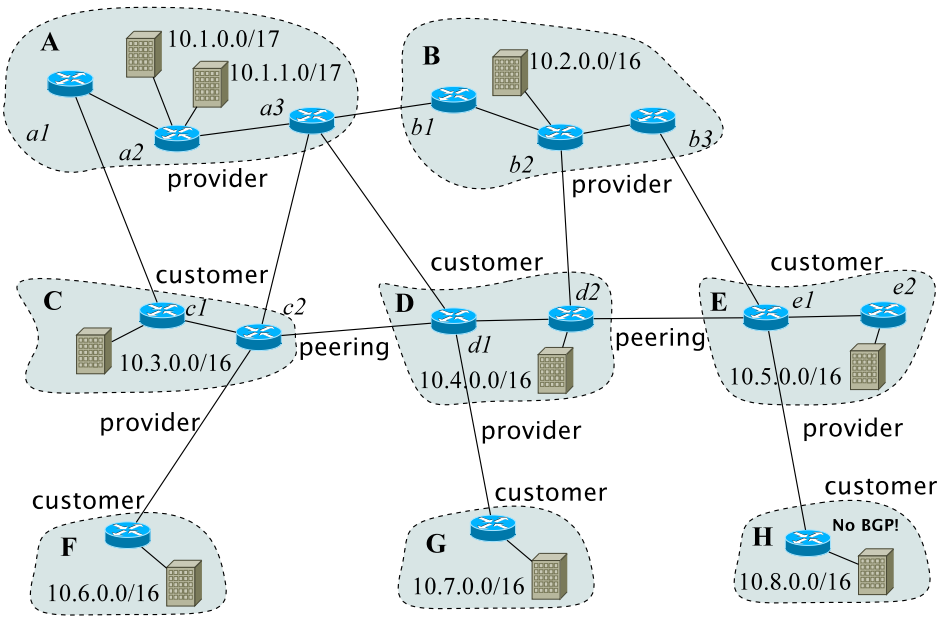

Apart from running some kind of intra-AS routing protocol, every router runs a BGP session to each of its neighbouring routers — with the **exception** of AS H, which does not run BGP on its routers and simply has a *default route* (i.e., prefix `0.0.0.0/0`) to its provider, AS E. ISPs *A* and *B* are Tier-1 providers. By only considering information from intra-AS and inter-AS routing protocols, determine the parts of the network topology (destination prefixes, ASes, routers, links between routers, abstract links between ASes, etc.) that are visible from the following ASes:

- *A*,
- *C*,
- *F*,
- *H*.

Draw those parts of the topology that they can infer from their BGP and Intra-AS routing protocols (i. e., individual routers, entire ASes, IP prefixes, etc.); aggregate/draw dotted lines for those parts that are only partially distinguishable; omit anything that is not visible. Should you feel unsure regarding the visibility of specific items, give a short explanation.

Exercise 11 — Network Address Translation

Goal: Understanding the mode of operation of different type of Network Address Translation boxes; understanding which techniques we can use to traverse them.

Consider a peer to peer network where every peer resides behind a Network Address Translation device. The goal is to establish a connection between two arbitrary hosts. We assume that we can exchange messages via a rendezvous point (RP) and that all NATs are of type Full Cone.

The following features *may* or *may not* be available:

- a STUN server that helps us to retrieve the external IP address and port of a mapping
- an ALG for the Session Initiation protocol (SIP) for port 5060
- UPnP
- static port forwarding entries
- a relay server in the public Internet

a) Please describe the advantages and disadvantages of the techniques listed above.

b) Find at least 4 possible combinations two peers can use to establish a UDP connection (to port 10000) between them. Please note who has to initiate the connection and what has to be exchanged via the RP before the connection can be established. (e.g. A uses nothing and B uses UPnP and A initiates the connection towards B. B has to send the allocated UPnP transport address to A via the RP).

c) What changes if both hosts implement a symmetric NAT and port prediction is not possible?

d) Try to figure out what type of NAT you have at home and describe how you figured it out. (hint: the STUN algorithm has been implemented not only as standalone tools (Java and C) but also as part of VoIP clients such as Ekiga.)

Exercise 12 — Hole Punching

Goal: Understanding how hole punching can be used to establish a direct connection between two hosts behind NAT.

Let's assume the following scenario:

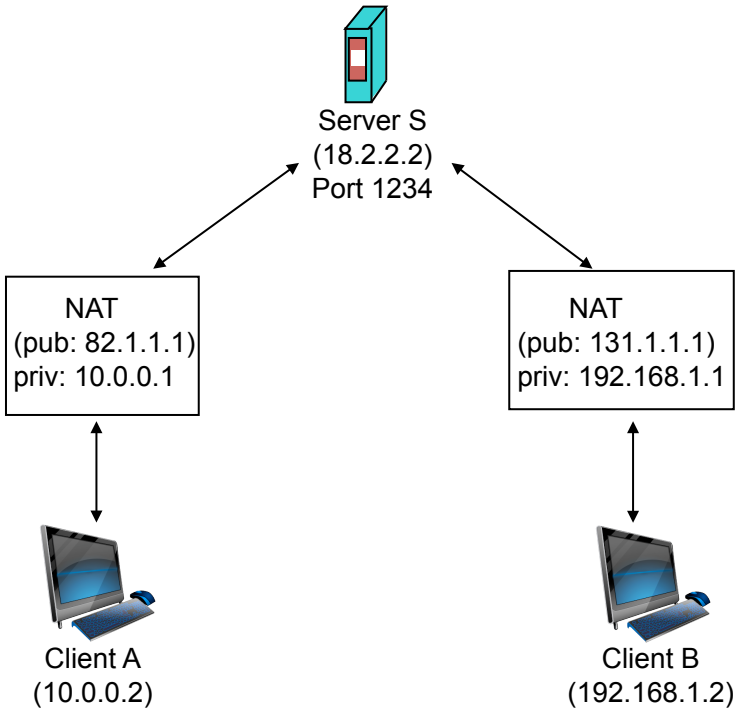

a) How do the NAT mapping tables look like after both hosts have established a connection to the server S (to port 1234)? You can use your own allocation algorithm for allocating source ports.

b) Host A and B try to establish a direct UDP connection between each other using Hole Punching. What has to be exchanged via the server S (rendezvous point) before the actual hole punching packets can be sent?

c) What is the IP 5-tuple of the packet that creates the "hole" in the NAT? And what is the IP 5-tuple of the packet that uses the "hole" to establish the actual connection?

d) If we assume a port address restricted filtering policy, who is allowed to use the created "hole"?

e) What changes if we assume a full cone NAT?