Computer Networks and Distributed Systems

Questionnaire no. 9

Context

- Lecture of 25/May: IP fragmentation and IP forwarding (Longest Prefix Matching)
- Exercise resolution session (25/May, in Lab B6): IP addressing with CIDR/VLSM and IP block partitioning
- Based on the following lecture presentation:
  - (Slides 11-19) http://paloalto.unileon.es/cn/ch3-part2-2015.pdf

1. Why is IP fragmentation necessary in IPv4?

2. Devise a simple example that illustrates IP fragmentation

An IP router receives an IP packet P whose Destination IP indicates it should be forwarded onto Ethernet interface vni3. P's header is 20 Bytes and its payload is 280 bytes and the MTU (Maximum Transfer Size) of interface vni3 is 104 Bytes, therefore, P does not fit vni3's MTU, then the router has to fragment P into several IP fragments. Calculate the fragmentation scheme and the parameters of each of the resulting fragments.

1. Each fragment has a 20 Byte IP header, therefore, the maximum payload size of each fragment is:
   Max size of fragment's payload = MTU – 20 = 104 – 20 = 84 Bytes

2. Each fragment transports a piece of the P's payload, the position where each fragment's payload fits into the whole P's payload is marked by the fragment's Offset field. This field, however, represents the number of an 8-byte word, not the number of a *1-byte* word as we usually assume. Therefore, we must check whether the Max size of a fragment's payload calculated above is divisible by 8:

   84 bytes mod 8 = 4

   Since the remainder is not 0, we adjust the Max size of each fragment's payload so that it is properly aligned on an 8-byte boundary:

   Aligned Max size of fragment's payload = Max size of fragment's payload – 4 = 84 – 4 = 80 bytes

   In summary: Each of the resulting fragments will have a **payload of a max size of 80 bytes**.

3. P's payload is 280 Bytes which we have to fragment into the following fragments:

   $N\ fragments = \frac{280\ Bytes}{80\ \frac{Bytes}{fragment}} = 3{,}5\ fragments$; A total of 3 full size fragments plus 1 partial sized fragment:

```
Fragment 0 = (Offset = 0; Payload size = 80 Bytes; MF flag = 1)

Fragment 1 = (Offset = 0 + 80/8; Payload size = 80 Bytes; MF flag = 1)
             (Offset = 10; Payload size = 80 Bytes; MF flag = 1)
```

```
Fragment 2 = (Offset = 10 + 80/8; Payload size = 80 Bytes; MF flag = 1)
             (Offset = 20; Payload size = 80 Bytes; MF flag = 1)

The preceding fragments contain a total number of bytes of 3·80 = 240, since the
total number of bytes contained in P's payload is 280, we have 280 – 240 bytes =
40 bytes remaining which we will include in the last fragment:

Fragment 3 = (Offset = 20 + 80/8; Payload size = 40 Bytes; MF flag = 0)
             (Offset = 30; Payload size = 40 Bytes; MF flag = 0)
```

3. Solve exercises no. 8, 9, 10 and 11 in pg. 11 in the following CN exam solutions document: http://paloalto.unileon.es/cn/CN-ExRefSol2013.pdf

4. Calculate whether IP 192.168.1.95 belongs to the IP block represented by network prefix 192.168.0.0/23

   1. The mask resulting from the CIDR prefix has a block of 23 bits 1:

      Mask = 255.255.254.0

   2. If the IP belongs to the IP block, then, the prefix resulting from (IP & Mask) should be equal to the prefix itself:

```
  192.168.  1. 95
& 255.255.254.  0
-----------------
  192.168.  0.  0
```

      Check that the network prefix given is correct:

```
  192.168.  0.  0
& 255.255.254.  0
-----------------
  192.168.  0.  0
```

      We conclude that the IP does belong to the IP block given, in other words, we might state that IP **matches the IP block with a length of 23 bits**.

5. Calculate whether IP 192.168.1.95 belongs to the IP block represented by network prefix 192.168.0.0/24

   1. The mask resulting from the CIDR prefix has a block of 24 bits 1:

      Mask = 255.255.255.0

   3. If the IP belongs to the IP block, then, the prefix resulting from (IP & Mask) should be equal to the prefix itself:

```
  192.168.  1. 95
& 255.255.255.  0
-----------------
  192.168.  1.  0
```

      Check that the network prefix given is correct:

```
  192.168.  0.  0
& 255.255.255.  0
```

```
------------------
  192.168.  0.  0
```

We conclude that the IP **does not** belong to the IP block given, in other words, the IP **didn't match** the prefix.