Politecnico di Torino

Exercises on IP Addressing

Fulvio Risso

December 7, 2017

License

Acknowledgments

The author would like to thank all the persons that contributed to those exercises. Particularly, special thanks go to Flavio Marinone, Guido Marchetto, Santo Vario, Dario Orfeo and Jon Brenas.

Contents

Part I.

Methodology

1. Classless addressing plans

The overall objective in this set of exercises consists in acquiring the capability to handle a classless IP addressing plan. To help students to overcome this problem, we present here a methodology that can be used to solve these exercises.

First, we define *Logical IP Network* (LIN)[1] the set of hosts that belong to the same IP network. Usually, this corresponds to a single data-link network (e.g., a LAN). Although in some cases we may have hosts belonging to the same IP network that span across multiple data-link networks, we do not consider this case, as it would need some additional configuration (e.g., proxy ARP). Vice versa, we will present the case in which multiple IP networks (multiple LIN) will be configured on the same data link network. However, in line of principle, the most common case is that a data-link network (e.g., a single Ethernet domain) maps to a single LIN. In the following, we will use the terms IP network and LIN as synonyms.

Given *(a)* the topology of the network that is the target of the addressing plan, *(b)* the number of hosts that are present in each link-layer network, and *(c)* the addressing space assigned for handling addressing, the steps that bring the student to the IP addressing plan can be summarized as follows:

1. identification of the list of the IP networks that are present in the given topology
2. determination of the number of IP addresses required in each IP network, followed by the number of addresses that need to be allocated for each LIN (taking into account that an IP network can not have arbitrary size)
3. verification that the address range available for the addressing plan is sufficient, or the determination of the address range required
4. assignment of the network address to each network
5. assignment of the address to the hosts (and to the routers) in each network

To show how this process works, we will use the sample network shown figure below, that should be configured using the address range 10.0.0.0/24.

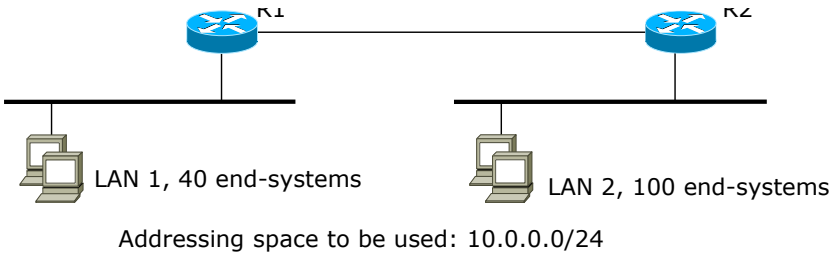

[1] Usually, the literature refers to this concept with the name *Logical IP Subnet* (LIS). However we prefer the term *Logical IP Network* as in modern IP the concept of *subnetwork* is no longer present.

1.0.1. Identification of the existing IP networks

The list of the IP networks in the above network includes the two local area networks (LAN1, LAN2) and the point-to-point connection between the two routers (please do not forget that two IP routers are always connected through an IP network).

1.0.2. Number of allocated/needed addresses

Each IP network needs a number of addresses equal to the number of end systems (thus 40 for LAN1 and 100 for LAN2), plus those needed for the correct behavior of IP, i.e., the two reserved addresses named *this net* (or *network*) and *directed broadcast*. Those reserved addresses correspond respectively to the first and the last address of the address space that will be assigned to this LIN. Further more, each LAN includes also an interface of a router, which brings the number of needed IP addresses to 43, 103 and 4, respectively for the networks LIN1, LIN2 and LIN3.

With respect to the addresses that we need to allocate to each network, each IP address can be split in a network part and an host part, whose size is given by the netmask (or the prefix length). As a consequence, an IP network cannot have an arbitrary size but must have a value equal to 2^n, e.g., 2, 4, 8, etc. The minimum size of an IP network will thus be equal to the number of needed addresses, computed at the previous step, rounded to the value of 2^n immediately equal or greater. As a consequence, we need 64 addresses for LAN1, 128 for LAN2 and 4 for the point-to-point link.

In the real life, the number of allocated addresses is also chosen by considering the foreseen future expansions of the various networks. For example, it would not be so clever to allocate 16 addresses to a network that already needs 15 address, because future expansions would be problematic as no other addresses will be available for the future hosts. On the other hand, however, in our example we allocate 4 addresses for the point-to-point link: this does not cause any problem because the number of hosts on this type of link would never grow because of the point-to-point nature of the connection.

It is worthy nothing that this steps gives immediately the netmask of each LIN. In fact, given the number of needed IP addresses, we can determine the size of the address block that has to be allocated to each LIN in term of netmask/prefix length. Particularly, prefix lengths /26, /25, /30 (respectively for the network LIN1, LIN2 and LIN3) will be used in the example, corresponding to the netmasks 255.255.255.192, 255.255.255.252 and 255.255.255.128.

The number of needed/allocated IP addresses is shown in the figure below.

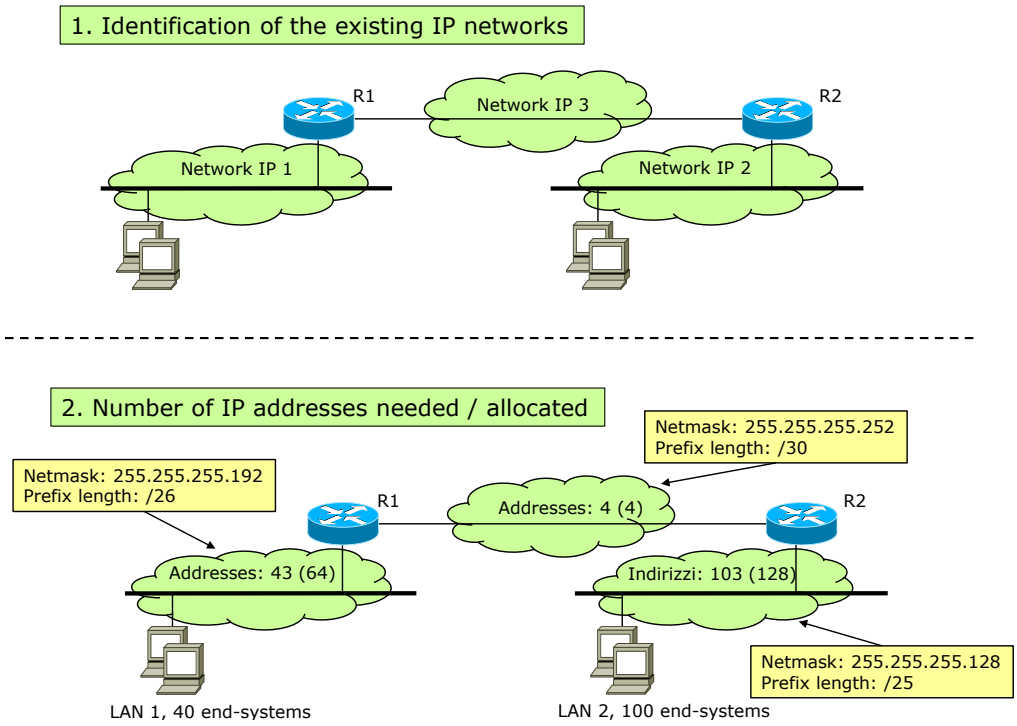

1.0.3. Validity of the address block

This points consists in determining of the number of IP addresses needed to handle the entire topology of the network.

In the best case we need just to sum all the addresses allocated in the previous point (in this case, $64 + 128 + 4 = 196$ addresses) and ask to the network administrator (or, in general, the responsible of the IP address allocation in the organization) an address range whose size is at least equal to the number of IP addresses we need.

However, in most cases the process is done in the opposite way. Giving a predetermined address range (in our case, the address range `10.0.0.0/24`), the network administrator has to verify if this is enough for its network; if not, it has to take some actions in order to reduce the consumption of IP addresses, e.g., by trying to reduce the allocated (but not used) addresses.

This second case is banal if the dimension of the address range allocated initially is greater or equal to the address range that we need in our network. In this case, one can directly go to the next step.

Vice-versa, in the case in which the assigned address range is not big enough to manage the entire network, we need to partitioning the existing IP networks in order to save addresses. This can be done by observing that some networks could have a (large) number of allocated addresses, but only a portion of those are used.

For example, the LAN1 (40 hosts) needs 43 addresses to handle its hosts but we need to allocate 64 because of the "predefined" size of the IP network. It follows that 21 addresses are currently unused. In this case, we could configure this network of 40 hosts as a set of two IP networks: the first including 32 addresses (29 hosts plus a router) followed by a second with 16 addresses (the remaining 11 hosts and one additional address for the default gateway, plus 4 additional unused addresses). This would result in the allocation of 48 addresses, with a saving of 16 addresses (48 allocated addresses compared to 64).

It is worthy noticing that in this case the number of addresses actually used increase compared with the original solution. In fact, while in the first case 43 addresses would have been used (40 hosts + router + network + broadcast), the second solution needs 46 addresses. In fact, in addition to the 40 hosts that need an IP address, each LIN needs one additional addresses for the router, one reserved for the network address and one for the broadcast, bringing the total number of used address to 46[2].

If we have to partition an IP network into multiple LIN, we need to return to the point *(1)* and restart our assignment process. In fact, we changed the number of LIN and therefore we need to compute again the number of addresses needed/allocated for each network and verify again that the requested addresses do not exceed the total number of available addresses (in our case, 256 addresses from the address range `10.0.0.0/24`).

1.0.4. Network address

This point is perhaps the most difficult (at least for a beginner), as the address range allocated to each LIN **must be in some precise positions** and **they can not be superposed**. For instance, if we allocate an address range of 32 addresses within the `10.0.0.0/24` block, we cannot uses the addresses `10.0.0.10 - 10.0.0.41`, as they do not belong to the same IP network.

Position of the address spaces

Being given an hypothetical block 10.0.0.0/24 and having the necessity to allocate a network of 128 addresses (i.e., a /25 network), the resulting network may only extend from address 10.0.0.0 to address 10.0.0.127, or from 10.0.0.128 to 10.0.0.255. It will not be possible, for example, to allocate a network /25 beginning with the address 10.0.0.100 and ending with the address 10.0.0.227, as clearly shown on the following table (next page) that reports the list of valid address range for networks ≤ 256 addresses; for networks of larger dimensions, it is possible to extend the table in order to be able to operate on address ranges with a larger number of addresses.

Th reason can be found in the way IP defines the splitting network/host, i.e., in the way the 32 bit addresses are partitioned. All the IP addresses belonging to the same address space must have the same network prefix and this forces the allocation of our address blocks in well-defined positions[3].

No overlapping of the address spaces

The address spaces assigned to the different LIN must not overlap in any way. For example, as the network LAN2 uses 103 addresses for serving 100 hosts, 25 addresses are currently unused. These addresses can not be reallocated to any other network because what it matters is not *how many addresses are actually used* but *how many addresses are allocated* for the whole network. For example, it will not be possible to assign a part of these addresses (for example, the block 10.0.0.104/30) to the point-to-point network because, even though the IP addresses are not duplicated (and thus one of the basic rules of IP, that requires that the addresses are unique, is respected), these addresses have been associated to the network LAN2 and there they should stay, even if not assigned to any host.

[2]Remember that each LIN connected to other networks needs a router whose IP address has to stay in the same address space of the served hosts.

[3]For example, assuming an address space of 4 bits (from 0000 to 1111), where the first two identify the network, a valid address range for a network will be 0000-0011 (whose network prefix will be 00), but not 0001-0100, even though both spaces include 4 ip addresses (0001, 0010, 0011, 0100).

When such an error is made (e.g., the address range 10.0.0.0/25 assigned to LAN2 and 10.0.0.104/30 assigned to the point-to-point network) we will have some routing problem, as routers would not be able to forward the packets direct to those host in the correct way. For example, any host in LAN2 knows that all addresses between .0 and .127 are directly reachable through the LAN, by sending an Ethernet frame directly to the destination, without delivering the packet to the router. That implies that addresses 10.0.0.104/30, that still belong to the 10.0.0.0/25 address range but have been assigned to hosts outside that link layer network, appear not to be reachable because they are not physically present on that LAN. Finally, note that a host cannot know how many addresses are actually assigned to other devices on its LIN: its knowledge is limited to its network address and to the netmask and thus it would consider all addresses in that range as directly reachable.

Using the table reported in next page, it possible to "color" the address range already used: this way, it would be possible to see, from the picture, which ranges are already used, i.e., the ranges that cannot be used by other networks.

The solution of points (3) and (4) is reported in the figure below.

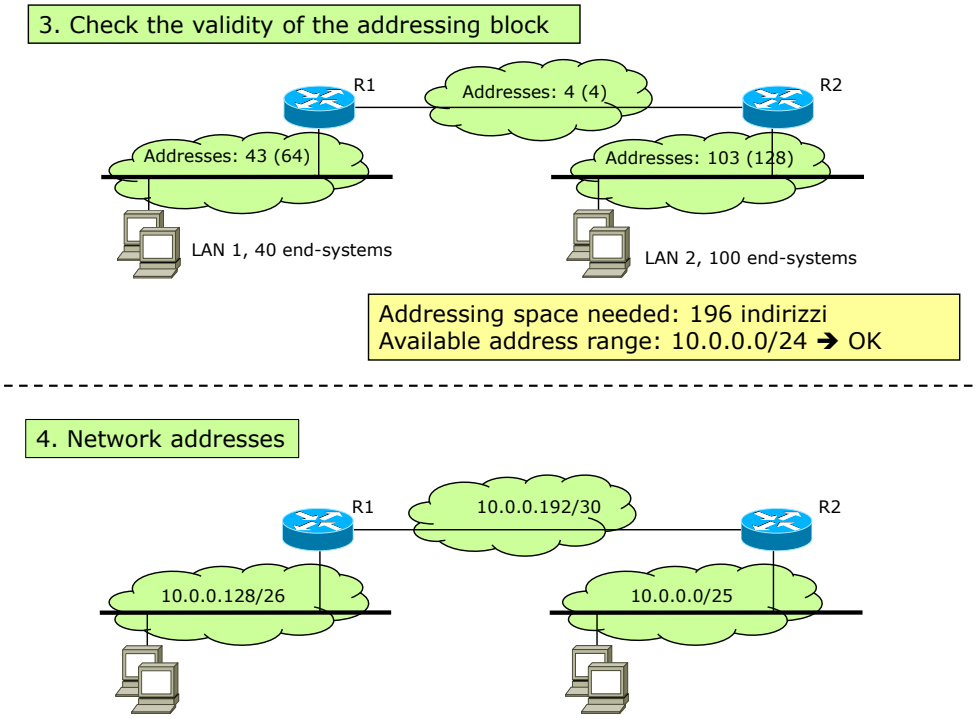

Contiguous address space

We suggest to complete the assignment of the address spaces to the LIN by starting from the largest IP network. In this way, the addresses assigned to all the LIS will result in a set of contiguous IP address ranges. This is due to the fact that the last address of a block with prefix length */N* is always followed by an address space with prefix length */(N+1)*, which can be clearly seen in the address range table. The assignment of the addresses with a diverse order often leads to the creation of "holes" in the allocated address space. Apart from looking a little chaotic, the problem is that this "random" assignment could lead to a waste of addresses, and in some cases to the impossibility to manage the network.

For example, let us assume that in our network we has assigned the addresses spaces as follows:

- LAN 1: 10.0.0.0/26 (addresses at the beginning of the block /24)
- point-to-point network: 10.0.0.252/30 (addresses at the end of the block /24)

In this case we would be in the impossibility to handle the LAN2, because all the possible address spaces (10.0.0.0/25 and 10.0.0.128/25) would be overlapped with another IP network.

This is the reason why we always suggest to assign address spaces from the largest to the smallest.

1.0.5. Hosts and routers address

The assignment of addresses to the hosts is an easy job once the previous steps have been completed correctly. It is possible to assign to the hosts (and the routers) any address belonging to the address range assigned to its LIS, excluding the first (*this net*) and the last (*directed broadcast*). In practice, network administrators tend to assign to router either the first or the last available address in the address range. This unwritten rule does not come from any theoretical prescription, but it represents an easy way to remember the address of the router. In fact, the addresses of the routers are often needed by the network administrator when connectivity problems occur in the network and he has to start some connectivity tests (e.g., *ping* or *traceroute*).

The solution of this point is reported in the next figure.

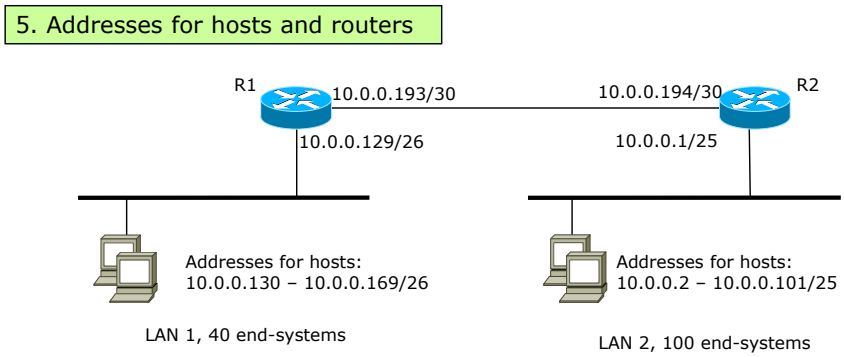

1.0.6. Some hints for calculating IP addresses

We would like to end this section by presenting some practical rules that are widely used when it comes to play with IP addresses, without having to deal with (human-not-so-friendly) binary computations.

- **Deriving the netmask**: Given a network of N elements (where $N = 2^M, M \leq 8$), the last decimal number of the netmask is equal to 256 - N. For example, in a network of 8 IP addresses the netmask will be 255.255.255.248 (265-8 = 248).
- **Network address**: A network address of an address range of N elements (where $N = 2^M, M \leq 8$) will always be 0 or a multiple of N. For example, in a network of 8 IP addresses the valid addresses will be x.y.z.0, x.y.z.8, x.y.z.16, etc.

Number of addresses available	256 addrs.	128 addrs.	64 addrs.	32 addrs.	16 addrs.	8 addrs.	4 addrs.
# if '1' bits in the last byte of the netmask	0 bits (/24)	1 bits (/25)	2 bits (/26)	3 bits (/27)	4 bits (/28)	5 bits (/29)	6 bits (/30)
Last byte of the netmask (dec/hex)	0 (0x00)	128 (0x80)	192 (0xC0)	224 (0xE0)	240 (0xF0)	248 (0xF8)	252 (0xFC)
Last byte of the netmask (bin)	00000000	10000000	11000000	11100000	11110000	11111000	11111100
Ranges of valid addresses	.0 - .255	.0 - .127	.0 - .63	.0 - .31	.0 - .15	.0 - .7	.0 - .3
							.4 - .7
						.8 - .15	.8 - .11
							.12 - .15
					.16 - .31	.16 - .23	.16 - .19
							.20 - .23
						.24 - .31	.24 - .27
							.28 - .31
				.32 - .63	.32 - .47	.32 - .39	.32 - .35
							.36 - .39
						.40 - .47	.40 - .43
							.44 - .47
					.48 - .63	.48 - .55	.48 - .51
							.52 - .55
						.56 - .63	.56 - .59
							.60 - .63
			.64 - .127	.64 - .95	.64 - .79	.64 - .71	.64 - .67
							.68 - .71
						.72 - .79	.72 - .75
							.76 - .79
					.80 - .95	.80 - .87	.80 - .83
							.84 - .87
						.88 - .95	.88 - .91
							.92 - .95
				.96 - .127	.96 - .111	.96 - .103	.96 - .99
							.100 - .103
						.104 - .111	.104 - .107
							.108 - .111
					.112 - .127	.112 - .119	.112 - .115
							.116 - .119
						.120 - .127	.120 - .123
							.124 - .127
		.128 - .255	.128 - .191	.128 - .159	.128 - .143	.128 - .135	.128 - .131
							.132 - .135
						.136 - .143	.136 - .139
							.140 - .143
					.144 - .159	.144 - .151	.144 - .147
							.148 - .151
						.152 - .159	.152 - .155
							.156 - .159
				.160 - .191	.160 - .175	.160 - .167	.160 - .163
							.164 - .167
						.168 - .175	.168 - .171
							.172 - .175
					.176 - .191	.176 - .183	.176 - .179
							.180 - .183
						.184 - .191	.184 - .187
							.188 - .191
			.192 - .255	.192 - .223	.192 - .207	.192 - .199	.192 - .195
							.196 - .199
						.200 - .207	.200 - .203
							.204 - .207
					.208 - .223	.208 - .215	.208 - .211
							.212 - .215
						.216 - .223	.216 - .219
							.220 - .223
				.224 - .255	.224 - .239	.224 - .231	.224 - .227
							.228 - .231
						.232 - .239	.232 - .235
							.236 - .239
					.240 - .255	.240 - .247	.240 - .239
							.244 - .247
						.248 - .255	.248 - .251
							.252 - .255

Valid address with prefix length between /24 and /30.

Notice that the above "practical" rules can easily be extended to networks that include more than 256 addresses, considering that each decimal number of the IP address and of the netmask is actually a group of 8 bits and thus can have 256 possible values. For example, a network of 512 elements is composed of *2* blocks of 256 elements, hence the netmask will be a number whose third decimal digit is 256-*2* (for instance, 255.255.254.0). With respect to the valid network addresses, they will be all those that have 0 in the last decimal digit and 0 or a multiple of *2* in the third one (ex. x.y.*0*.0, x.y.*2*.0, x.y.*4*.0, etc.).

1.1. Acknowledgments

The author would like to thank all the persons that contributed to those exercises. Particularly, special thanks go to Flavio Marinone, Guido Marchetto, Santo Vario, Dario Orfeo and Jon Brenas.

Part II.

Exercises

2. Classful addressing

2.1. Exercise 1

Assuming a classful addressing, determine if the following addresses are *network* or *host* addresses. Furthermore, determine also the class they belong to.

Address	Is it a network address?	Class (A/B/C)
130.192.0.0		
192.168.0.0		
80.45.0.0		
112.0.0.0		
198.0.1.0		
134.188.1.0		
224.0.0.3		
241.0.3.1		
235.0.0.0		

3. Classful addressing plans

3.1. Exercise 2

Define a classful addressing plan for the network depicted in the figure below. The chosen address ranges should belong to the private addressing space; use the first addresses available in classes A, B or C according to the size of each logical IP network.

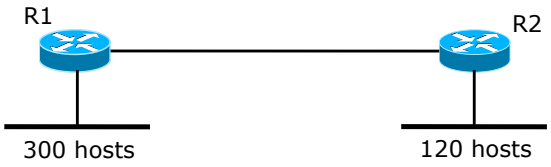

3.2. Exercise 3

Define a classful addressing plan for the network depicted in the figure below. The chosen address ranges should belong to the private addressing space; use the first addresses available in classes A, B or C according to the size of each logical IP network.

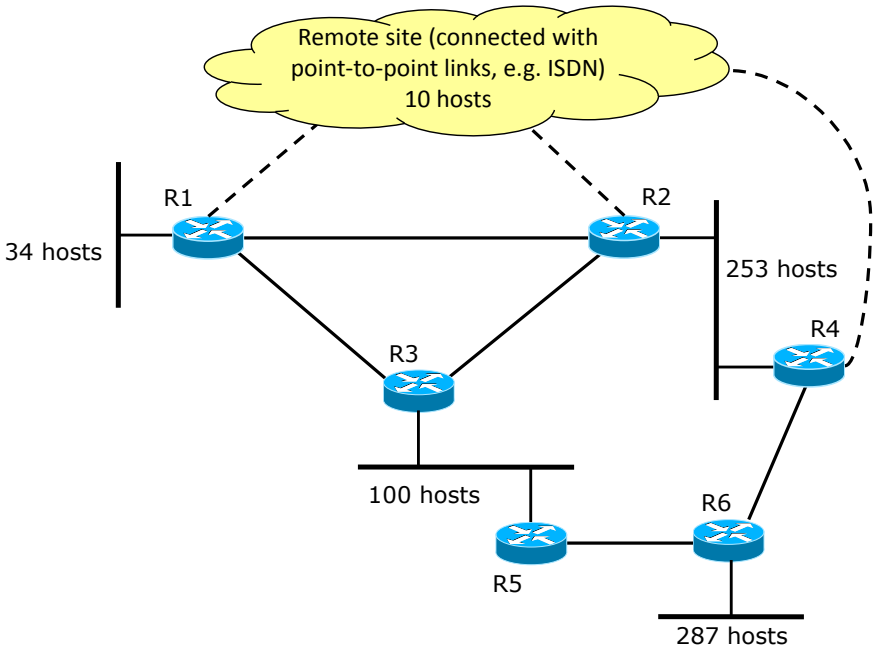

3.3. Exercise 4

Define a classful addressing plan for the network depicted in the figure below. The chosen address ranges should belong to the public addressing space; use the first addresses available in classes A, B or C according to the size of each logical IP network.

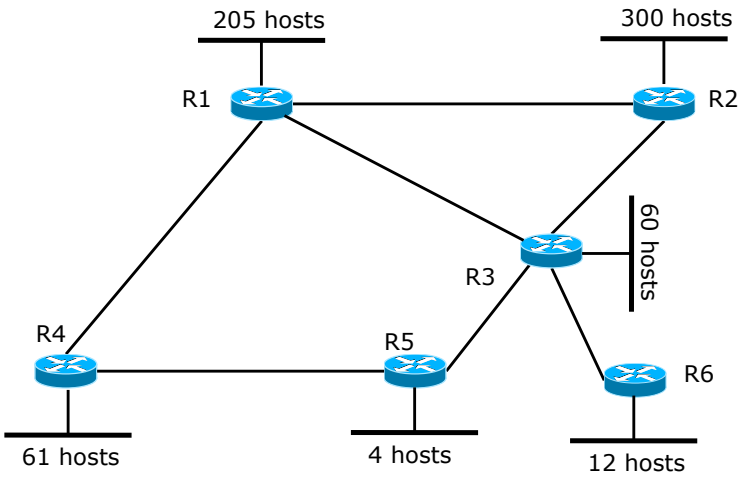

4. Classless addressing

4.1. Exercise 5

Assuming a classless addressing, define the netmask and the prefix length that should be assigned to hypothetic networks composed by the given number of hosts.

Number of hosts	Netmask	Prefix length	Number of available IP addresses
2			
27			
5			
100			
10			
300			
1010			
55			
167			
1540			

4.2. Exercise 6

Assuming a classless addressing plan, define the address ranges (in the form "network address/prefix length") that can be used to handle a set of IP networks that include the number of hosts shown in the table below. The address spaces assigned to the networks should be assigned in order, one immediately following the other, within the address range 192.168.0.0/16. Determine also the broadcast address for each network.

Number of hosts	Network address / prefix length	Broadcast address
2		
27		
5		
100		
10		

300		
1010		
55		
167		
1540		

4.3. Exercise 7

Assuming a classless addressing plan, define the IP networks that can be used to handle a LIN with the specified number of hosts (first column) within the given address range (second column). The student should specify the address range (in the form "network address/prefix length") that is the most appropriate to handle each IP network, considering that (a) no expansions (in terms of number of hosts) are expected in the future, and (2) each network is connected to the Internet and therefore a router is required. Furthermore, write also a possible address for the router and for the hosts.

Finally, in case the address range assigned to the network leads to a large waste of addresses, propose an alternative addressing based on the partitioning of the given network.

Number of hosts	Address range	Network address / prefix length	Router address	Hosts addresses
2	192.168.0.0/24			
27	192.168.0.0/24			
30	192.168.0.0/24			
126	192.168.0.0/24			
140	192.168.0.0/24			
230	192.168.0.0/24			

4.4. Exercise 8

Determine which couples of "network address/prefix length" identify a valid network.

Network address / Prefix length	Is it a valid network address?
192.168.5.0/24	
192.168.4.23/24	
192.168.2.36/30	
192.168.2.36/29	
192.168.2.32/28	
192.168.2.32/27	
192.168.3.0/23	
192.168.2.0/31	
192.168.2.0/23	
192.168.16.0/21	
192.168.12.0/21	

4.5. Exercise 9

Determine which networks can be used in private addressing, which ones can be used in public addressing, which ones are reserved for other purposes. For the last category, write why they are neither private nor public addresses.

Network address / Prefix length	Is it public?	Is it private?	What is it then?
1.1.1.0.24			
8.8.8.24/30			
10.10.10.0/24			
10.8.8.0/22			
20.2.2.0/24			
70.2.3.0/27			
127.0.0.0/30			
127.1.1.0/24			
130.192.0.0/16			
172.9.0.0//23			
172.31.0.0/24			
172.32.0.0/24			
180.12.4.0/22			
192.168.12.0/21			
192.168.16.0/24			
192.169.0.0/24			
200.200.200.0/24			
220.10.20.0/24			
224.0.0.0/24			
230.2.3.64/27			
241.0.0.0/24			
248.2.3.0/24			

5. Classless addressing plans

5.1. Exercise 10

Define a classless addressing plan for the network in the figure below, using first the address range 192.168.0.0/22, and then the address range 192.168.4.0/23. The address ranges assigned to the LIN should form a contiguous space; consider also that no expansions (in terms of the number of hosts) are required in the future.

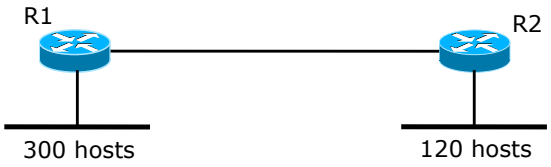

5.2. Exercise 11

Define a classless addressing plan for the network in the figure below using the address range 192.168.0.0/23. The address ranges assigned to the LIN should form a contiguous space; consider also that no expansions (in terms of the number of hosts) are required in the future, except for the network at the bottom (as shown in the figure).

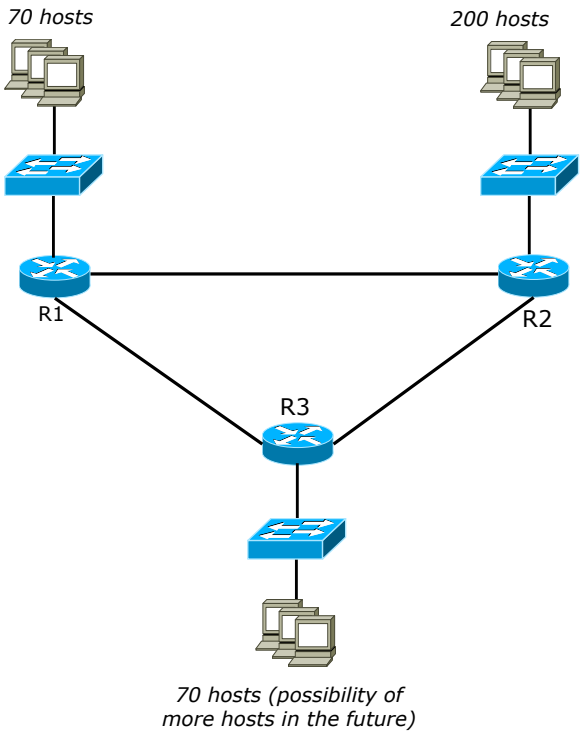

5.3. Exercise 12

Define a classless addressing plan for the network in the figure below, using first the address range 192.168.0.0/21, and then the address range 192.168.4.0/22. The address ranges assigned to the LIN should form a contiguous space; consider also that no expansions (in terms of the number of hosts) are required in the future.

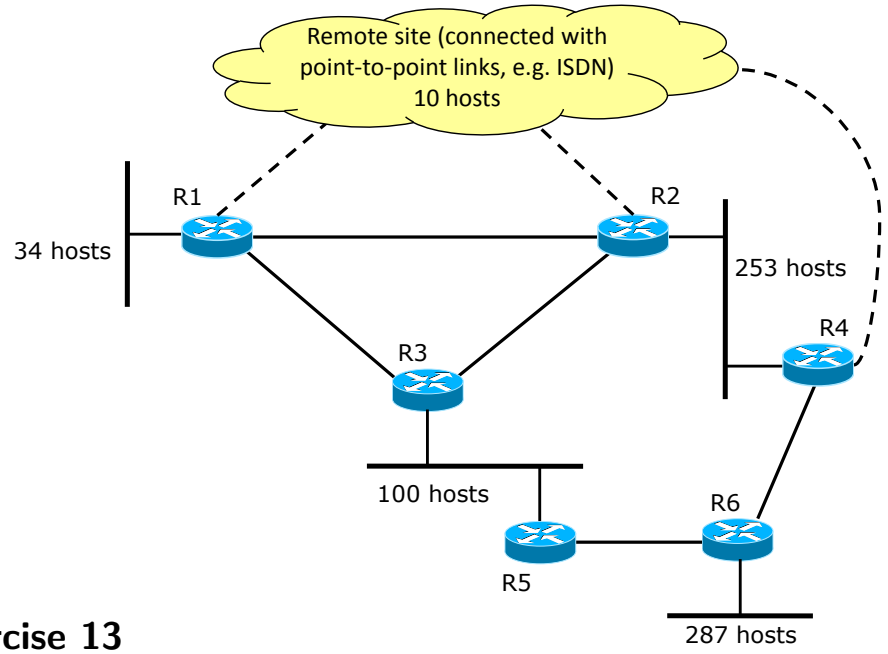

5.4. Exercise 13

Define a classless addressing plan for the network in the figure below using the address range 192.168.0.0/22. The address ranges assigned to the LIN should form a contiguous space; consider also that no expansions (in terms of the number of hosts) are required in the future, except for the network at the bottom (as shown in the figure).

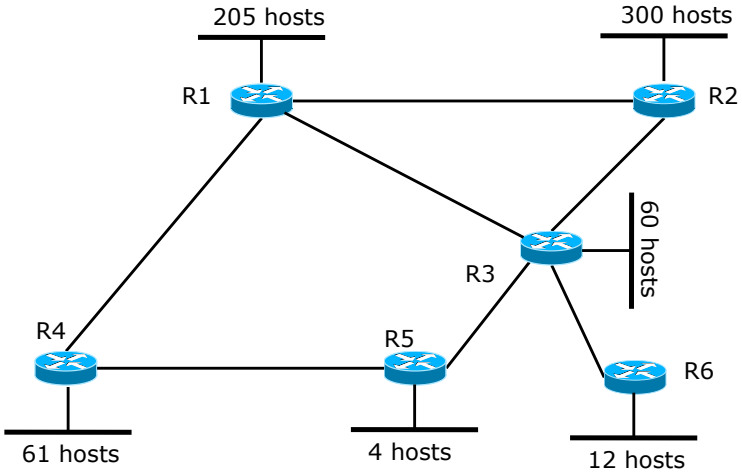

5.5. Exercise 14

An Internet provider has to build a backbone network to transport the traffic of four customers that requested an access to the Internet. Define a possible addressing plan, considering that only the networks of the final customers should be visible from the Internet.

The address spaces to be used are the ranges 192.168.0.0/21 for private addresses and 192.169.0.0/21 for public addresses.

The address ranges assigned to the LIN should form a contiguous space (within each address range); consider also that no expansions (in terms of the number of hosts) are required in the future.

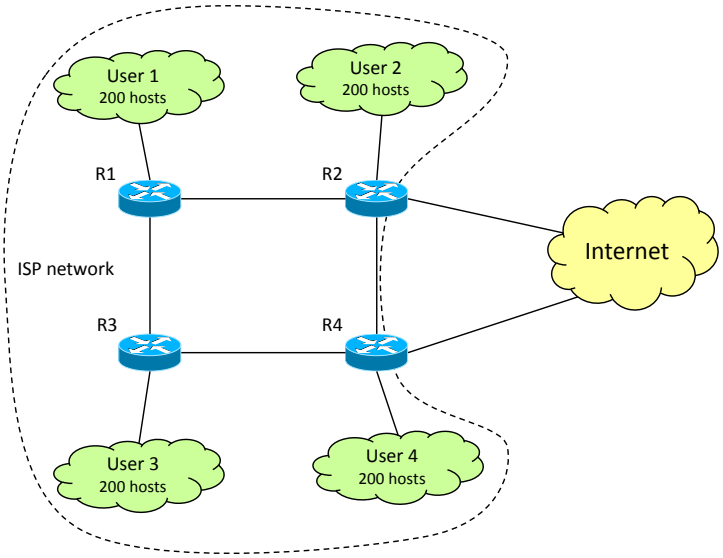

5.6. Exercise 15

An Internet provider connects a customer that requested a set of addresses to handle a network of 200 hosts, as shown in the figure; the provider assigns to that customer the address range 192.169.2.0/24.

Is the user allowed to handle its network with the "triangle-based" topology show in the figure (with the 200 hosts partitioned in two LANs of 110 and 90 clients), or is being forced to define a single LAN with all the hosts?

In case the customer is allowed to create its preferred "triangle-based" network topology, defines a possible address plan for that network.

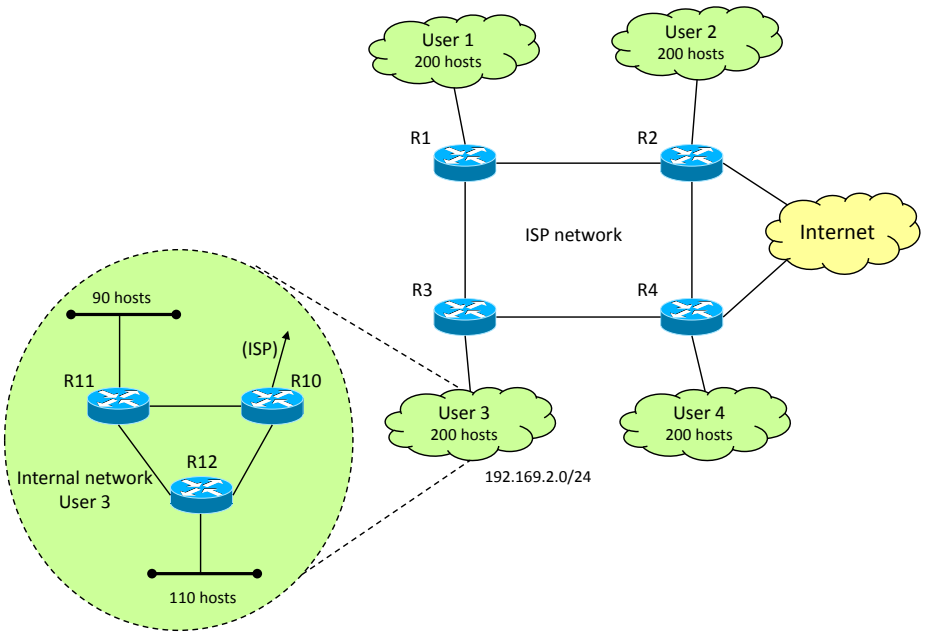

6. Troubleshooting

6.1. Exercise 16

Find the configuration error in the network depicted in the figure below and explain why such that error can prevent the hosts to work correctly.

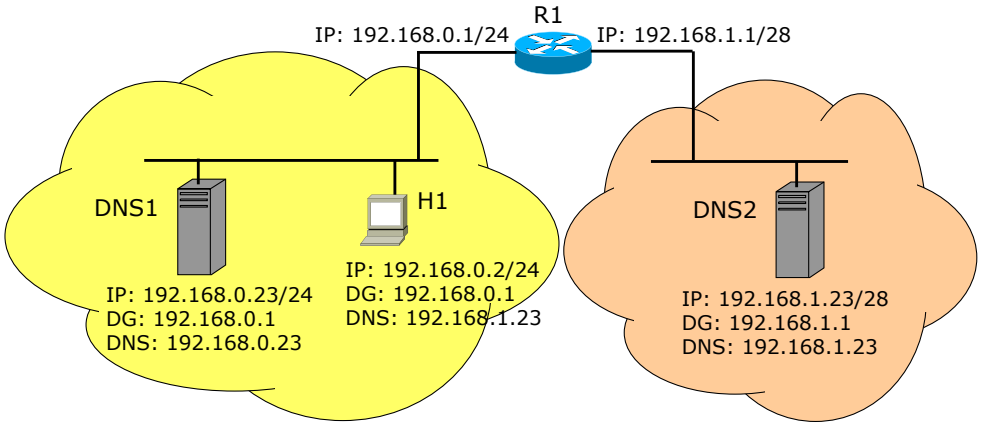

6.2. Exercise 17

Find the configuration error in the network depicted in the figure below. Assuming that host H1 would like to send an IP packet to host DNS 2, determine when the packet exchange breaks because of this error and explain the reason.

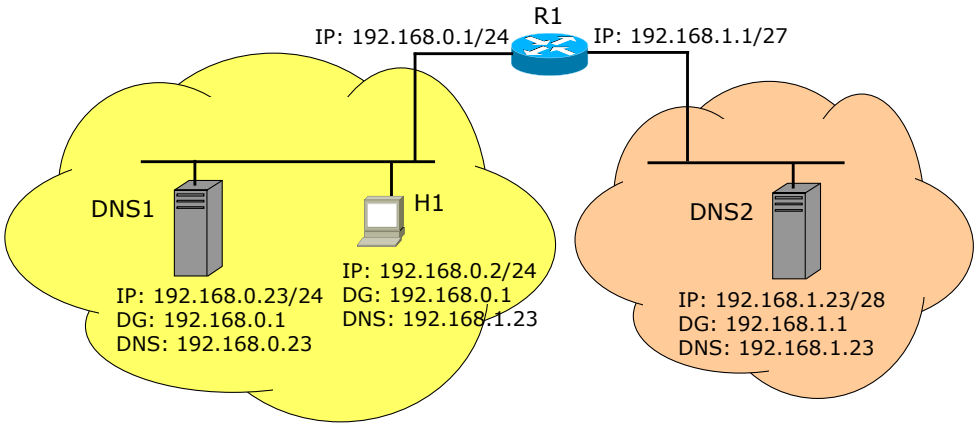

6.3. Exercise 18

Considering the network depicted in the figure below, is the command "`ping www.myserver.com`" typed on host H1 succesfull? Why?

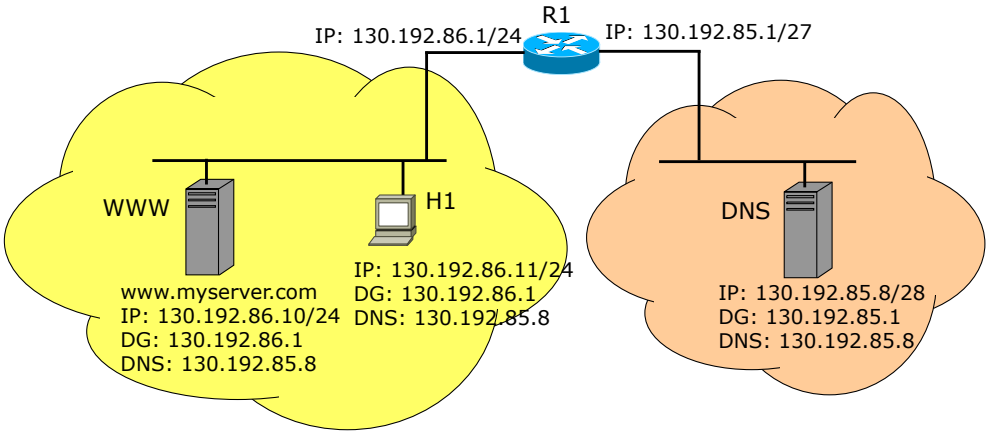

6.4. Exercise 19

Assume, as shown on network depited in the figure below, that the owner of host H1 mis-configured the IP addresses on its device; particularly the values of default gateway and IP address have been inverted. Describe the behavior of the network when such this error is present, supposing that all the hosts present in the topology would like to generate traffic.

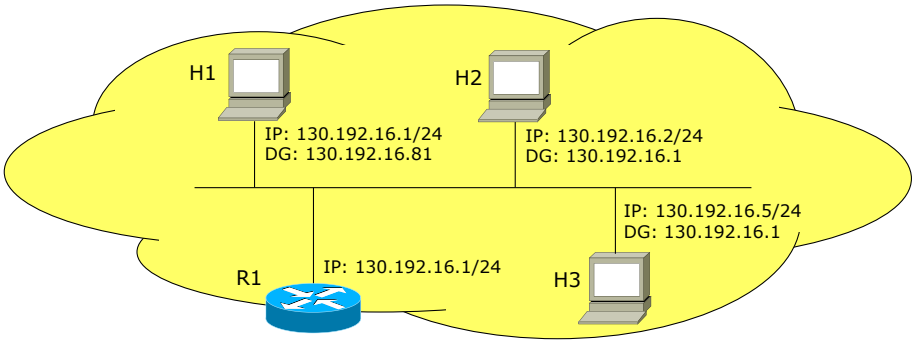

6.5. Exercise 20

Given the network depicted in the figure below, assume that the network administrator receives a phone call by the owner of host H1, which tells him that host 192.168.2.1 is unreachable, even though other hosts (e.g., 192.168.1.2) are working fine. The network administrator will then start some investigations, but its findings are that host H3 is perfectly working (e.g., 192.168.2.1 is reachable from host 192.168.1.2) and no problems are apparently there.

Given this situation, the student is asked to help the network administrator to envision some possible causes of this misbehaviour of the network.

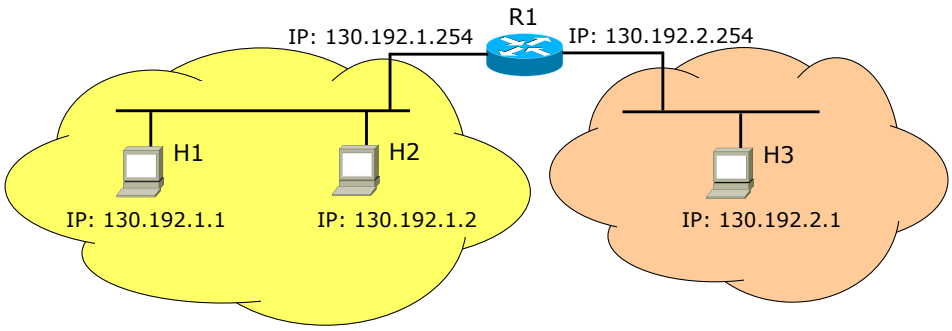

6.6. Exercise 21

Given the network depicted in the figure below, list all the configuration errors present in topology.

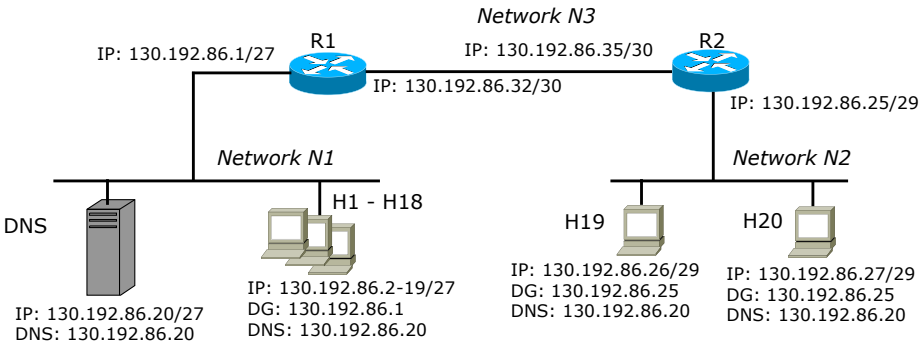

Part III.

Solutions

7. Classful addressing

7.1. Solution of exercise 1

The solution is:

Address	Is it a network address?	Class (A/B/C)
130.192.0.0	YES	B
192.168.0.0	YES	C
80.45.0.0	NO	A
112.0.0.0	YES	A
198.0.1.0	YES	C
134.188.1.0	NO	B
224.0.0.3	/	D
241.0.3.1	/	E
235.0.0.0	/	D

8. Classful addressing plans

8.1. Solution of exercise 2

The topology includes three IP networks, which can be configured using a single class B (the network with 350 hosts) plus two class C blocks. As the text of the exercise mandates the use of private addresses (the first available in each block), we must use the class B address 172.16.0.0 for the first network and the class C addresses 192.168.0.0 and 192.168.1.0 for the other two networks.

The solution is shown in the figure below.

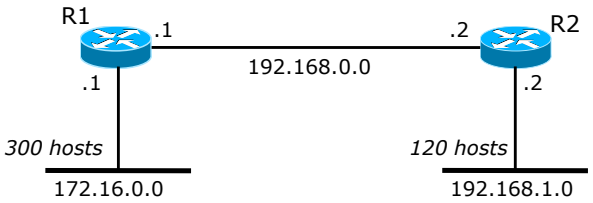

8.2. Solution of exercise 3

In this exercise, all the LIN but two can be configured with class C addreses. The first exception is the network with 253 elements, whose size exceeds the number of addresses available in a class C block because of the necessity to allocate space for the two reserved addresses (*this net* and *broadcast*) and the two routers. The second exception is the network with 287 elements. Both those networks have to be configured with class B addresses.

The solution is shown in the figure below.

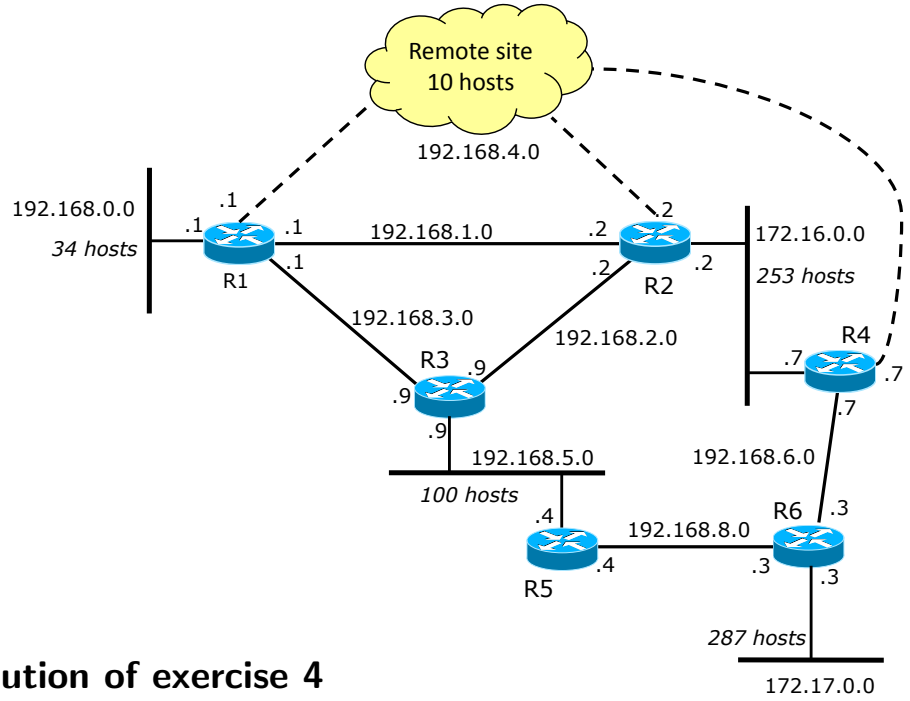

8.3. Solution of exercise 4

The topology includes only class C networks, excluding the LIN with 300 elements that has to be handled with a class B block. The solution is shown in the figure below.

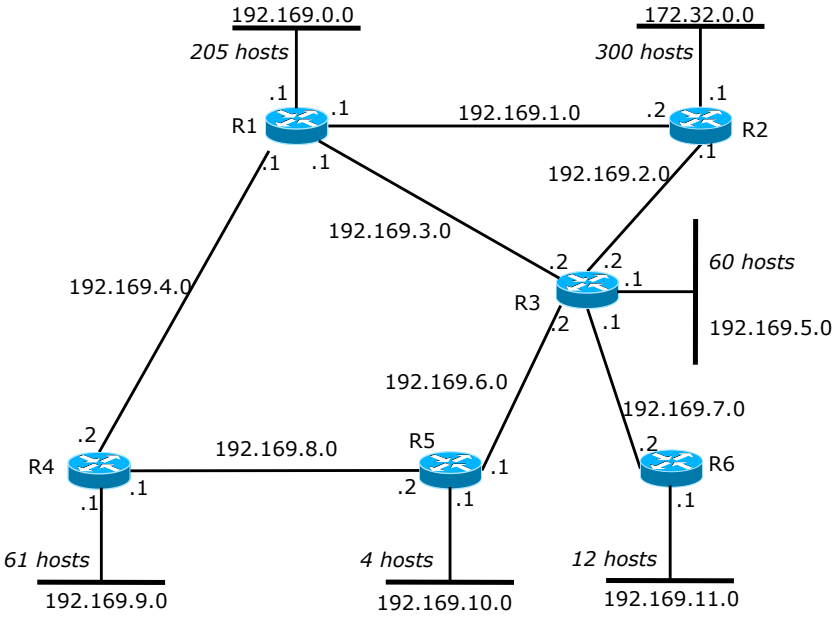

9. Classless addressing

9.1. Solution of exercise 5

The solution is:

Number of hosts	Netmask	Prefix length	Number of available IP addresses
2	255.255.255.252	/30	4 (-2)
27	255.255.255.224	/27	32 (-2)
5	255.255.255.248	/29	8 (-2)
100	255.255.255.128	/25	128 (-2)
10	255.255.255.240	/28	16 (-2)
300	255.255.254.0	/23	512 (-2)
1010	255.255.252.0	/22	1024 (-2)
55	255.255.255.192	/26	64 (-2)
167	255.255.255.0	/24	256 (-2)
1540	255.255.248.0	/21	2048 (-2)

9.2. Solution of exercise 6

The solution is:

Number of hosts	Network address / prefix length	Broadcast address
2	192.168.0.0/30	192.168.0.3
27	192.168.0.32/27	192.168.0.63
5	192.168.0.64/29	192.168.0.71
100	192.168.0.128/25	192.168.0.255
10	192.168.1.0/28	192.168.1.15
300	192.168.2.0/23	192.168.3.255
1010	192.168.4.0/22	192.168.7.255
55	192.168.8.0/26	192.168.8.63
167	192.168.9.0/24	192.168.9.255
1540	192.168.16.0/21	192.168.23.255

9.3. Solution of exercise 7

In this exercise we need to consider a number of addresses which is three more than the number of hosts of the network, because of the two reserved addresses (*this net* and *broadcast*), and the router. For instance, a network with two 2 hosts would need 5 addresses, bringing to the necessity to use an address space of 8 addresses (/29).

With respect to the assignment of the IP addresses to hosts and routers, it is worthy remember that those values are arbitrary, provided that they belong to the address range assigned to the given LIN (and that are not the reserved addresses *this net* and *broadcast*). In this solution we decided to assigne the first available address to the router and the others, following the natural order, to the hosts.

The solution is the following:

Number of hosts	Address range	Network address / prefix length	Router address	Hosts addresses
2	192.168.0.0/24	192.168.0.0/29	192.168.0.1	192.168.0.2-192.168.0.3
27	192.168.0.0/24	192.168.0.0/27	192.168.0.1	192.168.0.2-192.168.0.28
30	192.168.0.0/24	192.168.0.0/26	192.168.0.1	192.168.0.2-192.168.0.31
126	192.168.0.0/24	192.168.0.0/24	192.168.0.1	192.168.0.2-192.168.0.127
140	192.168.0.0/24	192.168.0.0/24	192.168.0.1	192.168.0.2-192.168.0.141
230	192.168.0.0/24	192.168.0.0/24	192.168.0.1	192.168.0.2-192.168.0.231

It is evident that the networks with 30, 126 and 140 hosts have a large number of unused addresses. For those networks we can propose an alternative addressing based on the splitting of the network into two distinct logical IP networks. It is worthy remember that the splitting of the hosts in two portions requires also one more address for the router, i.e., each resulting LIN must include also one address assigned the router (in addition to the two reserved addresses *this net* and *broadcast*). In other words, the router will become the so called *one-arm router*, i.e. a device with one interface and two IP addresses associated to it, the first that belongs to the address space of the first LIN (and that serves that LIN) while the second belongs to the address space of the second LIN.

A possible solution for those networks is the following:

Number of hosts	Address range	Network	Router address	Hosts addresses
30	192.168.0.0/24	192.168.0.0/27 + 192.168.0.32/30	.1 + .33	.2-.30 + .34
126	192.168.0.0/24	192.168.0.0/25 + 192.168.0.128/30	.1 + .129	.2-.126 + .130
140	192.168.0.0/24	192.168.0.0/25 + 192.168.0.128/27	.1 + .129	.2-.126 + .130-.144

9.4. Solution of exercise 8

The solution is the following:

Network address / Prefix length	Is it a valid network address?
192.168.5.0/24	YES
192.168.4.23/24	NO
192.168.2.36/30	YES
192.168.2.36/29	NO
192.168.2.32/28	YES
192.168.2.32/27	YES
192.168.3.0/23	NO
192.168.2.0/31	NO!!!
192.168.2.0/23	YES
192.168.16.0/21	YES
192.168.12.0/21	NO

Please note that the network 192.168.2.0/31 is not allowed, as this address range includes only 2 address that would be assigned to the reserved addresses *this net* and *broadcast*, leaving no room for any hosts.

9.5. Solution of exercise 9

The solution is the following:

Network address / Prefix length	Is it public?	Is it private?	What is it then?
1.1.1.0.24	YES	-	-
8.8.8.24/30	YES	-	-
10.10.10.0/24	-	YES	-
10.8.8.0/22	-	YES	-
20.2.2.0/24	YES	-	-
70.2.3.0/27	YES	-	-
127.0.0.0/30	-	-	Loopback (cannot be used)
127.1.1.0/24	-	-	Loopback (cannot be used)
130.192.0.0/16	YES	-	-
172.9.0.0/23	YES	-	-
172.31.0.0/24	-	YES	-
172.32.0.0/24	YES	-	-
180.12.4.0/22	YES	-	-
192.168.12.0/21	-	YES	-
192.168.16.0/24	-	YES	-
192.169.0.0/24	YES	-	-

200.200.200.0/24	YES	-	-
220.10.20.0/24	YES	-	-
224.0.0.0/24	-	-	Multicast (cannot be used)
230.2.3.64/27	-	-	Multicast (cannot be used)
241.0.0.0/24	-	-	Reserved (cannot be used)
248.2.3.0/24	-	-	Reserved (cannot be used)

10. Classless addressing plans

10.1. Solution of exercise 10

10.1.1. Address range /22

As per the first part of the esercise in which the address range 192.168.0.0/22 (i.e., 1024 addresses) has to be used, the soltion is rather simple. In fact, we can use 512 addresses for the LIN with 300 hosts (a /23 network), 128 addresses for the LIN with 120 hosts (a /25 network) and 4 addresses for the point-to-point link. The total number of allocated addresses is equal to 644, which is far smaller than the total number of adresses available in our /22 range. Because of this high number of available (and not allocated) addresses, it would have been possible to allocate a /24 network to the LINs with 120 hosts; however, the text of the exercise suggest that no future expansion are needed, it should be better to stay with the /25 address blocks.

The solution is shown in the figure below.

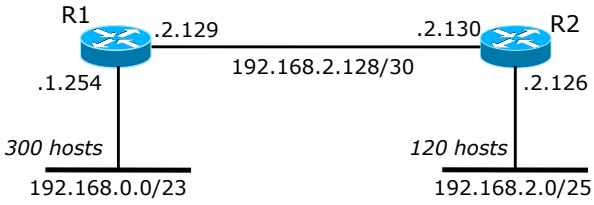

10.1.2. Address range /23

In this case the number of addresses available are not enough to handle the network with the same addressing plan defined in the previous solution. In fact, we have only 512 addresses, while the addressing plan required 644 addresses.

A possible solution is to partition LIN with 300 hosts in two distinct LINs, one keeping 253 hosts and the other with the remaining 47. The first LIN will be configured with a /24 network (253 hosts, one address for the router plus the two reserved addresses). For the second LIN, we can use a /26 network (64 addresses), in which 50 addresses will be actually used (47 hosts, one address for the router and the two reserved addresses *this net* and *broadcast*).

The number of addresses needed by this solution will be equal to 256+64 (for the network with 300 hosts), plus 128 (for the network with 120 hosts) and 4 (for the point-to-point link), which results in 452 addresses. Such this allocation is enough for us to be able to configure all our network with the assigned /23 address space.

As the exercise requires that the total occupied address space should be contiguous, we suggest to assign the address spaces to the LINs in (inverse) order of size, starting the allocation from the biggest network firsts. The LIN with 300 hosts will thus be handled by two non contiguous blocks: the

/24 block assigned to the first portion of the 300 hosts network should be followed by the /25 block assigned to the LIN with 120 hosts, further followed by the /26 block used to complete the coverage of the 300 hosts LIN.

The solution, making use of the address range 192.168.4.0/23, is shown in the figure below.

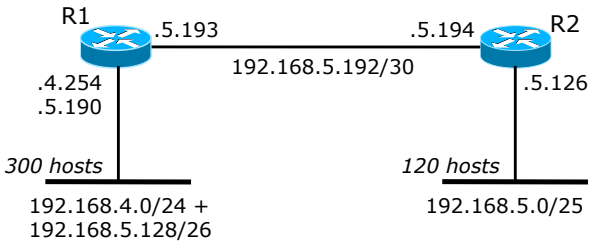

It is worthy noting that, even if it would be possible to partition the original LINs into even smaller portions (for example the network with 300 hosts could be partitioned into 253 + 29 + 13 + 5, i.e., a /26 + /27 + /28 +/29 network), this is not a good idea as it adds more complexity while a coarse partitioning was enough to make everthing working.

In fact, the partitioning of a single LIN into multiple LINs must be seen as a last-resort technique to be used when we have an insufficient number of available addresses; however, when possible, it should be avoided. One of the problems relates to the efficiency when forwarding data from one host to another: hosts in two different LINs cannot communicate directly with data-link frames (even though they are physically located on the same LAN level), hence the IP packet from one host belonging to LIN1 has to be sent to the router, which will then forward that packet to the second host belonging to LIN2. As a consequence, the packet traverses the LAN twice, with the effect of doubling the traffic on that network segment.

Another problem is that the IP broadcast address in the two LINs are different, hence services (e.g., some network discovery applications) that are based on sending IP (local) broadcast packets cannot operate across the boundary of the LIN, even if multiple LINs are in fact present on the same LAN.

10.2. Solution of exercise 11

The allocated address space is not large enough to handle the network without any further partition of at least one IP network in multiple LINs in order to save addresses. The two networks with 70 hosts are the best candidate for such an operation because they "waste" a large amount of addresses (in fact, only 73 out of 128 addresses are occupied, leaving 55 addresses unused). However, the network below the router R3 may need future expansion; therefore it should be better not to partition this network in order to leave room for future expensions. For this reason we will begin with the partition of the first network with 70 host (the one attached to router R1), verifying later if this is enough or we need to implement some other partition in order to save addresses. In our case no further partition is needed, as the number of needed IP address after the above partition of the network with 70 hosts fits within the available address space.

Notice that this exercise replaces the traditional shared (Ethernet) LANs with "switched" networks. This does not affect the IP addressing in any way, as the IP protocol works independently of the particular technology used in the data-link network. Therefore the solution with respect to the IP addressing plan will be the same, either with a shared or "switched" network.

The solution is shown in the figure below.

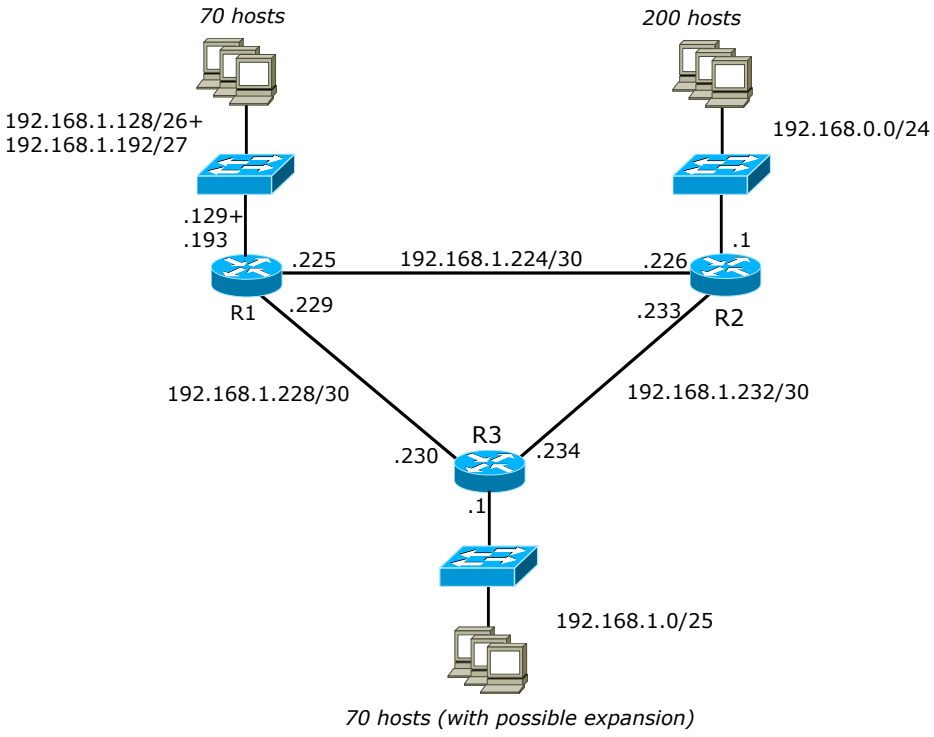

10.3. Solution of exercise 12

10.3.1. Address range /21

The solution in case the 192.168.0.0/21 address range is used is very simple, as the number of available addresses is enough to handle all the networks present in the topology. The solution is shown in the figure below.

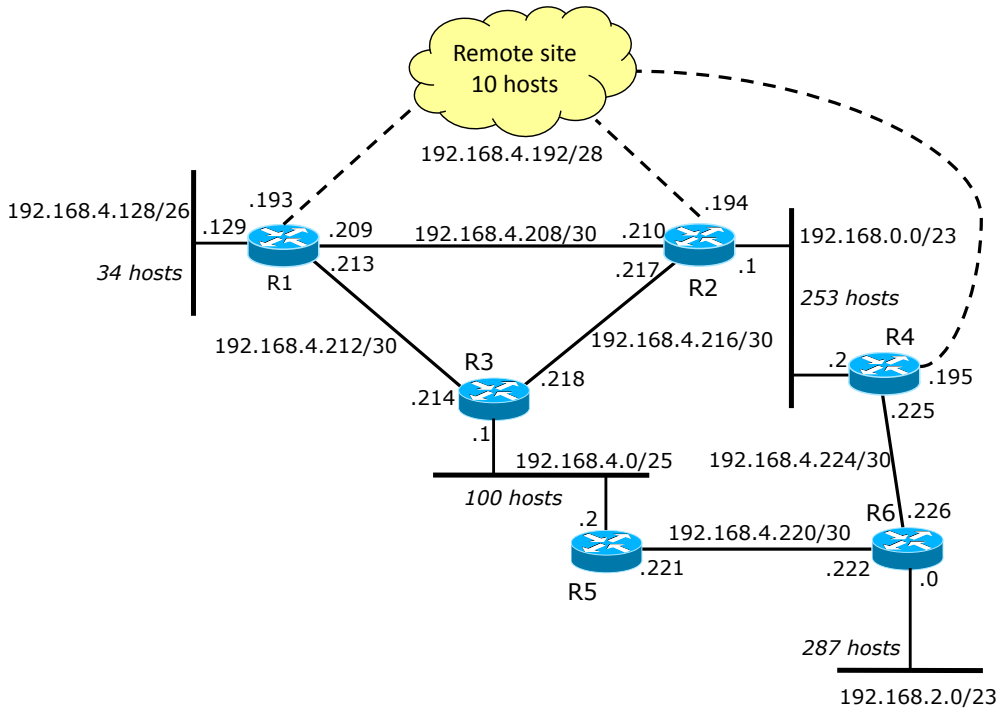

10.3.2. Address range /22

In case the 192.168.0.0/22 address range is used, we need to partition some LINs into multiple pieces in order to save addresses. In our case, we can partition the two biggest networks (287 and 253 hosts) in multiple LINs, as usual reserving 3 addresses (one for the router, plus the two reserved addresses) in each LIN because of the splitting.

The solution is shown in the figure below.

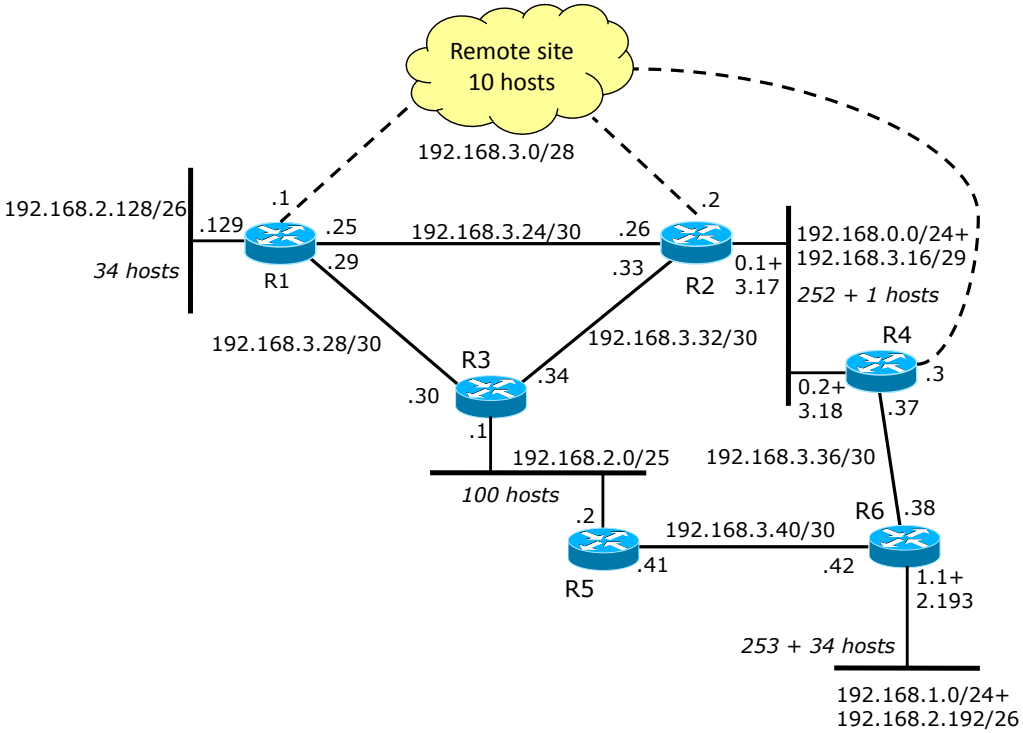

10.4. Solution of exercise 13

The solution is shown in the figure below.

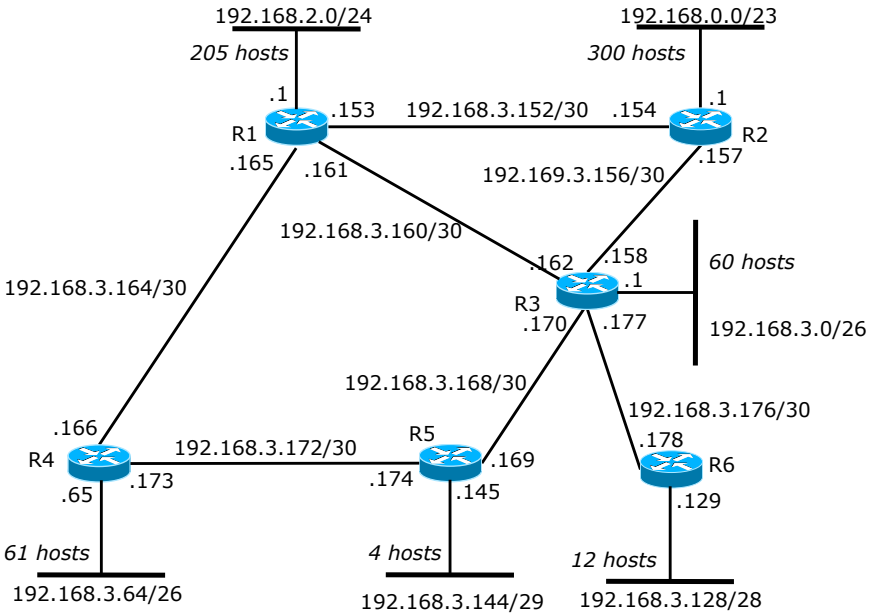

10.5. Solution of exercise 14

This exercise can be solved by noting that only customer networks need public addressing. The internal network of the provider (i.e., the point-to-point links that connect routers between themselves) does not have to be reachable from Internet and thus private addresses can be used. In this case, the routing annoucements inside the ISP network will be different from the one present outside that network (e.g., on the Internet): the ISP will obviously have to know all the existing IP networks (including private LIN) in order to be able to forward the traffic properly; however, outside its domain only networks with public IP addresses will be announced.

The fact that the addresses of the internal point-to-point links are not reachable from Internet (i.e., a user on the Internet cannot *ping* one of the addresses configured on point-to-point links) does not represent a big limitation. This technique can be used by the provider avoid attacks coming from the Internet and targeting its internal network (e.g., take control of the routers)[1]. Of course, the provider can reach (e.g., *ping*) the addresses configured on those links, as the routing on the internal network propagates all the network addresses, both public and private. Vice-versa, the two point-to-point links that connect the provider to Internet needs public addressing.

Supposing to use a /30 address space for the point-to-point links and a /24 address space for each customer, a possible solution for this exercise is reported in the figure below.

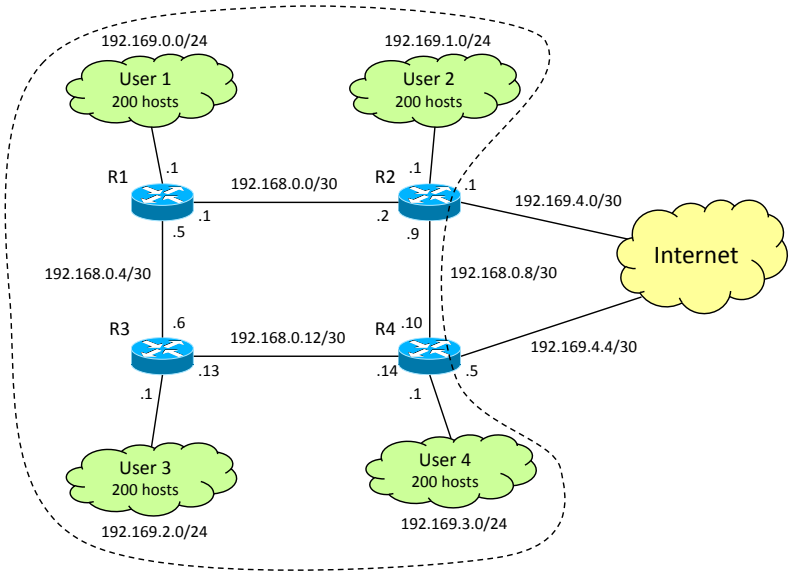

[1]For the same of precision, we should mention that all the routers in our topology *may* have also a public address on the network that connects them to the customer, which depends on the topology below (for instance, if no other routers are present in the network of the customer, this address must be public as it would belong to the LIN of the customer, which uses public addresses). In this case, this public address can be used from the external world to attack the router.

10.6. Solution of exercise 15

In this exercise, the difference between an *address space* and a *network address* becomes evident. An *address space* refers to a set of IP addresses that share the same value in the network portion of the address. Although those addresses can be assigned to a single IP network, in fact they represent just a set of contiguous and available addresses. A *network address* is an address space that is assigned to a logical IP network. As each address space can be further partitioned in multiple (smaller) address spaces, the most general case results in a large address space that is further partitioned in several address spaces, which are then assigned to LIN and used as network addresses.

Returning back to this exercise, the customer can use the 192.169.2.0/24 address space according to its needs. For instance, it can either assign the entire space to a single IP network directly connected to router R3, or partition the space in order to handle a more sophisticated typology inside its own domain. In fact, the 192.169.2.0/24 address range represents the set of addresses that the customer obtained from the provider to handle its own internal network, but this does not specify how those addresses can be used[2].

A minor difference that appears when the customer network includes additional routers is related to the IP address configured on the bottom interface of router R3. In fact, this interface must be configured with the netmask related to the LIN assigned to this link. For instance, the link that connects R3 to R10 in our topology is a point-to-point link and it is handled with a /30 LIN. In this case the bottom interface of router R3 will have to be configured with that netmask rather than the /24 that should be used in case the customer network is just a flat LAN.

The internal addressing space will be managed as shown in the figure below. Notice the necessity of partitioning the LAN connected to R11 into two IP networks (with respectively 61 and 29 hosts) in order to be able to handle the whole infrastructure with the /24 address range assigned by the provider.

The solution of this exercise is reported in the figure below.

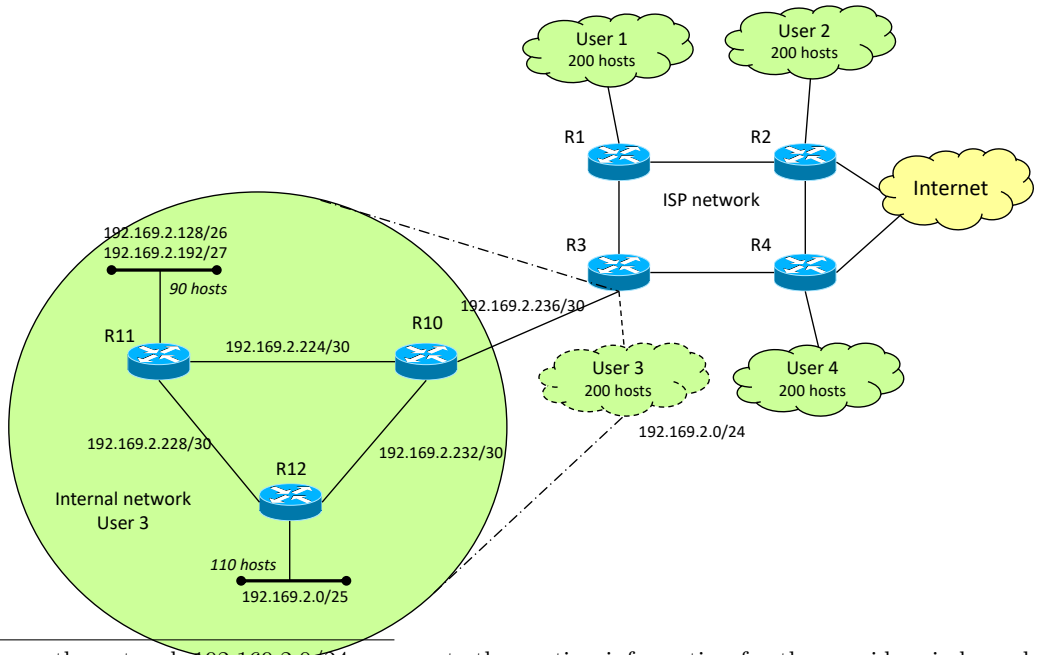

[2]For instance, the network ~~192.169.2.0~~/24 represents the routing information for the provider: independently of how this space is configured in the customer's network, the provider knows that all addresses in this range are reachable through the bottom interface of router R3. This concept will be more clear in the IP routing exercises.

11. Troubleshooting

11.1. Solution of exercise 16

The topology presents a configuration error between the router and DNS2 because the two machines do not belong to the same IP network and thus are not directly reachable. With the current topology, the router cannot send IP packets to DNS2 and vice-versa (and thus DNS2 is prevented from communicating with the outside world).

The router may reach DNS2 if both addresses belong to the range 192.168.1.2 - 192.168.1.14, or if the netmask of both devices refer to a network with a larger size. For instacne, a /27 netmask would identify the range 192.168.1.0 - 192.168.1.31, that would include both IP addresses currently configured on the machines.

Notice that having configured DNS2 like a DNS server of host H1 is not an error. The use of a DNS server that stay outside the IP network of the host is definitely common, as it would allow many LINs to make use of the same DNS server, which can be unique for the entire organization.

11.2. Solution of exercise 17

The topology includes a configuration error between the router and DNS2 because those machines have different netmasks.

With the current configuration, the router R1 can reach directly the server DNS2, as the 129.168.1.23 belongs to the same address space configured on the righmost interface of the router (192.168.1.0/27). Unfortunately, the opposite communication fails: DNS2 is configured as belonging to the network 192.168.1.16/28, a network that does not include the address 192.168.1.1 (that is associated to R1). DNS2 is thus not able to reach its default gateway and, obviously, to communicate to the external world.

11.3. Solution of exercise 18

By looking at the output of the *ping* command on the console of host H1, the command appears to execute succesfully.

In fact, the host H1, the router R1 and the WWW server can communicate together since they belong to the same LIN (associated to address range 130.192.86.0/24) and are located in the same data-link network, which guarantees e.g., that an Ethernet packet send by one of this host can be received by any other host on the same LAN.

Moreover, the router R1 can communicate to the DNS because both have an interface connected to the same LAN and have IP addresses that are "compatible" (130.192.85.8 belongs to the network 130.192.85.0/27, which is the address space configured on the router). Even the other way is possible: the DNS can communicate with the router because the address 130.192.85.1 of R1 belongs

to the IP network 130.192.85.0/28 configured on DNS2. Thus, host H1 is able to solve the name `www.mioserver.it` using the DNS, and then it can reach the WWW server with an ICMP Echo Request packet.

It is worthy noting that, even if in this particular case the network acts correctly for the indicated traffic, the network does present a configuration error: actually the rightmost interface of the router R1 and DNS2 have different netmasks. While in our case the network does not appear to experience any problem, malfunctioning ma appear for a different set of destinations; for example, an hypothetic host H2 with address 130.192.85.25/27, connected to that LAN, would be reachable from R1 but not from DNS2.

11.4. Solution of exercise 19

Given the configuration as in the exercise, host H1 appears isolated from the external world. It can reach any local destination (i.e., IP addresses belonging to its address range), but it cannot reach other destinations outside 130.192.16.0/24 as it has a wrong address for the default gateway. For instance, the address 130.192.16.81 (which represents the default gateway for host H1) should not be present at all on that network, as it should have been assigned by the network administrator to host H1.

For hosts H2 and H3 the situation is more complex. The problem is that the address of the default gateway (130.192.16.1) is present twice in the network: each ARP request for address 130.192.16.1, which is needed to discover the MAC address of the router, will receive two answers, one from R1 and the other from H1, because both of them are at the address 130.192.16.1. Unfortunately it is not possible to know a priori which answer will be used by the host that sent the ARP Request, because the choice is completely random (for instance, the host should use the MAC address written in the first ARP Reply received). In this case, two situations may occur:

- the host chooses the MAC address of the router: the Internet traffic for and to this host will flow without problems.
- the host chooses the MAC address of H1: the outgoing traffic toward the Internet will be sent to H1, which will drop those packets. In fact, hosts do not have any routing function and if the destination IP address of the received packet is not the IP address of the host, the packet is simply discarded. It is nevertheless interesting to note that the opposite direction (traffic from Internet to H2-H3) will be delivered without any problem.

Notice that this behavior may vary with time: when the ARP cache related to the address 130.192.16.1 expires in hosts H2-H3, the ARP process begins again, leading to one of the two possible outcomes shown previously. Thus an host may be able to communicate intermittently with the Internet, making the diagnosis of the failure much more problematic.

Finally, we should mention that many operating systems are able to detect a conflict related to a duplicated IP address; this is done by sending an ARP request with its own IP address as a target, before activating the IP protocol stack. If an answer is received, a duplicate address is present on the network; in this case, most operating system disable the IP protocol stack and prompt an error message to the user.

11.5. Solution of exercise 20

In this exercise, there may be a lot of reasons that can cause the described misbehaviour. We can begin by pointing out what should not be at the origin of the failure:

- DNS configuration is OK: as the failure is detected even when dealing directly with the IP addresses, the DNS cannot be the cause of our fault. This does not mean that the DNS configuration is correct, but simply that there must be something else wrong, as the error comes out before the DNS is involved in the process.
- No routing problems on the router: since host H2 sends and receives traffic from H3, and those hosts belong to different IP networks connected by the network R1, this suggests that R1 is beheving correctly.
- Configuration of the Default Gateway on host H3 is OK: the configuration of the default gateway on host H3 should be correct for the same reason shown in the previous point (host H2 sends and receives traffic from H3 which suggests that the value of the Default Gateway on H3 is correct).
- No errors in the configuration of the netmasks on H3 and on R1: the possibility to exchange IP packets between H2 and H3 seems to exclude also this problem.

At this point, we can suggest the following reasons for the fault:

- Wrong (or missing) configuration of the Default gateway on host H1: the host H1 would be able to reach only the destinations inside the network 130.192.1.0/24 (in fact, the text of the exercise confirms that H1 is able to send/receive traffic to/from H2).
- Wrong configuration of the netmask on host H1 (1): if the netmask configured on H1 defines a network whose size is smaller than 256 elements (prefix length $>$ /24), the address of the default gateway would not be reachable (it would be in a network non directly reachable at the IP layer); this would prevent H1 from reaching all the external destinations.
- Wrong configuration of the netmask on host H1 (2): if the netmask configured on H1 defines a network whose size is larger than 256 elements (prefix length $<$ /24), some remote addresses would be seen as directly reachable and the host H1 would try to reach them with direct addressing. In this case, host H1 would generate an ARP request in order to obtain the MAC address of the destination host, but it will never receive the answer because the requested host is in another LAN (broadcast messages, such as ARP Request, cannot traverse a router, hence the host on the other network will never receive the ARP Request).
- Access lists on the router: in some cases, routers are configured with some traffic filtering functionality to prevent some hosts from reaching some given destinations (or viceversa). In presence of an appropriate "deny" rule (for instance, a rule that tells the router that address 130.192.1.1 is not allowed to send packets to the network 130.192.2.0/24, or that address 130.192.1.1 is not allowed to send packets to the address 130.192.2.1, etc.), the traffic would not be forwarded by the router to the destination H3.
- Personal firewall on host H3: some personal firewalls prevent the generation of the ICMP Echo Reply when an ICMP Echo Request packet is received. The text of the exercise does not specify which application has been used to test the connectivity between the stations: if the application used was `ping`, it may be a valid reason for the lack of communication between H1 and H3. In this case, a different tool (e.g., the opening of a TCP connection through the `telnet` command line tools) may give different results.
- ICMP packet suppression on R1: in some networks, routers are configured not to forward the ICMP packets. The text of the exercise does not specify which application has been used to test the connectivity between stations: if the used application was `ping`, this may be a valid reason for the lack of communication between H1 and H3.

We can find even more esoteric reasons, but this is left to the imagination of the student (for example, the case in which two hosts with address 130.192.1.254 exist on the leftmost LAN, the case of a MAC spoofing attack on LAN1, etc.).

Finally, it is worthy noting that this exercise is much simpler than real world troubleshooting. The main difference is in the knowledge of the topology: here the student knows everything about the topology of the network, while in most of the real cases the topology is (partially) unknown, e.g., because it is under the control of different entities.

11.6. Solution of exercise 21

The errors present in this exercise are the following:

1. The DNS server does not have any configuration for the Default gateway; this prevents the host from being able to contact any host outside the range 130.192.86.0/27.
2. The addressing spaces of the networks N1 and N1 overlap, even if IP addresses associated to the hosts are not duplicated. In fact, the address range 130.192.86.24/29 (used in N2) belongs to the address range 130.192.86.0/27, used in N1.
3. The IP addresses assigned to the devices of network N3 are wrong, because they represent the addresses *this net* and *broadcast* of the address range 130.192.86.32/30. The correct addresses are 130.192.86.33 and 130.192.86.34.