

Project 1/Sprint 1 follow-up. LAN and virtual LAN. IPv4 addressing. ARP tables. IPv4 packets routing. Classful IPv4 addressing. IPv4 static routing. Cisco Packet Tracer activities. Practical exercises.

1. Project 1/Sprint 1 follow-up

1.1. Network outlets.

Once the area of each room was estimated, accordingly, the number of required network outlets for each room has been established.

Concerning network outlets for indoors wireless access-points, the maximum reach is about 30 meters, slabs, and columns have a significant impact in the signal propagation. To maximise the indoors coverage and avoid signal propagation to outdoors, access-points should be placed in floors central locations.

For a fair coverage of an area by a set of access-points, they ought to be closer than 50 meters from each other. As far as possible, access-points in adjacent floors ought to have different positions to avoid cells overlapping. Later, different Wi-Fi channels will be assigned to each access-point to further avoid the interference between neighbour cells.

Each individual outlet must be pinpointed in floor blueprints.

1.2. Cross-connects, cable pathways and cable types.

Having all outlets locations fixed, the next step it's deciding where to place each cross-connect. Several previously studied standards and guidelines, apply here.

Concerning horizontal cabling subsystems, remember no outlet can distance more than 80 meters from the horizontal cross-connect in a straight line, also cable length cannot be above 90 meters. If required consolidation points may be created.

Cable pathways outlining comes next. As far as possible, the maximum number of cables ought to share the same pathway, or at least part of it.

Now, each cable length can be calculated. A recheck on horizontal cabling lengths is prudent for the longer cables. For lengths above 90 meters, the only available option is the optical fibre, all the same, backbones, in general, should use optical fibre. Backbone redundant connections must not be forgotten.

1.3. Patch-panels and telecommunication enclosures.

Every cable entering a cross-connect is wired up to a patch panel. Being that the number of cables and cable types entering each cross-connect is now known, consistently, the number and type of required patch panels in each cross-connect can be settled.

The number of required connections at each cross-point must then be matched with patch-panel manufacturers' data, the vertical space occupied in the telecommunication enclosures is specified in U rack units. Typical 24 ports CAT6 copper patch panels take 1U and 48 ports CAT6 copper patch panels do occupy 2U on the telecommunications enclosure. Current optical fibre patch-panels have similar densities to copper patch-panels, older models are more modest: as low as four optical fibre connections for 1U.

Enforcing the previously suggested oversizing strategies, we can infer the telecommunications enclosures size by multiplying by four the amount of space required by housed patch-panels and round it up to the next commercially available size.

1.4. Inventories.

In fact, the structured cabling hardware inventory is mostly done, it's just an accounting matter to establish total numbers for network outlets, each type of patch-panels, and telecommunication enclosures. One key element yet missing from the inventory are cables themselves.

Building the cable inventory over the pathways blueprint can be a fairly significant effort, previously discussed simplifications and approximations should be used.

Patch cords are not regarded as structured cabling hardware, but they are ultimately required. Copper and optical fibre patch-cords are commercially available ranging from 0.5 meters up to 5 meters long. Inside telecommunications enclosures 0.5 meters models are most appropriate to connect patch-panels to active hardware.

1.5. Global inventory (sprint master).

It's up to the sprint master creating a global inventory for all structured cabling hardware required, this is just a matter of picking each team member's inventory and sum it all. The sprint master should also start preparing the **review.md** document for the incoming sprint review meeting.

2. LAN and virtual LAN (VLAN)

A LAN (Local Area Network) is a set of end nodes connected to the same shared layer two transmission technology. Within a LAN, nodes are able to directly transmit frames to each other with no restrictions. The LAN concept is also directly related to the **broadcast domain** concept. When a node sends a frame to the broadcast address, the frame will reach every other node within the LAN but never goes beyond the LAN limits.

Layer two communication between nodes belonging to different LANs is impossible, nodes belonging to different LANs are required to use a layer three protocol to be able to communicate. Under the point of view of layer three protocols, each LAN is a different network, routers operate at layer three and have the mission of transferring layer three packets between different LANs.

As we have already seen, more than one layer three network can be defined over a single LAN, nevertheless, **a single layer three network cannot be split over more than one LAN**.

A virtual LAN (VLAN) is a subset of a LAN we want to be operating as if it was a separate LAN, in other words, we can split a LAN into several VLANs in such a way each VLAN is equivalent to a separate LAN. Therefore, **each VLAN is a broadcast domain** and nodes belonging to different VLANs can't communicate directly using layer two technologies.

VLAN operation is based on layer two devices, notably switches, end nodes' network interfaces also support VLAN, so routers, servers and even workstations can use VLANs as well.

Port-based VLANs can be set up on switches by assigning switch ports to different VLANs, as result, the switch will then forward frames only between ports belonging to the same VLAN, including broadcast frames. Under operations point of view, such a switch becomes equivalent to several independent, and unconnected switches (one for each defined VLAN).

Let's take as an example an 8-port switch as represented in Figure 1. We can, for instance, define the following VLANs in this switch:

VLAN A: Port 1; Port 4; Port 8

VLAN B: Port 2; Port 5; Port 7

VLAN C: Port 3; Port 6




Figure 1 - Eight ports switch

The switch will be now equivalent to three unconnected switches (Figure 2, Figure 3, and Figure 4).




Figure 2 - VLAN A




Figure 3 - VLAN B




Figure 4 - VLAN C

For instance, a frame received in port 1 will only be retransmitted on ports 4 or 8, even if the destination address is FF:FF:FF:FF:FF:FF (the layer two broadcast address).

2.1. VLAN frame tagging

The major operation principle for a VLAN-aware switch is: never mix frames belonging to different VLANs. Yet, for the sake of hardware optimization, it's possible to assign several VLANs to the same switch port.

To be able to do so, frames must carry a tag identifying to which VLAN they belong, by looking to the tag of a frame incoming from a port assigned to multiple VLANs, the switch will then know to which VLAN it belongs to.

The IEEE 802.1q standard describes how to place 12-bits identifiers, called VLANID, in Ethernet frames. By using IEEE 802.1q assigning several VLANs to the same port of a switch is possible, in Cisco devices this is called **Trunk-mode**, in opposition to **Access-mode** where only one VLAN is assigned to the port.

Frames sent and received through a Trunk-mode port are not standard Ethernet frames, they carry IEEE 802.1q VLANIDs. Ports on both ends of a cable connection must be configured the same mode and using the same VLANIDs.

Trunk-mode ports can save a lot of hardware, imagine we have two switches, each with four VLANs: A, B, C and D, we usually want VLANs to be the same on both switches, so VLAN A on one switch must be connected to VLAN A on the other switch, and so on. With access mode ports only, this would require eight ports and four cable connections (one for each VLAN). By using trunk mode, the problem can be solved with a single cable connection and two trunk-mode ports. Now imagine if we were to interconnect 20 VLANs between two switches.

Switches identify VLANs by VLANID (12-bits number), however, network administrators can also assign them convenient arbitrary names for easier management. VLAN names are local to each switch, in what concerns frame transaction between switches what really matters is the VLANID. VLANIDs must be set accordingly in all switches.

2.2. Layer two infrastructures with VLANs

One advantage of VLANs is they can be remotely managed; the network administrator can remotely access each switch and therefore change VLANs assigned to each port.

If we want to achieve a flexible layer two platform capable of being adapted to any layer three networks layout requirements, we must create a **continuous layer two network**, based on interconnected VLAN-capable switches. Over this **continuous layer two network**, we can then enforce and manage whatever VLANs are required to meet the needs (Figure 5).




Figure 5 - Layer two infrastructure with VLANs

For instance, to change the network (VLAN) to which a node is connected it's just a matter of remotely accessing the switch to which the node is connected and change the VLAN assigned to the port. Without VLANs, accomplishing the same purpose would require going the appropriate cross-connect and physically change a patch cord connection at the patch panel.

Between different VLANs there's no layer two connectivity (as it happens with different LANs). Communication between different networks, being them LANs or VLANs, is provided by routers at layer three that are connected to several LANs or VLANs.

Layer 3 nodes like routers and servers can also use network interfaces in trunk-mode, for instance, a single physical connection (physical interface) of a router to the infrastructure may be used to connect it into several VLANs.

Still, it must be said, VLANs are not entirely equivalent to physically independent networks, this happens because the hardware is shared between different VLANs. For instance, a traffic overload in one VLAN will affect the shared infrastructure and will, therefore, disturb other VLANs.

3. Use the Packet Tracer tool to create the following layout

End node colours represent the VLAN each is connected to. Mind that, colours are just for the sake of this diagram's clarity and not to be set on Packet Tracer.




Figure 6 - Network layout with three switches and two VLANs

Although all nodes share the same layer two infrastructure, we aim at achieving two distinct and independent networks.

Yellow nodes network: PC0; PC1; PC5 and PC8

Red nodes network: PC2; PC3; PC4; PC6 and PC7

3.1. Configure the three switches to achieve the desired purpose (using VLANs).

- For the yellow nodes network use VLANID=5
- For the red nodes network use VLANID=6
- These end nodes are not VLAN-aware, so they don't recognise IEEE 802.1q frames, therefore, the switch ports they are connected to should be in access mode, assigned to the single desired VLAN.
- We want both VLANs to be global to all switches, thus ports used to interconnect switches should be set to trunk-mode, with both VLANs assigned to them.

3.2. Test each VLAN extent by broadcasting ICMP echo requests

To be able to use ICMP we must first set some nodes IPv4 addresses:

- Assign to PC1 the 192.168.20.1/24 IPv4 private address.
- Assign to PC4 the 192.168.30.1/24 IPv4 private address.

Now use the Add Complex PDU tool to send an ICMP echo request to the IPv4 broadcast address (255.255.255.255) every 5 seconds, from PC1 and then from PC4.

- Check that each VLAN is equivalent to an independent LAN (broadcast domain). Traffic, even when sent to the broadcast address, is never propagated from one VLAN to the other VLAN.

4. ARP – Address Resolution Protocol

When IPv4 operates over Ethernet (a layer two technology), each IPv4 packet is transported inside an Ethernet frame, we say the IPv4 packet is **encapsulated** in an Ethernet frame (Figure 7), in other words the IPv4 packet is the payload of the Ethernet frame. In turn, the IPv4 packet itself transports as payload data belonging to upper-level protocols, like for instance ICMP, UDP or TCP.

IPv4 addressing is independent of Ethernet addressing (as it's supposed to be), this presents a challenge. IPv4 32-bits node addresses mean nothing to Ethernet, Ethernet uses 48-bits node addresses (MAC addresses). The point is, when the IPv4 layer wants to send an IPv4 packet it must encapsulate the packet in an Ethernet frame, to do so successfully it must also set the correct **Destination MAC Address** for the frame, however, it only knows the **Destination IP Address**.

Somehow, the IPv4 layer is required to learn to which **Destination MAC Address** does that **Destination IP Address** corresponds (Figure 7). In other words, what is the **MAC Address** of the network node that is using that **IP Address**?




Figure 7 - IPv4 packet transported as payload of an Ethernet frame

ARP (Address Resolution Protocol) was designed to solve this issue, IPv4 can't operate without the help of ARP. We can see ARP as a function like this:

$$\text{MAC node address} = \text{ARP (IPv4 node address)}$$

ARP layer works side by side with IP and manages the so-called ARP table (**not to be confused with the MAC table of a switch**)

An ARP table (Figure 8) holds equivalences between IPv4 addresses and layer two MAC addresses. This is a dynamically managed table, while there are no communications, it is empty. Entries are created as needed, they have a short time to live and are removed if not refreshed.

When IPv4 asks ARP for a MAC address, it may be already at the table and in that case it's immediately returned.




Figure 8 – An ARP table

If the required MAC address is missing from the table, then the ARP layer uses the ARP protocol itself: an ARP request with the desired IPv4 address is sent in broadcast. Every IPv4 node has an ARP layer listening for ARP requests, they check if the requested IPv4 address is their own local IP address, if so, they reply with their own MAC address. Once the requesting node receives the ARP reply, it will add it to the ARP table.

Because ARP operates by using broadcast it will only work within a broadcast domain, in other words nodes are required to be on the same LAN or VLAN.

5. IPv4 packets routing

Two IPv4 nodes can communicate directly (without using routers) only if two preconditions are met:

A – Both are connected to the same LAN/VLAN (same broadcast domain)

B – Both nodes IPv4 addresses belong to the same IPv4 network.

We can see precondition **A** arises straight from ARP that only works within a broadcast domain.

An IPv4 sending node must somehow know if a given destination IP address is reachable directly or not. If reachable directly it just needs to use ARP and then encapsulate it inside a layer two frame, otherwise the packet must be sent to a router.

The way a sending node decides this is by matching the given destination IP address with the **local IPv4 network** if the network prefix is the same, we assume direct communication is possible. This is in fact condition B.

5.1. Routers

Routers (aka gateways) are layer three intermediate nodes; they forward layer three packets and not layer two frame like switches do. The mission of an IP router is receiving IP packets from one LAN/VLAN and retransmitting then on another LAN/VLAN.

5.2. Using routers

When an IP sending nodes checks the IP destination address of the packet doesn't belong to the local IP network, it knows a router must be used, so the **router IP address** must be known.

One IP node may be aware of several routers around it, but end nodes are usually aware of only one router they can use, this is an additional required configuration parameter usually called **default-gateway** (or **default-router**). If a node is not aware of any available router it will never be able to communicate with nodes beyond the local IP network.

So, if the destination IP address doesn't belong to the local IPv4 network, the IPv4 packet must be sent to the default-gateway instead. Sending to the default-gateway works the same as before, the IPv4 packet must be encapsulated into a layer two frame, and ARP must be used to set the appropriate **Destination MAC Address**. The only difference is that now, we will be requesting ARP for the default-gateway MAC address, and not the destination IP node MAC address.

Because communications with the router use layer two encapsulation and ARP, one condition must be met for a router address to be valid:

A router address is a valid next-hop only if it belongs to a local IPv4 network

5.3. Routing tables

The difference between a router and an end node is a router is supposed to retransmit IP packets, so it must be connected to more than one IP network. Thus, the router mission is more complex, while an end node has only two options (local destination or nonlocal destination) the router has more options.

A typical end node only needs to know the local IP network, if the destination address does not belong to the local network, then the packet is sent to the default-gateway. This means, all other networks are reachable through the default-gateway.

A typical router is connected to several networks and has several routers available around it to be used. Each neighbour router around it will provide access to some IP networks, the router must know which networks access is provided by each of its neighbour routers. This is the role of the routing table.

The routing table is a list of IP networks (IP and prefix length), and for each, the IP address of the neighbour router that should be used, this neighbour router is called the **next-hop**.

Take for instance a router connected to networks 192.168.10.0/24 and 192.168.20.0/24, the routing table could be something like what is presented in **Table 1**.

Table 1 – A routing table

Destination (network)	Next-hop
192.168.5.0/24	192.168.10.7
192.168.8.0/24	192.168.10.7
192.168.34.0/24	192.168.20.170
192.168.38.0/24	192.168.20.200

When this router receives an IPv4 packet for forwarding it will look at its destination IPv4 address to see which network it belongs to, options are:

- 1° Belongs to network 192.168.10.0/24 => direct sending (layer two)
- 2° Belongs to network 192.168.20.0/24 => direct sending (layer two)
- 3° Belongs to network 192.168.5.0/24 => send to 192.168.10.7 router
- 4° Belongs to network 192.168.8.0/24 => send to 192.168.10.7 router
- 5° Belongs to network 192.168.34.0/24 => send to 192.168.20.170 router
- 6° Belongs to network 192.168.38.0/24 => send to 192.168.20.200 router

Once a match is found the packet is sent and processing ends. If no match is found the packet is discarded, if this happens, it means the router doesn't know the destination network.

Checking if an address belongs to a network is achieved by applying the network mask (bit-to-bit and) to the address and see if the network address is obtained.

Routing tables can be manually set, this is called **static routing**. There are also routing protocols that are able to build routing tables and keep them updated, this is called **dynamic routing**.

5.4. Default route

It's impossible to store in single routing table all networks being used around the internet, however, if a network is unknown to a router, it will be unreachable.

The workaround is defining the special network 0.0.0.0/0 in each routing table. Because the mask is zero, it will always match any IP address, so it should always be the last entry at the routing table.

This entry is called the **default-route**, and the corresponding next-hop is called the **default-gateway** or **default-router**. Because it matches every IP address and is placed at the end of the table, the result will be that any packet addressed to an unknown network will be forwarded to the default-gateway.

Usually the default-route allows routing table simplifications, generally speaking, any routing table entry with a next-hop equal to the default-gateway can be removed. This is true because in the absence of that table entry, the table processing will continue until reaching the default-route, and the result is the same, meaning the packet is sent to the same router.

6. Use the Packet Tracer tool to create the layout shown in Figure 9

There are four IPv4 networks interconnected by three routers (Figure 9).




Figure 9 - Networks layout with three routers

6.1. Define all layer three nodes IPv4 addresses (both for routers and for end nodes)

Used IPv4 addresses must belong to the represented IPv4 networks.

Check that, for now, ARP tables are empty (use the Inspect tool – Magnifying Glass).

6.2. Check IPv4 connectivity.

Use the Add PDU tool to send ICMP echo requests between nodes.

Check that tests within each LAN are successful (namely from routers to local end nodes), but tests between different networks fail.

Also check that ARP tables are not empty anymore.

6.3. Define end nodes' (PCs) default gateways.

Notice that PC1 and PC2 have two alternative routers, in both cases use Router1 as default-gateway.

6.4. Test again communications, now only between end nodes (PCs)

Check that between PC1 and PC2 everything works fine, but elsewhere it fails.

Try again in simulation mode to try understanding what is happening.

6.5. Define each router routing table to solve the issue

For each router, check the remote networks it is not aware of. For each, add a static routing entry to inform to where the packets should be forwarded to reach that network.

6.6. Test again IPv4 connectivity between all end nodes

Check that now every node can communicate with every other node.

Check that ARP tables include only local addresses.

7. Keep the former layout and add an internet connection to Router 2

As shown in Figure 10, connect Router2 an ISP (Internet Service Provider) using a DSL (Digital Subscriber Line) across the PSTN (Public Switched Telephone Network).




Figure 10 - Networks layout with ISP connection

You will notice the Cisco 2811 Router2 has only two ethernet interfaces, and you need a third one to connect the DSL modem. Because this router has several free slots where we can plug the required hardware modules as we would do in a real device.

As with a real device, these hardware modules should never be plugged or unplugged with the power on, so we must first turn the power off, **but not before saving our configuration in Router2.**

To save the configuration click the **Save** button at the Router2 configuration window (Figure 11), you can see below that this is the same as running the **copy running-config startup-config** at the router's command line.




Figure 11 - Saving the current configuration

Now we can power the router off without losing the configuration. Switch to the physical view of Router2 (Figure 12), and click the power button, below on the right side of Figure 12.




Figure 12 - Physical view of the 2811 router in Cisco Packet Tracer

Now, with the router powered off, we can drag hardware modules, listed on the left side, and insert them into the free slots.

Pick the **NM-1FE-TX** module and plug it into the **NM** slot (the bigger one).

You can now power the router back on (click the power button again).

To establish the layer two connection between the modem and the ISP router, configure the cloud by associating the DSL connection to the Ethernet connection

Set appropriate IPv4 addresses for the new router and for the new interface of Router 2. Because the internet connection mask has 30 bits, there are only two valid node addresses, if the ISP Router is using 190.0.0.1/30, then the only available valid node address is 190.0.0.2/30.

Before advancing, check if there is IPv4 connectivity between Router2 and the ISP Router.

7.1. Change the routers routing tables to represent the new reality

Now there is an internet connection, therefore you must add a default-route to each routing table, whenever the destination address of a packet is locally unknown, it should be forwarded towards the ISP Router.

Add the necessary default-route entries in each local router.

In simulation mode use the Add Complex PDU tool to create an ICMP echo request addressed to a locally unknown address, say **100.10.10.10**. Send the request from PC0.

Check that all routers are forwarding unknown address packets in such a way they reach the ISP Router.

7.2. Simplifying routing tables.

Check that some entries in Router0 and Router1 are now useless because their next-hop is the default-gateway. **Remove them.**

Check that all works as before.

8. Classful IPv4 addressing.

An IPv4 node address is a 32-bits number, for the sake of human handling, it's represented as four sets of eight bits (octets). Each octet is represented as a decimal number and octets are separated by a dot. This is called the **dot-decimal notation**.

In Figure 13 we can see how from its 32 bits binary representation we attain the dot-decimal representation for the **192.168.0.5** IPv4 address.

Each octet in the dot-decimal notation can have a value from **0** (eight bits with the zero value) up to **255** (eight bits with the one value), however, when an IPv4 address is used to identify a network node, there will be additional restrictions to the values each octet may take.

An IPv4 node addresses are meant to uniquely identify nodes, there shouldn't be two nodes using the same address.




Figure 13 - The dot-decimal notation

To make the routing task easier, nodes are grouped into networks, and each IPv4 network is also identified by a unique IPv4 address (the network address). In fact, an IPv4 node address (32-bits) identifies both the node and the network the node belongs to.

Nodes belonging to the same IPv4 network can communicate directly with each other's, for that to be possible they must also be connected to the same LAN (same broadcast domain).

If two IPv4 nodes belong to different IPv4 networks they are not able to communicate directly, even if they were connected to the same LAN (same broadcast domain). If two IPv4 nodes are not connected to the same IPv4 network, the communication between them requires the use of IPv4 intermediate nodes (IPv4 routers).

8.1. Classful IPv4 addresses and network masks.

An IPv4 node address (32-bits) has two parts, the most significant bits (represented on the left side) are the **network prefix**, they uniquely identify the IPv4 network, and the remaining bits are used to uniquely identify nodes within that network, thus the full 32-bits uniquely identifies the node.

The number of most significant bits used to identify the network is called the **network prefix length** and it's not the same for every network. The network mask represents the network prefix length by expressing in dot-decimal an IPv4 address where network bits have value 1 and node bits have value 0.

Classful addressing (Figure 14) uses only three possible network prefix lengths: 8-bits, 16-bits, and 24-bits, corresponding to what are called respectively class A, B, and C networks. Therefore, corresponding network masks are 255.0.0.0, 255.255.0.0 and 255.255.255.0.

For classful addresses, the first most significant bits identify the class, and thus the network prefix length and the network mask.

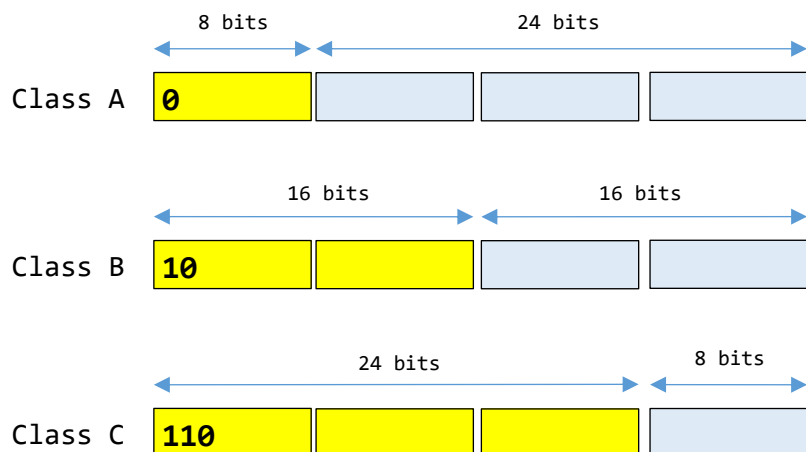


Figure 14 - Classful IPv4 addressing

The Figure 14 the three classes of classful IPv4 addresses are represented, yellow bits are used to identify the network, and the blue bits are used to identify nodes within the network.

Given an IPv4 node address, we can determine the network address it belongs to by zeroing all bits beyond the network prefix length, this is also equivalent to perform a bit-by-bit **and** operation with the network mask

Within every IPv4 network, two addresses are reserved, they are the address with all node bits having the zero value and the address with all node bits having the one value. The first represents the network's address itself and the second is the network's broadcast address. The number of different available networks and the maximum number of nodes each network can hold depend on the network prefix length, in Table 2 we can see the number of possible networks and maximum valid nodes in each network, depending on the class.

Table 2 - IPv4 classful networks

First bits	Values for leftmost octet	Class	Network mask	Possible networks	Maximum valid nodes on each network
0...	0 – 127	A	255.0.0.0	$2^7 = 128$	$(2^{24} - 2) = 16777214$
10...	128 – 191	B	255.255.0.0	$2^{14} = 16384$	$(2^{16} - 2) = 65534$
110...	192 – 223	C	255.255.255.0	$2^{21} = 2097150$	$(2^8 - 2) = 254$

Left octet values above 223 are reserved for special purposes like for instance multicast addresses.

9. IPv4 static routing

IPv4 intermediate nodes operate at the network layer and have the mission of transferring IPv4 packets between different IPv4 networks, they are usually called routers, and sometimes referred to as gateways.

IPv4 routers are usually connected to several local networks and make use of these local networks to forward IPv4 packets to other **neighbour routers**, known as next hops. Routers must know which networks are available behind each neighbour router, this knowledge is provided by the routing table. Table 3 represents an example of an IPv4 routing table.

Table 3 - An IPv4 routing table

Destination		Next hop
Network Address	Network Mask	
10.0.0.0	255.0.0.0	190.20.5.8
195.20.30.0	255.255.255	130.0.0.1
0.0.0.0	0.0.0.0	190.20.5.10

Each router only handles a single step (hop) in the path a packet must follow to go from the source node to the destination node. A single bad routing table, in a router along the path, makes the destination unreachable.

For every packet the router processes, the packet's destination address is matched with each entry at the routing table, the idea is checking if the address belongs to the network in the table entry, if so, the packet is forwarded to the table's corresponding next-hop, and the router mission is accomplished.

On Table 3, the last entry is called the **default-route**, and the corresponding next-hop is called the **default-gateway** or the **default-router**. Due to the network address and the network mask values that

are used, the default-route matches any IPv4 address, it's therefore used to represent all remote networks that are unknown to a router. In other words, any packet whose IPv4 destination address is not listed in the known destinations at the routing table will end up being forwarded to the default-router.

Defining a default-route in all routers is always required if the infrastructure has an internet connection, this is because it's impossible to add all the networks that exist over the internet to a routing table.

10. Practical exercises

10.1. Given the following classful IPv4 node addresses, determine the IPv4 network address it belongs to, the first valid node address on that network and the broadcast address on that network.

- | | |
|-----------------|----------------|
| a) 195.34.56.30 | b) 120.10.50.3 |
| c) 170.17.23.8 | d) 190.0.0.8 |

10.2. See the following diagram representing several IPv4 networks interconnected by routers.




Figure 15 - Networks layout and requirements

- Assign arbitrary classful network address to each network (A; B; C; D and E). Each network must be capable of supporting the specified number of nodes, and also, addresses wasting should be avoided as far as possible.
- Accordingly, set the IPv4 node addresses of each router's interface.
- Define each router's routing table.