Your Guide to the Home Network

www.homenethowto.com

Index

Revision: 20161022-1
www.homenethowto.com

Introduction

This document is written as a guide to present how computer networks work on a small scale such as a typical home network. The material was originally presented at www.homenethowto.com. There you can find the same material as in this document, but displayed in another format since this document has been edited for a text flow more suitable for documents.

One of our goals is that the information in this document should gather all of the basic knowledge that is required to understand and troubleshoot issues in a home network. If you understand how a network should behave then it is also much easier to figure out why things aren't working the way they should.

The document should be easy to read through and should not require any specific skills other than some general computer knowledge and an interest in figuring out how things work. Even so each chapter has a lot of information and the document handles a big variety of different networking areas.
Sometimes we have chosen to simplify things a bit by excluding more advanced parts of a subject. But both the text and the examples are never incorrect even when some bits might have been temporarily left for later on.

The chapters of this document are written in succession. You can read through the document in any given order, but if you find some area that you do not grasp completely then you might have to go back to an earlier chapter to find out more about previous subjects.

If you are teaching a course or if you are interested in using this material for any type of work related tasks then please contact us by using the form on the web page, www.homenethowto.com/contact.

Good luck with your network!

Best Regards,
the team at www.homenethowto.com

License Information

The content of the website at www.homenethowto.com is free of charge for private use. This document however is provided by us for a fee for both private and commercial use. See www.homenethowto.com/webshop-licensing/ for any changes to the License model.

The material, both online and in the PDF leaflet, is available for commercial use. It is well suited for courses and classes that introduce networking concepts and does so from a perspective that is familiar to most participants.

The material could be handed out to students or employees as a self study material, or could form the main curriculum in Computer Network Basics training classes.

Any form of commercial or professional use requires that each participant obtains a license for the material. This includes, but is not limited to, anybody who receives a PDF copy of the material or who is using the material via the web pages as part of a course or within a professional setting. Any instructor led training that is based on the material is also covered by the license requirement, where a license must be obtained for each participant regardless of whether or not the participant receives a physical copy of the material.

A license can be obtained by buying a copy of the PDF leaflet. The license model is primarily based on a gentleman's agreement, where we simply trust you and your organisation to obtain the correct amount of low cost licenses based on how many people that use or access the material.

Perhaps you would like to make the material available to a bigger group of people, for example all of your employees or every student in a higher education? Then please contact us via www.homenethowto.com/contact to discuss possible commercial licensing options.

Please also note that the license does not permit you to make any derivative work based on the material. If you want to create such derivative work (for example by re-branding the material to fit your corporate design templates) feel free to contact us to discuss possible solutions.

homenethowto.com, 2016

Introduction to home networks

A typical home network consists of a few devices where some devices have a more central role than others. Almost all home networks have a router as their central device. It is the router that connects to the Internet connection and then shares that connection to one or more computers on the Local Area Network or LAN.

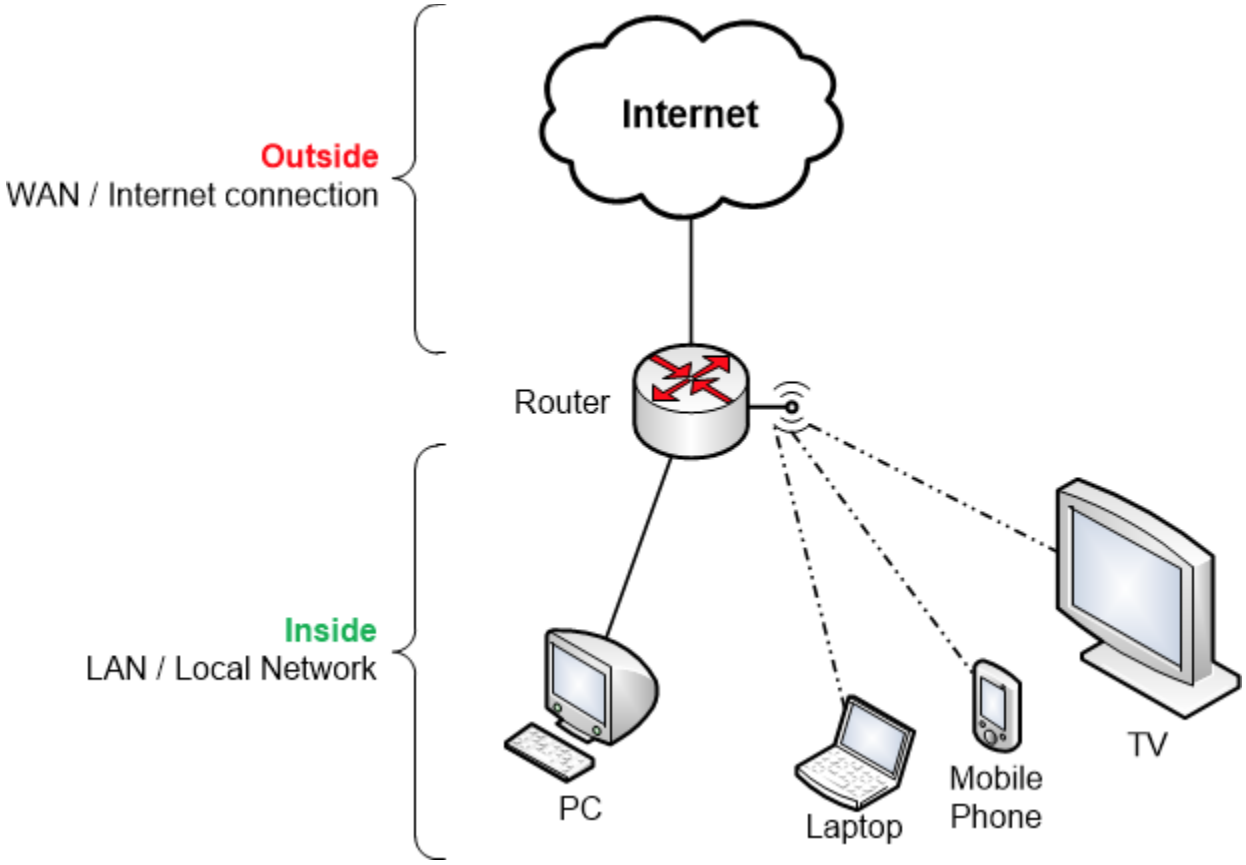

A home router typically has got an “outside” and an “inside”

- The outside is the Internet connection, which connects to a port on the router which is often called the WAN port
- The inside is the local network, or LAN. A router often has multiple ports that belong to the inside where you can connect different home computers, printers and other devices

Many routers also have built-in Wireless network. The Wireless network also belongs to the inside LAN of the router.

Related info

WAN means Wide Area Network, or a network connection that stretches out over a wider geographical area, such as an internet connection between a home and the Internet Service Provider (ISP)
LAN stands for Local Area Network, a network with limited coverage area, such as a network

within a home or a single company, often with a single owner for the whole LAN network

A closer look at the Home Router

Most people simply call their Home Router a *router* for short, but the term "router" is technically a bit misleading in this case. A Home Router actually consists of many different components which have been combined into a single box.

This guide will discuss these components and what they do throughout the guide, spread out over multiple sections. But the sooner you learn about what a "Home Router" actually is and how it works the better.

These are the main components of a Home Router:

- An integrated Switch with a number of LAN ports where you can connect computers and other devices using network cables
- An Access Point with one or more antennas to which wireless devices can connect
- Often there is a built-in modem, at least if the Home Router is meant to be connected to Cable or DSL-based Internet connections. The WAN port of the router is then connected to the integrated modem.
- The actual Router function which forwards traffic between the inside and outside networks mentioned previously. The Router normally also handles a number of other features:
    - Handing out IP addresses to devices on your home network
    - Handles Address Translations and Port Forwards
    - Takes care of any firewall rules
    - Replies to DNS-queries
    - Includes a Web interface that you can connect to via your web browser to configure the Home Router

Many of the features above can either be run in the actual integrated Router component, or they can run as separate services in a mini server in the Home Router.

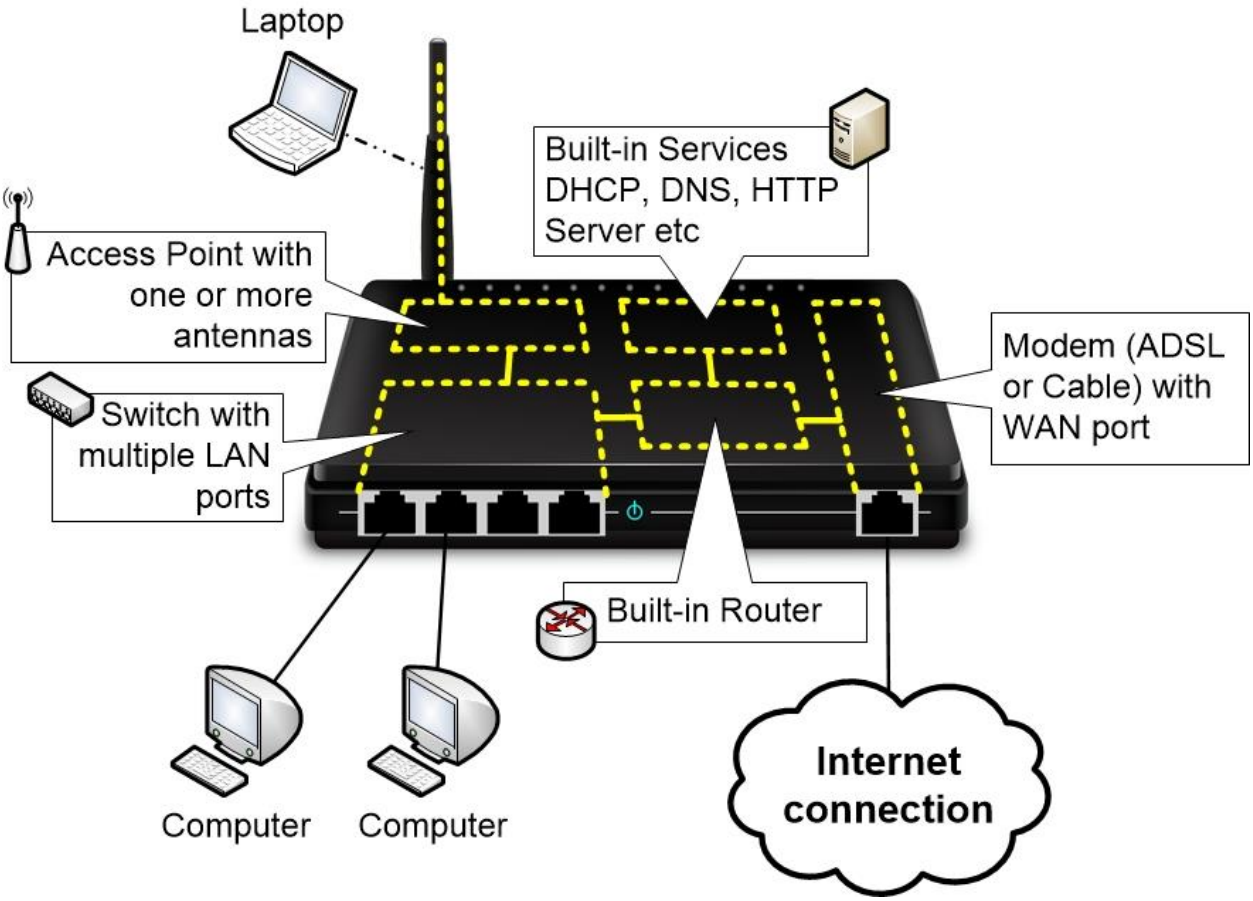

By looking at the picture above you can also see which of the integrated components that are typically connected to each other. You can also easily see that network traffic doesn't actually always have to pass through the built-in Router function. Traffic going between two computers on the LAN only has to pass through the integrated Switch. Similarly, traffic passing between the LAN ports and the wireless network just have to pass through the internal Switch and Access Point.

Traffic only has to be handled by the internal Router if the traffic is either going to the Internet or if the computers are requesting data from one of the Services that are running on the Home Router.

Each of the internal components will be discussed in more details later on in the guide.

Routing, an introduction to IP addresses

The main role of a *Router* is to route data traffic. The router knows in which direction different destinations are, and when it receives data traffic it will forward the traffic in the direction of the destination. A router always picks the best route it knows for the data traffic.

A normal home router often has a very simple task to perform because it doesn't have many possible routes to choose from. It knows that a local LAN network exists on the inside of the router, and that every other possible network thus must exist on the outside via the Internet connection. As long as the destination is not known to exist on the inside, then the home router can safely assume that it should send the traffic to the outside instead, knowing that the traffic will be heading in the right direction.

As long as two computers on the inside of the router are talking to each other, then the router passes the traffic on directly between those computers. If however a computer on the inside tries to browse to a web page on the Internet, then the router will send the traffic to the Internet Service Provider (*ISP*) on the outside WAN interface. Then the ISP's more powerful routers will take over responsibility for passing on the traffic towards its destination.

In fact that is all the Internet is - thousands of routers that are interconnected with each other. Each ISP has their own routers which find their way towards different destinations. The ISPs agree to connect their routers to each other so that traffic from one ISP's customers can find its way toward destinations (web servers for example) that are connected to other ISP's routers.

So in a very simplified way, the Internet isn't much more than a lot of interconnected routers which are owned by different ISPs. Companies who want to run a web server will contact their ISP and ask for an Internet connection so that they can connect their web servers to the ISP's routers. Home users will also lease an Internet connection from an ISP to connect their home routers and their computers to the Internet.

By understanding a bit more about how your home network is functioning you will also gain some understanding in how Internet as a whole works. It's basically the same thing, just on a larger scale!

Routing introduced

IP addresses are the type of addresses that computers use to send data traffic between each other. IP addresses can be compared somewhat to street addresses in real life. Imagine that a person wants to send a letter to a friend. The sender must then know the street address of the receiver and must print it down on an envelope, and must then put the message inside the envelope.

The sender doesn't have to know exactly where the receiver's address is located. It is enough that the sender knows the exact address of the receiver. Even so the sender can drop the envelope off in a post box and thereby send the message. A mailman will pick up the letter and bring it to a main sorting office.
Then the postal service will take over responsibility for transporting the letter closer and closer to the receiver until the letter finally has been delivered to the correct address. So the sender is delegating responsibility for handling the letter to the postal service, which through its chain of delivery will transport the letter to its final destination.

In the picture below you can see that different parts of the address on the envelope have different importance depending on where in the mail delivery chain the letter is currently at.

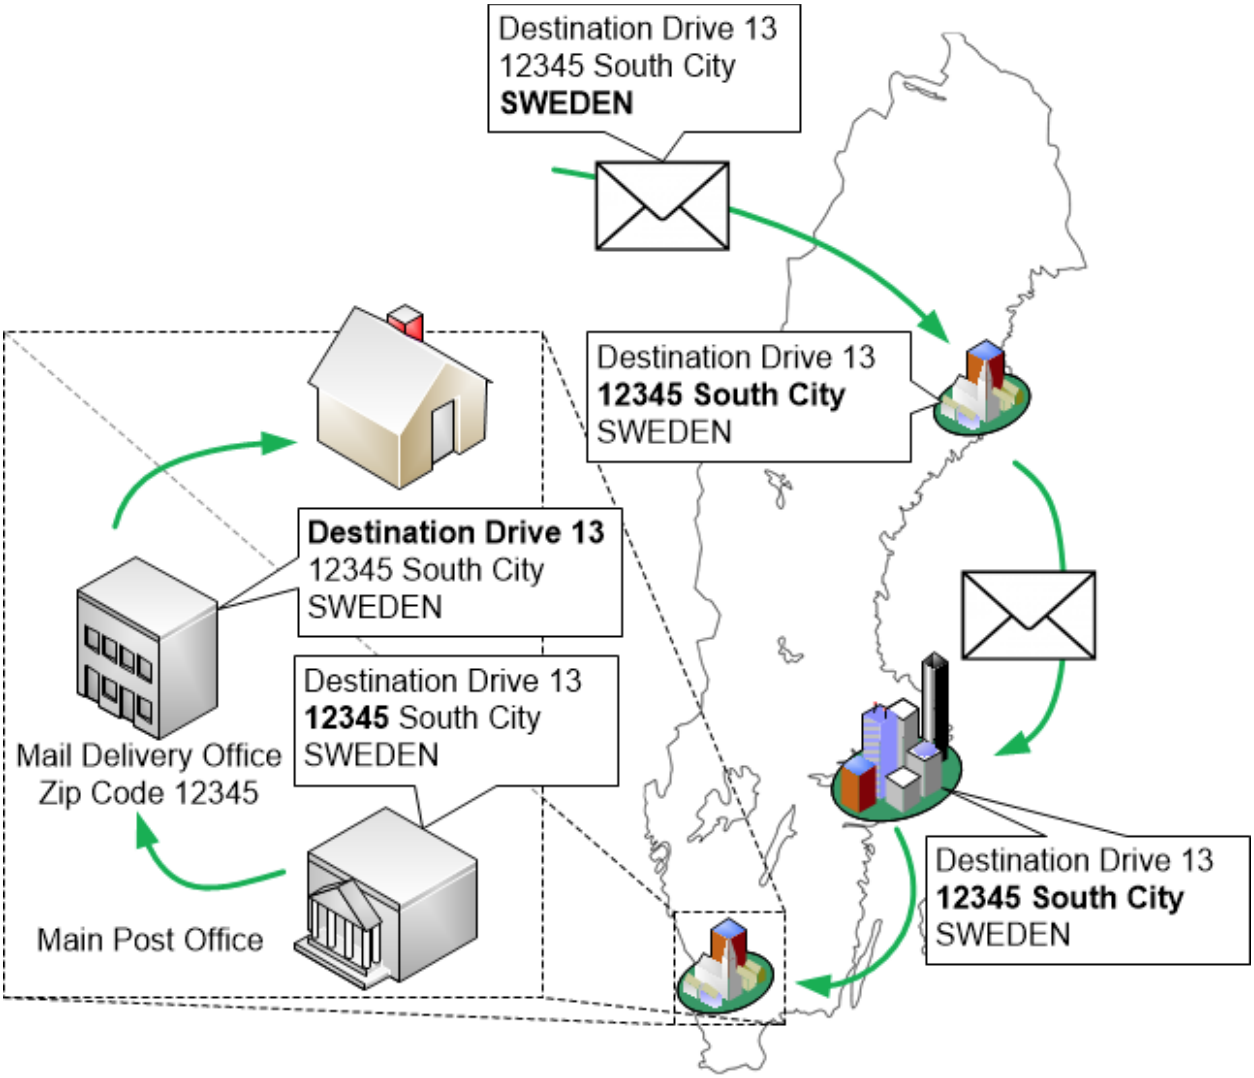

At the first sorting office the address on the envelope is investigated to see where the letter is being sent. Then the first sorting office sends the letter towards its destination.
If the receiver is far away, or perhaps even in another country, then the letter will first be sent to another central post office closer to the receiver. This could even be repeated multiple times. So the first sorting central might not even know where the exact destination street address is. But all they have to look at is the more general direction that the letter should be sent toward so that the letter ends up in the correct country, region or city.

The closer the letter gets to the final destination the more important all the details of the destination street address becomes. When the letter eventually ends up in a main sorting office

in the right town, then that sorting office will have to give the letter to the correct mail delivery office based on the postal code. Finally the letter is handled by a mailman who knows exactly where the final destination is, and he delivers the letter to the receiver.

The mail sorting offices in the example are represented in the world of computer networks by routers. A home router who is sending out traffic to the Internet doesn't have to know exactly where www.homenethowto.com is located. All it needs to know is that the general direction it needs to send the traffic towards is the outside Internet connection, because that is where all networks except for the inside LAN network are located.
Then the task of sending the traffic towards the destination is delegated to the ISP router that the home router is connected to. The ISP router will also look at the traffic without knowing exactly where the destination is. But it will have more detailed knowledge than the home router and will also have more possible routes to choose from.

In this way computer traffic is sent or routed over the Internet closer and closer to its destination via various ISP's routers. The closer the traffic gets the more detailed knowledge of the destination the routers will have. Eventually the traffic will end up at a router that has a direct connection to the destination, which could be a web server or similar. That last router will then deliver the traffic to the web server.

IP addresses

Similar to how full street addresses consist of multiple parts, IP addresses also have multiple parts. Where a street address can consist of the increasingly specific parts Country, City, Postal Code, Street and House Number, an IP address consists of two main parts: IP network and addresses within the IP network. The addresses within the IP network are also called the *host addresses* because different hosts or computers have different unique addresses within the IP network.

- The **IP network** corresponds to the more general parts of a full street address, such as City, Postal Code and Street.
- The **Host Addresses** within the IP network correspond to specific house numbers on a street.

On a computer network, all devices that are connected to the same local network can talk directly with each other, as long as they have IP addresses that belong to the same IP network. If the computers belong to different IP networks then they have to communicate with each other via a router. The main purpose of a router is to be able to forward traffic to different destinations. Within computer networking those destinations are different IP networks.

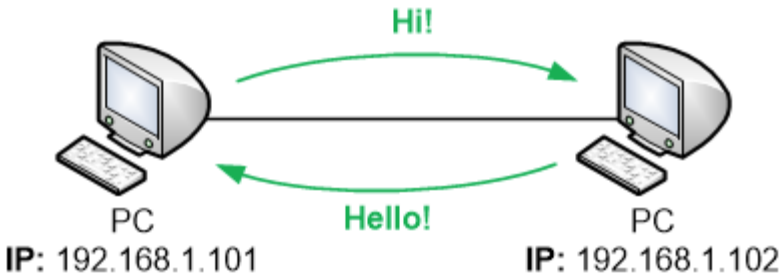

Two computers on the same IP Network

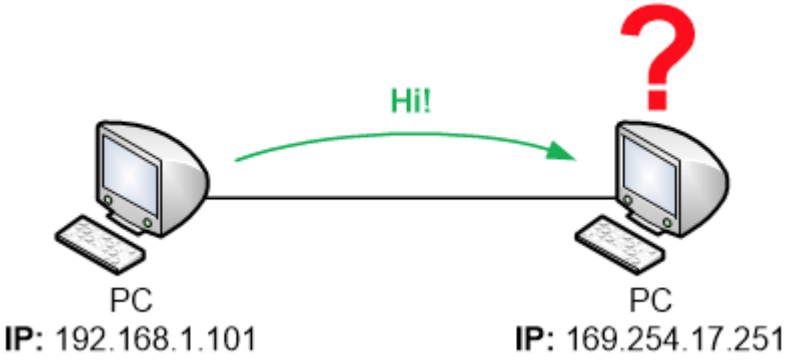

Two computers on different IP Networks

So how do you know if two IP addresses are on the same IP network or on two different IP networks? Unfortunately it isn't as easy as it is with street addresses to see which part of an IP address is the IP network and which part belongs to the host addresses.

The answer is within something called a *Subnet Mask*. An IP address is always combined with a Subnet Mask, and it is the Subnet Mask that determines which part of the IP address that belongs to the IP network and which part that belongs to host addresses.

Subnet Masks and IP Networks

To really understand exactly how the Subnet Mask works you would have to study and learn about binary numbers and a few other more advanced topics. Luckily you don't need any of that deeper understanding unless you work with IT or computer networks.

Both an IP address and a Subnet Mask consists of four parts separated by periods. Each part of an IP address and a Subnet Mask can have a value between 0-255

	Part 1	Part 2	Part 3	Part 4
IP address	113.	211.	197.	5
Subnet Mask	255.	255.	255.	0

In its simplest form, each part of the Subnet Mask is either the number 255 or the number 0 (zero).

- 255 means that the corresponding part of the IP address is part of the IP network.
- 0 (Zero) means that the corresponding part of the IP address belongs to the Host Addresses.

Here are a few examples of IP addresses and Subnet Masks in combination:

- If the Subnet Mask is 255 in the first part then the first part of the IP address shows what IP network the address belongs to.

	IP network	**Host Addresses**		
IP address	**50.**	211.	197.	5
Subnet Mask	**255.**	0.	0.	0

- If the second part of the Subnet Mask is also 255, then the second part of the IP address is also part of the IP network.

	IP network		**Host Addresses**	
IP address	**50.**	**211.**	197.	5
Subnet Mask	**255.**	**255.**	0.	0

- If the third part of the Subnet Mask is 255, then the third part of the IP address also belongs to the IP network.

	IP network			**Host Addresses**
IP address	**50.**	**211.**	**197.**	5
Subnet Mask	**255.**	**255.**	**255.**	0

Network Address, the name of the IP network

Each IP network has a so-called *Network Address* which is the "name" of the IP network. If you wanted to tell somebody which IP network a computer is on you would always tell them the first (lowest numbered) address on the IP network, which is the Network Address.

Example 1:
In the picture below, the first three parts of the IP address belongs to the IP network. This is determined by the Subnet Mask.

IP: 101.102.103.5
Subnet Mask: 255.255.255.0

0 (Zero) is the lowest address that is available in the fourth part of the IP address. The computer therefore belongs to the IP network 101.102.103.0
The fourth part (.5) of the IP address shows which host address that the computer is using on the IP network.

Example 2:
The next computer below belongs to the IP network 211.139.157.0
It is using the host address 9 on the IP network, and its IP address is 211.139.157.9

IP: 211.139.157.9
Subnet Mask: 255.255.255.0

Example 3:
This computer belongs to IP network 192.168.1.0, which is one of the most common IP networks that you can find on home networks.
It has received host address 7 on the IP network, and its full IP address is 192.168.1.7

IP: 192.168.1.7
Subnet Mask: 255.255.255.0

IP Networks in your home network

In a normal home network, the Subnet Mask is usually “255.255.255.0”. What that means is that the three first parts of the IP address determine which IP network that the IP addresses belong

to. The last part of the IP address determines which unique address within that IP network that each individual computer has got.

As we mentioned earlier, each part of an IP address can have a value between 0-255. So the fourth part of the IP address permits for 256 different addresses (zero up to 255) that can be used for computers, IP phones, routers, laptops, printers and other devices in the home network. These type of devices are commonly referred to as *hosts* och *clients* in the world of computer networks.
On a normal home network those addresses are always more than enough to cover the devices that are connected to the network.

The first address (zero) and the last address (255) can not be used for computers. They are reserved for special functions. The first address is the Network Address which has already been discussed. The last address is used for broadcasting, which is discussed in the Switching part of this material.
So in the end, there are 254 available addresses that can be used. Normally the home router will use one of those addresses, and the rest are available for your computers and other devices.

In a home network it is usually the router that determines which IP addresses that the computers can use. The home router hands out IP addresses, the Subnet Mask and other details to the computers. This will be discussed later on in the section about how addresses are assigned to devices.

Related info

When speaking about the computers, wireless phones, printers etc that are connected to a computer network the term ***client*** is often used. A client is any device that acts as an end user device, something that a user is interacting with. This could be a PC, your mobile phone, the wireless printer you have, your Smart TV or your game console.

The term ***device*** is a bit more generic and can normally include clients but sometimes also the network equipment itself such as your router or switch.

Another term that is sometimes used is ***host*** which more commonly refers to computers and servers on the network - things that are running operating systems such as Windows, Linux, Android or Mac OS X.

But none of these terms are really set in stone, and they are often used interchangeably. The important bit is that since most of these devices behave in more or less the same way on a network it is often simpler to refer to them as clients, hosts or devices instead of having to specify exactly what types of devices the information is about.

Default Gateway, finding other IP networks

A computer that has an IP address and a Subnet Mask can talk directly with other computers that share the same Subnet Mask and have IP addresses within the same IP network.

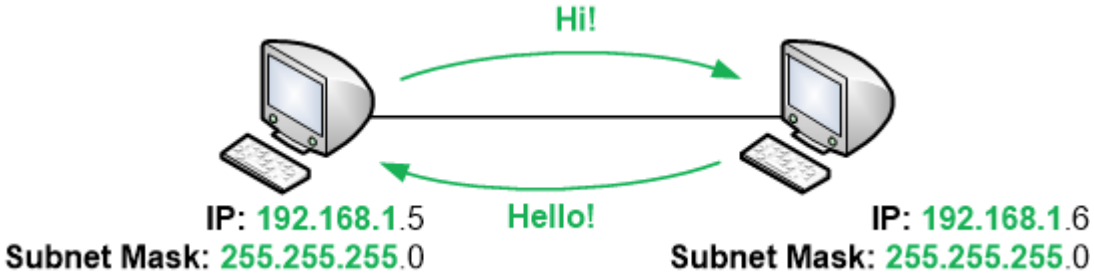

But what about a computer that wants to talk with something that has an IP address on another IP network? In the below example the left computer belongs to IP network 192.168.1.0 and the right computer belongs to IP network 192.168.2.0 - two different IP networks.

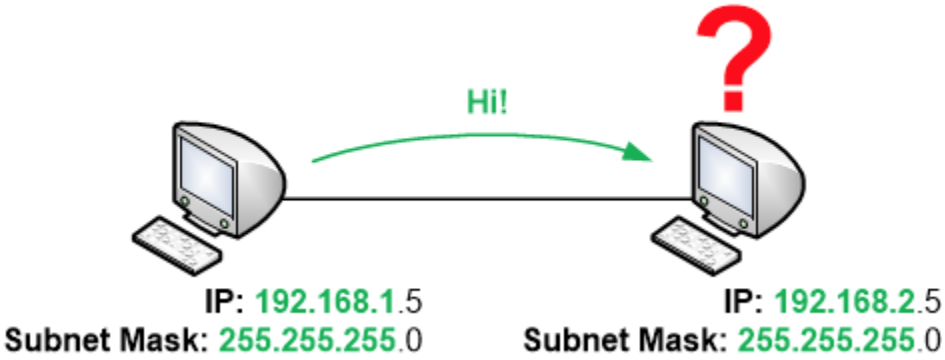

In previous parts, we concluded that to be able to communicate with other IP networks a router is used. The computers above need to communicate via a router. But how does the traffic actually end up in a router?

To find its local router the computer needs a so-called *Default Gateway*, sometimes simply called a *Gateway*. A Default Gateway is always a router that can connect to more than one IP network and can route traffic between those IP networks. The router will have its own IP address on each IP network that it connects to.

A computer can only talk with other IP addresses within its own IP network, the local IP network that the computer belongs to. So for the computer to be able to communicate via a router, at least one of the router's IP addresses must belong to the same IP network as the computer.

When the router hands out an IP address and a Subnet Mask to the computer it also sends out its own IP address to the computer and tells the computer to use that address as its Default Gateway.

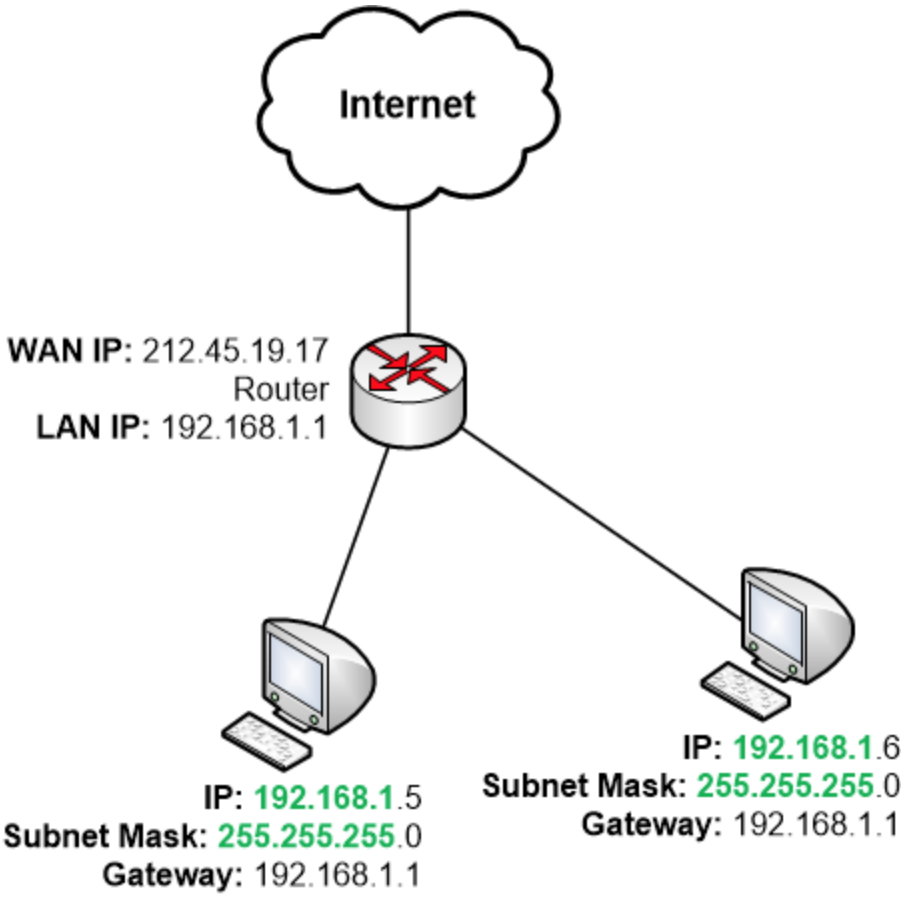

With the help of the Default Gateway address, the computers can then find their way to all other IP networks in the whole world. The computer knows that it can talk directly with any other IP address on the same IP network as itself. But as soon as it needs to talk to any other IP network it just needs to send the traffic to the Default Gateway, which is the router. Then the router will take over responsibility for routing the traffic towards the destination on the Internet.

So you could say that by sending the traffic to the default gateway the computer is delegating responsibility for forwarding the traffic to the router.The computer can trust that the router can find the destination. In turn, the home router will then trusts the ISP's routers to take over responsibility for forwarding traffic further along the path toward the destination.

DNS, names and IP addresses

In previous parts, we have so far talked about how IP addresses work, how Subnet Masks determine which IP network that a computer belongs to, and how a computer can send traffic to other IP networks via a Default Gateway, which is the router's IP address.

But as a user, you rarely need to bother much about IP addresses, which is somewhat lucky since they can be hard to remember. Instead you are used to relying on something called

Domain Names to connect to things on the Internet, for example www.homenethowto.com or www.apple.com
Those names are called Domain Names and make it easier for us people to remember how to access the services on the Internet that we want to browse and connect to.

But when computers send traffic to each other they must use IP addresses as destinations, even if the traffic is going to a web server that has an associated domain name.
Therefore the computer must have a way to translate the names that we use to IP addresses that the computer can use.

The translation is made by DNS, which stands for Domain Name System, a system for translating back and forth between IP addresses and DNS names or domain names.

DNS is served by a large number of servers on the Internet which can reply to queries about domain names. The DNS servers are owned by different companies and organisations, for example ISP's, web hosting companies and similar.
When you buy or lease an Internet connection from an ISP they also always provide DNS as a service to you, and your home router will also automatically learn about and use those DNS servers.

For a computer to be able to look up which IP address that a particular domain name has got, the computer must first find its way to a DNS server. Luckily the computer can obtain this information from the home router. When the computer gets its IP address from the home router, then the router also passes along information about which DNS server that the computer should use.

Additional Info

A **server** is really just a regular computer. The difference is that the server is specifically made to be good at hosting one or more services to other computers on the network. A DNS server, for example, could just as well be any regular home PC with an installed program that can reply to any incoming DNS queries.

But most often, servers on the Internet are more expensive high-end varieties of computers that are made to be durable even when powered on indefinitely. They have hardware components that are made to be able to handle many queries simultaneously from many different users. In addition, they are made to be mounted in special Rack mounts in data centers and as such have different appearances to a normal home computer.

But you can more or less think of a server as a regular computer made for a specific purpose that it is really good at.

A lot of home routers tell the computers on the LAN to send their DNS queries to the home router itself. The home router will then either answer directly if it knows the ansewr, or it will forward the query to a DNS server on the Internet.

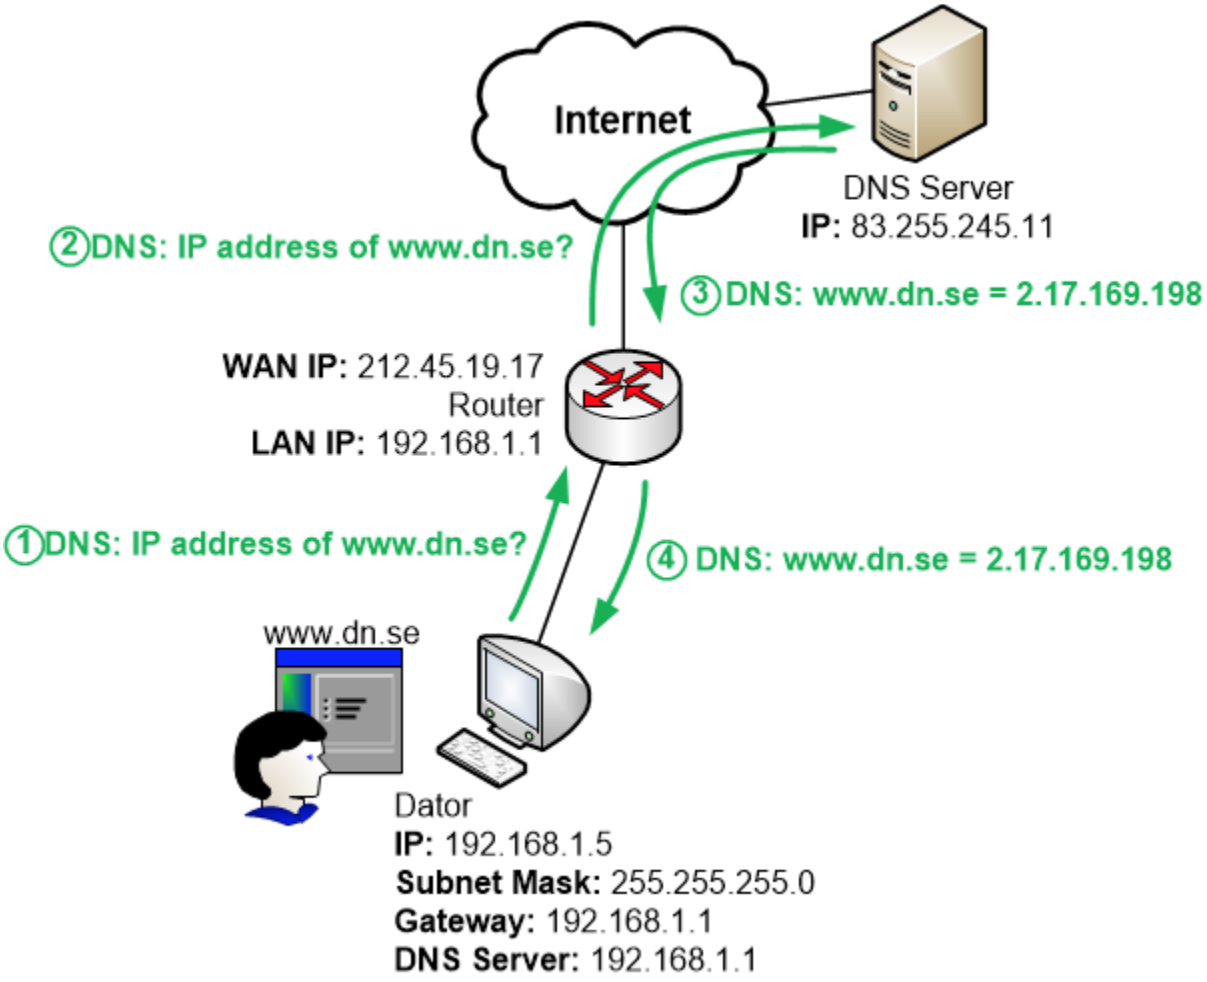

When a computer wants to browse to a domain name it queries the DNS server for what IP address that domain name has got. Once it gets a DNS response back containing the IP address of the domain name it can use that IP address as destination for the traffic.

There are also home routers that simply tell the computers on the LAN to send their DNS queries directly to the DNS servers of the ISP. This works just as well. Since the DNS servers are on the Internet on another IP network the computer cannot ask its DNS questions directly to the DNS servers. Each query must pass via the Default Gateway which is the router. So the DNS queries still goes from the computer via the home router to the DNS servers on the Internet. The difference is that the home router just forwards the traffic and does not have to handle the actual DNS query itself. It just passes along the query without looking at the contents.

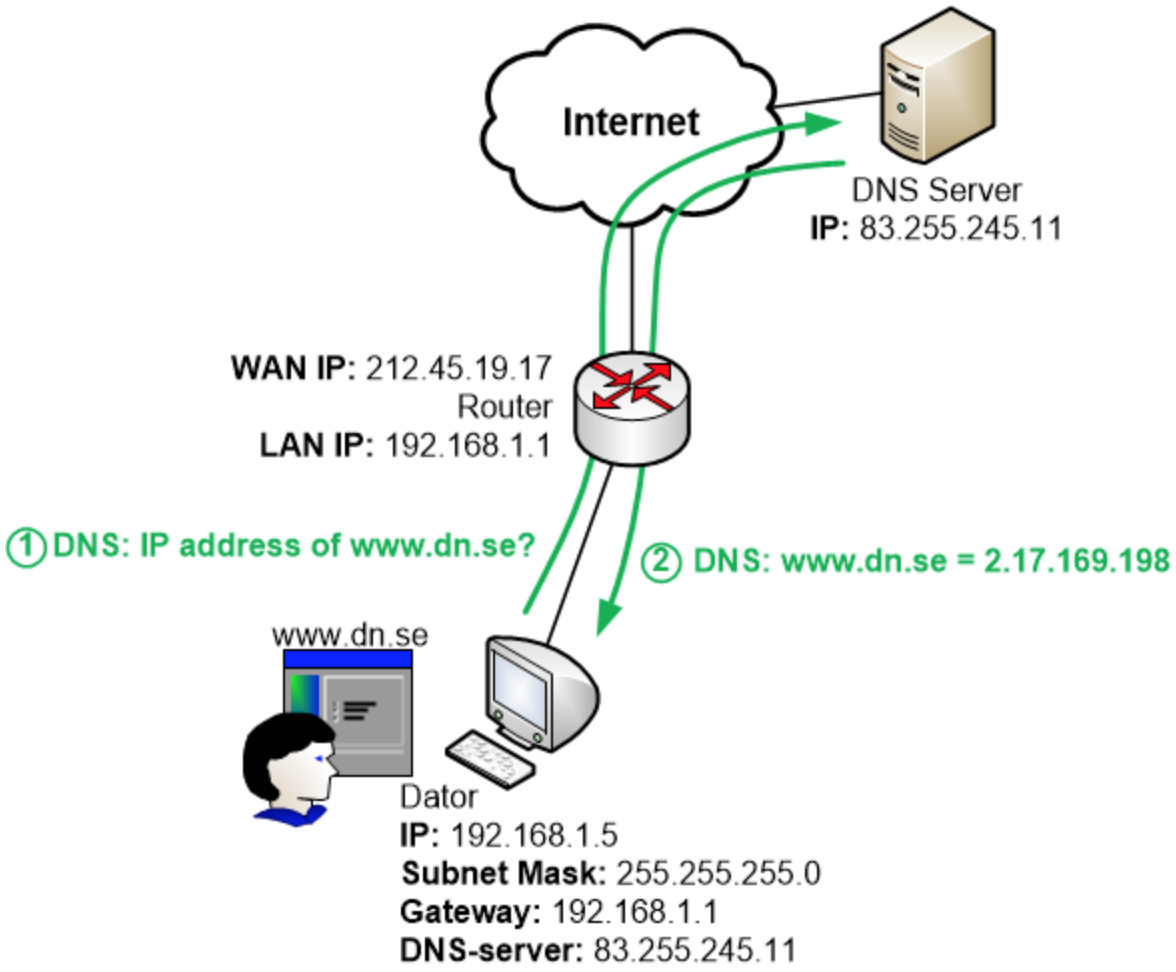

How the computer obtains an IP address

A computer needs the following information to function normally on a computer network:

- IP address
- Subnet Mask
- IP address of a Default Gateway (router)
- IP address of a DNS server

There are two ways that a computer can obtain those details. Either automatically, or via manual configuration.

DHCP - automatic assignment of IP addresses

In a home network, the router normally decides how the LAN should work. The router will forward traffic between the clients on the LAN and also between the LAN and the Internet.

With that in mind, it is only natural that the router also hands out IP addresses and other necessary information to the computers on the network. This is done via *DHCP*, which stands

for *Dynamic Host Configuration Protocol*. In other words it is a protocol to automatically hand out configuration to computers and other devices on the network.

Usually when you receive your home router it is already pre-configured with a DHCP server to hand out configuration to your computers and other devices. The router is also prepared so that the addresses that the router hands out via DHCP is on the same IP network that the router is configured to use for its own LAN IP address. This is necessary for the clients to be able to use the router as their Default Gateway.

When a computer connects to a network it will try to ask for an IP address. This is done by sending out a DHCP request where it asks if there are any available DHCP servers on the network. If any DHCP server responds then the computer will use DHCP to ask for an IP address and all the other necessary information it needs from the DHCP server.

So when your router sees this DHCP request it will hand out an available IP address from its pool of free IP addresses, together with the other details that the computer needs.

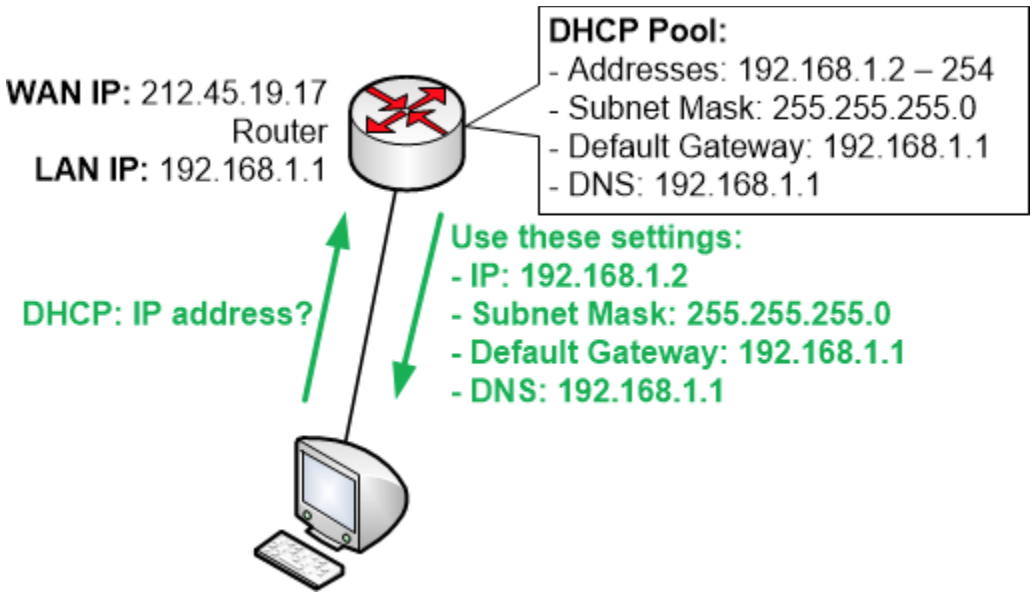

In the above example the router's DHCP server has a pool of available IP addresses starting with 192.168.1.2 and going all the way up to 192.168.1.254. The router will hand out the first available IP address from that pool and will mark the address as "leased" so that it does not hand out the same IP address to any other client on the network.

All clients on the LAN will receive the same Subnet Mask, Default Gateway and DNS Server settings from the DHCP server, since those details are common for all clients.

Manual configuration of an IP address

Instead of letting the computer obtain its IP address from the router via DHCP you can choose to manually configure the IP settings on the computer. Normally this is avoided since it can

cause a few different problems unless it is handled handled properly by the administrator, which is you.

When and why would you need to manually configure an IP address on a client?

If a computer obtains its IP address automatically via DHCP then it is not certain that the computer will obtain the same IP address each and every time you start the computer. The DHCP server remembers which computer that has gotten which IP address, but only for a certain amount of time. If a computer is powered off for too long (often a day or two, depending on how the router is configured) then the DHCP server will forget which IP address that it handed out to the computer. Also if the router is powered off for any reason then it will typically forget about any DHCP leases it has handed out.

In some special cases, this could lead to problems. One such example is if you have had to make a Port Forward (a subject which is discussed in further defail later on in this guide). Port forwards often point to an internal LAN IP address of a computer. As long as the computer keeps the same IP address the Port Forward will work. But if the computer changes IP addresses every so often, then after each IP address change the Port Forward must be updated in the router configuration.

In that situation, it is often recommended to configure the computer that should receive the Port Forward with an IP address manually. That way the IP address stays the same and the Port Forward keeps working.

When you configure an IP address manually on a computer you need to configure the same settings that a computer normally receives via DHCP:

- An available IP address on the same IP network as the router
- The same Subnet Mask that the router is using
- Default Gateway, which should be set to the LAN IP address of the router
- DNS Server address - either the router LAN IP address or another DNS server on the Internet. You may use the same address that the router normally hands out via DHCP

IP address conflicts

If you choose to manually configure an IP address on a computer, then you also should make sure to exclude that IP address from the pool of DHCP addresses in your home router. Otherwise the router might hand out the same IP address to some other computer on the network.

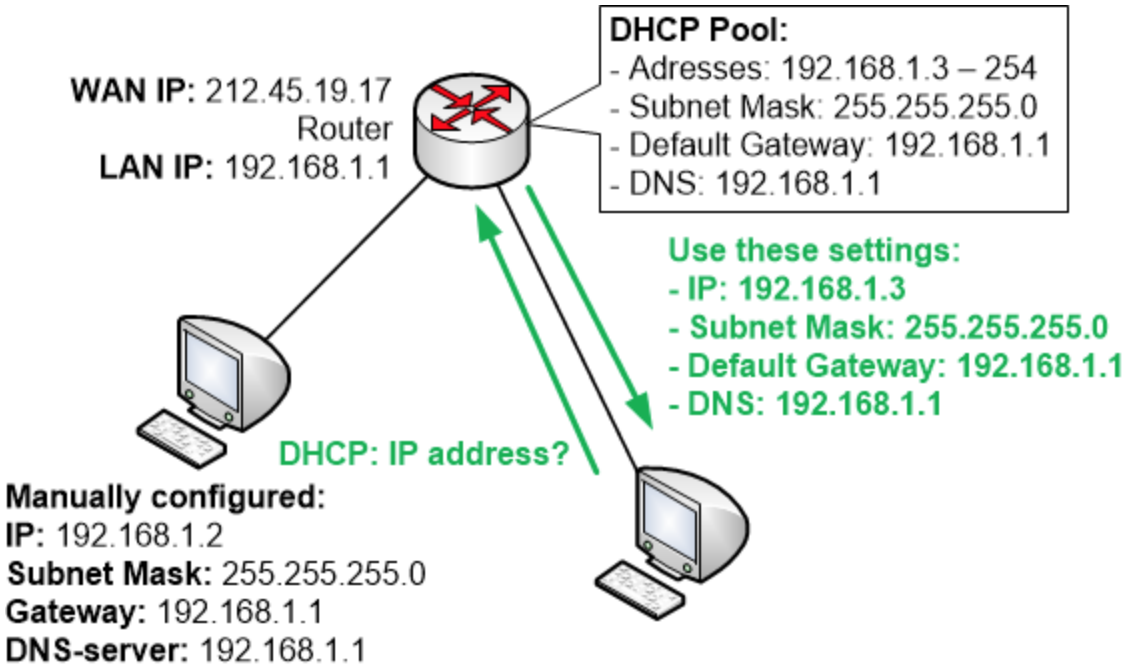

Using the street address analogy again, if two houses on the same street for some reason had the exact same house number, then the confusion would be great. Some packages and letters would of course end up at the correct house whereas others would end up at the wrong place. It would very much be hit and miss with a big random element to it.

The same thing would happen on a computer network where two devices were configured to use the exact same IP address. You then have an *IP address conflict* on the network, and the result is basically that communication stops working for the involved clients. Network communication simply does not work if only approximately half of the traffic ends up in at the correct place.

In modern networks and with newer operating systems the computers will try to avoid IP address conflicts by checking first if the IP address seems to be taken already. But still even then only the first computer that obtains the IP address will work correctly. The second computer that accidentally is given the same IP address as the first one will notice the IP address conflict and will then simply avoid talking on the network until it has been given another IP address.

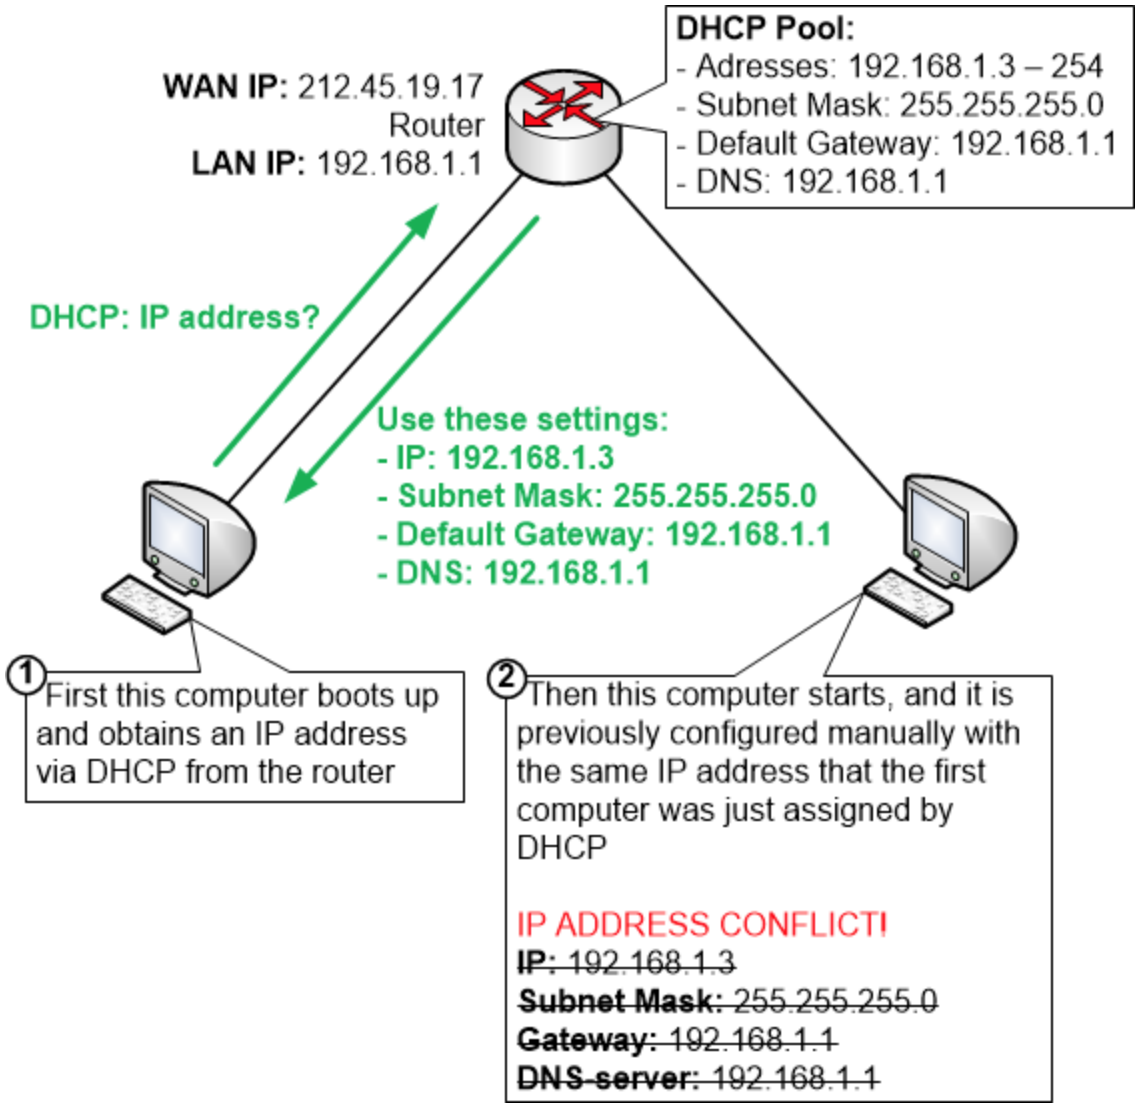

Public and Private IP addresses

On the Internet, a limited number of IP addresses exist. We have already mentioned IP address conflicts in the previous chapter. Each computer that wants to communicate on a network needs a unique IP address to function.

This is also true for the Internet in general. All IP addresses in use must be unique and must not be used anywhere else on the Internet.

However, there are exceptions to this rule. One such exception is all IP addresses beginning with “192.168” which belong to a special category of IP addresses that are commonly called *Private IP addresses*.
These Private IP addresses are reserved for local use within LAN networks. They can be reused in many places. Actually, you will find IP addresses from the 192.168. address range at almost any home and even at many corporate networks.

Two other such Private address spaces exist

- Any address beginning with “10.”
- All addresses beginning with “172.16” up to “172.31” for example 172.16 or 172.23 or 172.31

Private IP address space	
From	**To**
10.0.0.0	10.255.255.255
172.16.0.0	172.31.255.255
192.168.0.0	192.168.255.255

Most often home routers use IP networks that begin with “192.168” on the inside LAN. The two absolutely most commonly used IP networks on the inside LAN of home routers are “192.168.0.0” and “192.168.1.0” with the Subnet Mask 255.255.255.0

As mentioned earlier, the above special addresses are called Private addresses. They cannot be used on the Internet, they can only be used within local networks. If you try to use Private addresses on the Internet then your Internet Service Provider will block your traffic automatically, sensing that the traffic is coming from a Private IP address. This automatic block is being done to avoid any IP address conflicts on the Internet. These addresses are used in so many places that without the block we would have guaranteed and constant IP address conflicts all over the Internet.

So how can the home network function and how can you browse the Internet from a computer which is configured with one of these Private IP addresses? Obviously, almost everybody is using such private addresses, so what is preventing those constant IP address conflicts?

The secret is in something called Address Translation, or NAT (Network Address Translation) which is discussed in the next section.

Address Translation

As mentioned in the previous section the number of Public IP addresses (all IP addresses that can be used on the Internet) is limited. In total, there are about 4 billion IP addresses, much fewer than the number of people on the earth. IP addresses are needed not only for people but for all things that need to communicate on a computer network. This includes servers, web services, network equipment, cars and so on.
In practice, each Internet-connected person uses many public IP addresses. Not only simply because each person has got multiple devices (mobile phone, computer etc) but also each

person has multiple roles where they use Internet-connected devices, for example at work and at home.

The one big solution to this problem is the Private IP address space that can be reused over and over without limitation. But private addresses cannot be allowed to be used on the Internet. Traffic can't be sent over the Internet *to* private addresses, nor be sent *from* private addresses.

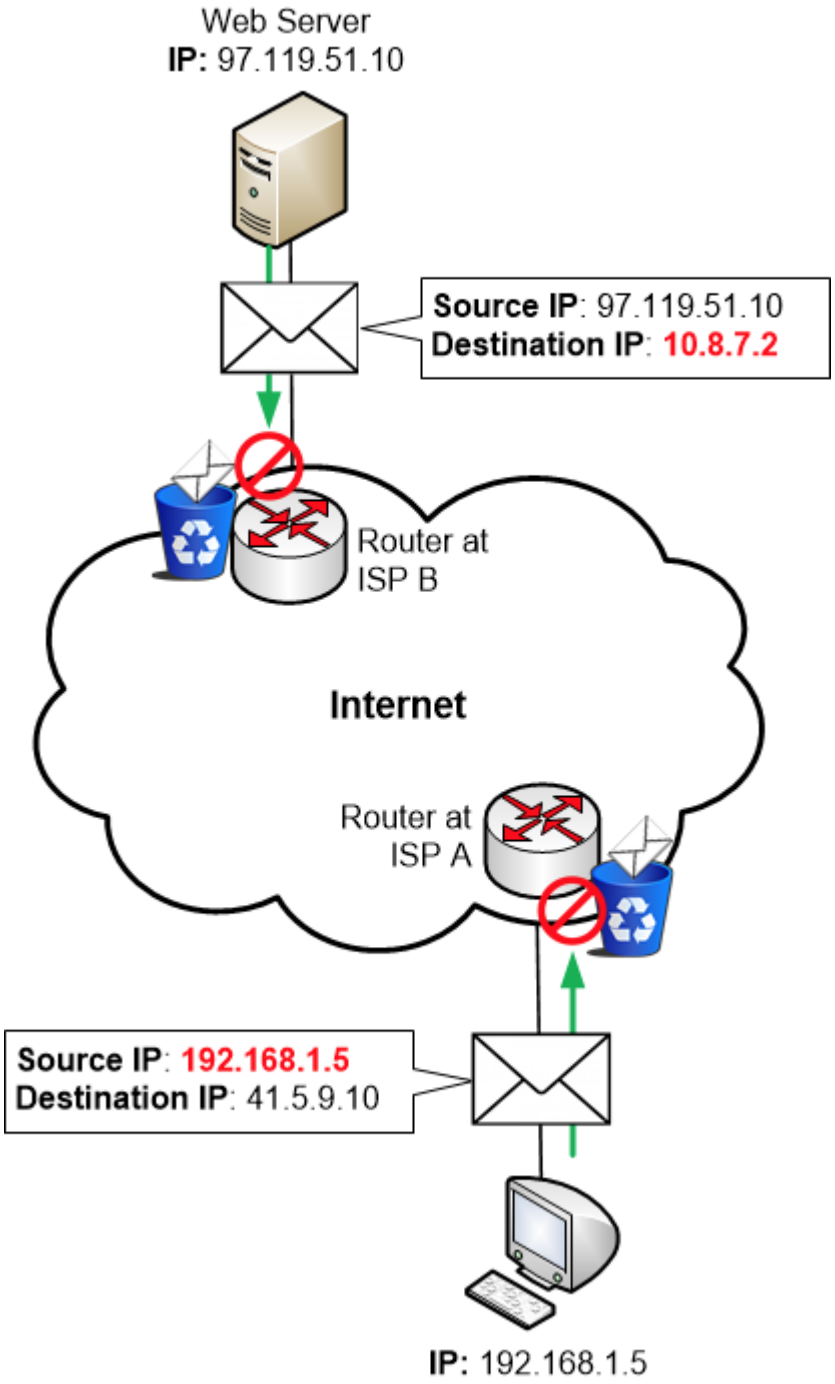

So the Private IP addresses create another problem that in turn also requires a solution. That solution is to let many private IP addresses share one single public IP address. This is accomplished using *Address Translation* or *NAT* (Network Address Translation), which is described below.

When a computer wants to communicate it sends off a packet with data. The packet always has two IP addresses inscribed in the envelope or *header* of the packet.

- *Source Address*, which is the IP address of the sender. This has to be entered into the packet so the receiver knows where it should send its replies, like a "return address"
- *Destination address*, the IP address of the receiver that the packet is being sent to

In a home network where a computer wants to talk to something on the Internet, the source address will be a Private IP address on the LAN. The destination address of the packet will be a Public IP address of a server on the Internet.
If that packet is sent to the Internet then the ISP will block and throw away the packet since it has a private IP address as its source.

To fix this problem the home router steps in and translates the source address from a private address to a public IP address.
The router itself has a public IP address on its outside WAN interface. It got that public IP address from the ISP. The router will simply let every client on the inside LAN share that single public IP address.

The following picture shows (somewhat simplified) how the address translation works in the home router:

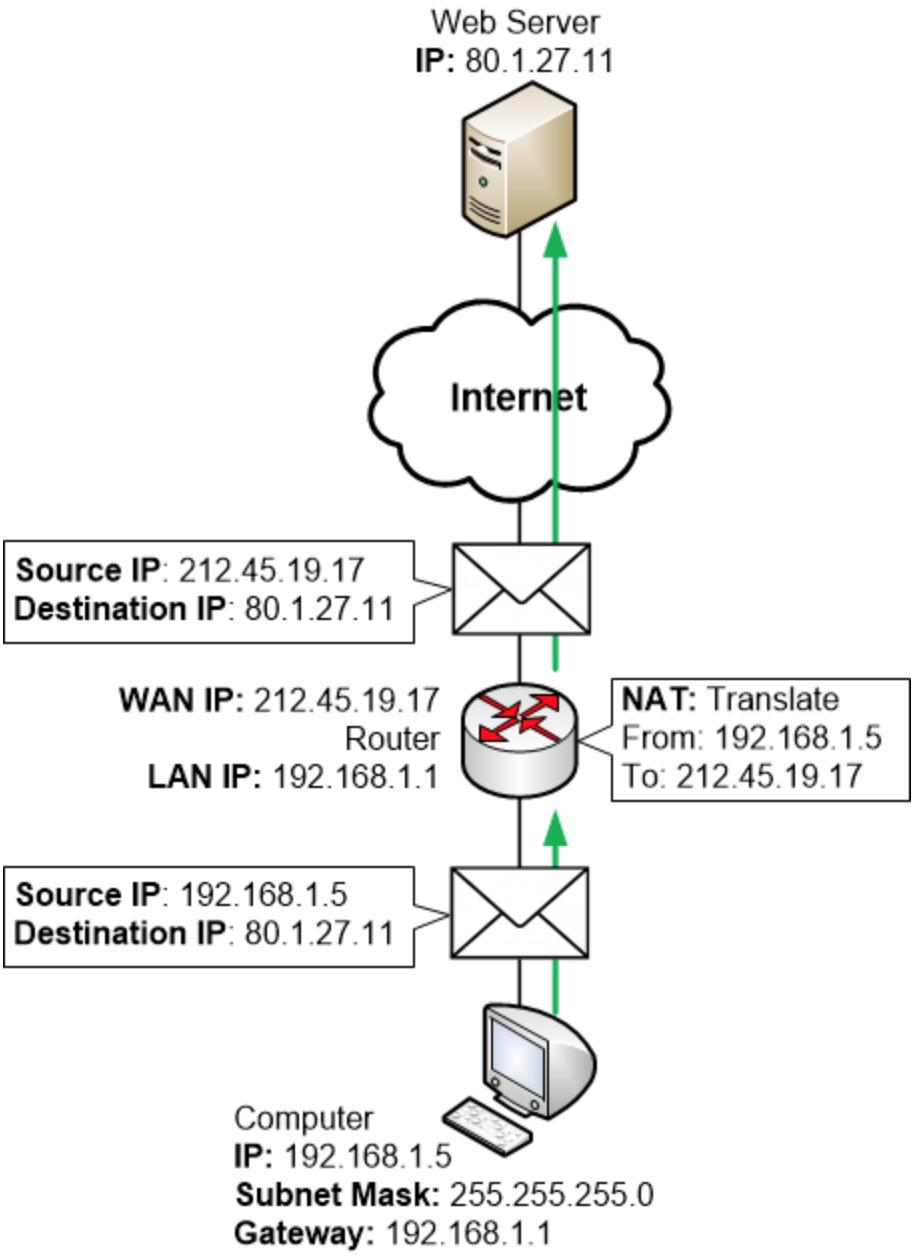

When the router sends the packet on to the Internet the packet will appear to come from the home router's public IP address. From the perspective of the ISPs and the web server, the packet is coming from the public IP address of the home router. When the web server replies back to the computer it will send its reply back to the public IP address of the home router, and the ISP finds its way back there without any trouble.

If several clients are operating at the same time on the LAN then it is both possible and likely that several computers want to browse at the same time to the Internet. Then the home router will keep track of what traffic that belongs to which computer. This lets the router know which inside LAN computer that the returning reply traffic should be sent to.

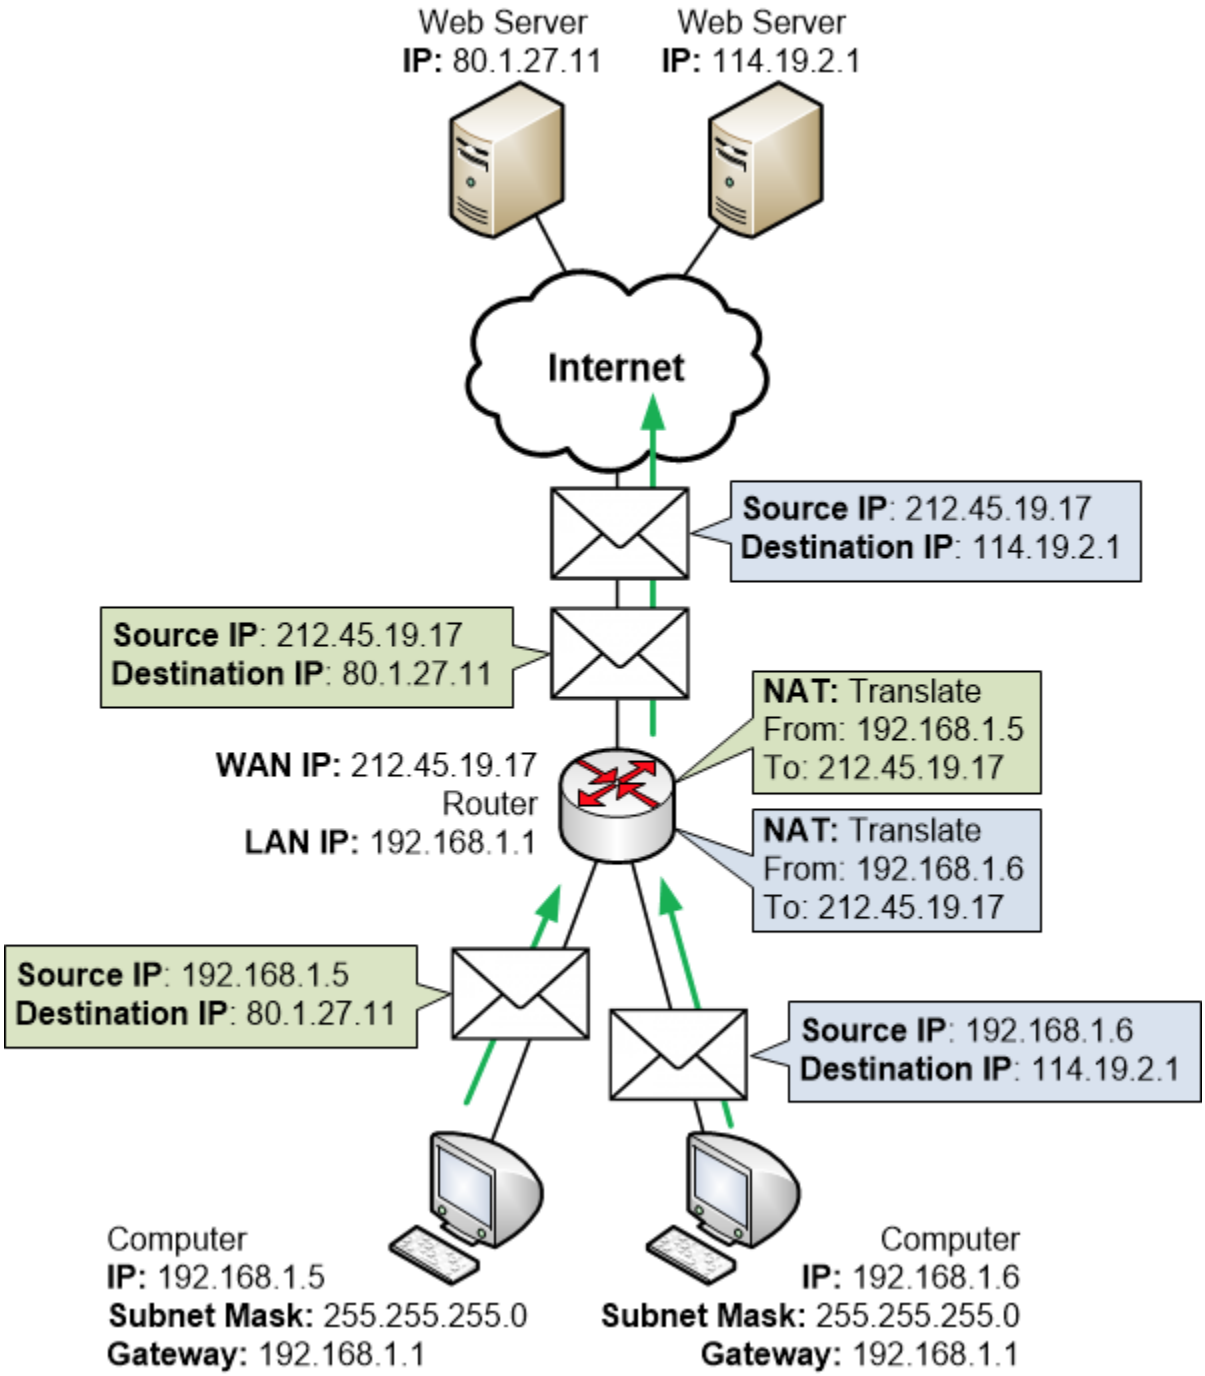

This type of NAT or address translation is often called *Hide NAT* because you "hide" your LAN computers behind a shared public IP address. It is a function that all home routers have built in and which is enabled right from the start. It is also very rarely something that you need to care much about because it simply works.

However as you can probably tell by now there are a lot of things going on in computer networks that most people don't know about, and that you might have to learn about if you want to make any changes to your home network.

Ports - addresses for programs and services

IP addresses are addresses that are assigned to computers and which can be compared to street addresses. If you want to send a letter to somebody then you would print down your message on a piece of paper. Then you would put the paper in an envelope, put the address of the receiver on the envelope and finally send it off. If a program on a computer wants to talk with a server, then the message is prepared by the program, and then the OS will put the message in a packet and send the packet to the IP address of the server.

But several people could share a single street address. Maybe a whole family lives in the same house. Often you don't want to address the whole household with your letter but rather a single person living in that household. So a letter is normally also addressed to one specific person living at the address.

In a similar fashion, a single server can run multiple services or programs at the same time. For example, the server could run a Web Service, a DNS Service, an FTP Service and many other services. All of those services could be running simultaneously. And since they are running on a single server they also share the single IP address that the server is configured with.

So the server must have a way to know which service that each packet is meant for so that it can look at incoming data traffic and hand over the traffic to the correct service.

The solution that is used in network communication is something called *Ports*. These ports are used to give addresses to different services.

Additional Info

A ***service*** is really just a program like any other, except that a *service* is usually a program that is constantly running in the background on the computer. There are many examples of services.

Some are more obvious, such as a Web Server program that is serving Web Pages to people who are browsing in their Web Browser to the computer.

Others are more hidden and might be included in other programs or even the OS that you have installed. For example, Windows is running a Windows Update service in the background, which will look for new updates for your computer every now and then, and might prompt you to install any updates.

A computer can manage just fine with a single IP address, but it can listen on many ports simultaneously. Each program that wants to be able to receive network connections and data packets will require its own unique port.

The programs and services that you start on a computer will tell the OS (Windows, Mac OS X etc) which ports that the service would like to listen to. The OS will then start listening for traffic to that port, and if any such traffic arrives at the computer then the OS will forward that traffic to the correct program.

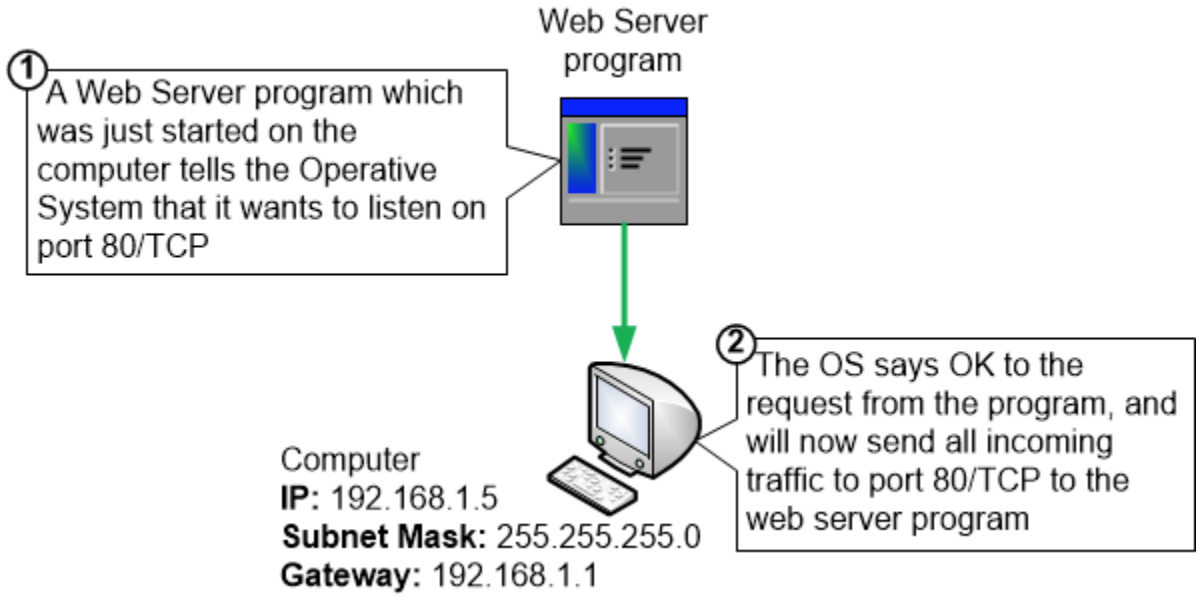

There are Internet standards that determine which ports that common programs and services should use. This makes it a bit easier to know which ports that different services are probably using. A web server for example almost always listens for traffic on port 80, and a secure Web Server listens on port 443.

Here is a table showing a few examples of Ports that some common programs can listen to:

Service	Port and Protocol
FTP server	21/TCP
DNS server	53/UDP
HTTP Web server	80/TCP
HTTPS Web server	443/TCP

The table above also displays two different so-called *Protocols* that different services use, TCP and UDP. Those protocols will be discussed in the next section.

UDP and TCP - two ways of sending traffic

When a computer is sending out traffic it needs to send the data packets from its own IP address (source address) to a destination IP address. So IP will handle the addresses of different devices on a network.

But different programs and services have very different requirements regarding how they prefer that traffic should be sent over the network. To some programs, it is extremely important that not a single data packet is lost, and that no packets are received in the wrong order. Other programs might not care if some errors or packet losses occur. Instead, they might prefer that the traffic is just simply being sent as quickly as possible.

This is where TCP and UDP come into play. They are *Transport Protocols* that govern how traffic is sent over the network.

Additional Info

A protocol is a compilation of information regarding decisions on how something should be performed. One obvious example of a protocol is a simple meeting protocol that describes what was discussed during the meeting and which decisions that were made.

Networks and computers use loads of different protocols that govern how communication should be handled, how data should be sent and so on. Here are two examples:

- IP addresses and how network traffic should be addressed with IP addresses is governed by the *IP protocol*. IP actually stands for *Internet Protocol*
- TCP and UDP are protocols governing how traffic should be sent

The protocols are produced by different standard bodies where experts from many companies and organisations are working together to agree on how the communication should function. This is to guarantee that equipment such as routers from different vendors can work together when you interconnect them.

Without these standards, no equipment from any vendor would be able to work together with equipment from any other vendor.

TCP is absolutely packed with functions that make sure that traffic will arrive in the correct order, that no packets are lost, that any lost packets are automatically sent again, and so on.

UDP, on the other hand, is designed with speed in mind. It has very few safeguards and it doesn't care if packets are lost or arrive in the wrong order. UDP is made to be fast and simple without any extra controls.

TCP	UDP

Keeps track of lost packets. Makes sure that lost packets are re-sent	Doesn't keep track of lost packets
Adds sequence numbers to packets and reorders any packets that arrive in the wrong order	Doesn't care about packet arrival order
Slower, because of all added additional functionality	Faster, because it lacks any extra features
Requires more computer resources, because the OS needs to keep track of ongoing communication sessions and manage them on a much deeper level	Requires less computer resources
Examples of programs and services that use TCP: - HTTP - HTTPS - FTP - Many computer games	Examples of programs and services that use UDP: - DNS - IP telephony - DHCP - Many computer games

It is the person who creates a computer program that chooses which transport protocol that the program should use when it communicates. The choice is made based on how the program should act and which requirements the program will have for its network communication.
Many programs might use both UDP and TCP for different types of traffic. For example, a lot of online computer games could use TCP for player logins and a lot of other features but might use UDP to transfer live continuous events in the game world between the player and the servers.

Most programs that you use on a daily basis will utilise TCP. This is because the program can then rely on TCP to always deliver all packets correctly. With UDP, the programmer who is creating the program must make a lot of decisions about how the program should detect and handle lost packet and packets that arrive out of order.

Of course, there are a lot of programs and services that prefer UDP. DNS is such an example. Each computer needs DNS to be able to browse the Internet and make domain name lookups. After all, that is how computers find out which IP addresses that different domain names correspond to.
One reason for why DNS uses UDP is that it doesn't matter much if a DNS query disappears every now and then. The computer will notice that it doesn't get any DNS reply, and if no reply arrives then the computer will simply send out another DNS query. The only effect is a slight delay in the communication, but once the computer gets the DNS response then the communication can proceed.

Another example where UDP is commonly used is IP telephony, also known as VoIP (Voice over IP), which means making phone calls over the computer network instead of placing calls

over the regular telephony network. When you speak with somebody over IP telephony the phone or computer will transform speech into data packets that are sent over the network to the person you are talking to.

Humans don't have any particular issues with understanding what somebody else says even if a few milliseconds of speech were to disappear here and there. So even if a data packet would be lost every now and then we can still understand what the other person is saying.

If however TCP would be used for IP telephony and a data packet would disappear, then TCP would make sure that every other data packet that contains the speech is put on hold in a queue while the lost packet is re-sent. So the whole conversation is put on hold while TCP is waiting for any lost packets to be retransmitted. Imagine this happening over and over throughout a conversation.

For that reason, it is better for a VoIP call if a few packets here and there are lost (UDP) than if you would have to wait for any lost packets to be re-sent and arrive in the correct order (TCP).

So because different programs have different communication needs and requirements, the various programs will use either TCP or UDP when they want to send out network traffic.

Back to the ports

Let's get back to the Ports again now that we have mentioned what TCP and UDP are and how they are different from each other.

The ports that are used are not a part of the Internet Protocol or IP address. Instead, they are a part of the TCP and UDP transport protocols. So any data traffic can be said to be sent to an IP address - meaning which computer or server that the traffic is sent to, and to a Port - meaning which service or program that the traffic is sent to.

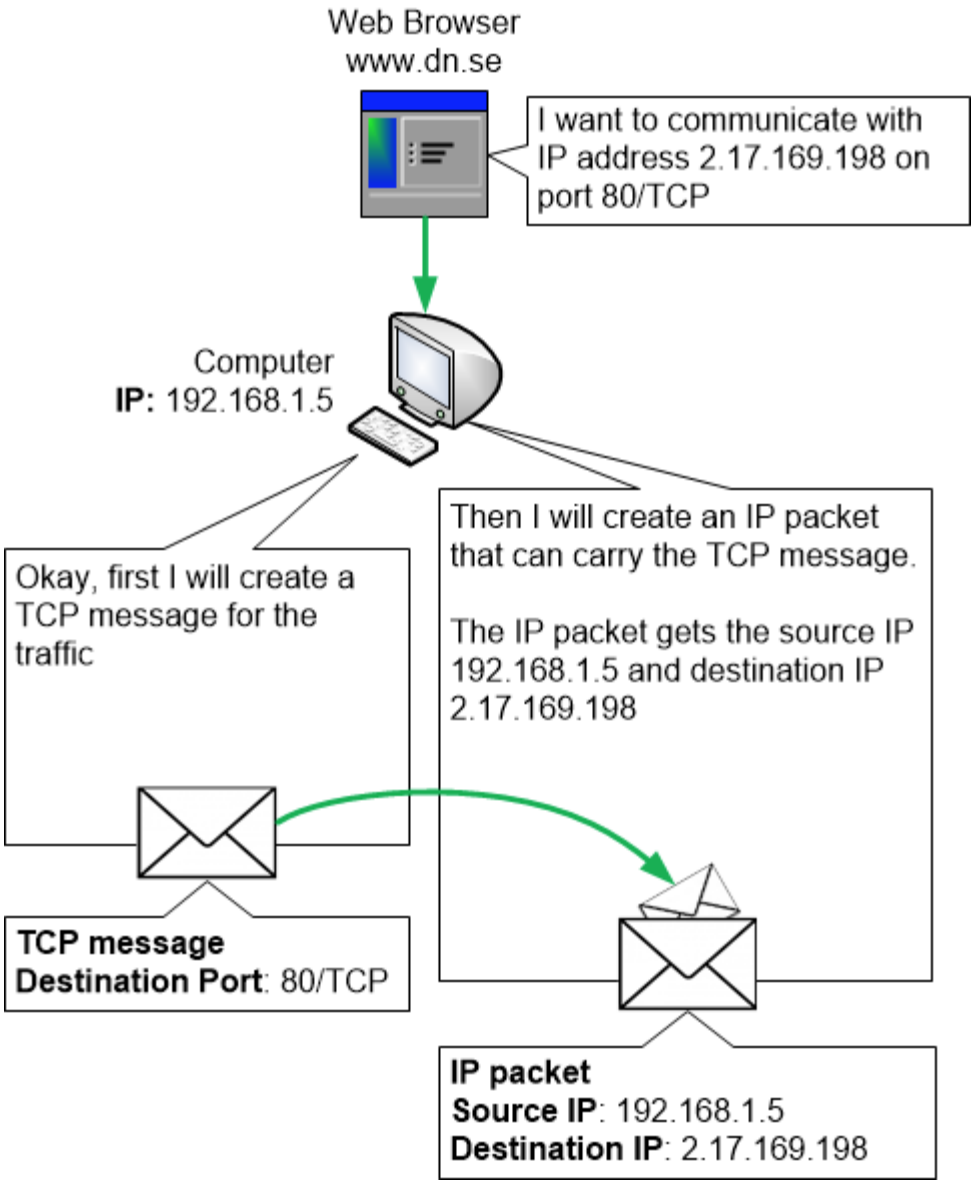

When the computer has some data that it wants to transmit it puts the data inside a UDP or TCP message, depending on which protocol that the application wants to use. The computer enters the Port addresses into the UDP or TCP message.
Then the UDP or TCP message is put inside of an IP packet, which is in turn addressed by the computer using IP addresses.

Since applications always listen to a specific Port but also a specific Transport Protocol (UDP or TCP), then you always have to know not just which port that is used, but also which Transport Protocol that the program uses.

A lot of people simplify by saying for example “Program X listens on port 1337”. And even though that statement is not exactly incorrect it also doesn’t provide the full picture.

Instead, you would have to say “Program X listens on port 1337 TCP” or “The program listens on port 1337 UDP”. That extra piece of information becomes important when you want to do a Port Forward, or if you were to permit the traffic through a firewall. Of course, you could create a Port Forward for both UDP and TCP just to be sure, but sometimes this could cause problems. Also, it is always best practice in computer networks not to permit more traffic from the Internet than necessary to reach your internal LAN network.

Specialisation: Ports, TCP and UDP in depth

This section goes a bit deeper into how network communication functions regarding Ports.
If you want to understand how Address Translation works in depth then you have to understand a few more things about ports than what we previously discussed.

IP packets always have a source IP address (who sent the package) and a destination IP address (who is the recipient of the package). The same thing goes for the Transport Protocols which also require both a source and a destination port.

Each UDP or TCP message has a destination port, meaning which service or program that the message is intended for. Maybe the message is intended for a Web Server listening on port 80/TCP, or perhaps a DNS Server listening on port 53/UDP.

But all UDP and TCP messages also have a source port. The source port is most often randomly picked by the computer. The source port can be used by routers and firewalls in the network to distinguish between different communication flows or sessions since each session will have a different random source port associated with it.

For example, let’s say that you open two web browsers at the same time on your computer. Then you browse simultaneously from both web browser to the same web page on the Internet. You will now have two sessions from your computer’s IP address to the IP address of the web server. Both sessions are also going to the same destination Port, 80/TCP.
There will be two replies coming back from the Web server, one reply for each web browser on your computer. But how will your computer know which browser that should receive which reply?

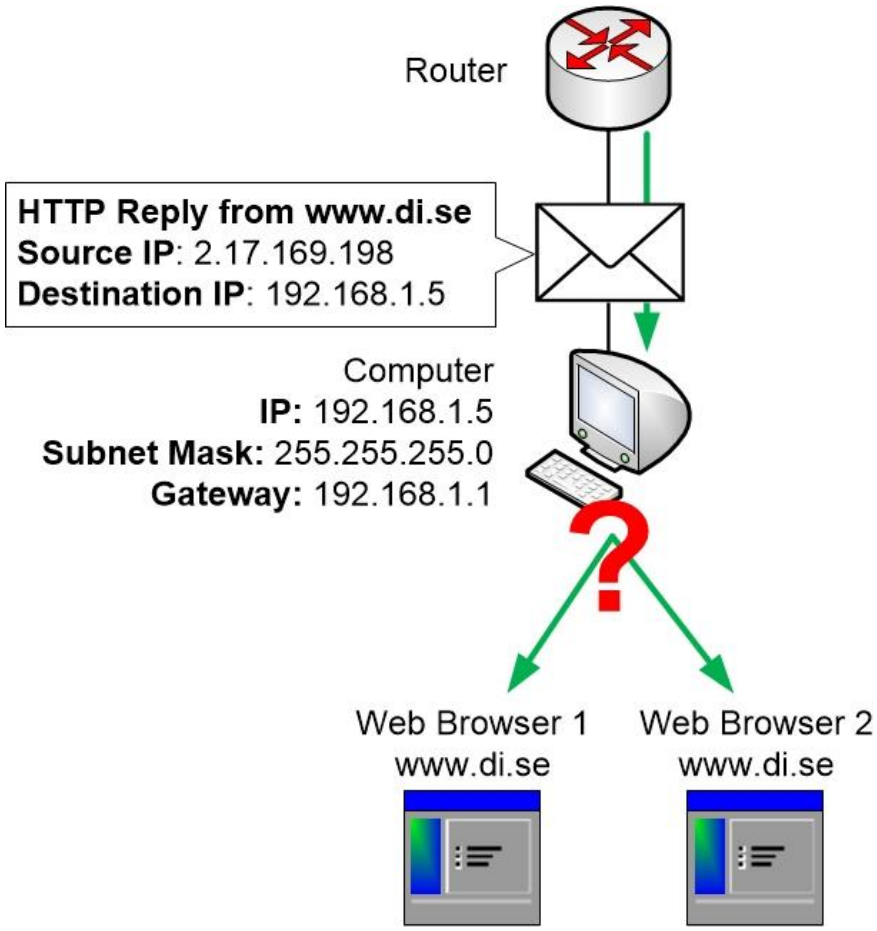

The answer is in the random source port that the computer picks for each traffic flow. An ongoing communication between two devices is called a session. The OS of the computer will remember some details for each session, including which program that was involved, what IP addresses that the communication is going between, and what source and destination port that is used. In this case, two web browsers have a session each.

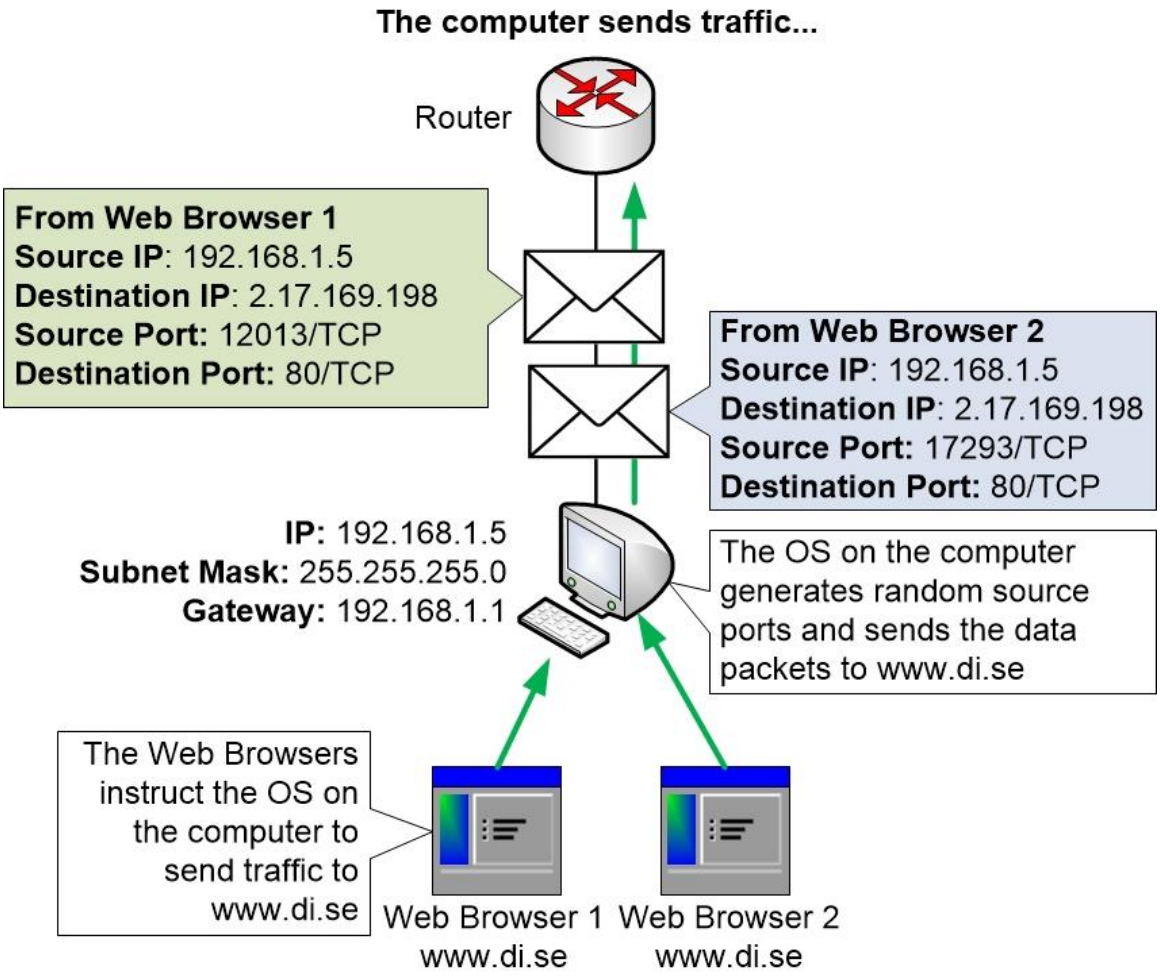

When the replies are coming back from the web server the computer can look at those ports and compare them with the session table in memory to see which web browser that should receive which reply.

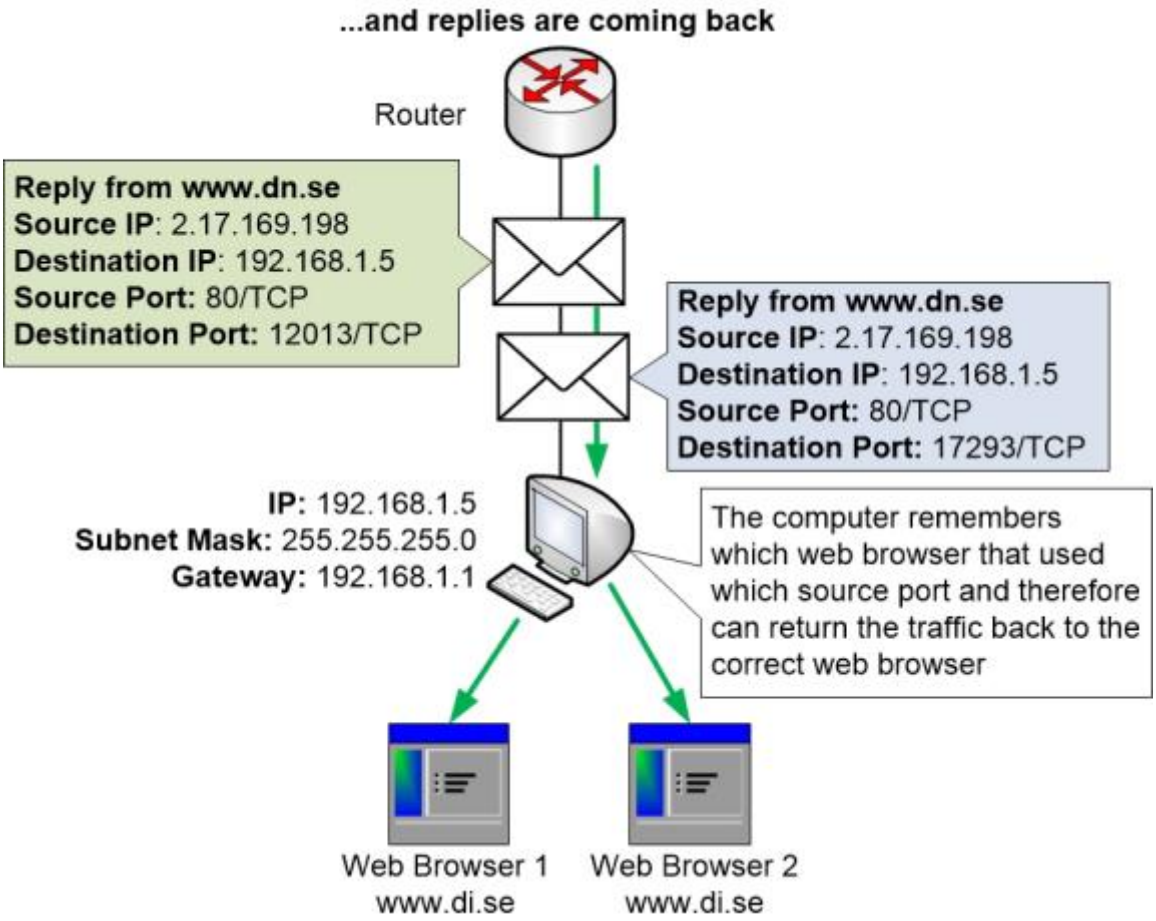

Also note in the two pictures above that the source and destination IP addresses and Ports swap places for the return traffic. This is because the return traffic is coming from the web server and is going to your computer.

Your computer is browsing from a random TCP source port to the destination port 80/TCP of the web server. The reply is coming from port 80/TCP on the web server and is going to the random destination port that your computer picked.

Specialisation: Address Translation in depth

This section is going through almost the complete picture of how Address Translation works and is more advanced than many of the other sections. The main concept of Address Translation as it has been discussed in other sections are still valid, but by adding more knowledge about source and destination ports it becomes possible to delve deeper into Address Translation.

As long as a computer on your home network is initialising the traffic (in other words your computer is starting up the communication by sending the first message of the communication) then your home router can keep track of the replies coming back and send those replies to the correct computer on your home network.

For example, if you open a web browser on your computer which is connected to your home network, and you browse to a web page on the Internet, then your computer is initialising the communication by sending out a message to a web server on the Internet. The web server replies, and your home router makes sure that the answer is sent to the correct internal computer.

But what happens if two computers browse simultaneously to the same web page on the Internet? How can the router know when the replies come back which computer on the LAN that it should send each reply to?

In practice, the router keeps track not only of which IP addresses that the traffic is going to, but also which source and destination ports that the traffic is using. The source ports are randomised by the computer, which means that each session thus will have a unique combination of IP addresses and ports that the router can remember and associate with each ongoing traffic session.

By combining the information about IP addresses and ports that the router has about each communication flow it can distinguish different sessions from each other. So when the replies come back the router can determine just by looking at the IP addresses and ports that are in use which session that the traffic belongs to. The router can then check a *session table* where it saves information about these ongoing flows to see how the traffic was Address Translated, and to which computer on the internal LAN that the replies should be sent to.

This is how the complete flow of traffic through the router would appear if we also bring in the ports and the address translation.

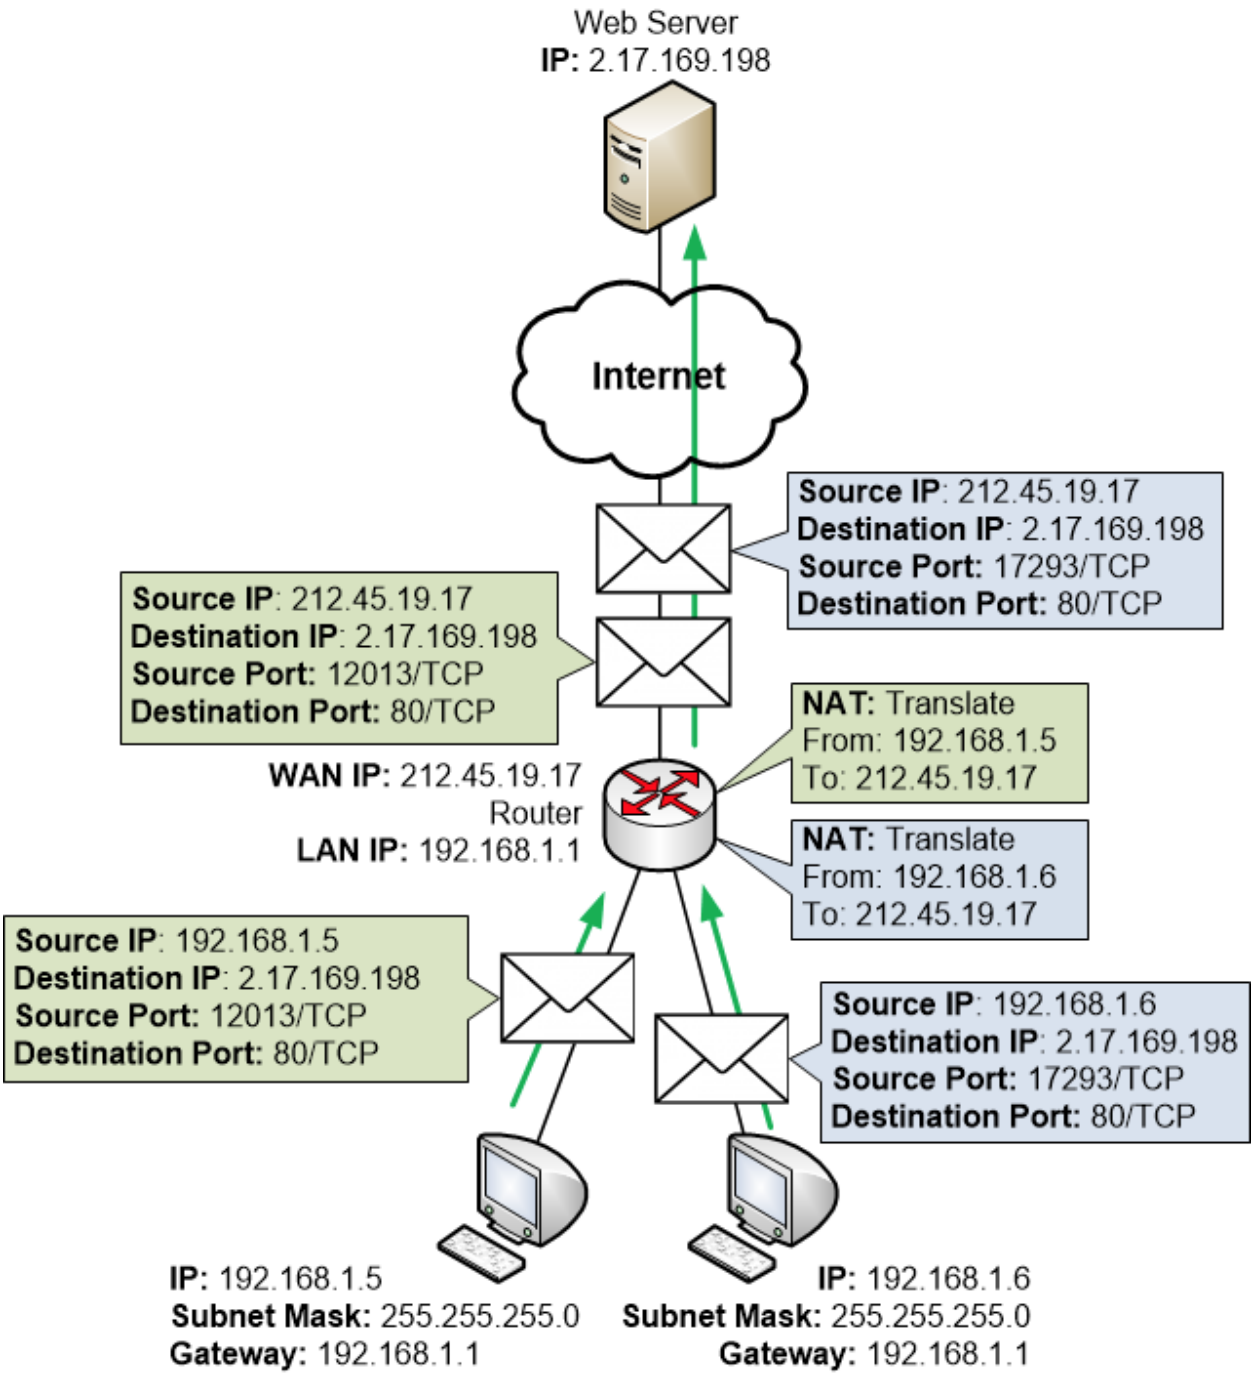

For each traffic flow or session that is going through the router, the router will remember the following information and save it in a table in memory:

- From what IP address on the LAN is the traffic coming
- To what IP address on the Internet is the traffic going
- From what port is the traffic coming
- To what port is the traffic going
- How did the router Address Translate that particular traffic

This is what the table that the router builds up would look like based on the information from the previous picture:

Source IP	Destination IP	Source Port	Destination Port	Source IP translated to:
192.168.1.5	2.17.169.198	12013/TCP	80/TCP	212.45.19.17
192.168.1.6	2.17.169.198	17293/TCP	80/TCP	212.45.19.17

When replies are coming back the router can look in its table and see what session that the traffic belongs to. Then it knows exactly how it should handle the replies, how the replies should be reverse address translated and to which computer it should send the traffic.

Port Forward

Earlier sections in this guide brought up NAT, address translation, and how it lets multiple computers on a LAN share the home router's public IP address.

If however the opposite scenario is taking place - something on the Internet wants to initialise the communication toward a computer on the internal LAN - then that would be impossible without some type of special solution such as a Port Forward. Remember that all computers on your LAN are "hidden" behind the public IP address of your home router.

Let's say we install a Web Server on our home network, and we want people on the Internet to be able to browse to our Web Server. We also have several other computers on our home network, and we want both our computers and the Web Server to share the same public IP address that the router has on its outside.

When somebody is browsing from the Internet to the public IP address of the home router, how should that router know that it should pass the traffic to the Web Server?

The router cannot do this automatically. Instead, you would have to configure the router to do what we want it to do in this particular scenario.

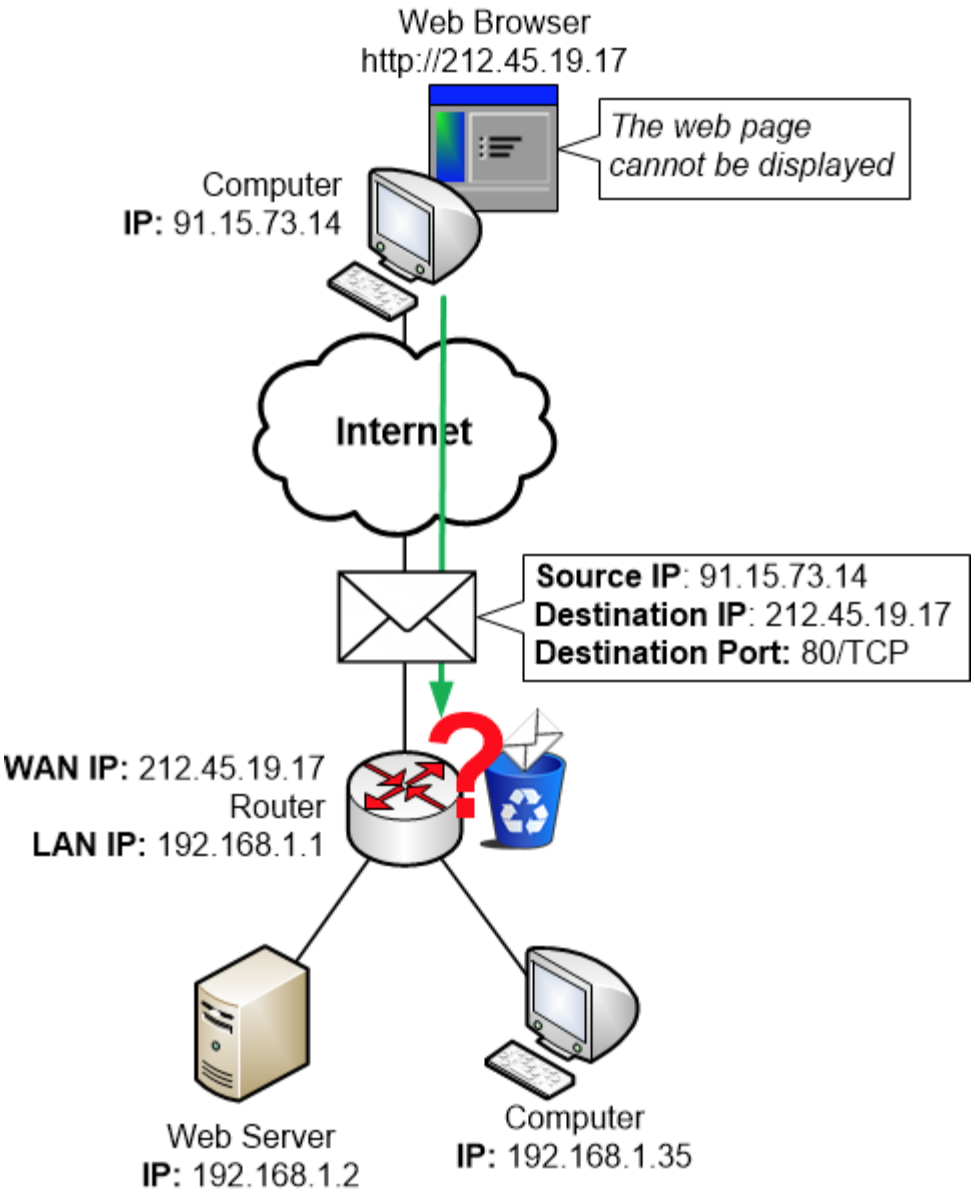

If you do not configure the router for this scenario, then if somebody on the Internet is browsing to the router's public IP address, the router wouldn't know what to do with the traffic. The router cannot find any matching pre-existing session in its memory, so the router doesn't have any other choice but to discard the traffic.

The result is that the person on the Internet who was trying to browse to our Web Server simply doesn't get any replies back. Their web browser will eventually time out and display an error information message.

The solution to this problem is to create a Port Forward. You as the administrator of the home router will have to investigate which ports that the web server on the inside LAN wants to listen

to. Then you make sure that any traffic from the Internet that is sent to those ports are forwarded in the router to the correct device.

Luckily these days there is an easier and completely automatic way of doing Port Forwards. It is handled by a protocol called UPnP which is described in a section of its own within this guide. But worth noting is that UPnP won't always work correctly.

So if you would like to understand the theory behind how Port Forwards work then this section is for you. If however you just want the simplicity then start by looking at the section about UPnP and come back to this section only if UPnP did not work in your case.

Understanding Port Forwards

Port Forwards are among the most advanced things that a typical home network owner will deal with. Sometimes depending on your router it might be simple to perform the actual configuration in your router. But the underlying theory behind Port Forwards is more complex.

First of all, you have to know why Port Forwards would sometimes be needed. You must also have heard about Port Forwards to even have a clue about any problem that could be associated with running a server on your home network.

Once you have gotten that far it is often rather simple for an experienced computer user to look for guides and try to follow those guides to perform a Port Forward in their home router. But the problem is that many guides are simply badly written, and might even be written by people who themselves do not have a clue about how Port Forwards actually work.

Also, if you don't know the theory behind the configuration change that you are performing then it gets really difficult to try to troubleshoot why it is not working if something doesn't go according to plan.

All programs or services that you can connect to always listen on a specific port. A Web Server for example always listens for TCP traffic on port 80. So if we install a Web Server on our home network then we know that it will listen for traffic on port 80/TCP by default.

That is all the information we need to set up a Port Forward rule in our home router. We can configure the router in such a way that if anybody browses to the public IP address of the router on port 80/TCP, then the router will forward that traffic to our internal LAN Web Server. That way we can "publish" our internal Web Server to the public Internet.

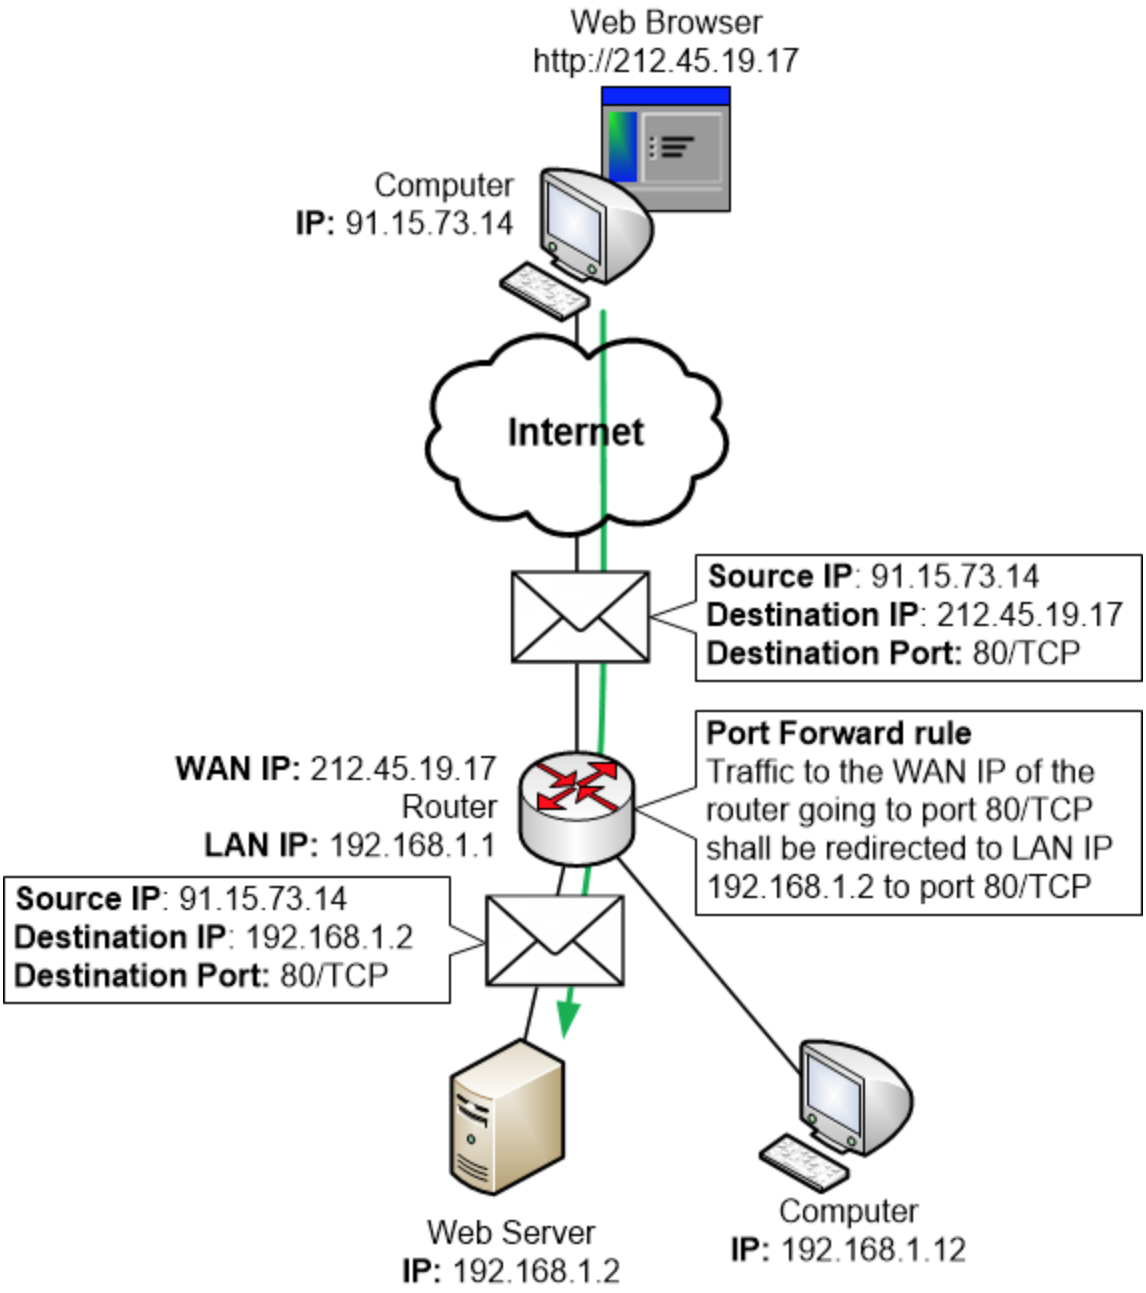

As you can see in the picture above an Address Translation is performed by the router on the IP packets as they pass through the home router. The destination IP address is translated in the IP packet.
In fact, Port Forwarding is actually just a special type of NAT or Address Translation. But since it is used for a specific purpose it has gotten its own name, “Port Forward”.

Many computer games also require Port Forwards to function. This is often true for multiplayer games where one player can start a Game Server within the game and the other players connect to the Game Server. Since the Game Server is started on a computer which sits on a

local LAN behind a home router, it might be necessary to configure that home router with Port Forward rules to make it work properly. Otherwise, when a player on the Internet wants to connect to the game server the home router doesn't know where it should send the traffic.

Different games will require different Port Forward rules. Most games will actually require multiple Port Forward rules before they start to work, and there might be a mix of both UDP and TCP ports that must be forwarded to the computer which is running the game. Sometimes a whole range of ports must be forwarded.

To figure out which ports that must be Port Forwarded you have to either google for the game's name and the keywords "port forward", or you could try to find the information on the homepage of the game.

In all honesty, the game publishers are often absolutely incompetent regarding Port Forwards. They commonly list far too many ports that they tell you must be forwarded, and they are often confused themselves as to which ports are actually required to run the game. So to be on the safe side, the publishers often list loads of ports in their Port Forward help articles on their websites.

The most common mistake they make is that they cannot distinguish between outgoing traffic (from the gaming computer to the Internet) and incoming traffic (from the Internet to the gaming computer). The result is that they might list all ports in both directions and tell you to forward all of them.

Unfortunately, there is no general rule that can be applied to the problem of incompetent game publishers. You could try to enable UPnP if possible, but if that doesn't work then you might have to search the Internet to find others who have solved the puzzle of making a certain game work with Port Forwards and copy what they did.

Example game requirements

Here is one example of a game which has some Port Forward requirements listed. This particular game is *Titanfall* for PC:

- UDP port 8125
- TCP port range 25000 - 25099
- TCP port range 30000 - 30099
- UDP port range 25000 - 25099
- UDP port range 30000 - 30099

The game publisher also lists port 80/TCP and port 443/TCP. However, ports 80/TCP and 443/TCP (Web Server ports) should never have to be Port Forwarded to your computer unless you are running a Web Server on your computer. Your game is not a Web Server. So you can probably safely assume that you do not have to Port Forward port 80/TCP or 443/TCP to an internal computer unless you actually want to run a Web Server on that computer.

What the game *actually* uses those two ports for is to let the game connect to the publisher's Web Servers on the Internet to download information and updates. In other words, they are only required for outgoing traffic to the Internet.

Configuring Port Forwards

How you configure a Port Forward depends completely on what router you have. Often there is a setting available in the router called “Port Forward” or something similar to it. But no matter what the exact name of the function is on your particular type of home router, the main idea is that you have to first pick which ports that should be forwarded to an internal computer, and then you have to pick which internal computer or which internal IP address that the traffic should be forwarded to.

Remember when we discussed earlier in the guide how you can configure a computer with a static IP address? When you need to set up a port forward rule to your computer, then it would normally also be wise to configure that computer with a static IP address. Otherwise, your new port forward rule will stop working if your computer changes its IP address later on.

UPnP - automatic Port Forward

UPnP stands for *Universal Plug n Play*. It is a standard for letting computers, routers and other devices that are connected to a network share information with each other about which services they are running and then automatically connect to those services when they have a need for them.
For example, a computer that has a bunch of movies or lots of music stored on itself can announce that to other devices on the network. If you then boot up your Smart TV which is connected to the same network you could then use your UPnP enabled Smart TV to browse those movies, pick one and start playing it on the TV over the LAN network.

Another function of UPnP is that a router could announce to the LAN computers that it has the capability to do automatic Port Forwarding. Once the computers hear about that they can then automatically tell the router to perform any Port Forwards that they require. For example if you boot up a game and the game tells the computer that it wants a couple of Port Forwards to function properly, the computer can then ask the router to perform those Port Forwards.

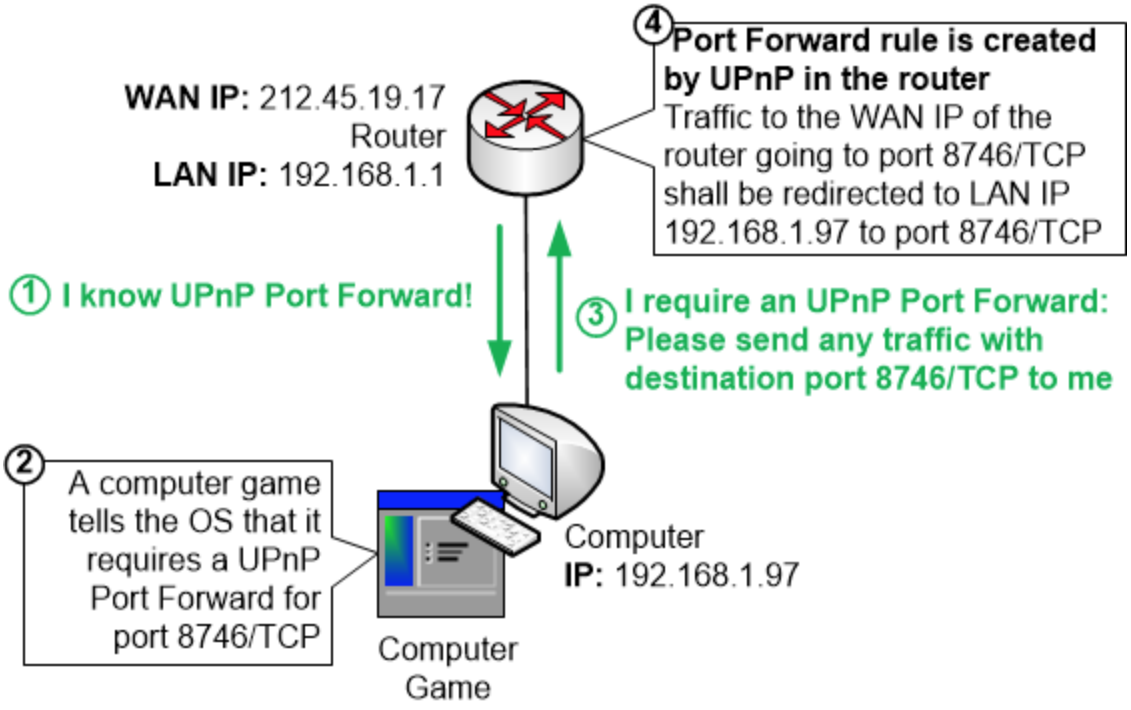

There are a number of prerequisites for UPnP Port Forwards to work at all:

- The router must have support for UPnP Port Forward, and UPnP Port Forward must be enabled in the router
- The computer OS must have UPnP Port Forward support and UPnP Port Forward must be enabled in the computer OS
- The program which requires the Port Forward must have support for UPnP Port Forward so that it can tell the OS about the Port Forwards it want

Additional Info

As you will see below, most devices come with UPnP disabled and you have to manually enable it. The major reason is that UPnP has a negative impact on the overall security of your home network.

If you look at UPnP from another perspective, what it does is that it lets programs on your computer tell your home router to open up holes through its protection mechanisms. Malware and viruses that might be running on your computer are programs too. If UPnP is enabled in your home network then malware could use UPnP to open up holes through your home router's security features at will to let attackers on the Internet straight through into your home network.

Newer home routers often have UPnP Port Forward support, but most likely the functionality is disabled for security reasons. You might have to manually enable it in your router. The function

is often called something like “UPnP Port Forwarding” or just “UPnP” but it is also possible that your router manufacturer has chosen to call it something completely different.

Modern operating systems have support for UPnP, but it is often disabled and must be enabled before it can be used. In Windows, the UPnP function is called Network Discovery.

Finally, not all programs have built in support for UPnP Port Forward. If the program does not have UPnP support then there is no fully automated way of making the port forwards work for that particular program. So the first thing you might want to check is whether or not the program has UPnP support.
Luckily most modern games have UPnP support, but you still have to google it to find out for sure if your particular game has UPnP support or not.

Switching

Switching is the task that is performed by Switches in a network. Whereas Routers are forwarding network traffic based on IP addresses, switches instead forward network traffic based on something called MAC addresses.

The subject of Switching contains a wide scope of information, and most people who start to learn about computer networks think that Switching is the easiest part to learn, since "*it just works*". You connect a couple of computers to a switch and they can then immediately communicate with each other. No configuration required! So most people assume that the Switch and the tasks it performs must be really simple.

The truth is that a switch and the theory behind switching is rather complex. If you want to understand what switching actually is and what a switch really does, then you need to acquire a lot of basic knowledge about how network communication works.

For that reason, the section about Switching is divided into several subsections. We encourage you to go through the sections in the intended order of presentation unless you already know the contents of one or several sections.

A reminder about Home Routers

We would like to remind you that Home Routers consist of several integrated parts, including a Switch.

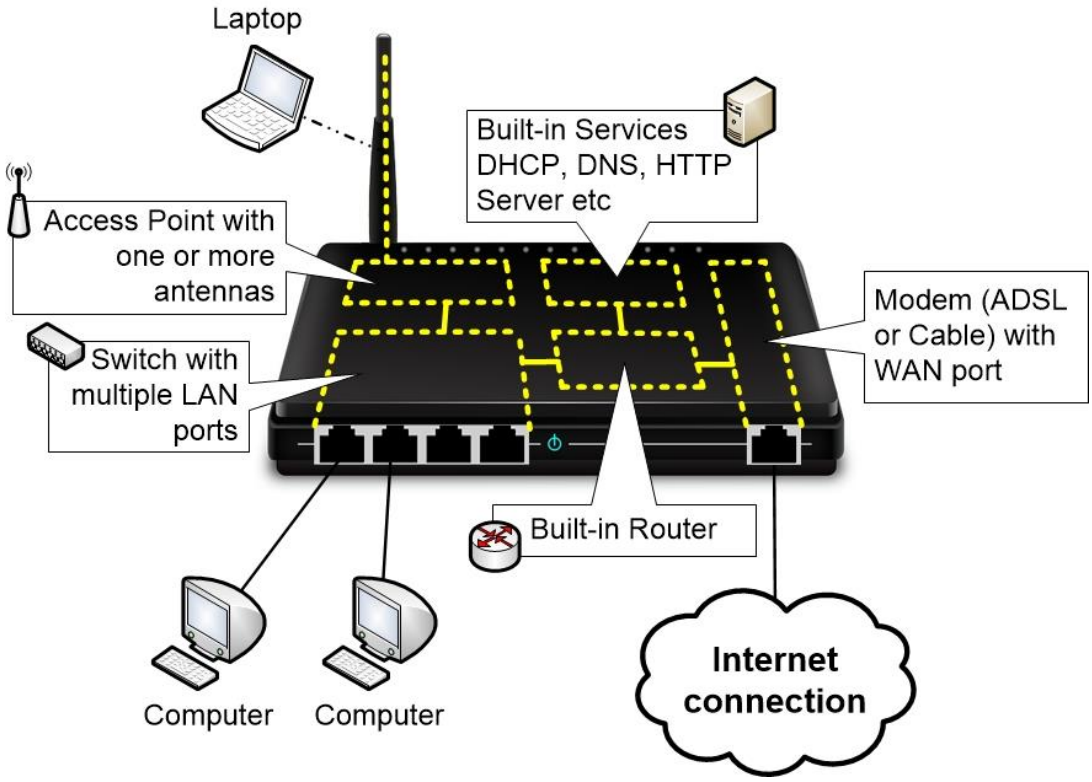

A lot of the pictures and examples in the Switching section showcase computers that are connected to a Home Router. Have in mind that as long as the computers on the LAN network communicate directly with each other the traffic will not be handled by the integrated Router part of the Home Router. The traffic will instead be handled only by the integrated Switch.

Hubs

A hub is a network equipment with several interfaces (also called ports, not to be confused with TCP or UDP ports since these are physical ports where you can connect cables). You can connect computers, printers, routers etc to those ports and they will be able to communicate with each other.

A lot of people confuse Switches and Hubs and mistakenly call their Switches “hubs”. Let’s start by simply saying that switches and hubs work in completely different ways. Hubs are antique legacy devices that no longer belong in any computer network. Switches, on the other hand, are modern day devices.

Why do we even mention hubs in this guide then?

There are still a few hubs to be found out there in the world. Also, hubs serve as a great introduction to switching, because by using hubs in our examples we can showcase several fundamental networking principles that relate to switching. Last but not least, since people sometimes still use the word “hub” when they talk about switches, why not start off by clarifying the differences?

All a hub will ever do is to copy electrical signals that are entering one port to all other ports. So whatever a computer sends into the hub will be copied by the hub to all other devices on every other port.

This means that a hub is completely unintelligent. It doesn’t care about network traffic or addresses at all. All it does is to copy electrical signals. Hubs are more or less never used these days, and you can barely buy them in stores.

One problem with a hub is that many messages that a computer sends out are meant for just one single receiver. But if the computer sends that message into a hub then the hub will copy the message to every other connected computer. So not only does the intended recipient receive the message, all other computers must listen to the message too.

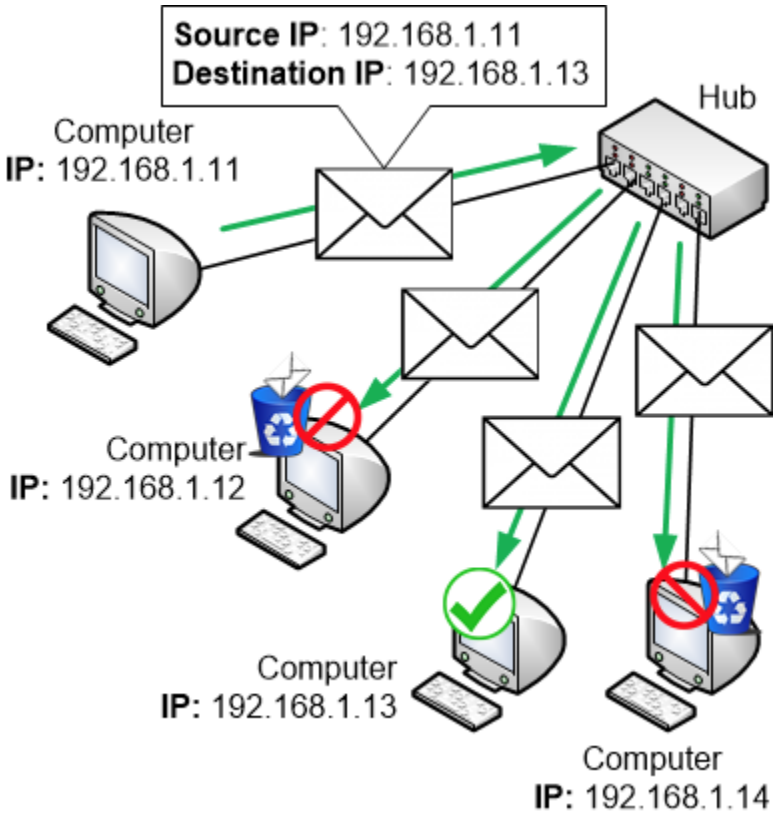

It is then up to each computer to evaluate the destination address of the packet to see if the message was intended for them or not. If the message was intended for another recipient then the computer must throw the packet away. This leads to a lot of unnecessary work for all involved computers and much of the network traffic on the network is thrown away because it ended up at the wrong destination.

Another huge downside of hubs is that they are very limiting to the network traffic. Only one computer can talk at a time through a hub. When one computer is talking every other computer must stay completely silent. This is because the electrical signals of multiple computers that talk simultaneously will mix together in the hub, creating disturbances to the signals so that neither signal can be interpreted.

If two computers talk at the same time a *collision* occurs. When this happens all computers will notice the disturbances and must stop talking for a while before trying again.
This also makes a hub really slow. Only one device can communicate at a time, and when a collision accidentally occurs every device must be silent for a while.

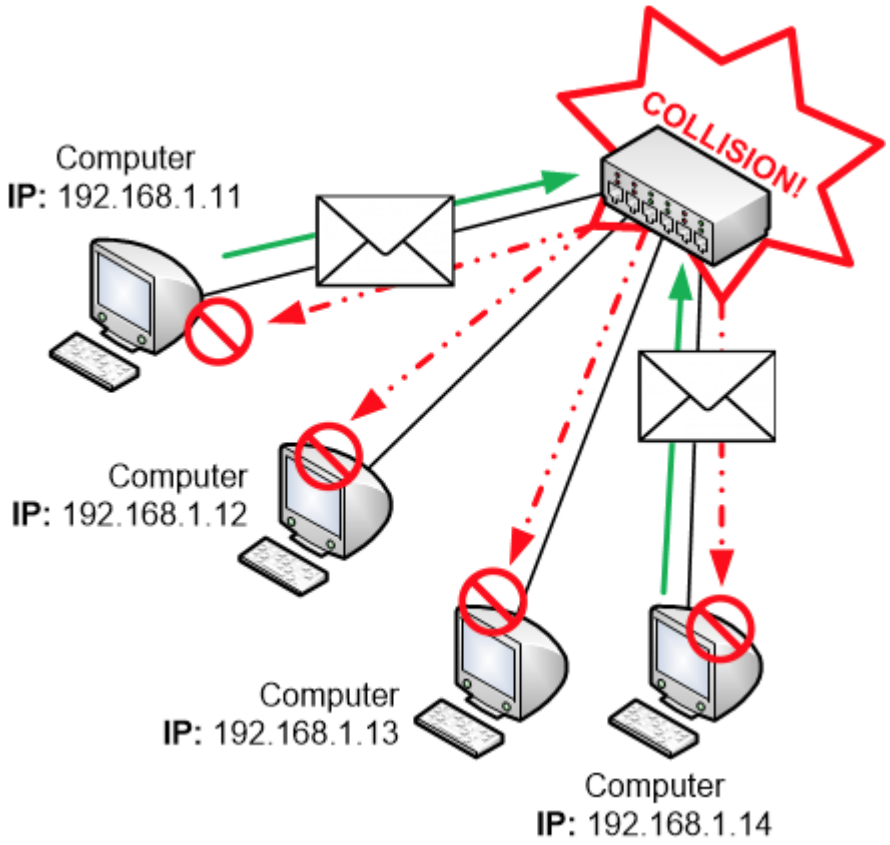

The more computers you connect to a hub the bigger the risk gets for collisions to happen since more involved devices will indirectly compete with each other for the available communication time slots.

It is commonly said that with a hub all computers are sharing the available bandwidth. In practice, the bandwidth situation is even worse than just shared. If four computers are connected to a 100Mbit hub and if all computers want to communicate then each computer will end up getting less than 25Mbit bandwidth. This is because it is necessary to have some quiet time between each communication. Also, any collisions that do happen will interrupt the traffic flow for a while.

Additional Info

It is called “Half Duplex” when only one device at a time can communicate on a network. If more than one device can talk at once it is called Full Duplex.

- A hub is always Half Duplex, so only one device at a time can communicate.
- A switch can handle Full Duplex communication so multiple computers can talk with each other simultaneously through a switch.

If you have a hub at home you should really consider swapping it for a Switch instead. The cost is small but the difference in performance is often huge and will be noticeable right away.

Often the equipment says on the label if it is a hub or a switch. But since Hubs and Switches have basically the same physical appearance you might have to google the model name to find out unless it says right on the equipment what type of equipment it is.

MAC addresses

All equipment that can be connected to computer networks (computers, routers, servers, printers, smartphones and so on) have a MAC address. It is an address which is written into the network interface of the device during manufacturing.

A MAC address consists of 12 hexadecimal characters and could look like this:

- 01:23:45:67:89:ab
- 00:fe:19:2a:73:dc
- 02:0a:95:9d:68:16

Additional Info

Regarding Hexadecimal numbers:

In our decimal system that we use in everyday life each digit can have 10 values, ranging from zero to nine:

- 0, 1, 2, 3, 4, 5, 6, 7, 8, 9

In the hexadecimal system, each digit can have 16 different values. Starting from zero and going up to nine, and then continuing further using letters A through F:

- 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, A, B, C, D, E, F

A has the value 10, B means 11, C means 12, D is 13, E is 14 and F has the value 15.

MAC addresses and IP addresses are two completely different types of addresses, but both are used by computers that communicate with each other. Each time a computer sends out network traffic the traffic has both a source and destination IP address, but it also has a source and destination MAC address.

IP addresses are relevant on a global scale. They hold the final destination of the packet and can tell us which address the packet is originally coming from.
In contrast, MAC addresses are used on a more local scale, and hold information about the next hop destination in the local LAN network.

The difference can be compared to a journey to a final destination (IP address) where you can have multiple legs of the journey, each time stopping temporarily at interim stops (routers with local MAC addresses) before moving on.

An IP packet travelling over the Internet has many such temporary stops along the way in different routers that the packet must pass through. Each router forwards the packet to the next router on the path, until the packet has reached its final destination.

The destination IP address on the IP packet must always stay the same throughout the whole journey, in the same way that the delivery address of a letter that is being sent cannot be changed along the way. The delivery address stays the same until the letter has been delivered no matter where along the delivery path that the letter is currently at.

But even with regular letters, the postal delivery service must use temporary destinations or *next-hop addresses* where the letter is going to be delivered next. The mailman who fetches the letter from the mailbox is not going to deliver the letter straight to the recipient. Instead, the mailman will fetch the letter to a mail sorting office. So even though the letter has a final destination address, the mailman will take it to another more temporary stop along the way where the letter can be sorted for further delivery toward the destination.

This can be repeated multiple times, with the letter passing by several such temporary next hop addresses before it is finally delivered.

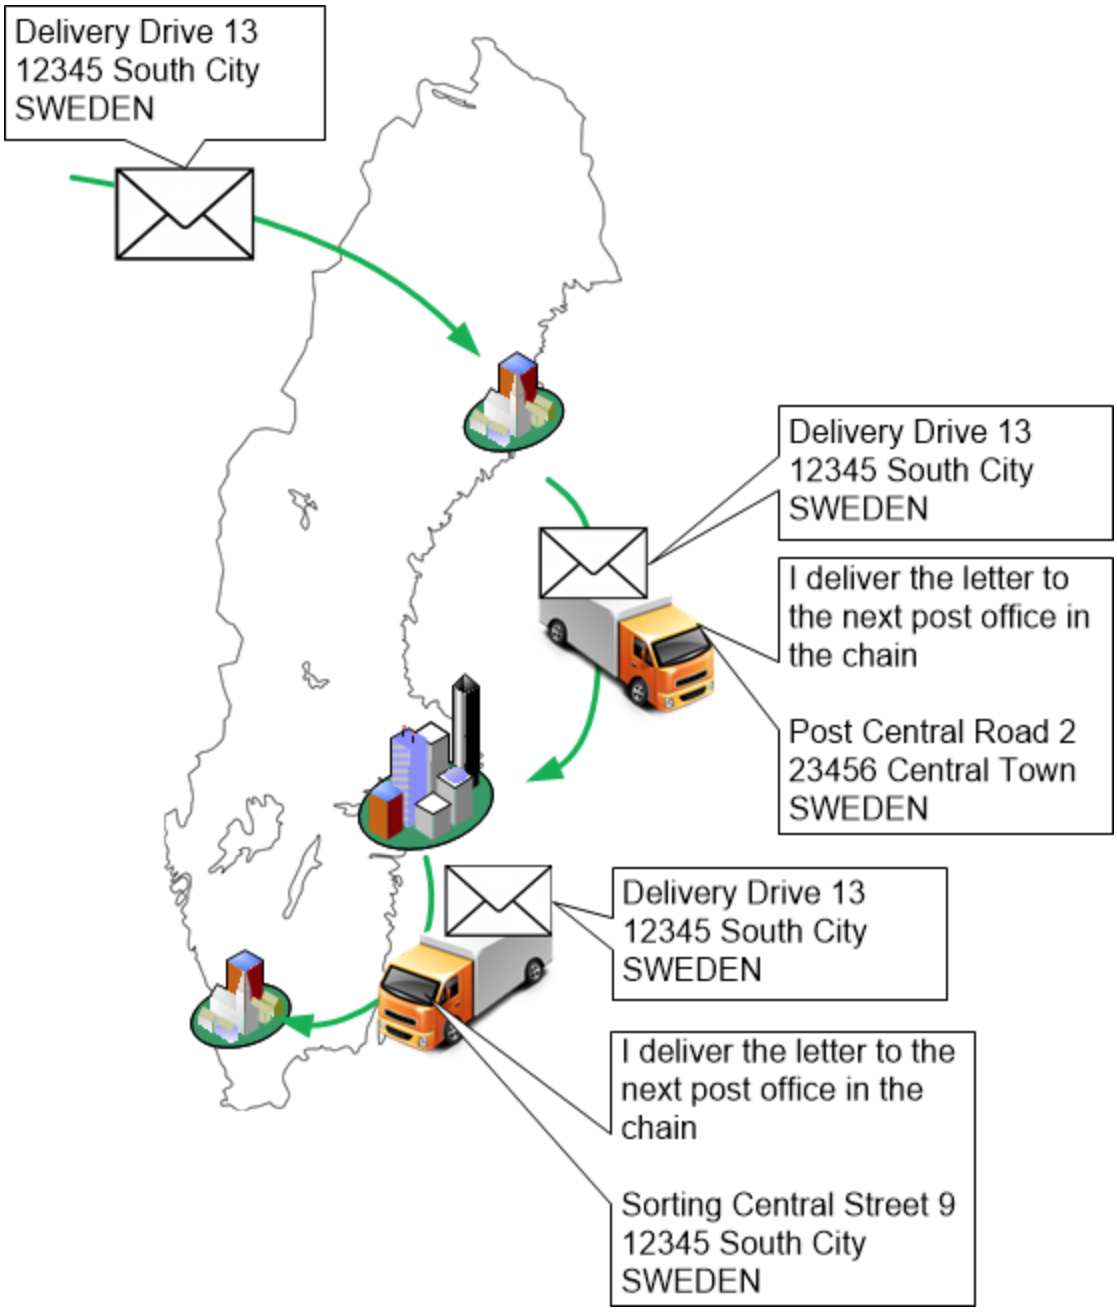

In computer networks, it is the role of the MAC address to act as the next hop address on the local LAN. The mailman is replaced by Switches in the network environment. These switches can be located in a Local LAN, or sit between routers on the Internet.
The switches do not deliver packets based on the destination IP address. Instead, the switches look at the destination MAC address to see where it should send the packet next.

When a computer has a packet to send it knows the destination IP address where it wants to send the traffic. But the computer must also make sure to add a destination MAC address to the

traffic which points to the next hop router. This is how the computer makes sure that the packet will end up at its default gateway.

When the router receives the packet it will look at the destination IP address of the packet. Based on that destination IP address the router will know which way to forward the packet and which the next hop router will be. Then it can strip off the old destination MAC address (which was the router's own MAC address) and replace it with a new destination MAC address that points to the next hop router.

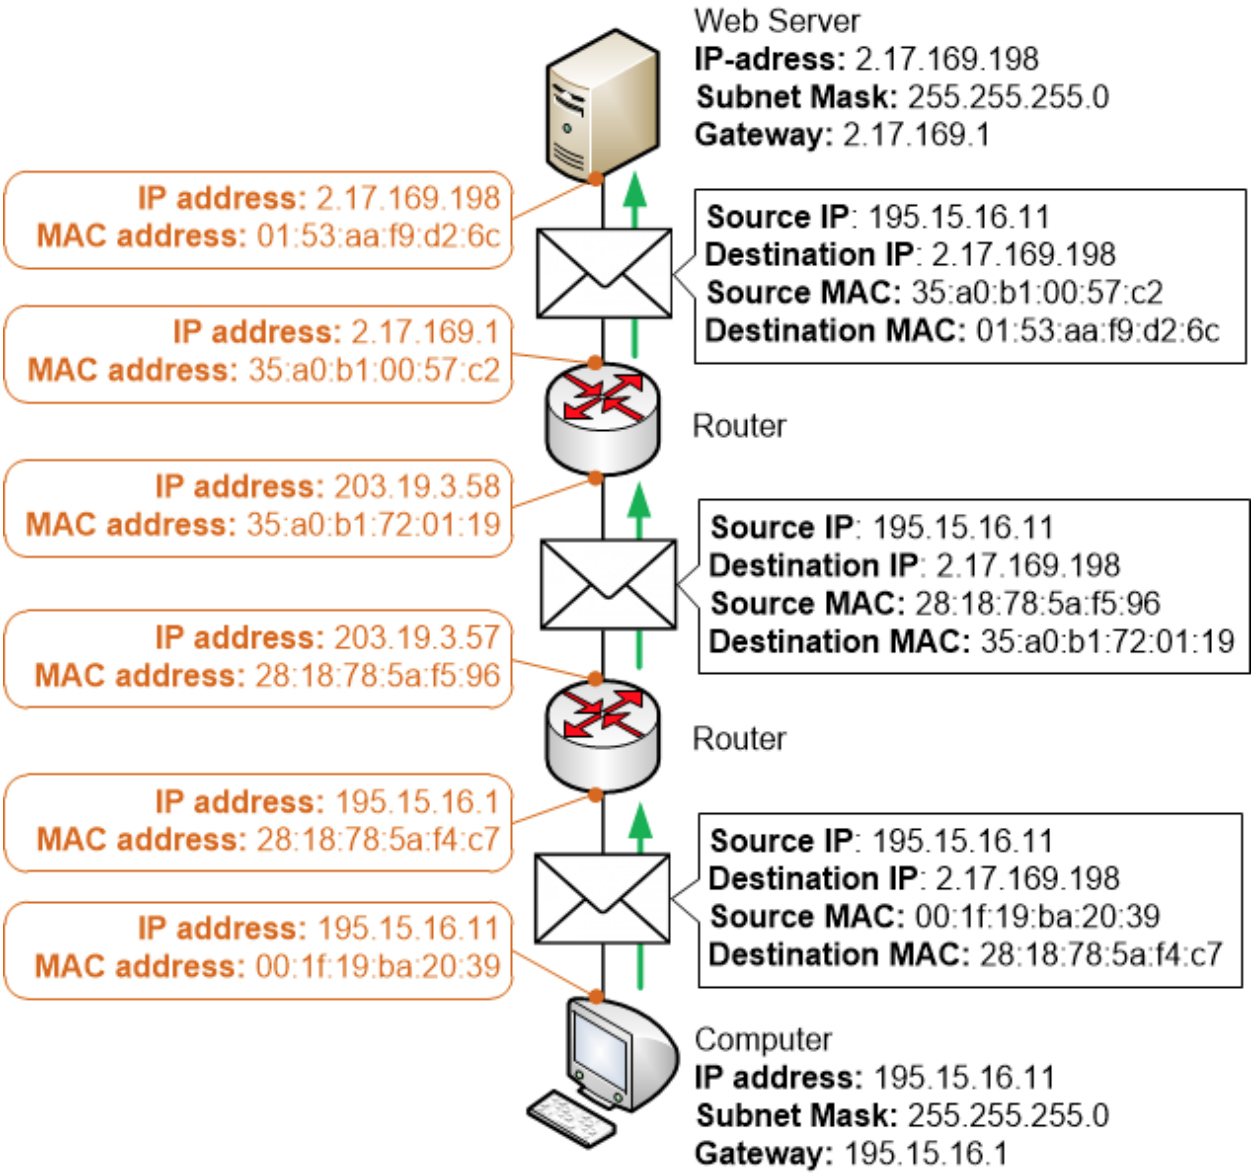

So Source and Destination IP addresses don't change as the packet is transmitted and routed over the Internet. The MAC addresses, however, are being changed for each new local network

that the packet traverses as it is being sent between each pair of routers on its way to the final destination.

It is worth repeating that MAC addresses are always used in combination with IP addresses, even within small LANs where two computers want to communicate directly with each other.

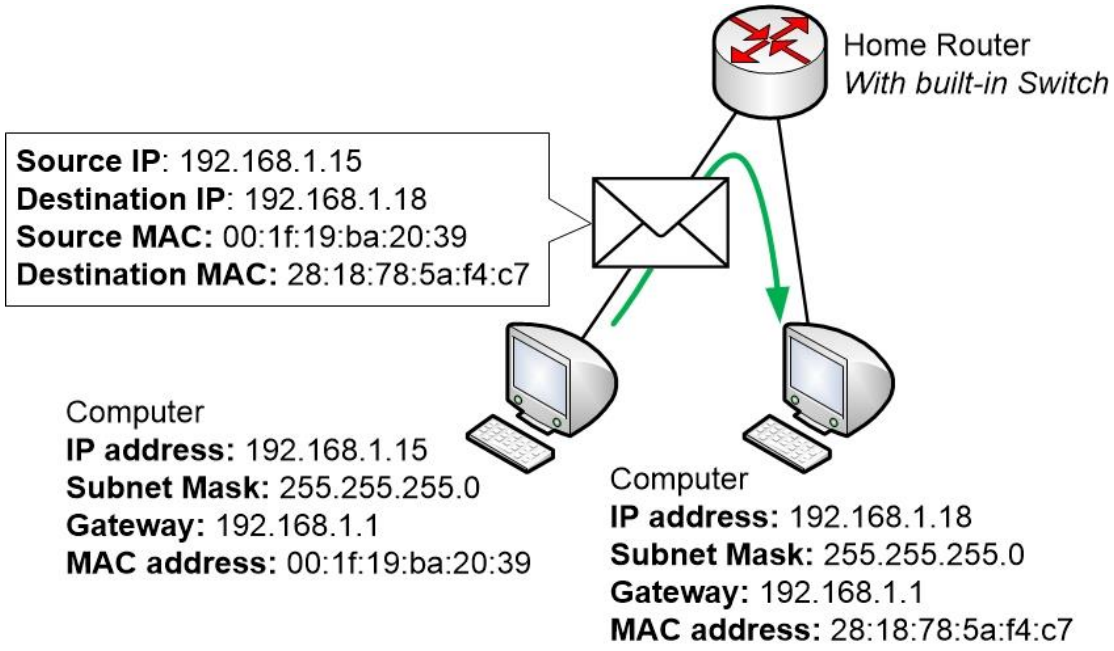

Broadcast

Broadcasting is a term that is used when talking about Radio and TV transmissions. A radio antenna or TV transmitter is sending out a single signal that can be received by anybody with a radio within reach of the signals. It doesn't matter if your radio is turned on or if it is tuned to listen to that particular radio channel. The signals are reaching your radio equipment no matter if you chose to listen to the radio signal or not.

The word Broadcast is used on computer networks too and basically means the same thing as it does for radio or TV broadcasts.
A device, such as a computer or a router, sends out a broadcast message on the local LAN that is intended to reach everybody else on that local LAN.

Here are two examples for when a broadcast can be used on a computer network:

- A computer just starts up and wants an IP address. It sends out a broadcast message to try to locate a DHCP server to ask for an IP address. As the computer has just booted up it doesn't know if there are any DHCP servers on the local LAN or what IP address any such DHCP server might have. Therefore, the computer sends out a broadcast which will reach every other device on the LAN to ask any available DHCP servers to reply back with an IP address.

- A windows computer wants to know which other Windows computers that are connected to the local LAN to be able to share files and folders between computers. It automatically sends out a broadcast on the LAN to locate any other Windows computers.

When a computer sends out a broadcast it will use a special destination MAC address, *FF:FF:FF:FF:FF:FF*. That address is called the Broadcast Address and is used specifically for this purpose. All other equipment on the LAN will then understand that the traffic is a broadcast that is directed at everybody else within the LAN.

Any computers, routers or other devices that receive a broadcast will pick up the message to read the contents. But not every device will be the intended recipient of the traffic. Any device that reads the message just to notice that the message was not aimed at them will simply throw the message away after reading it.

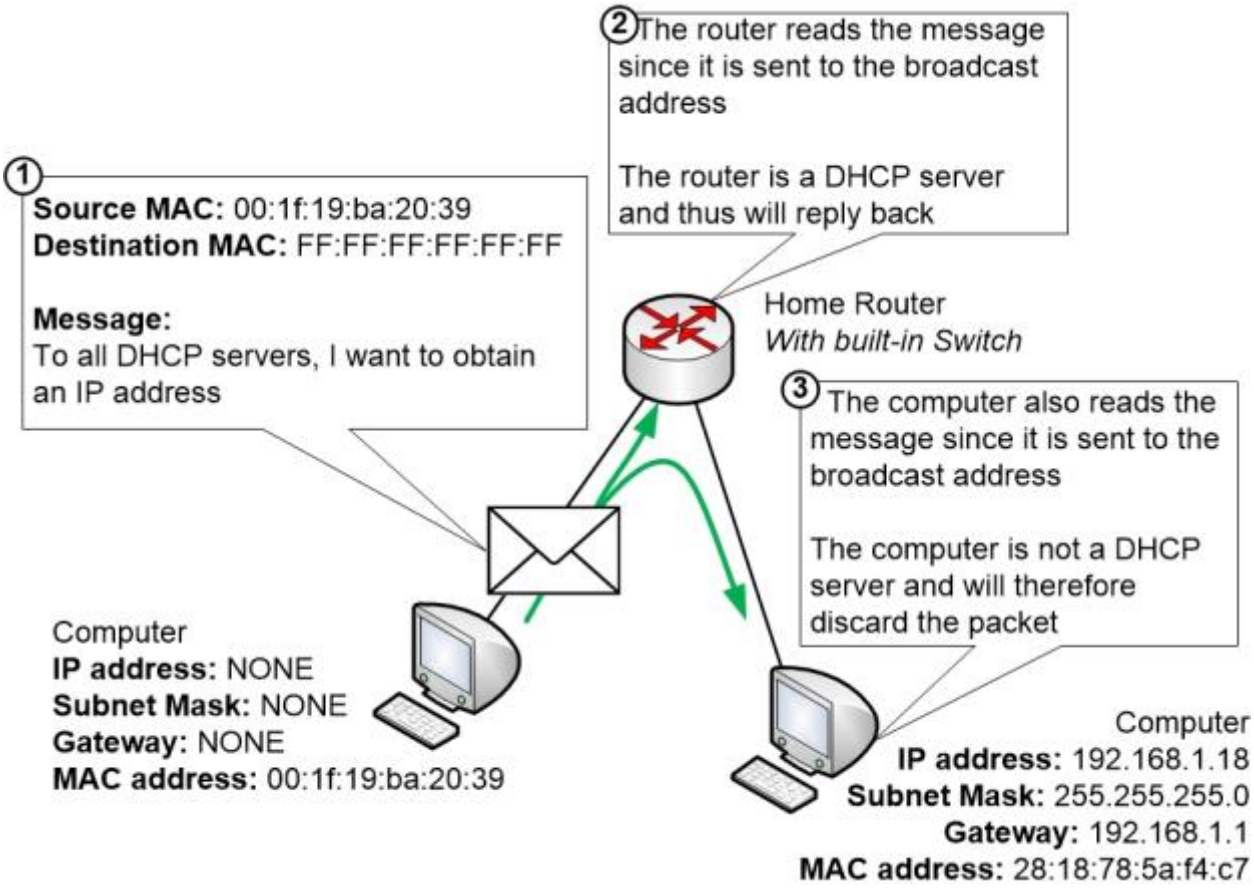

In the example above a computer is looking for a DHCP server to obtain an IP address. All the other devices on the LAN receive the message, but most of them will just simply throw the message away since they are not DHCP servers and cannot hand out any IP addresses.

The home router has a DHCP server built in and replies back to announce itself to the computer and to offer an IP address.

Switches

A switch is a network device with multiple interfaces or ports. The ports can connect computers and other devices, and any devices that are connected to the switch can communicate with each other. Switches aimed at the home network market segment often have 6-16 ports, but there are a lot of varieties available with different amounts of ports.

In fact, most people already have a switch at home, often without even knowing it. The LAN ports of your home router where you can connect your internal computers, printers etcetera are in fact built-in switch ports that act exactly like the ports in a standalone switch.

Whereas a hub just copies electrical signals between ports, a switch works intelligently with MAC addresses which have been mentioned previously to make sure that traffic that is sent between devices end up at the right place.

What a switch does is that it constantly monitors the traffic which is entering the switch from connected devices. It then learns about where the different MAC addresses of those devices are connected. It does this by looking at the traffic that arrives from computers to read the source MAC address of the traffic.

Additional Info

In network communication, data is packaged into different types of "envelopes" at different stages of the communication.

- When data from an application is going to be transmitted by a computer it is put inside a TCP or UDP *Segment* with port information.
- The Segment is put inside an IP *Packet* that contains the IP addresses.
- The IP packet is wrapped inside a *Frame* with MAC addresses.

Since a switch is working with MAC addresses it is looking at the information in the Frame, which carries the MAC address information. So when speaking about how switches work we will talk about how they transmit Frames instead of using the word Packets.

If a frame arrives on port 1 of the switch, and the frame is coming from Source MAC address 00:11:22:33:44:55 then the switch will automatically learn that a device with MAC address 00:11:22:33:44:55 is connected on port 1.

The switch will store this information in a *MAC address table* that it keeps in memory.

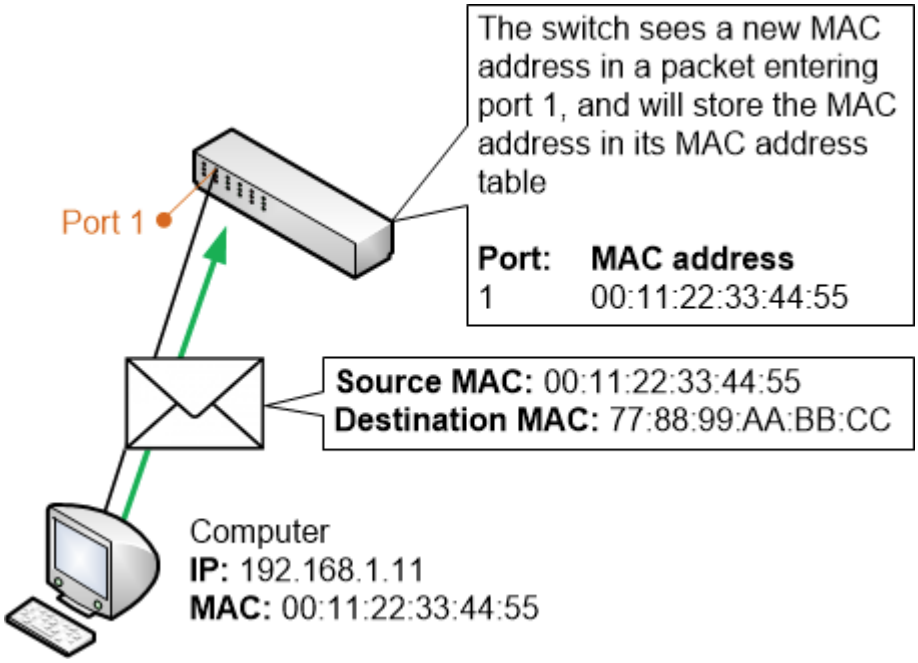

When the switch has seen at least one frame from each connected device it will know exactly which MAC addresses that are connected to what ports, and it will then also be able to forward traffic only to the correct destination ports.

The way a switch forwards traffic is based on always trying to send traffic only to the correct destination port. Whenever traffic enters the switch the switch will read what destination MAC address the traffic is being sent to, and then it will compare the destination MAC address to its own table of known destinations to find out if it knows where the destination is located.
If it can match the destination to a port in the table then it will forward the traffic only to that port.

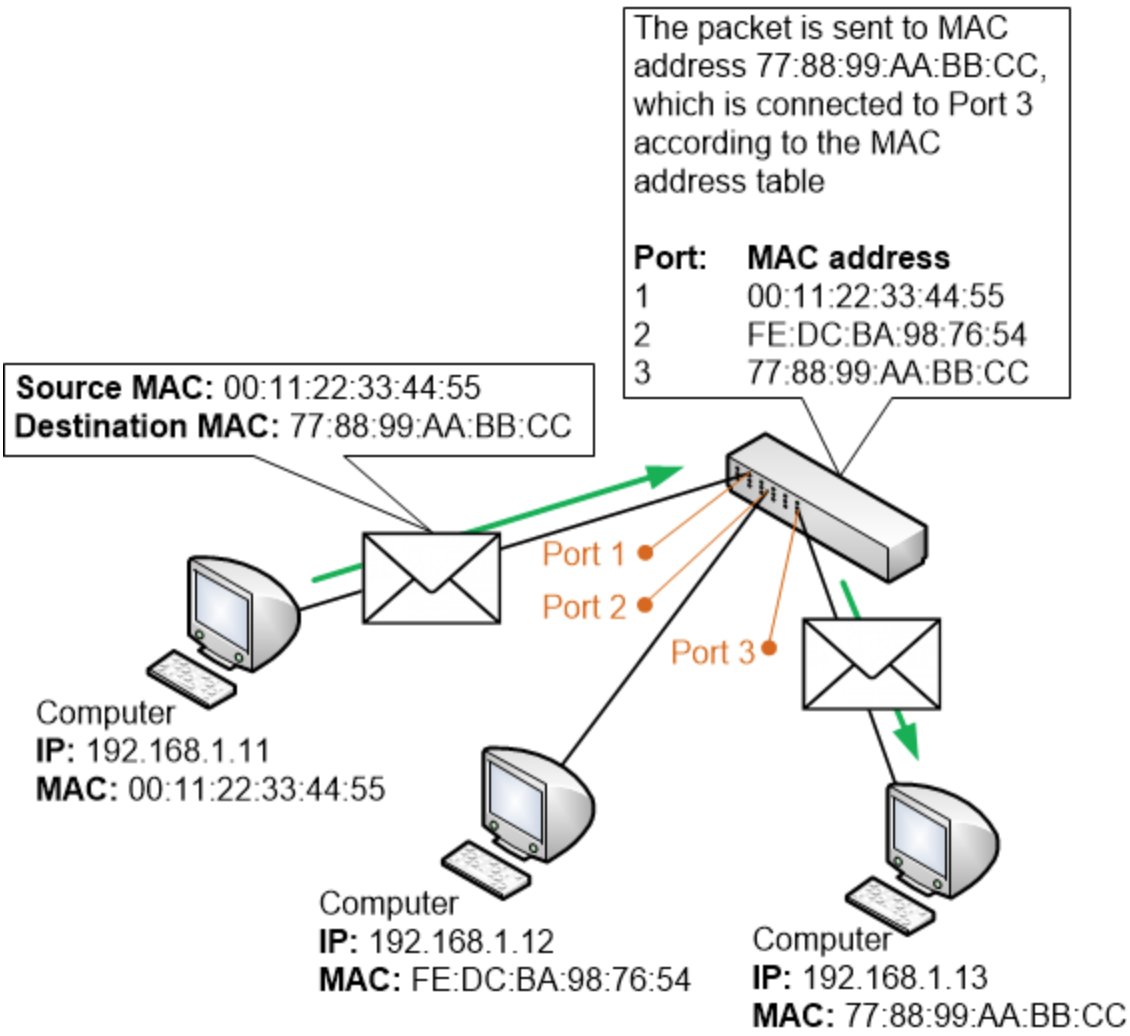

The most common exception is broadcast traffic. A broadcast is sent by computers when they want to send a message to all other devices on the same LAN. As mentioned earlier a computer could be looking for a DHCP server and uses a broadcast DHCP request to find out if there are any connected DHCP servers.

For broadcasts to work as intended the switch must handle broadcasts as a special case, and must send broadcasts to all other connected ports. It doesn't act quite like a hub, however, because the hub is Half Duplex. The switch, which is Full Duplex, can still handle other traffic at the same time as it is sending a broadcast, so there is no need for all other computers to be silent while the broadcast is being sent.

Broadcast messages are sent to the destination MAC address FF:FF:FF:FF:FF:FF. If a switch receives a message that is sent to that destination MAC address then the switch knows it is a broadcast and will forward the message to every other connected port.

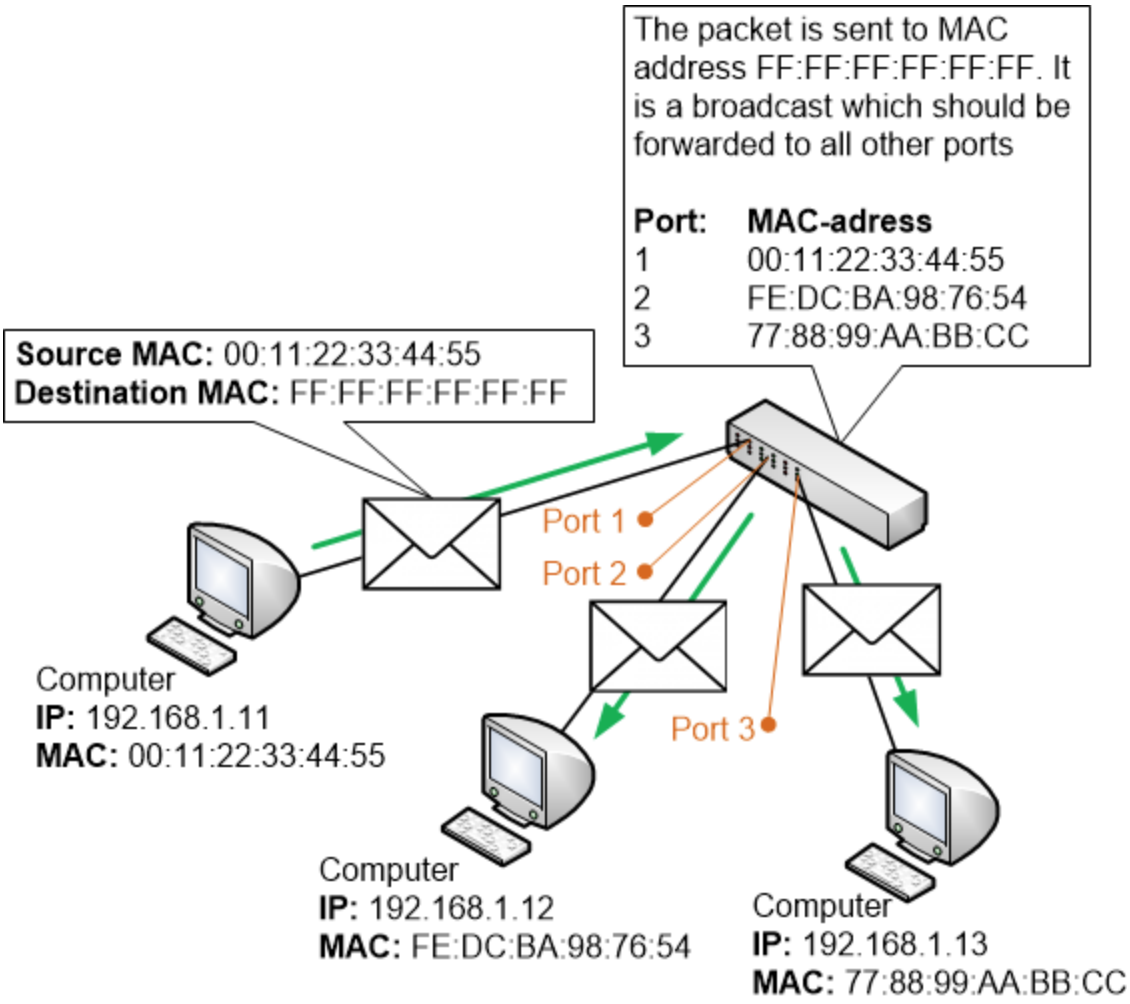

There is also another common scenario where switches send out traffic to all other ports, and that is when the switch simply does not yet know where the destination address is connected. Let's say a computer sends a message to another computer via a switch, but the switch hasn't learned where the destination MAC address is located. This means that the switch cannot know where the second computer is connected. Then the switch simply treats the traffic as broadcast traffic and sends it out all other ports.

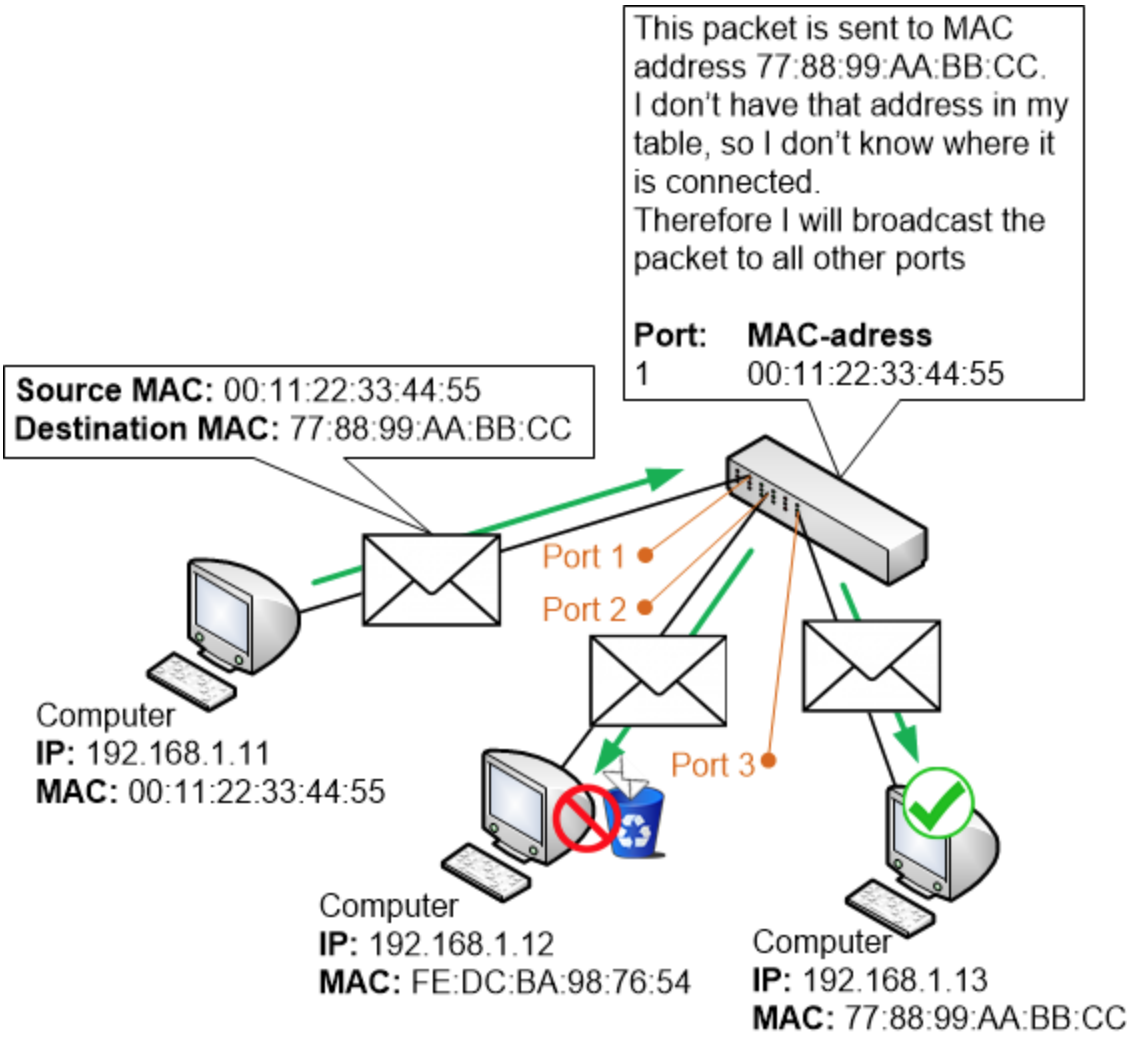

But the switch is also smart. Once the second computer replies back, then the switch can read the source MAC address in the traffic and will learn which port the MAC address is connected to.

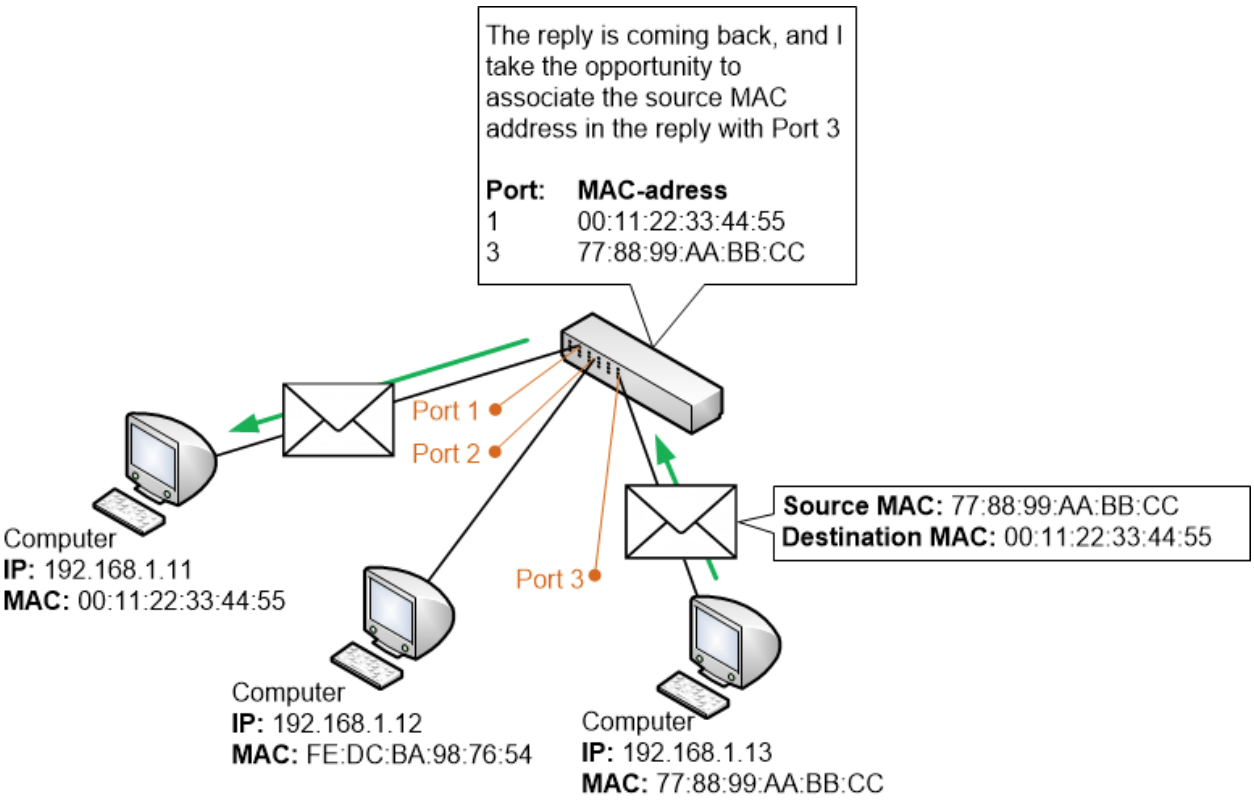

So as the computers continue to communicate with each other the switch will remember where the MAC addresses of those computers are connected. The rest of the communication will be forwarded by the switch only to the correct switch ports where the intended recipients are connected, without bothering any other computers on the LAN.

ARP - Associating MAC addresses with IP addresses

We have gone through how MAC addresses and Switching work, and other sections discuss the functionality of IP addresses and Routing.
What hasn't been discussed however is the glue that binds those together. How do the MAC and IP addresses interact?

There is another protocol called *ARP* which stands for *Address Resolution Protocol*. It is used to associate MAC addresses with IP addresses and is a way for a computer to look up an unknown MAC address for a device that it wants to communicate with.

Most commonly a computer knows what IP address that it is sending the traffic to. For example, if you are browsing to a web page on the Internet, you would enter an address in the browser. The computer would then use DNS to do a name resolution lookup to obtain the IP address of that web page. So the computer will easily find out what the IP address of the web server is, meaning that the destination IP address is known to the computer in one way or another.

But let's go back to that DNS message. You have just instructed the computer to browse to a web page in a web browser. The computer must use DNS to find out the IP address of the web page.

After the computer has created the DNS query it will put the query inside an IP packet and send the packet to a DNS server. Let's say the computer is configured to use the home router as its DNS server.

So the computer creates a packet with a DNS query that it wants to send off to the home router's IP address. But the computer must also enter the MAC address of the Router as the destination MAC address for the traffic. How will the computer know which MAC address that the router has?

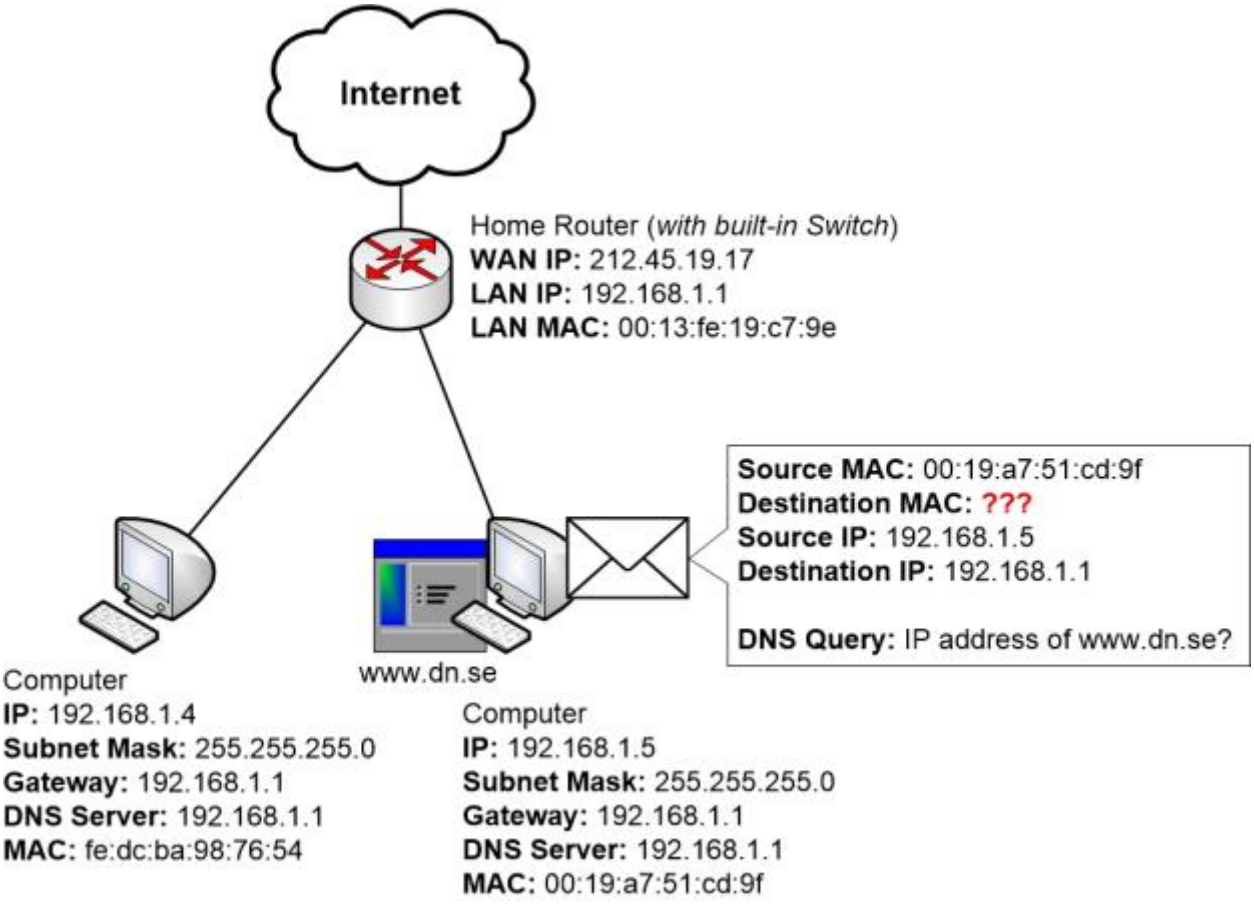

This is where the ARP protocol comes into play. ARP lets devices on the network ask each other which MAC addresses they have.

To find out what MAC address the router has got the computer will first put its DNS query on hold in a queue. Then it will create an *ARP request*.

The ARP request contains a simple question. In this case, the computer wants to find out which MAC address that the 192.168.1.1 device has got. So the request is basically as follows:
"Device with IP address 192.168.1.1, respond back with your MAC address"

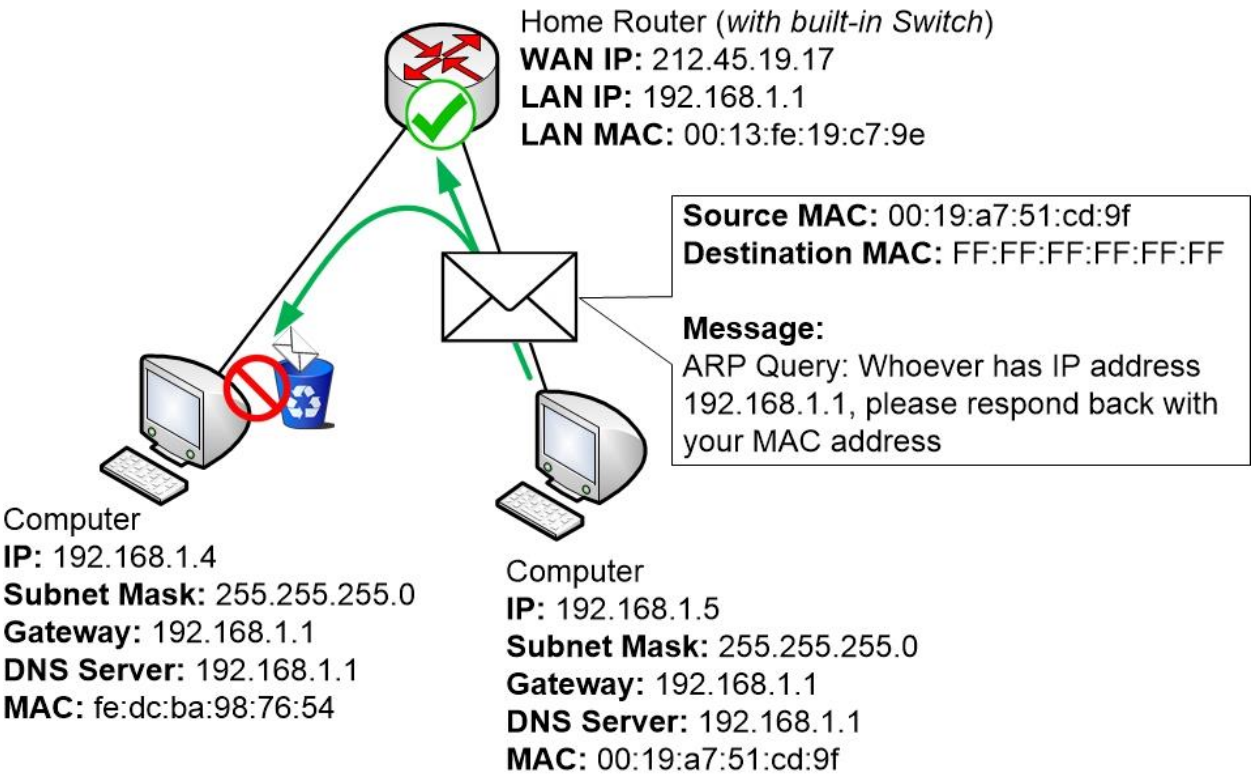

ARP requests are always sent as broadcasts because we don't know what MAC address we want to send the message to. Since it is a broadcasted message, every other device on the LAN will receive the message. This is because the integrated switch in the Home Router handles the message as a broadcast and forwards it to all other ports including the integrated router.
But all devices except one will notice when they read the contents of the ARP request that the message is intended for another device with IP address 192.168.1.1

The home router, which is configured with IP address 192.168.1.1, will read the message and will notice that the message is directed at itself. It will then construct an *ARP reply*:

"I have IP address 192.168.1.1 and my MAC address is 00:13:fe:19:c7:9e"

Every time a computer receives an ARP reply it will save the response for at least a few minutes in an *ARP table* (or *ARP cache*) in memory. This is so that the computer doesn't have to do an ARP request for each packet it wants to send. From now on and for as long as it keeps communicating with the router it will remember the router's MAC address.

If however they stop communicating for a while then the computer will clear out the router's MAC address from its ARP table.

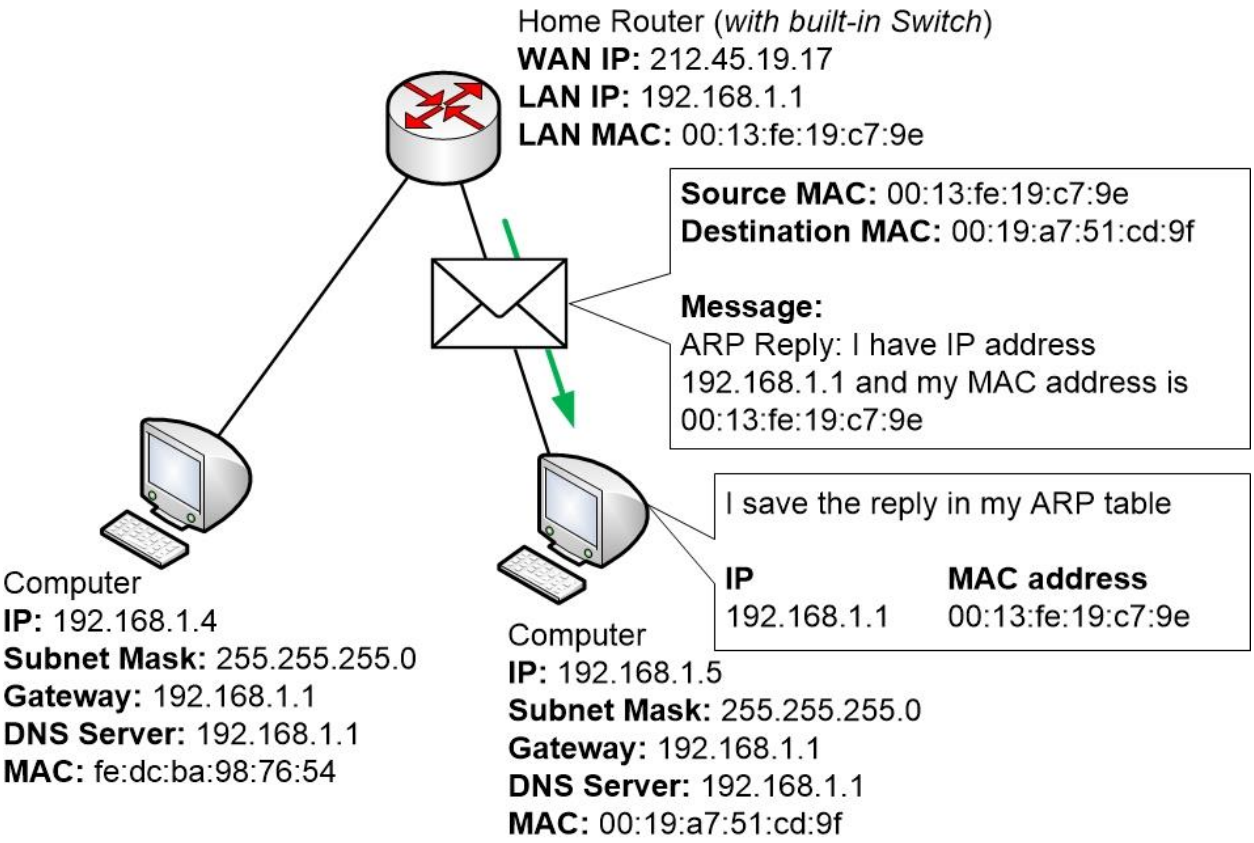

Each time a computer is sending a packet to an IP address it will look in its ARP cache to see if it already knows what MAC address that is associated with that IP address.

- If the address exists in the ARP cache then the MAC address in the table will be used.
- If the address does not exist in the ARP cache, then an ARP request must be created and sent out.

Finally once the computer has gone through the ARP request it now has all the necessary information to send the DNS query to the router to find out what IP address that the web page has got.

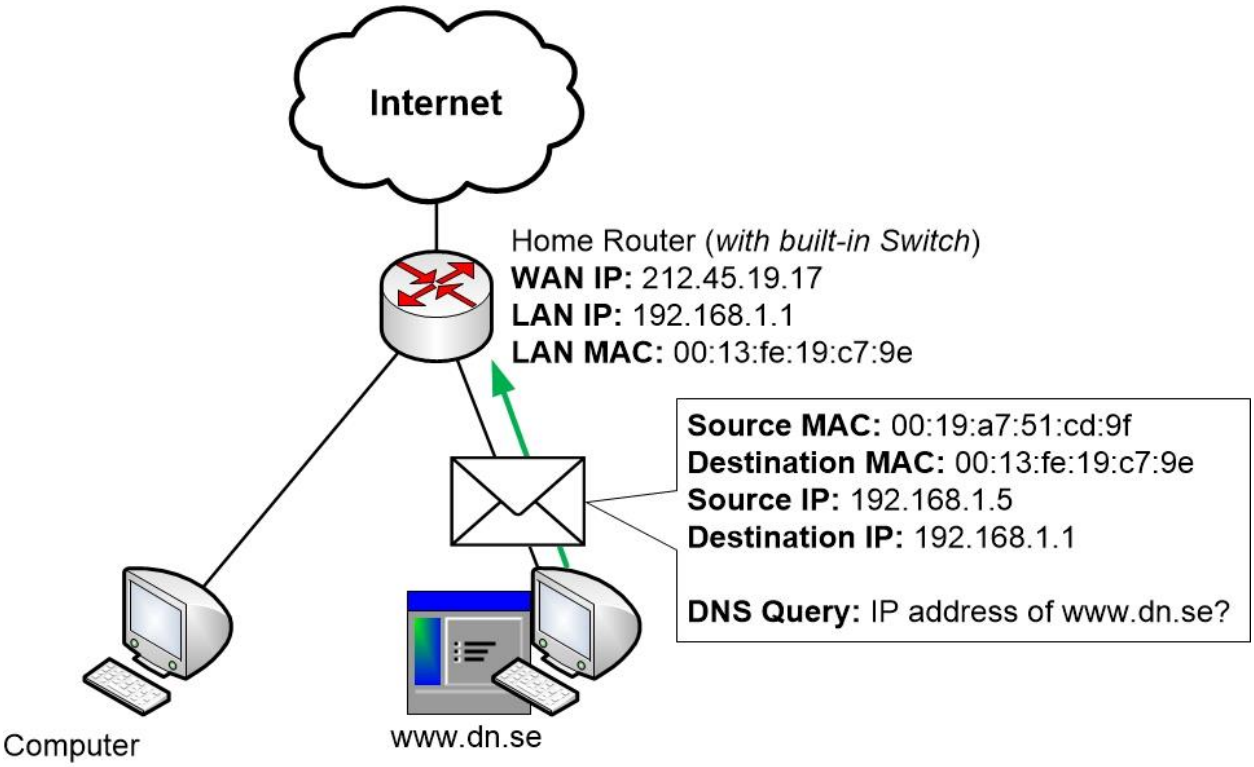

Specialisation: Traffic example, a step-by-step walkthrough

In this section we will do a complete step-by-step walkthrough of a traffic example, showing which steps that a computer will go through when it wants to communicate over a computer network. This example will incorporate a lot of information from the other sections and serves as an all-inclusive example.

Most people are surprised to see how many steps that are involved for a computer to be able to just send out a simple packet, and there are lots of steps that are completely hidden to a regular user and which you would never know about unless you knew about them beforehand.

If you have read through most of this website then you probably already have an idea about the number of steps that might be included in this walkthrough.

In the example below we will display a small home network where a computer has just booted up. The computer has a manually configured IP address and has not yet communicated on the network.

A user sits down by the computer and opens up a Web Browser and tries to browse to www.iis.se

First let's have a look at the whole picture. This image shows the network topology in the example. Within a lot of the steps, we will zoom in on just the most relevant part of the network to avoid having to draw the whole picture every single time.

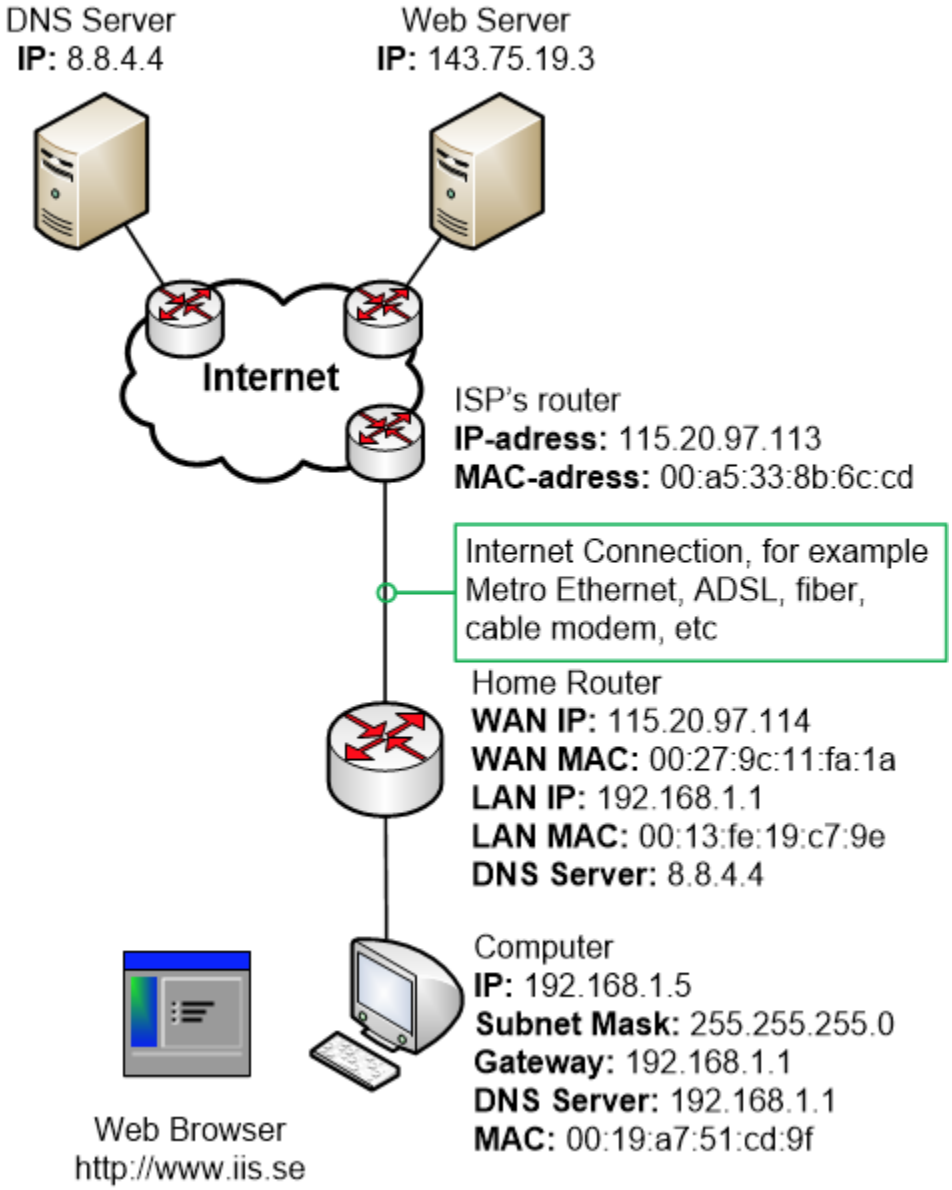

Now let's start with the traffic example walkthrough!

Step 1: The computer wants to send traffic

A computer connected to the home network has just booted up. The computer has a manually configured IP address, Subnet Mask, DNS server and a Default Gateway. Both the DNS server and the Default Gateway address is pointing at the LAN IP address of the home router.

The user of the computer opens a web browser and goes to www.iis.se

The first thing that happens is that the Web Browser instructs the OS on the computer to set up the communication between the computer and www.iis.se

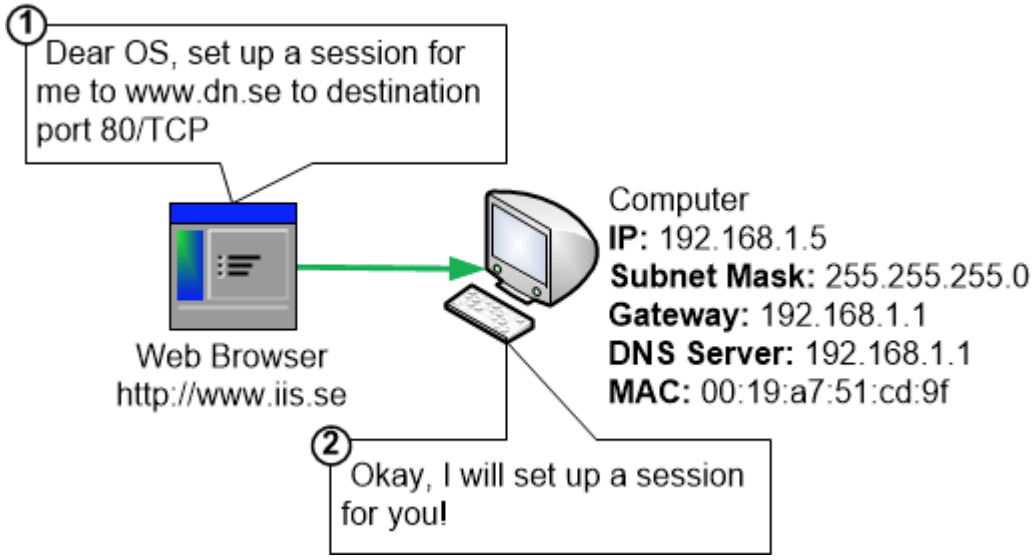

Step 2: DNS

This part is divided into numerous sub-steps

Step 2a: DNS cache

The computer OS checks its DNS cache to see if it already knows what IP address that www.iis.se has got. Since the computer just started up and it hasn't previously contacted www.iis.se the DNS cache is completely empty.

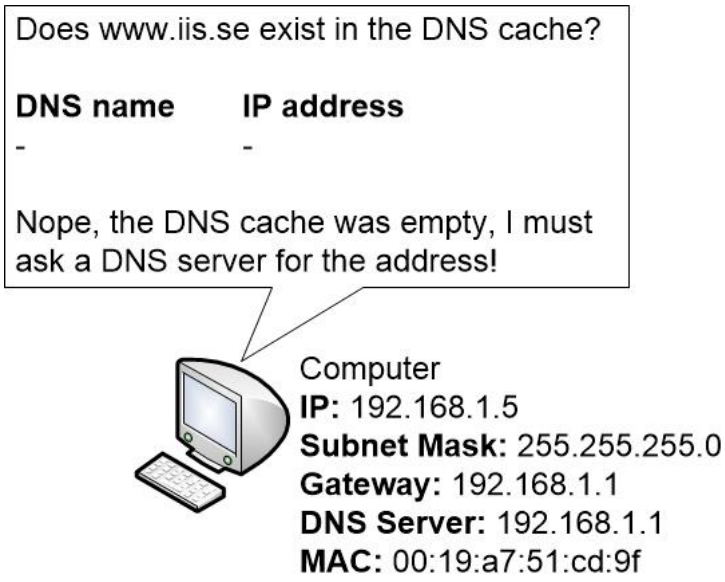

The computer must now ask its DNS server what IP address that www.iis.se has got.

Step 2b: Putting a DNS query together

The computer will construct a DNS query that it can send off to the DNS server, 192.168.1.1, that it is configured to use.
The destination address of the DNS query is 192.168.1.1, and the source IP address is the IP address of the computer itself, 192.168.1.5

DNS uses UDP as its transport protocol. The destination port for DNS queries is 53/UDP. Later when the DNS query reaches the DNS server the DNS server will be able to tell by looking at the destination port 53/UDP that the message is intended for a DNS server program, and can forward the message to the running DNS program.
The computer OS must also randomise a source port which is also written into the message.

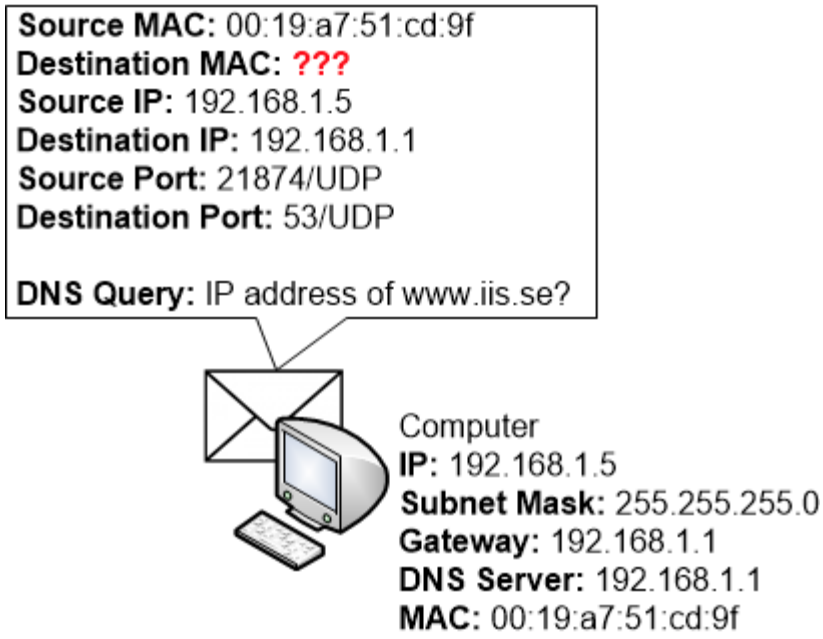

But when the computer puts the DNS query together it notices that it must check what destination MAC address that it should send the packet to.
So for now, the OS puts the packet in a queue in memory and then starts working on figuring out what destination MAC address to use.

Step 2c: Check the ARP table for a valid MAC address

The computer will now check its ARP table to see if it knows what MAC address that is associated with the IP address of the router, 192.168.1.1

But the computer has a completely empty ARP table since it just booted up and hasn't yet learned any ARP entries.

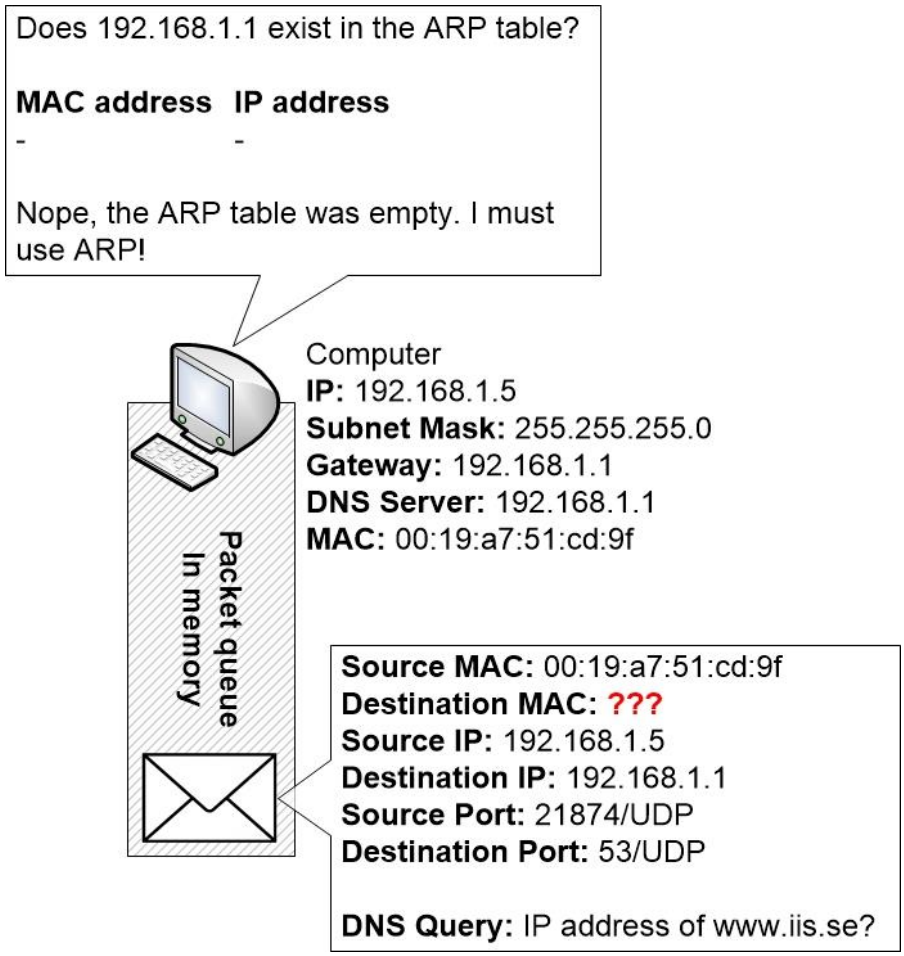

Step 2d: ARP request to the network

Now the computer must construct an ARP request to the rest of the network. The request will be sent to destination MAC address FF:FF:FF:FF:FF:FF which is the broadcast address. The result is that every other computer and device on the LAN will receive the request and read the contents.

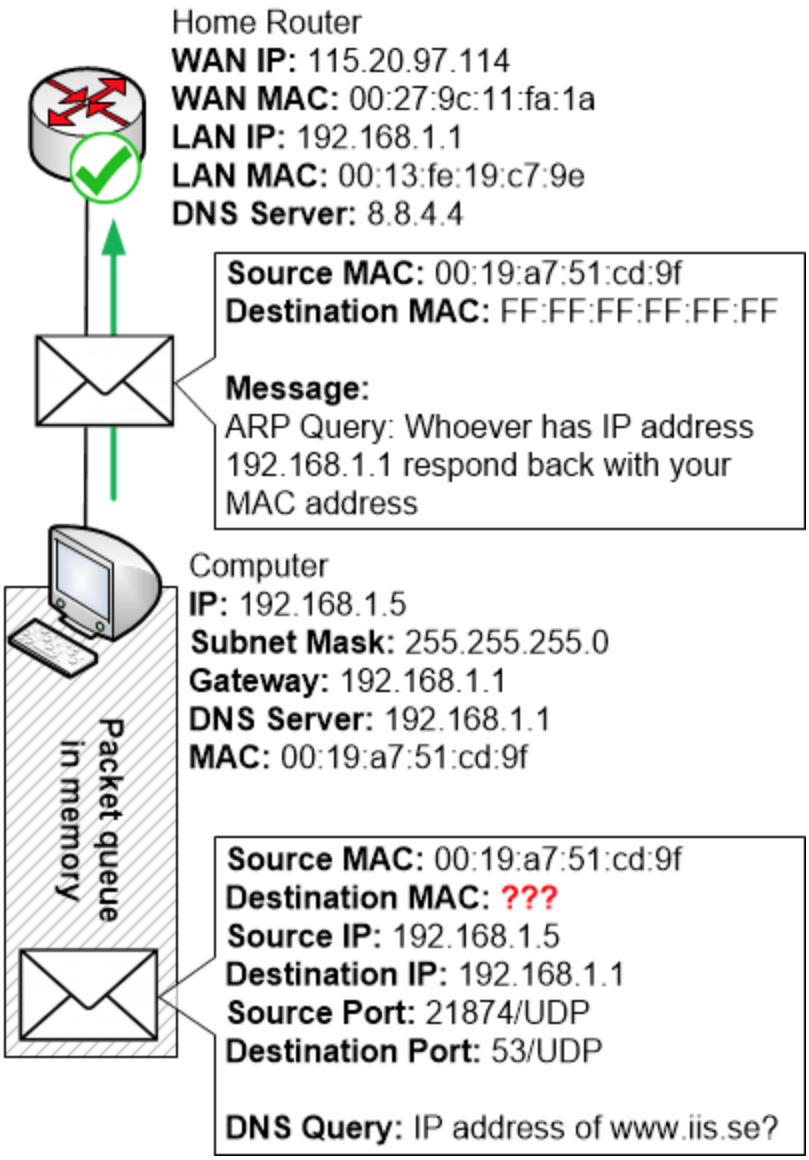

The home router receives the ARP request and reads the message since the request is sent to the broadcast MAC address FF:FF:FF:FF:FF:FF
The home router can see in the message that the computer is asking for the device with IP address 192.168.1.1. Because the router is configured to use that IP address the home router will respond to this message by constructing an ARP reply and sending it back to the computer.

Step 2e: ARP reply from the router

When the ARP reply is received by the computer the OS will read the reply. It will then enter the reply into its ARP table to remember for a few minutes which MAC address that is associated with 192.168.1.1

Now the computer finally has gathered all the information it requires to be able to send off the DNS message.

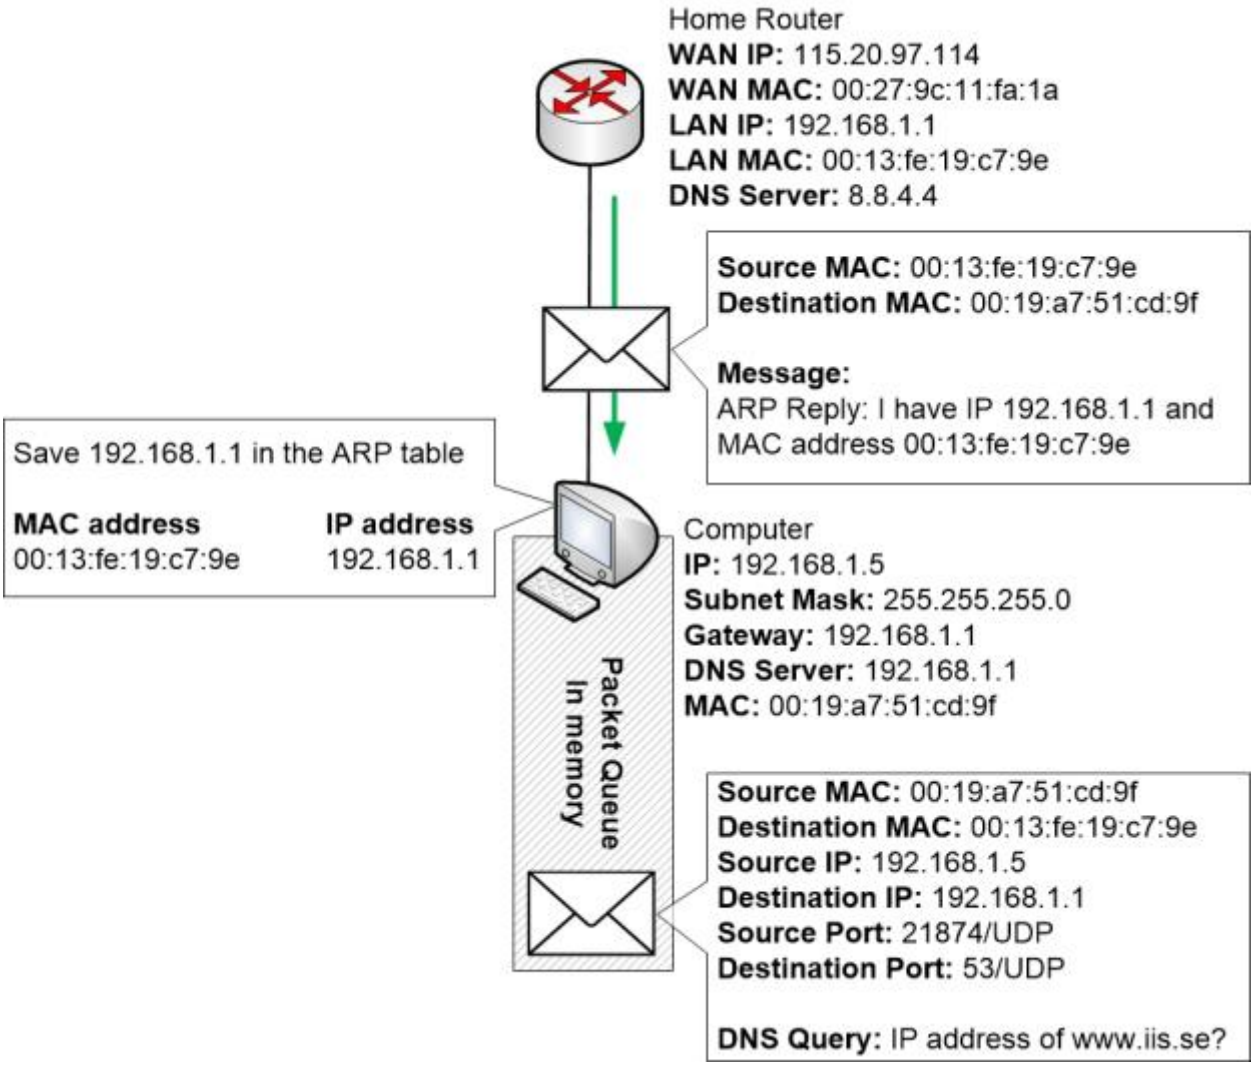

Step 2f: Send off the DNS query

The DNS query is now going to be sent from the computer to the DNS server, which exists as a service that is running on the home router 192.168.1.1

The home router receives the query, sees that it is a DNS query aimed at the router's own IP address and MAC address, and understands that it must handle this DNS query and send back an answer.

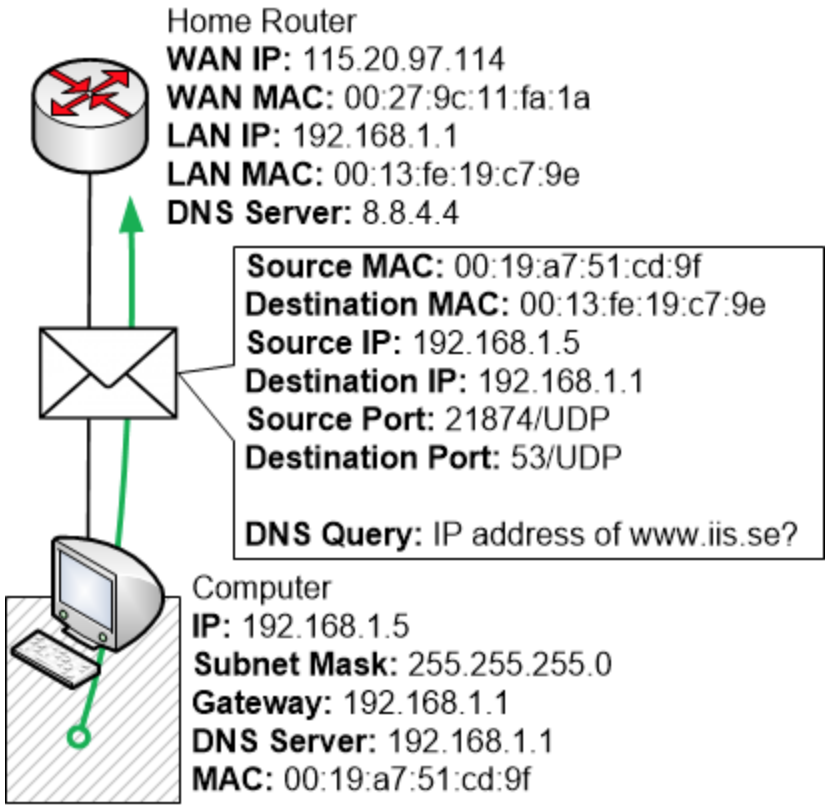

Step 2g: The home router checks its DNS cache

The home router is a DNS server, but it is also depending on other DNS servers on the Internet. The home router can't know every single DNS address on the Internet. Instead it will ask those DNS servers that are responsible for different domains (such as example.com) as needed.
The home router also has a DNS cache just like the computer. Every time the home router handles a DNS query from a computer it will also save the DNS reply in its own DNS cache for some time. This is to avoid having to handle the same DNS queries over and over and to speed up the response time.

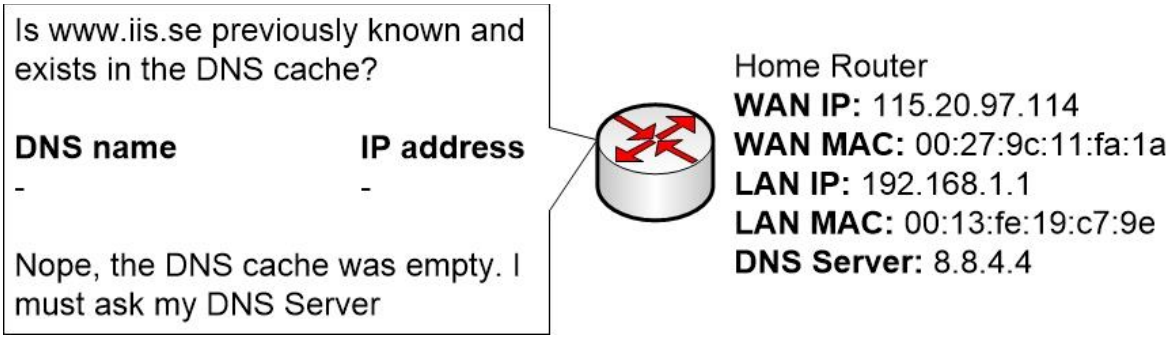

In this case, the home router hasn't gotten any question about www.iis.se in a long time so it doesn't exist in the DNS cache on the router. Therefore, the router must ask its configured DNS servers on the Internet to answer this DNS query.

Step 2h: The home router prepares and sends away its DNS query

Now the router prepares a DNS query that it will send to its DNS server. The router has learned about available DNS servers via DHCP from the Internet Service Provider, the ISP, when the home router first booted up and got its own public IP address from the ISP.

So the home router will prepare a DNS query for transmission by putting the query inside a UDP message with destination port 53/UDP and a random UDP source port. It will then put the message inside an IP packet. The IP packet is sent from the home router's public IP address to the DNS server address.

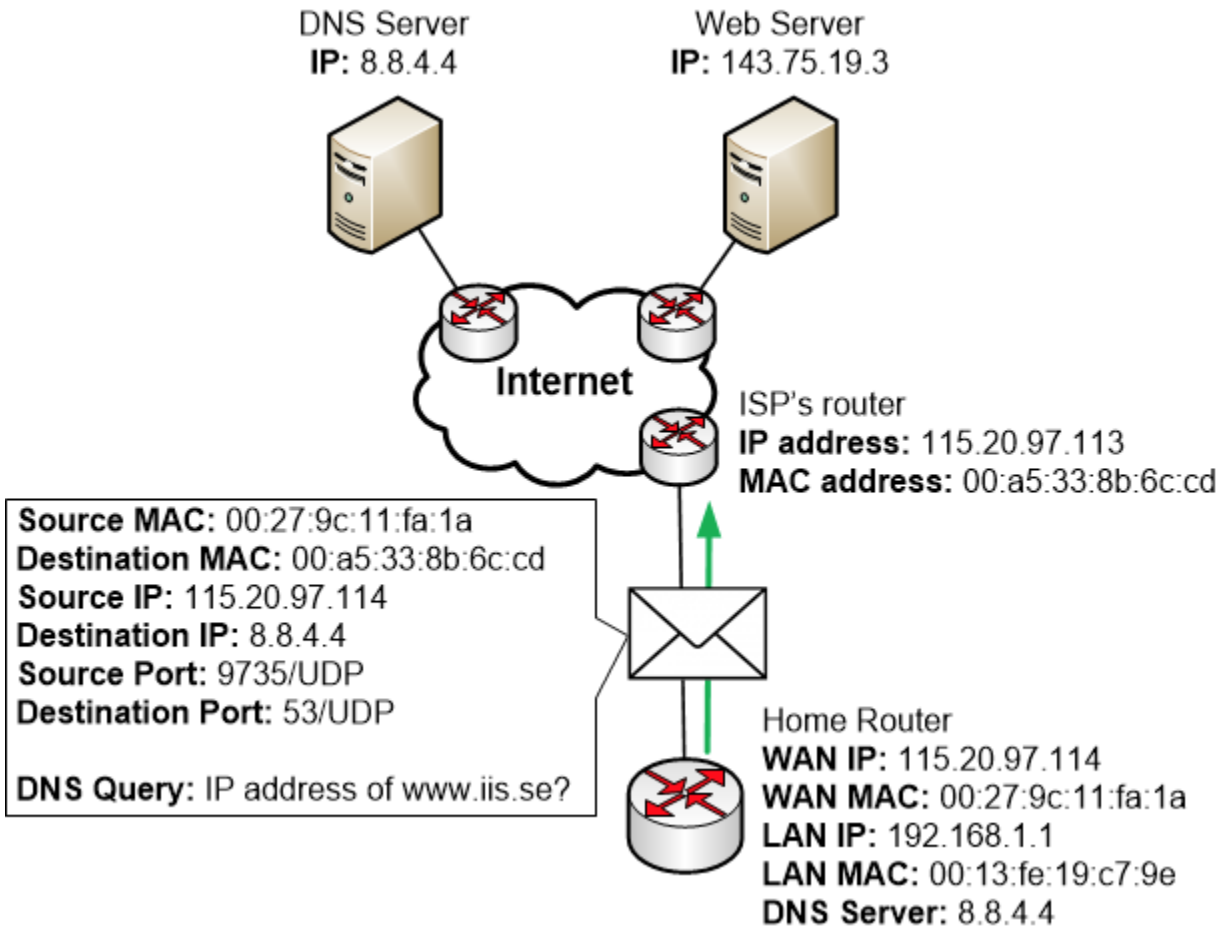

When the home router has prepared the packet and is ready to send it, then the home router will look in its routing table to see which way it should send the packet. It can see in the routing table that the best path to the inside LAN 192.168.1.0 is via the LAN ports, but this packet should be sent to another IP network on the Internet. So the home router picks the WAN port as the best destination.

At this stage of the process, it is possible that the home router would have to perform an ARP request to find out which MAC address of the next hop router 115.20.97.113, but we assume that the home router already has got this information in its ARP cache.

Step 2i: The DNS query is routed over the Internet

Here a number of steps have been somewhat simplified and shortened.

Each router on the Internet that receives the DNS request will perform the following:

- Receive the packet
- Looks at the destination IP address to see where the packet is going
- Looks in its routing table to see which path that is best for the packet
- Removes the old MAC addresses from the packet and adds new ones. It will use its own MAC address on the outbound interface as the Source MAC address for the traffic, and will put the next-hop router's MAC address as the destination MAC address
- Sends off the packet to the next hop router

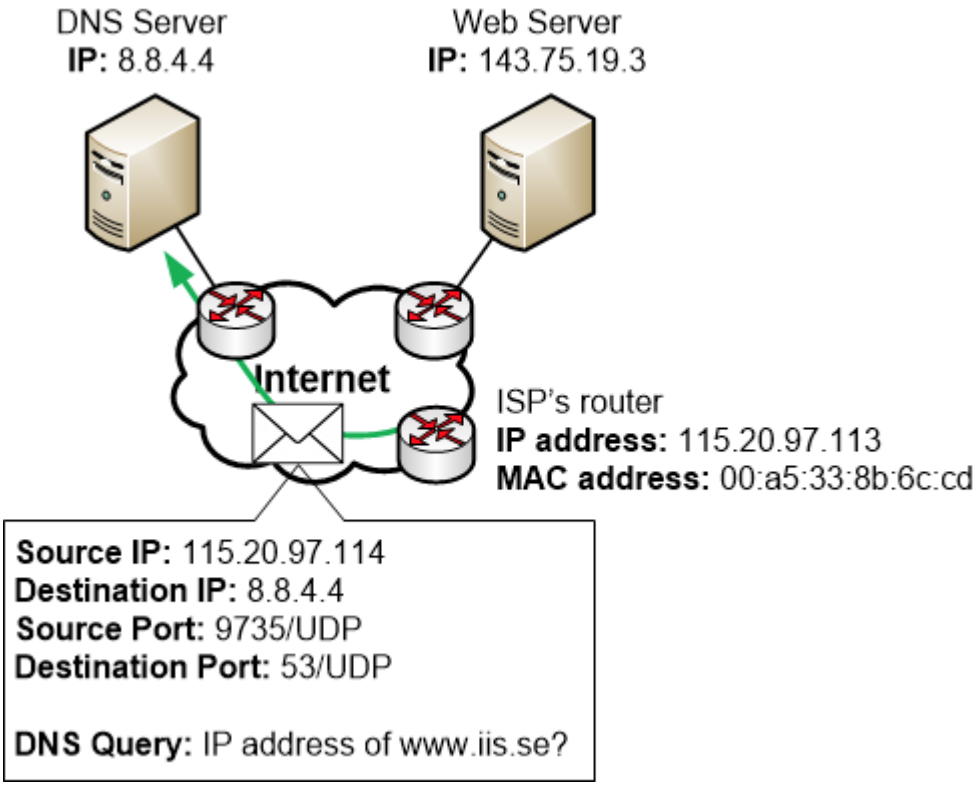

Step 2j: The DNS server responds

Eventually, the packet reaches the DNS server which will handle the packet and prepare a response.

Just like a regular computer the server has an IP address, a Subnet Mask and a Default Gateway. So it works in much the same way as a regular computer.

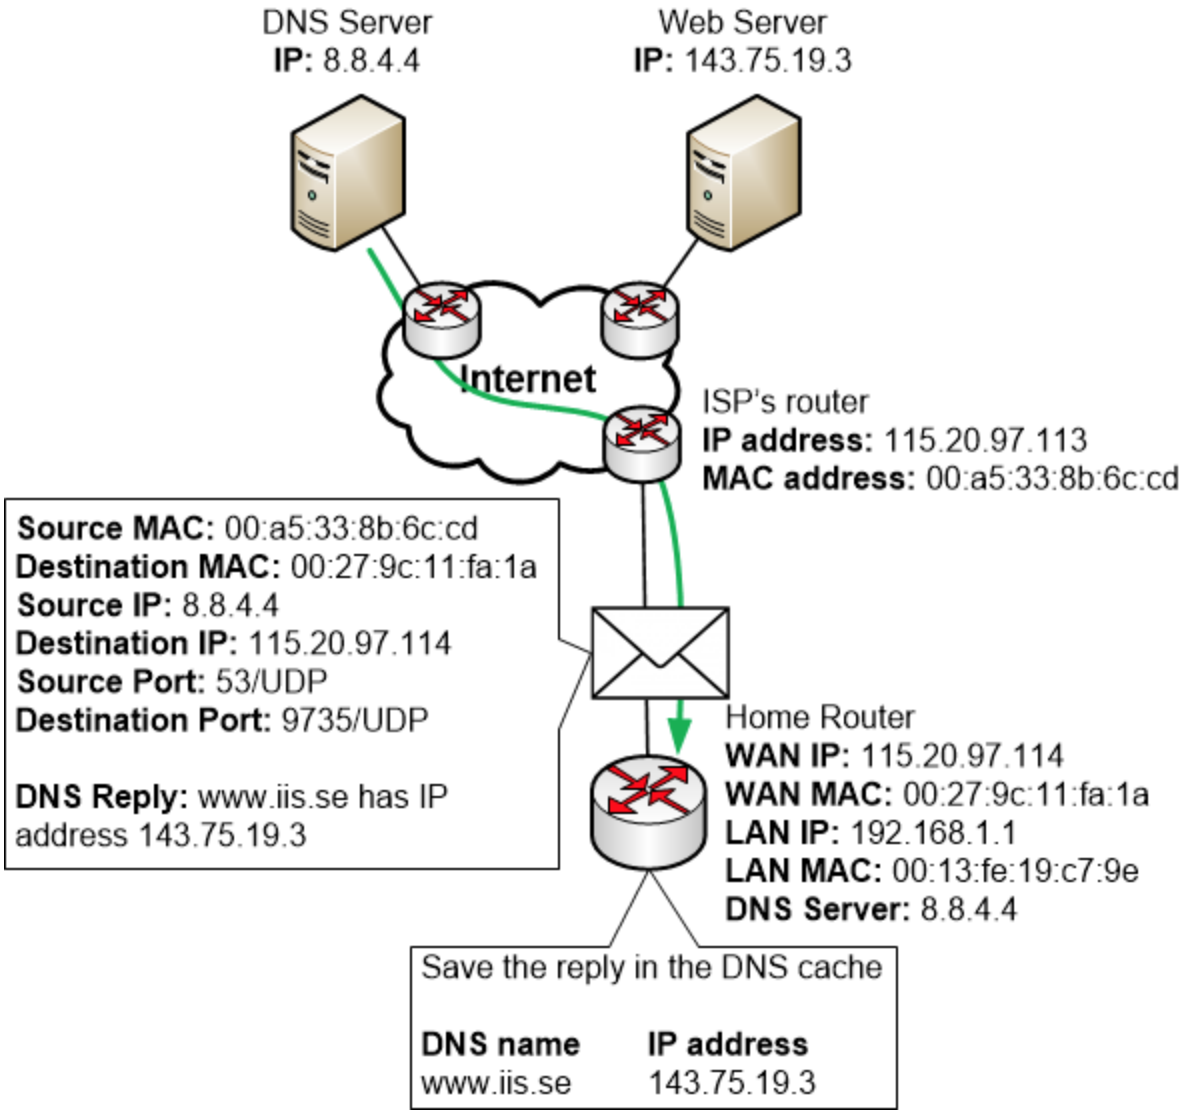

Step 2k: The home router can send a DNS reply to the computer

Now after receiving the DNS reply from the DNS server, the home router can finally create a DNS reply and send it to the computer to let the computer know which IP address that www.iis.se has got.

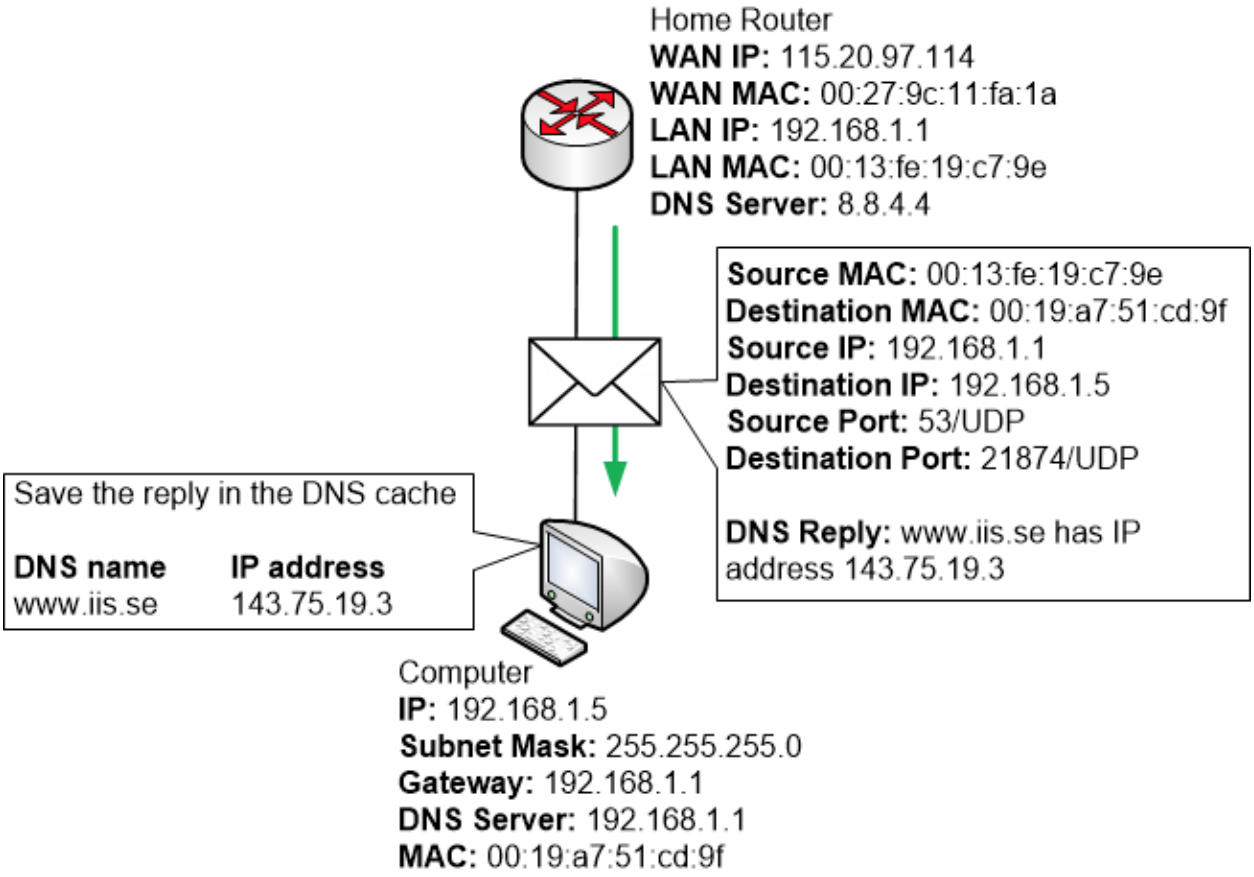

Step 3: The computer sets up a session to www.iis.se

During this part, a lot of things happen at once.

Mainly the computer will initialise something called a "TCP 3-way Handshake" which is a setup phase of TCP communication that consists of three messages between the computer and the server. When TCP is being used, as is the case with web browsing, TCP always tries to make sure that everything works as well as possible which includes setting up a session via a handshake. This is done to prepare the server for an incoming session and to decide which ports that should be used for the communication.

The TCP 3-way Handshake consists of three messages:

- The first one is sent from the computer and is called "SYN" which stands for Synchronise. It lets the other side know that we want to synchronise settings for a TCP session. The message also contains the random source TCP port that the computer has chosen.
- The second message is the reply back from the server and is called "SYN-ACK", which stands for Synchronise Acknowledgement. This simply means that the server acknowledges that it received the message and that it also is prepared to set up a session for communication

- The third message is sent by the computer and finalises the session by sending “ACK” or Acknowledgement. This means that everything is now fully prepared.

The computer has got the correct ARP information for the IP address of the home router in its ARP cache. So the computer can send off any packets it wants to the Internet via the home router without first having to go through ARP lookups.

But this is also the first time so far in this example that the computer wants to communicate directly with something that is located beyond the router. The TCP 3-way handshake will take place directly between the computer and the Web Server on the Internet.
Earlier during the DNS lookup process, the computer just talked with the Home Router. The home router, in turn, talked with the DNS servers on the Internet. But no communication was travelling directly between the computer and any IP address on the Internet.

The difference is that now that the computer wants to communicate directly with something on the Internet then the home router has to perform Address Translation for the traffic.

Step 3a: The computer sends a TCP SYN message

Here we take a closer look at what happens when the computer is establishing the TCP session by initialising the TCP 3-way handshake.

To do so, the computer OS will randomise a TCP source port that it will use for the communication. Then it assembles the TCP SYN message and sends it to the Web Server. The TCP message doesn’t contain any other data. It is just an empty TCP message.

When the TCP SYN message passes through the router the router will perform NAT on the message. The router will also save information about the performed NAT in its NAT table so that it can keep track of the session and perform reverse NAT on any replies.

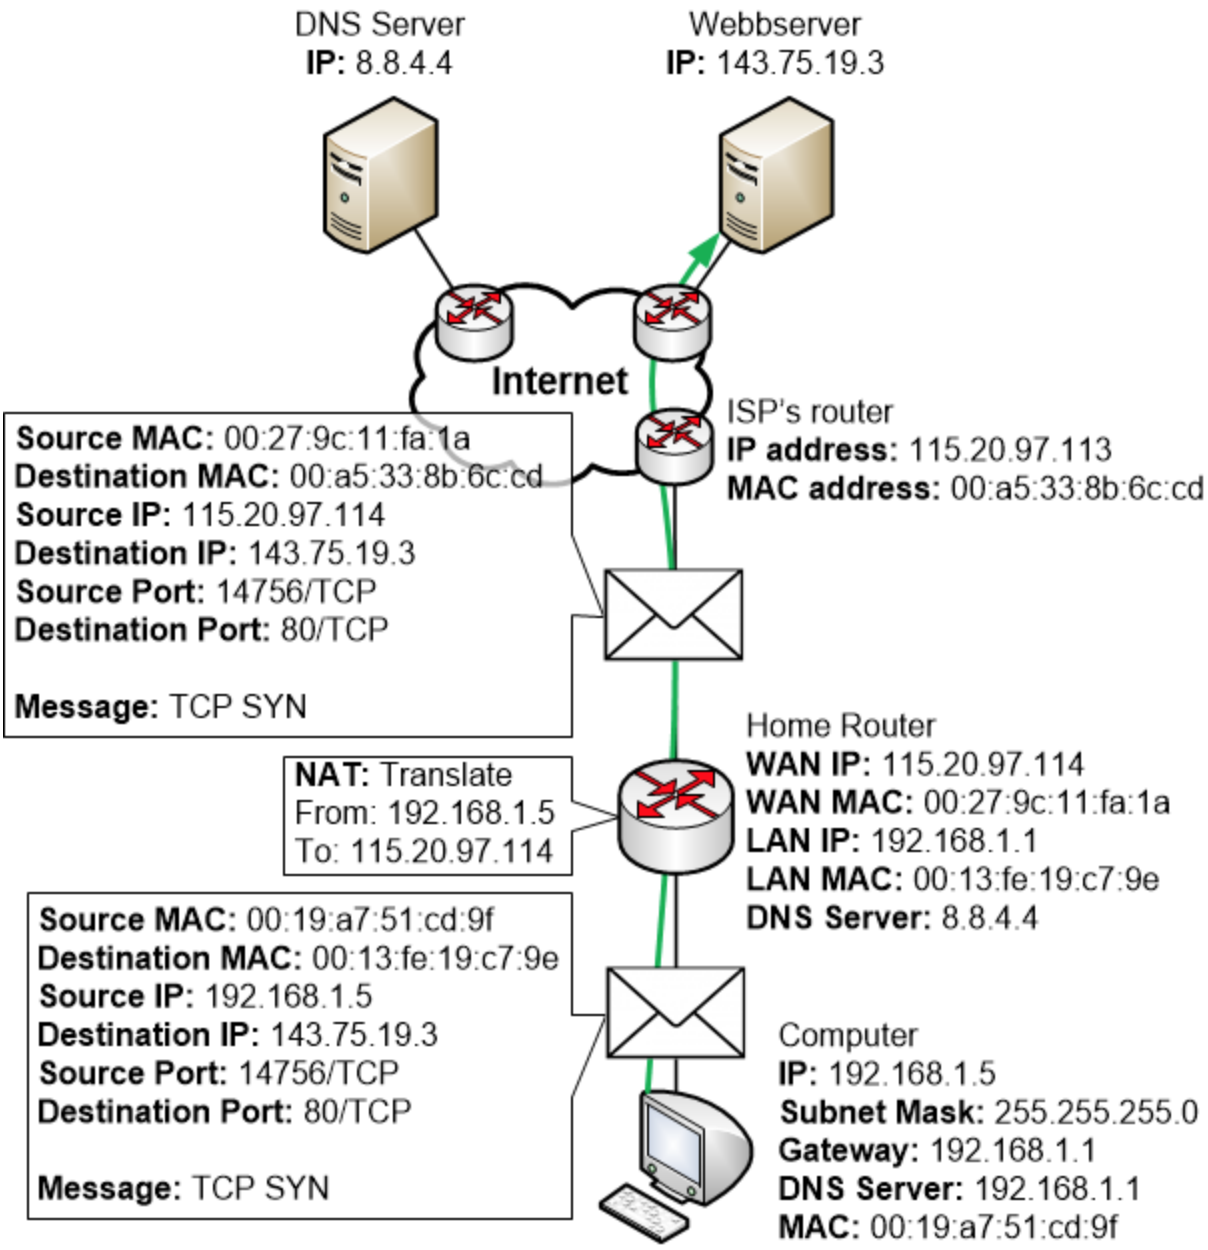

The picture shows a TCP SYN travelling from the computer to the web server

- "I would like to synchronise a TCP session with you"

Step 3b: The Web Server replies with TCP SYN-ACK

Here you can see the returning TCP SYN-ACK response from the Web Server to the computer:

- "Okay, I'm good to set up a session and I confirm that I got your message"

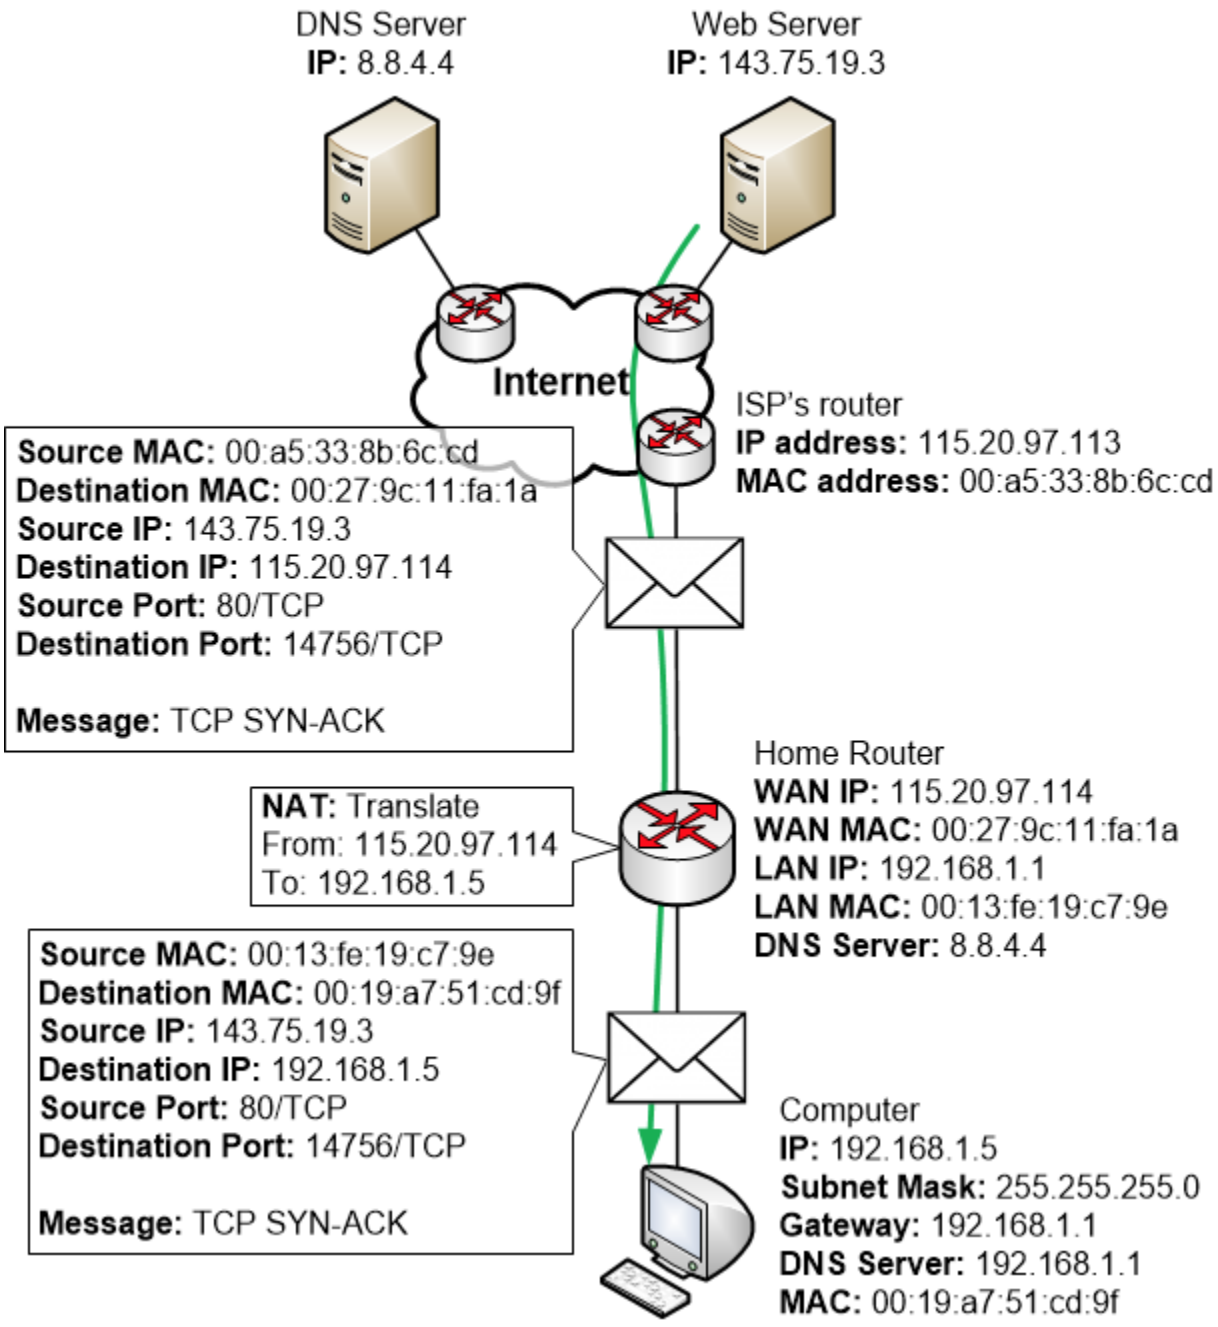

The packet matches the NAT table entry on the router so that the router can see which LAN computer it should forward the packet to and how it should perform NAT on the packet.

Step 3c: The computer sends a TCP ACK

Finally, a TCP ACK is sent from the computer to the Web Server

- "Then I confirm that from now on we have established a session!"

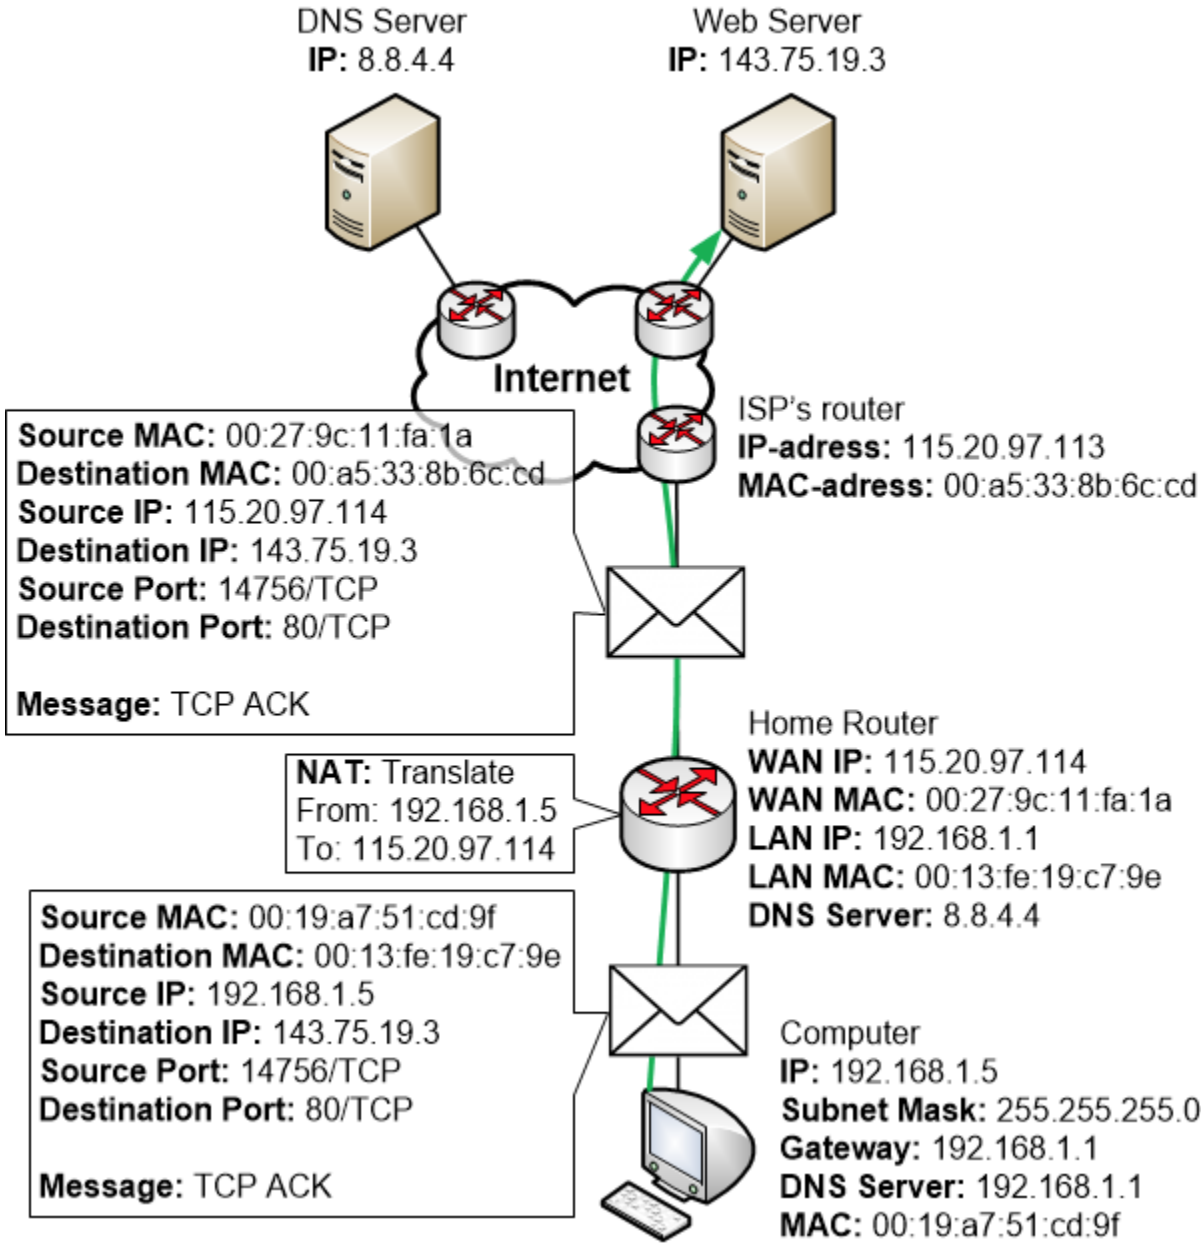

As long as the computer and the server keep on communicating with each other they will keep using the same session for the communication. This also includes using the same TCP ports,

which lets all devices along the way keep track of the session, the address translation and so on.
The session could last just long enough to download the web page, or the web server and the computer could choose to keep the session alive for longer in case the user wants to keep browsing around on the web pages.

Step 4: The web browser talks with the Web Server

Once the TCP session is established by the OS then the OS will let the Web Browser know that it is now alright for it to start communicating with the Web Server.

The web browser will do this by using the HTTP protocol, which is the standard protocol for transferring web pages on the Internet

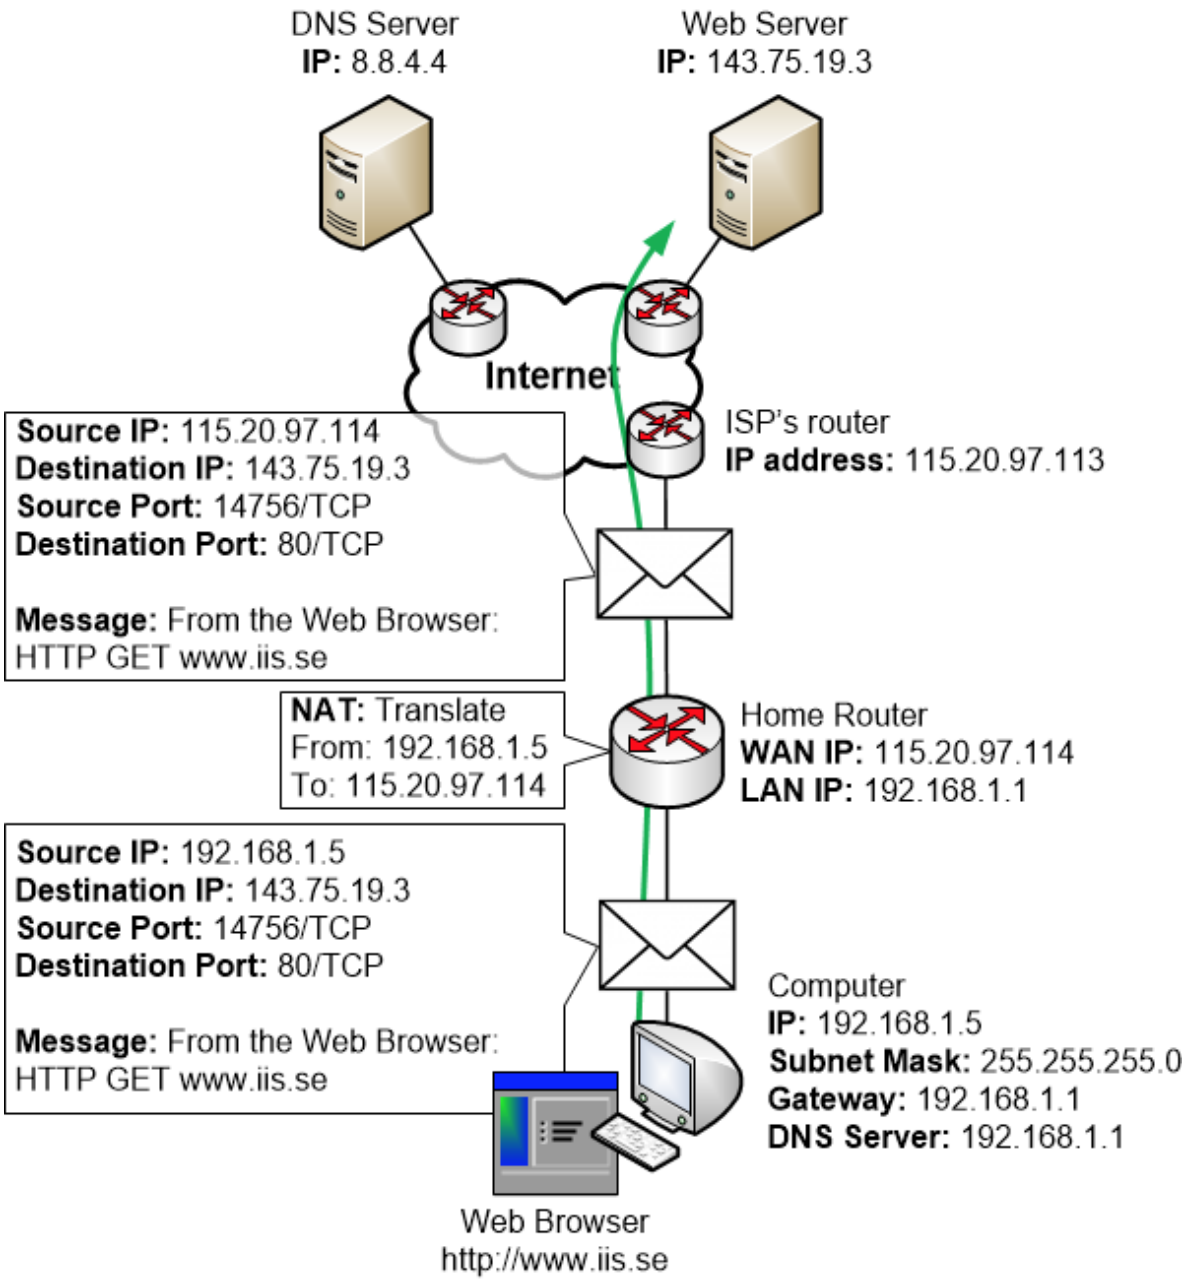

This also means that we have concluded the example including most steps that are involved with setting up the communication. From now on the computer and the server can communicate with each other to transmit the web pages until they are done. Then they can choose to end the session by sending a so-called TCP RESET message if they want, which lets all devices know that the session has now ended.

Specialisation: Speed and Size - Bits and Bytes

Transfer Speed (also known as *throughput* or *bandwidth*) is confusing in itself. But it gets extra messy when you also throw sizes into the mix ("how big is that file?"). The words for expressing throughput per second and file sizes are confusingly similar but are very different.
Often people use the incorrect term (bit vs byte) which is not surprising.

The background for the two terms is that the world of computers is binary. All information is broken down into binary digits called bits. A bit can either have the value 0 (zero) or 1 (one). Whenever you store something on your computer, such as a photo, then the picture is stored on the computer's hard disk drive in the form of binary digits or *bits*.

If you think about how we use numbers in the real world we don't always use the most basic form such as "gram". Instead, we say "kilos" or "Kilograms" (1,000 grams) which is much more convenient for most day to day use cases. The same goes for the terms score, dozen or gross. They are all used for convenience instead of the specific numbers they correspond to.

In the same way, we have the word *byte* in the computer world. A byte is 8 bits. One reason for why we started to use the word byte is that a lot of information that was stored on computers required 8 bits of data. For example, a normal typed character used to require 8 bits to be stored on a computer. It then makes sense to have a separate word to express the most common used number of bits.

So the basic formula to convert between bytes and bits is:

- 1 byte = 8 bits

1 bit is a very small value. Same thing with 1 byte. So we have to be able to add prefixes just like we do with weights (gram, kilogram…)

Prefix	Meaning	Unit
Kilo	1,000 (Thousand)	K
Mega	1,000,000 (Million)	M
Giga	1,000,000,000 (Billion)	G
Tera	1,000,000,000,000 (Trillion)	T

- 1 Kilobit is 1,000 bits
- 10 Kilobit is 10,000 bits

- 100 Kilobit is 100,000 bits, which could also be written as 0,1 Megabit
- 1000 Kilobit is 1,000,000 bits, which could also be written as 1 Megabit

But just to add some confusion - the above is just true for bits. For bytes the formula for converting between the prefixes is different. 1 Kilobyte is not exactly 1,000 bytes but rather 1024 bytes. We will show you further down how the conversion works for bytes.

Unit symbols

The unit symbols are really important when you talk about bits and bytes. It is one of the most common causes for confusion to use the wrong unit symbol.

- **bytes** commonly uses the unit symbol "**B**" with a capital B
- **bits** can use "**b**" as its symbol, but that is easily confused with the capital B for byte. So it is also common to use the full word "**bit**"

Since a byte is 8 times as much as 1 bit it is important to keep them apart and understand the difference.

Examples:

- MB means Megabyte
- KB means Kilobyte
- Mb or Mbit means Megabit
- Gb or Gbit means Gigabit

As previously mentioned, it is very common for people to be unaware of the difference between bits and bytes. A lot of the time the wrong term is used. This can lead to misconceptions about both files sizes and transfer speeds. It is more common for people to know file sizes than it is to know transfer speeds since you often work with files but rarely have to deal with bandwidth or throughput.

Sometimes you can see someone posting on the Internet because they are unsatisfied with their download speeds. They might have a 20Mbps Internet connection, but they can "only download files at about 2.4 MB per second!". Some programs that you use to download files report the download speeds using bits per second, whereas other programs might report the number of bytes per second. If you don't know the difference then you wouldn't even be able to spot the difference between "Mb" and "MB".

Based on the above, should they be dissatisfied with their Internet connection? Nope, on the contrary, they are getting really good results! Each Byte is 8 bits. So how many Megabit (Mbit) are 2.4 Megabyte (MB)? It is more or less as simple as calculating 2.4MB x 8 = 19.2Mb.

So in the example above the difference between being dissatisfied and satisfied is in the upper- and lower-case B!

But where did the last 0.8 Mbps go? We have a 20Mbps Internet connection, but we are downloading at 19.2Mbps. You will obtain the answer to that question further down.

Why this mix of bits and bytes?

We already touched on this subject, but for storage and for expressing space on hard disk drives it was much simpler to express file sizes using bytes. Files are big and most files were text files containing text. And since text consists of characters and each character took up 8 bits of storage it was easier to express storage space in the number of bytes that the storage could hold.

For data transfers and computer networks, however, it makes more sense to measure the number of bits per second that are being transferred. This is because the network equipment can typically only transfer one bit at a time.

Throughput, the simplified version

In this simplified version we skip explaining something called "overhead", but we bring it up further down.

Let's say we have an Internet connection handling 50 Mbps. Mbps stands for Megabit per second and is often written as either "Mbps", "Mb/s" or "Mbit/s". By now you will note the lower case b which stands for bit.

Now we are downloading a file that is 6.25 MB big.

- 6.25 MB is 50 Mbit (6.25 x 8)

So under optimal conditions, it would take a second to download the 6.25MB big file if your Internet connection is 50 Mbit/s.

How about if we instead download a file that is 4.5 GB big?

- 4.5 GB is about 4600 MB (4.5 x 1024). 4600 MB is about 36800 Mbit (4600 x 8)

In a best case scenario, it would then take 736 seconds (36800 Mbit / 50 Mbit/s) to download the file.

There are a lot of good online calculators out there on the Internet that can help you convert between Bytes and bits. Some of them can also help you calculate things like download speeds.

As you can tell there are good reasons for why confusion often arises around bits and Bytes. Not even the manufacturers themselves can agree on how they should perform the calculations.

Hard disk drives for examples are storage areas for files, and their capacity is measured in GB or TB today. But hard disk manufacturers measure hard disk capacity based on decimal base 1,000. So according to hard disk manufacturers, 100,000 MB equals 100 GB. However, the

computer OS uses binary base 1024 for calculating hard disk drive capacity. So when you try to store files on the hard disk drive you can't fit 100 GB on there.

This is how hard disk manufacturers do the math:

- 100,000,000,000 bytes = 100,000,000 KB (divided by 1,000) = 100,000 MB (divided by 1,000) = **100 GB** (divided by 1,000)

But this is how much storage space that you actually get as reported by the OS:

- 100,000,000,000 Bytes = 97,656,250 KB (divided by 1024) = 95,367 MB (divided by 1024) = **93,1 GB** (divided by 1024)

Throughput, the advanced version

Unfortunately, it is not enough to state that an Internet connection with 50 Mbps bandwidth could transfer 50 Mbit of data per second. Those 50 Mbit per second include all of the data that has to be transferred. Not only the data that you want to transmit but everything else as well.

There are several more things that need to be covered by that bandwidth, including for example the following:

- Overhead
- Session setup
- Application Data and Control Data
- Adaptive transfer rates

Overhead

Overhead means information that is not the data you and me actually want to send, but which has to be transmitted anyway.

When you send a letter to somebody by mail you write the letter on a piece of paper. Then you put the paper in an envelope and you type down the address of the recipient on the envelope.

The whole envelope has to be sent in the mail, even if the important message that you want to transfer is just the text on the paper inside of the envelope. The address that you put on the envelope has to be there, but it does not contain any of the important information that you wanted to transmit.

The envelope is the mail equivalent of Overhead. Required information which is necessary to transmit the message, but which does not belong to the message itself.

In computer communication, the overhead consists of for example IP addresses on the IP packets, MAC addresses, Port numbers, TCP or UDP information, and so on. All of which has to be transmitted with each packet but which is not the actual data you want to transmit.

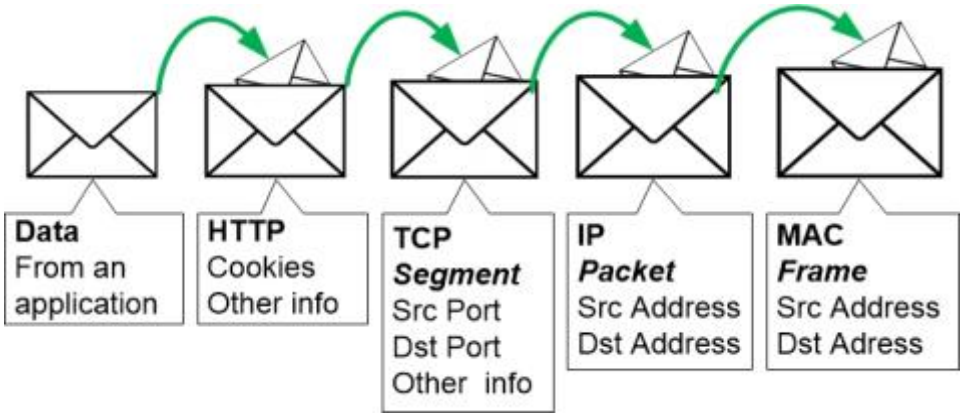

When the message is being sent to the network all parts of the message are being sent including the overhead information. A 50 Mbps Internet connection can only transmit 50 Mbit per second, including that extra overhead.

If the file that we want to transfer is big it also has to be chopped up in smaller parts, and these parts are then put in different IP packets. Each of those packets then need the same type of overhead information with addresses on it.

Normally each packet can contain a maximum of 1,460 Bytes of information. Then some additional 40 Bytes are used for overhead addressing and such. So about 2.7% (40/1500) of the available bandwidth is used to transfer overhead information, and this is just under optimal circumstances.

And if you have read the specialisation section with a full traffic example then you might remember the TCP 3-way handshake we showcased there. Network communication often has to be setup using sessions, which also uses bandwidth for the messaging.

Not to mention that a lot of applications are communicating in the background non-stop without you even knowing about it. If you perform a file transfer using Windows Explorer your computer will send hundreds of messages in the background to verify data, check the file system on the other computer, browse to the correct destination, figure out how much space is available and so on. All of this is information that you never see in Windows Explorer but which is still consuming available bandwidth. These messages contain Control Data and various types of background Application Data.

Adaptive transfer rate

Most file transfers use the TCP protocol. TCP always tries to help by being nice and not send more data than the receiver can handle, so TCP tries to find a sweet spot for how quickly to send the data. Not too fast and not too slow.

This means that often a computer that is going to send a file is starting out a bit slower while TCP is testing the territory of how quick transfer speeds that the connection and the other

computer can handle. The speed will then gradually increase until TCP has found what it thinks is the maximum transfer speed. But it is not a perfect system. The transfer speed will vary up and down because of little adjustments here and there.

The result is that it is rare for a single file transfer to be able to utilise all of the available bandwidth.

However there are some applications such as BitTorrent that work by using multiple simultaneous downloads from several different sources, and the different bits that have been downloaded are then assembled after they are fully downloaded. This makes it easier to reach higher speeds even over rather bad internet connections with high latency.

Wireless

Computer networks will function the same way on a higher level no matter if they are wireless or wired. The difference is on the lower levels of the communication, including how the signals are physically transported between the devices that are communicating. Nowadays wireless for home networks is almost exclusively Wi-Fi. Wi-Fi is a radio based wireless network following certain standards that we will discuss in this main section.

Wireless communication builds upon the same type of communication as wired LAN's. You still use IP addresses and MAC addresses, default gateways, DNS servers and so on. But instead of letting the signals and the communication travel over physical cables the signals will be transmitted using radio waves.

Despite the similarities wireless networks are still perhaps the most common causes for issues with computer networks, both in home networking and for corporations. There are bad coverage areas, people are constantly being disconnected, you have slow transfer speeds even though you seem to have a good signal strength and so on. Is it so hard to work with and set up wireless networks?

Yes, wireless networks are extremely complex!

Because wireless networks are so easy to enable they also seem to work by magic. You don't have to connect any cables, the signals are invisible and there are no issues that you can see or touch. No broken cables, and no connectors that are loose in their sockets.

But the "magic" in how wireless networks function is the very reason for why they also often work very poorly. Most people have close to zero knowledge about how radio works. If you don't understand what makes something tick then you also don't know what could cause it to **not** work. And once it starts working poorly you also don't know how to troubleshoot it and do not know about which underlying causes that could exist for the problems that you are noticing.

All wireless devices are also sold in a way to let you believe that it is really simple to hook them up and start using them, which can also be true...if you live out in the woods 500 meters away from any neighbors and you don't have any furniture indoors. You also live in a one-room apartment with only thin drywalls. You don't own a microwave oven or any Bluetooth devices. You also don't have any other "magic" wireless electronics such as wireless mice or keyboards, no Smart TV, no internet connected devices etc. In fact the only wireless devices within a kilometer of your house is that single wireless laptop you bought. Then it is almost guaranteed to work!

The wireless products themselves are often just bad. You should never underestimate the ability of manufacturers to produce bad quality wireless products without any comprehension for how wireless actually works.
If a wireless device is cheap then, unfortunately, it will almost always be noticeable. It might still work! But it will be more susceptible to the inherent problems of Wireless networks and won't have any of those more advanced features that are associated with higher quality equipment, features which help to alleviate those problems.

Why is wireless difficult?

Wireless communication is difficult because you more or less have to have a background in radio communication to understand radio based wireless networks. Most people simply don't know anything about radio communication other than perhaps that it consists of electromagnetic waves of some sort that can be used to relay information. There's also something about frequencies that you can tune into, oh and you can listen to different radio channels!

The biggest challenge in wireless networks are the radio signals themselves and how they affect and disturb each other. Among other things the radio signals:

- are disturbed by some electronics and electromagnetic fields
- are disrupted by the air itself and by microscopical contaminations in the air
- can travel fairly easily through some material but not through others
- are weakened by various amounts in different material
- spread out through the air in different patterns depending on the antenna that is used
- depend on the drivers in the computer and how those drivers instruct the computer to behave on the wireless network
- are blocked very effectively for example by concrete walls, steel doors and newer types of windows and glass panes with isolating qualities
- bounce around rooms to create standing waves, localised disturbances and areas with bad signal quality

There are simply a million different things that affect how your wireless network behaves.

As a matter of fact in a country such as Sweden which generally has a lot of extremely qualified and highly skilled and educated engineers you could only find a handful of people with enough skills, knowledge and experience to be able to set up a perfectly functioning wireless network in a complex environment such as a sports arena or a big shopping mall. It is simply too complex.

At the same time, there are hundreds of network technicians who work daily with rather big wireless network installations in office environments and who manage just well, but who still don't really understand exactly what they are doing. Many of those people could easily install a wired network on that same sports arena. That probably says a lot of how complex wireless really is.

Anybody could set up a small wireless home network, and since those networks are usually really small it almost always works. It might not be perfect, you might have some areas where you have a bad signal strength and so on, but it mostly works.
However if you let a skilled wireless network engineer perform a wireless network installation at your house with good quality equipment then the network could work so much better!

What does the section about wireless contain?

This section will provide information on how wireless networks actually work. It contains some basic radio knowledge to let you understand the concepts behind wireless communication and why it can be difficult.
Then we go through different building blocks of a wireless network so that we have the groundwork ready. This is also where we go through some of the terms that will be used in later wireless sections.

We also look at a few of the ways in which you can build up wireless networks and explain what the differences are and which Pros and Cons you get from using the different building blocks.

Finally, we suggest a few solutions on how you can build your wireless home network for better results compared to setting up the same network without any knowledge.

Specialisation: Basics of Radio Communication

This section contains a lot of theory and basic information. The information is important to understand how Wi-Fi works.

Wireless Wi-Fi networks work by radio. All Wi-Fi enabled devices have an antenna by which they can send and receive radio signals.

We will kick off this section by discussing radio waves and a little bit about how radio communication works.

About radio waves

Radio signals are transmitted by an antenna with a certain output signal strength or *effect* as it is called. As the radio signal is spreading out from the antenna it becomes weaker and weaker the further away from the sender it gets.
Eventually, the signal strength is so weak that the signal can no longer be understood by a receiver. This is similar to how you cannot hear somebody from too far away because the sound is weakened and spreads out with distance. The further sound waves travel the weaker they become.

There are many reasons for the weakening of the radio signal. Here are a few examples:

- As the signal is spreading out the effect of the signal is also spread out over a bigger area. Much like if you throw a pebble in a pond the height of the waves decreases the further out from the impact that the waves spread. The energy of the pebble's impact

which transferred to the water has to spread out over a bigger and bigger area the further the waves travel
- Absorption, that the effect or energy of the signal is simply absorbed by the air itself, the water in the air, particles or other things so that the energy of the signal is lost
- Scattering, that the signal is split up after colliding with particles, objects or edges. One example of scattering includes how a rainbow can be created when light is split up into different wavelengths after hitting water droplets in the air

Radio signals are electromagnetic signals just like light. And similar to light, radio signals can be reflected by or bounce on things. This can be both positive and negative depending on the circumstances. In a small home it might not matter much if radio signals bounce around a bit, and sometimes it can even be beneficial, but other times it can be a bother for the communication.

A radio wave can act much like regular waves in the ocean. If two waves meet up and their crests and troughs line up then they will form a new much bigger wave, much bigger than the individual waves were.

If two waves meet up and the troughs of one wave line up with the crests of the other wave then the two waves will cancel each other out. If there's a perfect match between the two waves for both their wave heights and their wave lengths then they could cancel each other out more or less completely.

Radio Waves act in the same way. Two radio waves that meet up will interact with each other to increase and decrease the signal strength. In the following picture we can see two radio waves (dashed lines) with different wavelengths and wave heights (amplitude) that meet up, and together they form a new combined radio wave (solid line) that appear very different to the original radio waves.

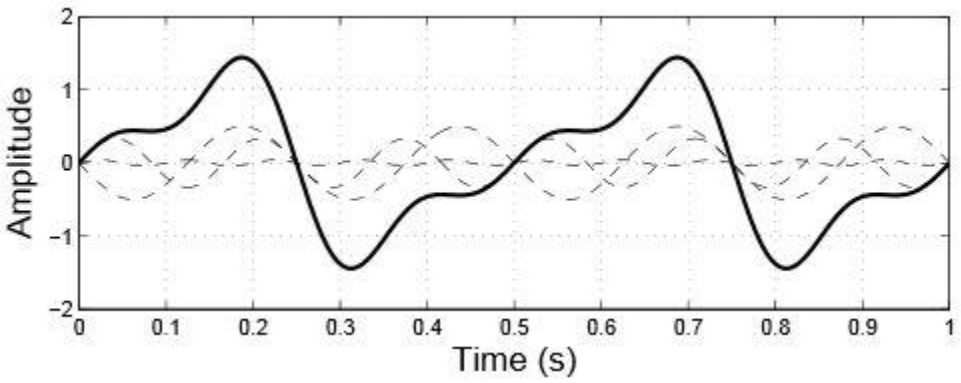

Even if you just have a single radio transmitter the radio waves are spreading out in all directions and will bounce around the room. Basically, everything you see around you is something that radio waves can interact with by being reflected, diffracted or scattered.

Extremely simplified you could say that radio waves pass through softer materials easier and are reflected against harder surfaces. The result is that the air is filled with radio signals bouncing around in all directions. So even with a single transmitter sending out a single signal you can still get interference between the radio signals that are bouncing around.

Imagine putting a Wi-Fi router in the room you are in. When the radio signal encounters walls, furniture, ceiling and all other things the radio signal will be affected by the materials it encounters. Some of the signal effect will pass through the items, some is absorbed and some of it will reflect to bounce back into the room. This happens on every surface you see, so the room is literally filled with bouncing radio signals.

The bouncing signals are often much weaker than the original signal directly from the antenna since not all of the signal will be reflected. If you were to put the Wi-Fi router right next to a concrete wall or steel plate then a lot of the signals would bounce on the hard surface right after leaving the antenna, and the bouncing signal is still strong enough to definitely interfere the antenna signals considerably.

You always get the strongest signal strength with the best signal quality if you have a completely free line of sight to the antenna of the Wi-Fi router. Anything that is blocking the signal will always affect it in some way. For that reason, it is always absolutely best to place a Wi-Fi router in an open space with a free line of sight to as many places that you want to cover with your Wi-Fi signal as possible.

If instead you were to place your Wi-Fi router in a closet, under a desk or in a box in the hallway then you've already ruined any chances of achieving perfect radio signal quality. You could still end up with a wireless network that works well enough, but it will be considerably worse than if you would move the Wi-Fi router to an open space.

About antennas and radiation patterns

A radio antenna can have different appearances depending on what purpose it is designed to fulfil. Antennas also have a different length depending on what wavelength that they are made to transmit and receive radio waves on.

Antennas are manufactured to radiate radio waves in different patterns. The manufacturing process and the design of the antenna can control in which particular directions (if any) that the antenna should focus its radio wave emissions.

The most extreme example is *parabolic antennas*, where the whole antenna is built to send (or receive, or both) as much of its effect as possible in one single narrow beam. This design concentrates all of the energy of the radio transmission in one single direction. Parabolic antennas are used for longer distance communication where you have another single antenna far away that you want to establish communication with.

But there are also many other variations on so-called directional antennas that direct radio waves in one or several directions.

The absolutely most common type of antenna for Wi-Fi routers and home networks is the dipole antenna which has an omnidirectional radiation pattern. Such an antenna will be more or less equally powerful in all directions, which is what you usually want in a home network environment. Also, the manufacturer couldn't know what type of antenna that you would otherwise require in your particular home, so they always default to including omnidirectional antennas.

This is another reason for why you should place your Wi-Fi router in an open area of your house. The router will radiate radio waves more or less equally in all directions. If you place your router in a corner then some of that radio effect will enter the walls and be absorbed right away without producing any positive results for your Wi-Fi network.

Many routers are equipped with internal Wi-Fi antennas with a particular *antenna alignment*, and the router is then meant to be placed in a particular way for the antenna to have the correct alignment. You should always look at the instructions manual for your router and place it in the correct way standing up or lying down. Otherwise, the radio communication won't be working under optimal conditions.

If the router has an external antenna that can be set in different angles, then you can often choose how you want to place your router, and then pick an angle for the antennas. Often the manual for the product will have instructions for antenna alignment. Basically, it is best if the antenna alignment of the Wi-Fi router matches up with the antenna alignment of the devices that you want to connect to the Wi-Fi router.

Matching antenna alignment is not always as straightforward as one might think. Most devices have built-in antennas that you can't see, so you don't know what alignment they have. You can't look at a mobile phone and see what antenna alignment it has got. Also, how will you hold your mobile phone when you are using it? If it is lying down then the antenna is lying down. If you hold it straight up then the antenna is upright. Or perhaps sideways. Luckily a phone rarely needs a high-speed Wi-Fi connection so you will often get away with phone Wi-Fi quality.

But look at your laptop instead. The antenna is often integrated into the screen and is upright when you are using the laptop, but not always. For example, a lot of MacBooks have their antennas integrated into the screen hinge, lying down.

If you have multiple antennas on your router you could face some of them straight up and others lying down. But once again whether or not this is beneficial depends on the type of router you have and what capabilities it supports. Many manufacturers recommend angling the antennas. For three antennas you would have the middle antenna somewhat straight up, and then angle

the two side antennas 45-60 degrees. But check with the router manual to see what their recommendation is for your particular home router.

Radio channels, an introduction

If you've ever listened to the radio then you know that there are different radio channels you can tune in to such as Kiss FM, AMP or Now FM.

Then you probably also know that those radio channels are operating on different Frequencies that they are allowed to transmit on. Different countries have different governing bodies that rule on how the radio frequencies can be used. In the USA for example the Federal Communications Commission (FCC) regulates radio communication, including who can use which radio frequencies for transmitting FM radio.

This is necessary because if two channels were to transmit on the same frequency within range of each other then they would disturb one another. There must also be some unused frequency space between the channels. This is because radio signals are never transmitted just on an exact frequency such as 99.3MHz. The signals are actually "smeared out" over a little range of frequencies surrounding a center frequency. This is called the *channel width*.

So if a radio channel is transmitting on the frequency 99.3MHz, then it is really just the center frequency which is 99.3MHz, but the radio signals can also be heard over the nearby frequencies such as 99.2MHz or 99.4MHz. You can even hear this effect if you have an old radio available. When you tune in a channel you can hear how the channel first appears with some distortions as you close in on the right frequency, and then as you get right on top of the center frequency you get perfect sound quality.

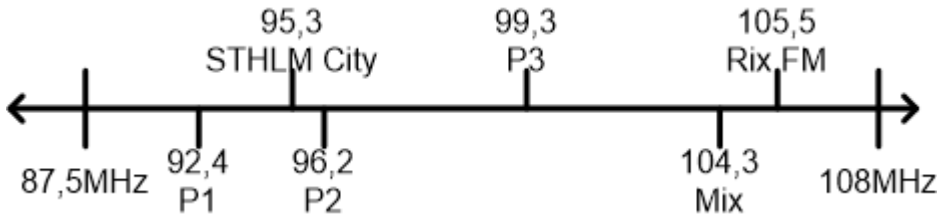

Additional Info

The **frequency** is simply explained how long or short the wavelength of the radio signal is, which affects how many times per second the wave swings up and down or oscillates. With higher frequencies the wave oscillates more often and the wavelength is also shorter. Lower frequencies mean fewer oscillations of the wave per second and the wavelength becomes longer.

When you listen to the radio in your car or kitchen you often listen to the so called FM band of radio frequencies between about 87.5 - 108 MHz. One Hertz (Hz) means one oscillation per second (one wavelength). MHz stands for Mega Hertz - million oscillations per second.

But let's say that a radio channel is transmitting in the Stockholm area on the frequency 96.5MHz.

Then nobody else can transmit a radio signal on 96.5MHz there because the radio signals would disturb each other. But how far away do you have to be to put up another transmitter on 96.5MHz without causing disturbances? The simple answer is that it depends on how strong the radio transmitter is, or which *effect* it has, and also on what type of antenna that is used and how that antenna focuses the radio transmission. The important bit is that if two radio senders use the same frequency then their radio waves will disturb each other *if they are within range of each other and their transmissions overlap.*

So as long as two transmitters send on the same frequency their signals will disturb each other if they overlap each other's coverage areas. In the picture below the two radio towers are sending on the same frequency, and there is a huge overlap in the middle of their signals. That whole overlapping area is then affected by the disturbances, so it will be very difficult in that overlapping area to listen to the radio signal that either tower is transmitting.

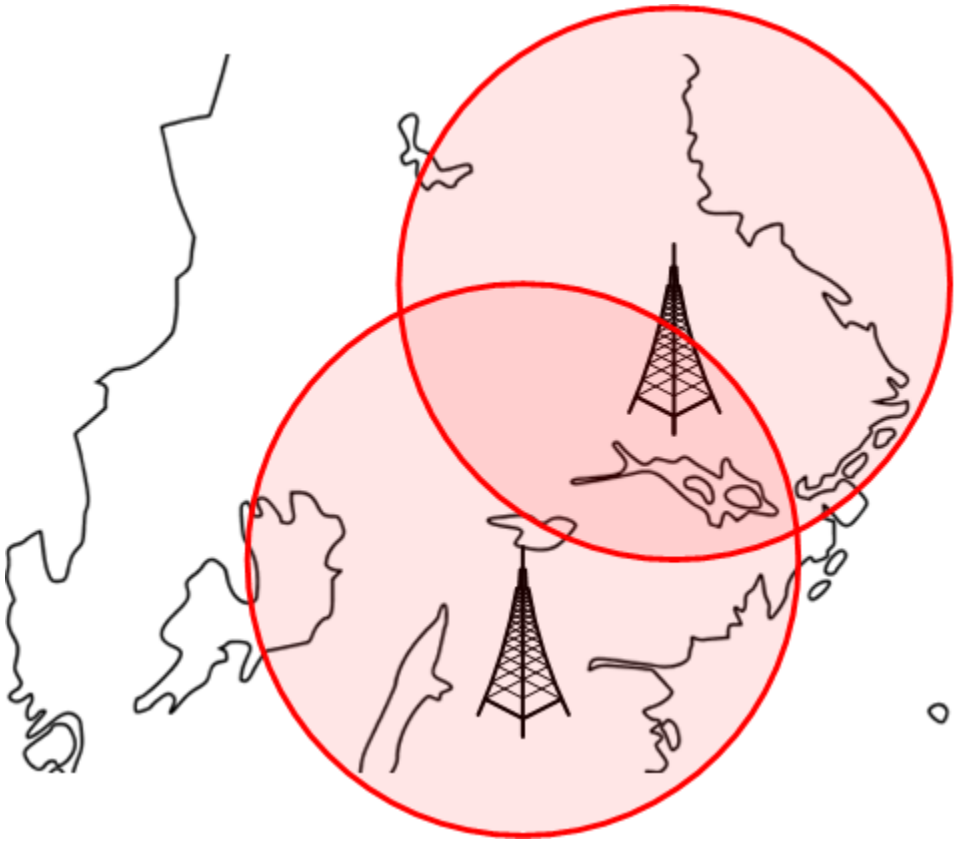

The frequency 96.5MHz can be used at many different places simultaneously. The important thing is that the transmitters are not within reach of each other.

But like we stated earlier the radio signals are only centered on a frequency. In reality, the signals are transmitted over a range of frequencies. How much the signals are smeared out over that frequency range depends on the exact technology that is being used to transmit the signal. For radio on the FM band signals are often about 100kHz "wide", for example from 96.45MHz up to 96.55MHz. Or from 105.4MHz to 105.5MHz. It is usually not allowed to use frequencies even close to those of other radio channels. Instead, the radio channels that FCC and other governing bodies hand out are well distributed with enough gaps between the different radio channels so that they do not disturb each other.

As long as you just have a single transmitter you won't have any issues. The problem will always arise when you want to cover a larger area with multiple transmitters that should transmit simultaneously for extended coverage, but you don't want the transmitters to disrupt each other. That's when the puzzle of radio channels begins!

A puzzle of channels and transmitters

Each radio transmitter on the FM frequency band has a reach ranging from that of a city, or sometimes a county or even somewhat bigger areas.

But the further the distance from the transmitter the weaker the signal gets, and eventually the signal is weak enough that you cannot listen to it without also hearing disturbances. That's when you have to put up another transmitter to increase the total coverage area. The new transmitter must be close enough to the first one so that the radio signals overlap. Otherwise, you couldn't go in your car between the two areas without losing the signal. So the signals must overlap.

The puzzle is in the fact that the two transmitters must send out the same radio show, but they must use two different frequencies so that they can overlap and be within range of each other without disturbing each other by transmitting on the same radio frequency.

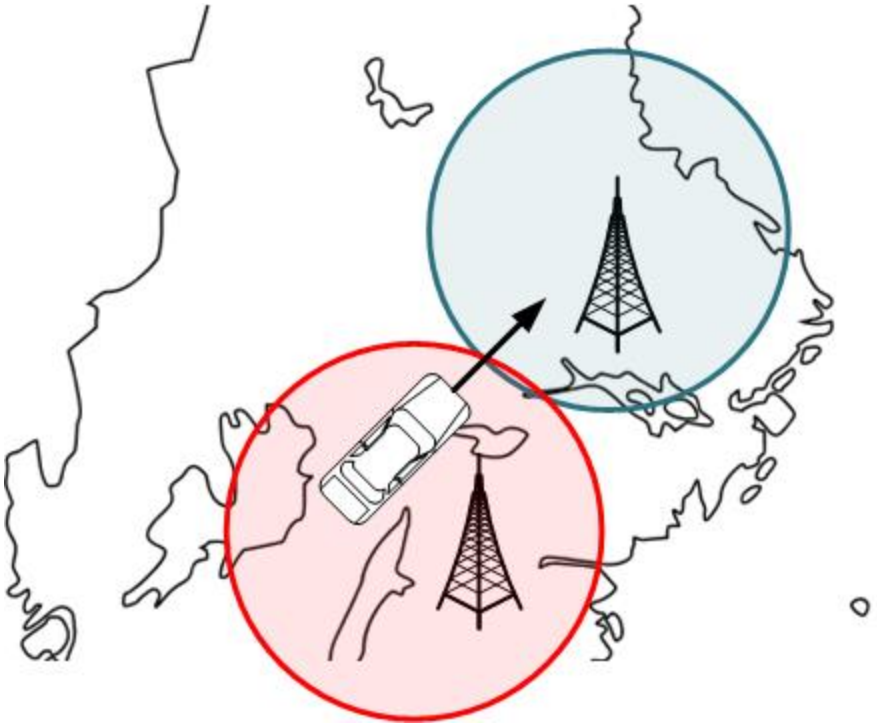

The two transmitters in the picture above are far away from each other. When going by car from one transmitter to the next, you will probably cross areas with bad coverage where of the transmitters have a good enough signal. To solve this problem the senders must be closer together, or the effect of the senders could be increased so that each sender gets a farther reach.

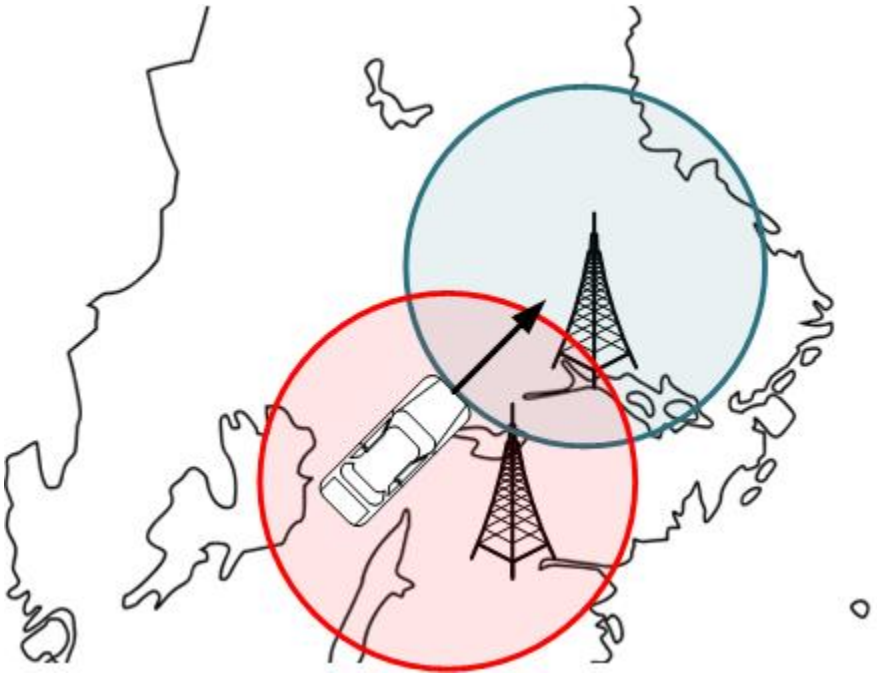

In the picture above the senders have been moved closer to each other, which means that the overlapping area between the two senders is bigger. Then it gets much easier to ensure that the overlap between the senders is big enough so that you always have a good signal from at least one of the towers when you go from one area to the next.

The transmitters above are using different frequencies to broadcast the same radio show. Car stereos handle this automatically these days by moving between frequencies without you noticing.

The puzzle I mentioned starts when you want to cover a big area and you also have many radio shows that you want to transmit simultaneously. Let's say each transmitter needs to broadcast 10 radio shows. That means each transmitter will work on 10 different frequencies, and none of the channel frequencies must collide with any other channel frequencies on any other radio tower within range. Then you have a much more challenging puzzle that requires planning from someone who can control exactly what frequency that is used for which show in what area.

Without that control, it would be up to each individual company to coordinate their radio usage with every other company. For North America, there are more than 9000 radio channels that would have to coordinate how they transmit their programs, which would quickly become an impossible task.

Here we have drawn up a simplified radio frequency planning puzzle. In this picture, each radio tower only transmits on a single frequency. Each frequency is represented by a unique color. What we need to do is two things:

1) We need to cover the whole area, so there must be no place on the map that is not covered by at least one radio tower
2) None of the towers must transmit on a frequency that overlaps with another tower within reach that sends on the same frequency

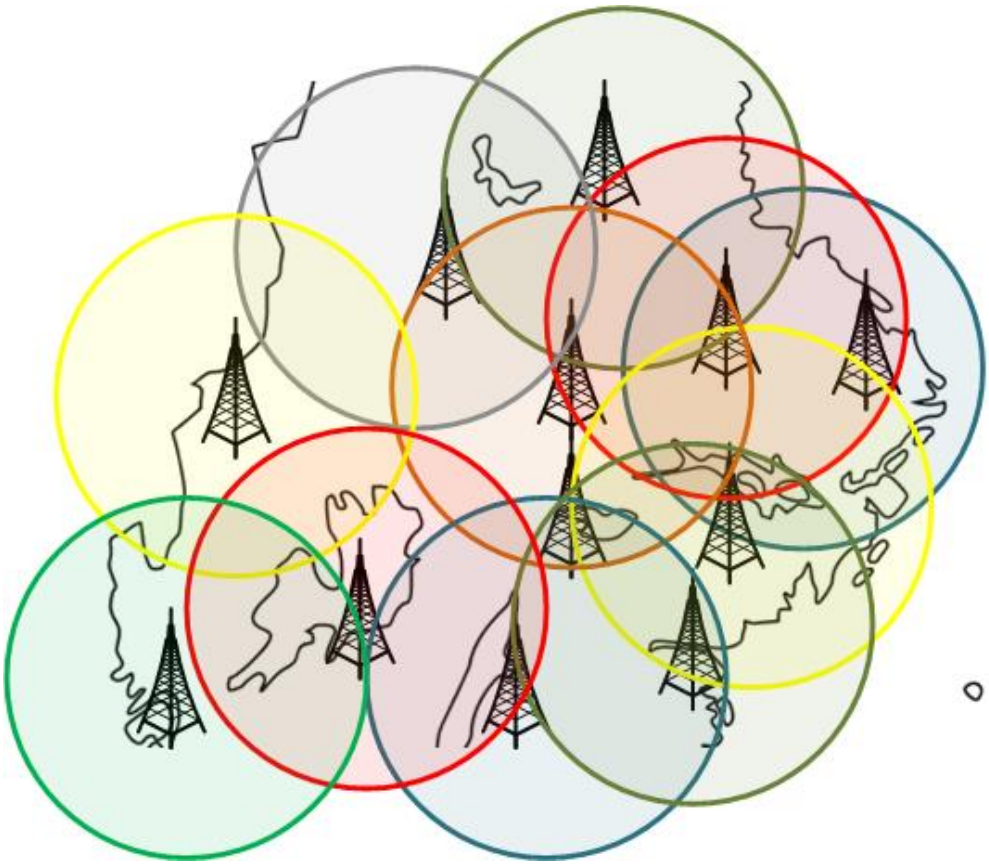

In the picture above we have succeeded since the whole map is covered and no same-color circles are overlapping.

Radio theory in Wi-Fi networks

This section contains a lot of theory and basic information. The information is important to understand how Wi-Fi works.

Wi-Fi is a radio based wireless network. Wi-Fi is often called *WLAN* (Wireless LAN), *Wi-Fi*, *WiFi,* "wireless networks" or just simply "wireless".

A wireless LAN needs an access point that computers, phones, and other devices can connect to. Those devices that connect to the wireless network are called *clients*. In a normal home network, the router usually fulfils the role of the Access Point. since most home routers have one built in. But you can also buy separate access points that you can connect to your network in case your router does not have built-in wireless LAN or if you want to upgrade your wireless network.

By applying the knowledge from the advanced section on radio theory on Wi-Fi networks, there are a few things we can conclude right away

- Having a free line of sight between the radio transmitter and radio receiver is always best. Put the access point in an open space. Not buried underneath things, inside cupboards or under furniture.
- The access points and their antennas are made to be placed in a specific way with a specific antenna alignment. Make sure you read the manual of your router or access point.
- Wi-Fi uses radio signals and functions in much the same way as FM-radio when it comes to overlapping frequencies and radio channels. We will look at this in more detail later.

Wi-Fi uses radio signals on two main frequency bands:

- 2,4GHz
- 5GHz

Both of these are so-called unlicensed frequency bands, which means that agencies such as FCC and other international and national governmental bodies have decided that anybody can use these bands freely without first applying for permission, as long as you follow a few rules on things like what effect your equipment can use.

But the two frequency bands are not only used for Wi-Fi networks. A whole range of other wireless things also utilise the same unlicensed bands. There are even devices which do not even communicate using radio waves that still emit radio waves on those frequencies:

- Wireless Headsets
- Wireless mice and keyboards
- Surveillance cameras
- Anything that has to do with Bluetooth
- Some baby monitors
- Many wireless phones
- Wireless speakers
- Microwave ovens
- Strip lights
- Radar

So there are not only a lot of things that communicate using radio over these frequencies. There are also a lot of non-Wi-Fi equipment that happen to use the same radio frequencies or might emit radio signals on those frequencies as a side effect when they operate.

A lot of people are used to just looking at their laptops or their phone to see what signal strength that they have. But the truth is that you can have full signal strength and still have a really bad signal quality. In fact, you can have full signal strength and still have enough disturbances from non-Wi-Fi appliances that the wireless network is completely unusable.

Channels on Wi-Fi networks

Just like the FM band is divided into radio channels the Wi-Fi network also uses channels. Exactly which channels that are available differ a little bit depending on what country you are in. The reason is that not all countries have defined the unlicensed radio bands in exactly the same way. Some have included a wider frequency range, and others have a bit smaller ranges available. This makes it possible to fit a different number of channels into the available frequency range depending on where in the world you are.

Channels on the 2.4GHz band

There are usually 13 channels available on the 2.4GHz band, channels 1 - 13. Unfortunately, there is a complication with how they can be used. Since Wi-Fi channels are very wide and take up a lot of frequency space, each channel actually overlaps with a lot of nearby channels. The width of a channel is counted in how many MHz that the channel is transmitting data over.The exact width of each channel can vary somewhat depending on exactly which Wi-Fi standard that is used, but the basic channel width is either 20 or 22 MHz.

In the specialisation section on radio communication basics, we stated that you should always avoid using overlapping channels within reach of each other. This is because two radio signals that collide (and they do so when their frequencies overlap) will disrupt each other. Since most of the 13 channels are overlapping you are only really left with 3 channels that can be used completely without any overlap.

This is referred to as there being 3 non-overlapping channels on the 2.4GHz band.

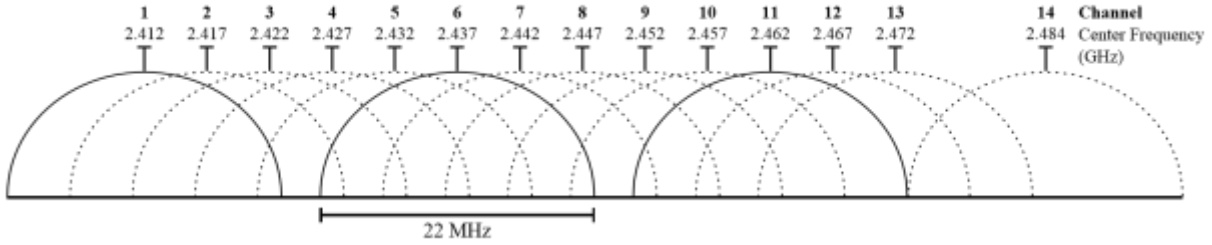

- Channel 1 is overlapping with every channel up to and including channel 5
- Channel 6 overlaps every channel from channel 2 up to channel 10
- Channel 11 overlaps all channels from 7 to 13

You can use channels 1,6 and 11 on three different Wi-Fi transmitters without them overlapping each other.

Of course you could also use channels 1,6,12 or 1,7,13 or 1,8,13. But there is a reason for why it is best to use 1,6,11. This is because if two Wi-Fi devices within reach of each other are running on the exact same radio channel then they will detect one another and can at least try to work together to minimise disturbances. But if they run on different channels that are overlapping then they cannot detect each other and thus cannot do anything about the interference.

If you and your neighbour both have an access point each running on channel 1, then your access points can detect each other and make sure that they try to disrupt each other as little as possible. But if your access point is running on Channel 1 and your neighbours have their access point on Channel 2, then there will be interference between the two that the access points cannot control.

By now it has more or less become a silent agreement among network professionals to simply stick with the channels 1,6 and 11 and only ever use those channels on the 2.4GHz band. This will ensure that any overlaps between nearby devices have the least possible negative impact on the wireless network communication.

Having said that, even if you use channels 1,6 and 11 you can and will run into problems with overlaps. If you live in a building with more than 3 households or if you want to mount more than 3 access points in your apartment then you might encounter channel overlap.

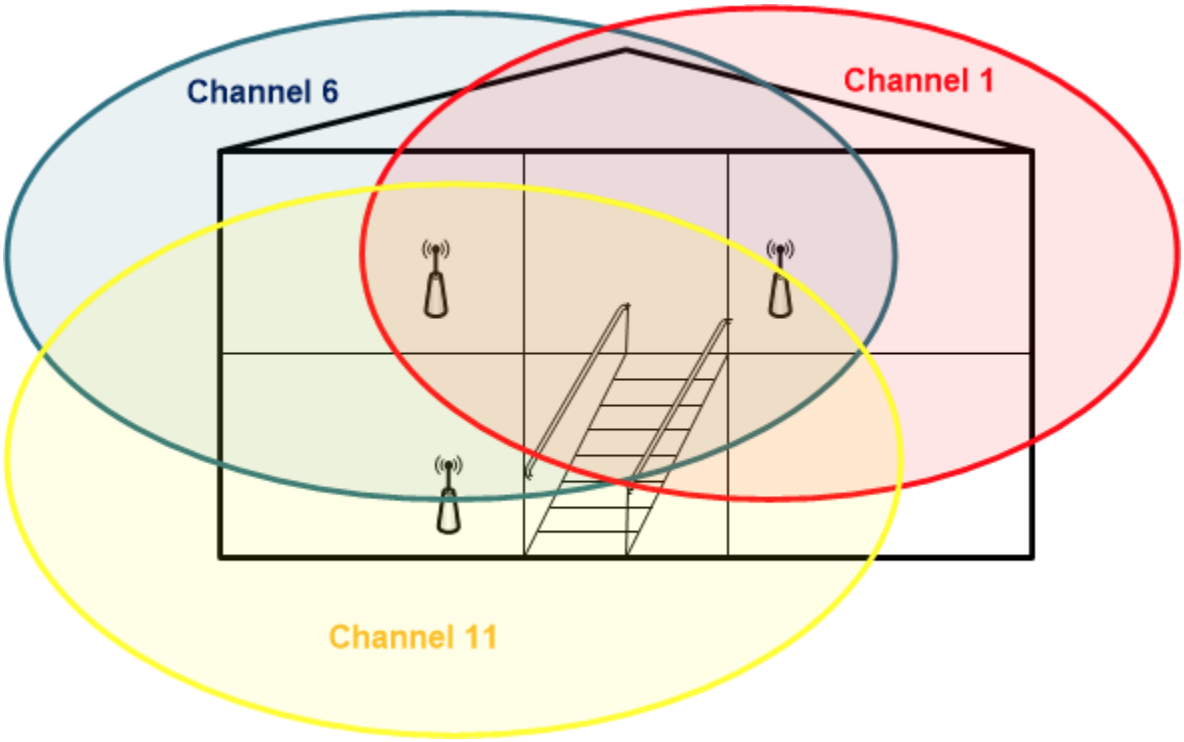

As long as you only have three access points and as long as every access point is transmitting on a channel of their own then the overlaps don't matter. Those channels are non-overlapping and thus do not interfere with each other.

But if a fourth access point is added then there simply aren't any non-overlapping channels left on the 2.4GHz band, so it is difficult to avoid interference completely.

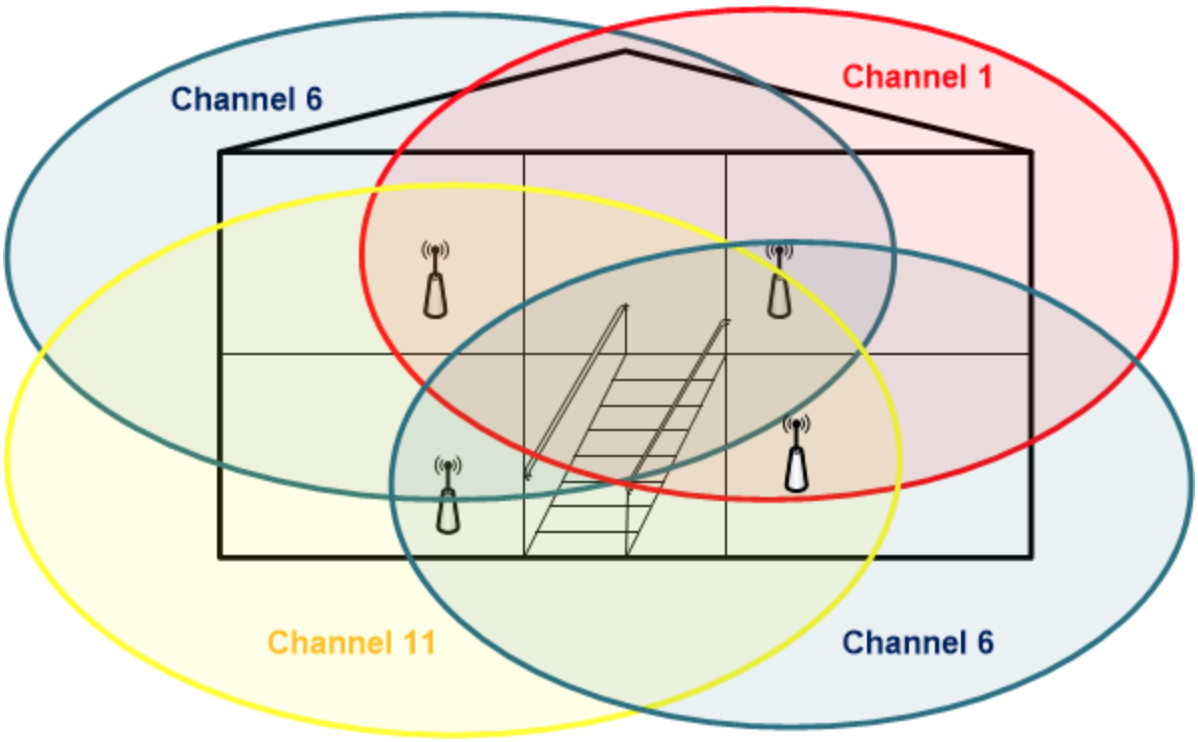

Note in the picture above that only the two access points that are using channel 6 will experience interference. The other two access points on channel 11 and channel 1 are completely unaffected.

The signals also overlap mostly in the central staircase where nobody really uses their Wi-Fi much. The signals will also have travelled through a number of walls and as such will be weakened enough that the interference might not be too bad.

There is a possible solution to the scenario if you have control over and can administer all four access points. You can then experiment by decreasing the output effect of the top left and the bottom right access points that are running on Channel 6. Decreasing the effect will shrink the radio cell sizes and might alleviate the problem.
This is an effective solution if you have a lot of access points that overlap. Within each room where you mount the access points you will still have good signal quality, but the wireless signal won't reach as far into the neighbouring radio cells.

Channels on the 5GHz band

There are many more non-overlapping channels available on the 5GHz band. The 5GHz channels are spread further apart from each other so none of the available channels are overlapping. The number of channels differs from country to country. In Sweden, for example, there are close to 20 channels, and since none of them overlap they can all be used simultaneously within reach of each other without interference.

This can be of huge benefit when you are building a wireless network. First of all, there still aren't as many people who are using the 5GHz frequencies so there is less competition for the available channels. Secondly, even when people do use the 5GHz band there are more available channels to begin with so you do not run out of free channels as quickly. And third, a lot of equipment still doesn't have support for 5GHz wireless which constricts that equipment to using the 2.4GHz band.

Modern routers and access points often have an automatic radio channel selection that they can and will use. They do this by listening in on the surrounding Wi-Fi environment to see which channels that seem to be taken already by other nearby equipment. Then they pick a channel that seems to be the least utilised.

Many newer Wi-Fi routers and access points can even run on the 2.4GHz and 5GHz bands simultaneously. This is called Dual Band and is discussed later on in the guide.

Another thing that is also explained more in detail later is that newer routers and access points support higher speeds by using wider Wi-Fi channels. Instead of using a normal 20MHz wide 2.4GHz channel a router can use a twice as wide 40MHz channel to double the amount of data that can be sent. Those wider channels will of course also take up more of the available frequency space. If you double the channel-width of 2.4GHz Channel 1, then the resulting wider channel will overlap most other available channels, leaving you with only two non-overlapping channels - Channel 1 and Channel 11.

As a final note, remember that these radio bands are unlicensed. Your neighbours might not have a clue when it comes to wireless networks, but if they choose to set up 13 access points that absolutely overwhelm every single available 2.4GHz channel with constant interference, then they have the right to do so. You could of course ask them to cooperate, but your best bet otherwise is to move to the 5GHz band where there is hopefully less contention!

Throughput on Wi-Fi networks

A *theoretical maximum throughput* is marketed for all Wi-Fi equipment. But what does it mean? Why is the maximum throughput only achieved in theory and not in practice? If the maximum theoretical throughput is never actually achieved, then what throughput should you be expecting?

In one of the specialisation sections, we discussed *overhead* which is extra data such as IP addressing, MAC addressing, UDP or TCP ports and other information that is required to send data packets over the network. The overhead information can consume several percent of the available bandwidth, but will do so regardless of whether or not your network connection is wireless.

Wireless networks in particular are very prone to disturbances from outer sources. Therefore, wireless devices have to include a lot of extra *verification data* or *correction data* to enable the

receiver to check for any errors that might have been introduced via interference during the transmission. This lets the receiver figure out what the original message contained even if something went wrong during the transmission and parts of the message are scrambled by disturbances.

Part of the correction data works in a similar fashion to the phonetic alphabet that the military uses. Instead of spelling out things as simple as possible (*R-O-U-T-E-R*) the military use a phonetic alphabet to replace each letter with a spoken word. Using the phonetic alphabet the word *Router* would be transmitted as *Romeo - Oscar - Uniform - Tango - Echo - Romeo*.

Even if small parts of the message end up being unintelligible due to interference, disturbances or surrounding noise levels, the risk of somebody mishearing a full word is significantly lower compared to if single letters had been transmitted instead. In fact, most of the message could be lost in transmission and the receiver could still figure out what the original message was.

Wireless communication works in a very similar way. Instead of sending a binary bit that is either “zero” or “one”, a wireless device might transmit a whole series of bits that together mean “one” or “zero”. Even if parts of the transmission is scrambled the receiver can still figure out the meaning of the message.

The devices will also automatically lower their speeds the worse the signal quality and signal strength gets. When two devices move further apart they communicate more clearly and slowly to overcome any disturbances. This means that to even have any remote chance of experiencing the best possible transfer speeds on a wireless network you have to be really close to the Wi-Fi access point. Possibly within just a few meters at most.

All the above combined means that you rarely can count on actually seeing any data transfers reaching more than half of the theoretical maximum throughput that is promised in any given wireless standard. If you reach 25 Mbps on a 54 Mbps Wi-Fi network you should probably be happy with your results. The same thing goes if you achieve 600-700 Mbps on a 1300 Mbps Wi-Fi network.

So why is the theoretical maximum transfer speed even marketed then? The simple reason is that the manufacturers cannot make any assumptions whatsoever about your radio environment or how you install or use the equipment. It doesn't make sense for the manufacturers to market numbers based on assumptions that will rarely be correct. Although one might argue that they could have avoided a lot of dissatisfaction and misconceptions over the years had they only chosen to communicate more clearly what those theoretical maximum download speeds actually mean.

Wi-Fi is (normally) Half Duplex

This is just a reminder that Wi-Fi communication is normally Half Duplex, just like communication through hubs. Only one device can communicate at a time on any given radio

channel. If two devices on the same channel transmit simultaneously they will disturb each other. So Wi-Fi networks behave just like any other typical Half Duplex network.

If you want to reach higher throughput speeds it is almost required that you are the only one who is actively using the wireless network. It doesn't matter too much if a few more devices are connected to the network, as long as they are silent and do not communicate. Disruptions will only occur between devices that are actively transmitting data onto the wireless network.

However, please note that computers continually communicate in the background no matter if you are actually using them or not. The problem with the shared bandwidth also increases with the number of computers and other devices that you hook up to the wireless network.

There is one exception to the Half Duplex Wi-Fi rule. The newest 802.11ac standard with support for something called *MU-MIMO* lets the access point act more like a switch rather than a hub under certain conditions. We will explain MU-MIMO and other subjects in the next section.

MIMO explained

MIMO stands for *Multiple Input, Multiple Output*. It is a function that uses multiple antennas which communicate on the same radio channel, but it does so using clever radio techniques to avoid disturbances. One such technique is called *Spatial Multiplexing*, which means sending multiple data streams or *Spatial Streams* at once. These Spatial Streams are precisely engineered so that they do not cause interference even though they are sent on the same radio channel.

To be able to use MIMO communication both the Access Point and the Client must have multiple antennas and support for the same MIMO standard. Sending multiple streams at once means that the total throughput can be increased. This is what the *multiplexing* term means. You send multiple separate data streams that are then combined or *multiplexed* to form a total combined higher throughput.

Different Wi-Fi standards can use a different number of antennas to achieve higher throughput by using MIMO. Support for MIMO first started with 802.11n, and then was further improved with 802.11ac.

SU-MIMO

SU-MIMO stands for Single-User MIMO. It means that all the available antennas can be used to communicate with one client at a time to increase the bandwidth to that client. So even though there are multiple antennas they can only be used to communicate with a single wireless device at a time.
SU-MIMO is used in the 802.11n and 802.11ac standards for devices that support it.

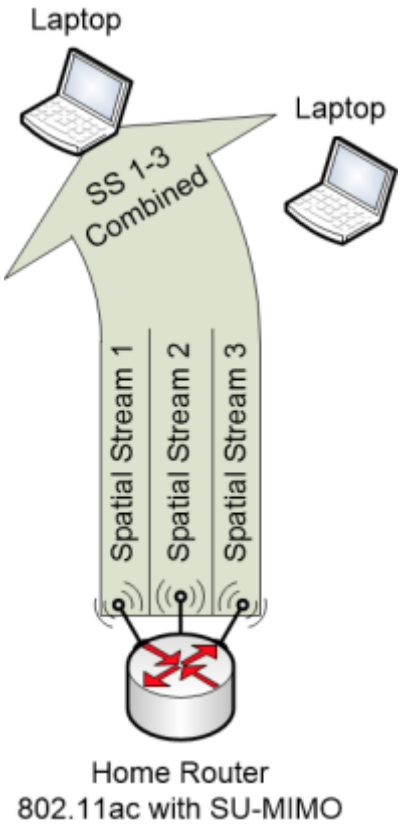

MU-MIMO

MU-MIMO stands for Multiple-User MIMO, and as the name suggests it is an expansion of SU-MIMO that is used in newer 802.11ac devices. It lets the access point talk to multiple clients at once using the MIMO technology. This actually makes the access point behave a little bit more like a switch than a hub, and if you have several wireless devices that often want to communicate at once then MU-MIMO can greatly improve the overall quality of your wireless communication. Of course, this requires that both the access point and the clients adhere to the same 802.11ac and MIMO standard.

MU-MIMO can thus both improve throughput by using several antennas to talk to a single client, or it can divide the attention of those antennas to multiple clients and let several clients communicate wirelessly simultaneously. It can also do combinations where it uses some antennas to increase throughput to one client while using another antenna to simultaneously communicate with another client.

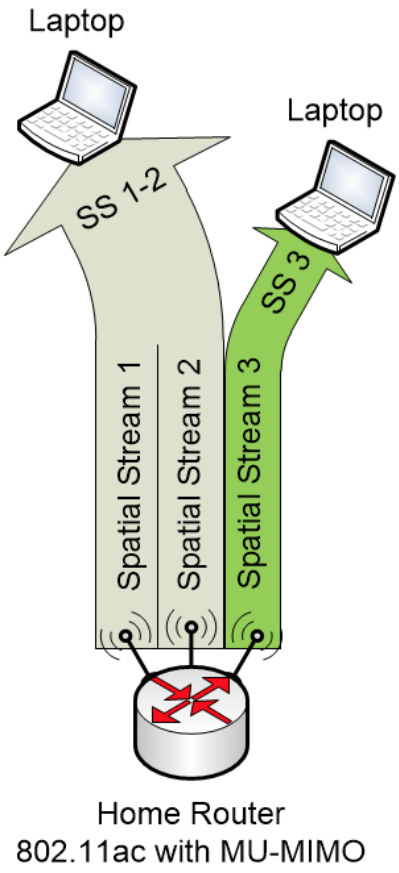

Variable Channel Widths

Normally a Wi-Fi radio channel is about 20MHz wide. One trick used by some 802.11 Wi-Fi standards is to increase the width of the radio channel from 20 to 40 MHz, or even further all the way up to 160MHz wide channels. By doing that the radios can fit more data into each channel, since the channel is now much wider and can transmit much more data at once.

- 802.11n supports either 20MHz or 40MHz wide channels.
- 802.11ac has even more options available: 20MHz, 40MHz, 80MHz, 80+80MHz (two separate 80MHz channels at once) or 160MHz

However while the throughput is increased by using wider channels, the radio signals will also take up much more space in the frequency range. Your router will disturb and be disturbed by other nearby devices more easily when it is using wider radio channels. This includes both Wi-Fi devices and other types of devices that use those same unlicensed radio bands.

Therefore, access points supporting wider channels have functionality built in to detect disturbances and try to work around them automatically, for example by moving to another radio channel or by disabling the wider channels to improve the quality of the radio communication.

This is one reason for why it is incredibly useful if the wireless devices have support for Dual Band and communication over the 5GHz band. On the 5GHz band there are usually no problems running twice as wide radio channels because there are so many available channels that do not overlap. In the newest standards, the access points will be able to use channels as wide as 80 or 160MHz. Then the overlapping channels can become a problem even on the 5GHz band, simply because each Wi-Fi radio uses up a big chunk of the available unlicensed frequencies.

Wi-Fi standards

There are a lot of standards for Wi-Fi networks and how they should operate. The standards for Wi-Fi communication are always named beginning with “802.11” and end with different characters such as “a”, “b”, “n” or “ac”.

Even though standards are developed by standards groups where experts from manufacturers and other organisations agree on how things should function a somewhat common problem with standards is that different equipment manufacturers might then not interpret the finalised standards in exactly the same way. This could prevent equipment from different manufacturers from working perfectly with each other.

To fix that problem for Wi-Fi networks almost all manufacturers agreed a long time ago to create an organisation called the “Wi-Fi Alliance”. All manufacturers within this organisation work together to test each other’s equipment against the equipment of other manufacturers so that they work together. Any equipment that fulfils the requirements get the official stamp of the organisation which looks like this:

The stamp consists of the “WiFi” brand in the middle, surrounded by a bunch of characters representing which Wi-Fi standards that the product can support. The product has also been tested within Wi-Fi Alliance for those supported standards.
So if you find two products such as a Wi-Fi router and a Wi-Fi network card for a computer and both have the WiFi logo which includes the “ac” label then those products will be able to communicate with each other.

802.11b

802.11b is one of the oldest Wi-Fi standards. It supports a theoretical max throughput of 11 Mbps over short ranges using the 2.4GHz band.

802.11a

802.11a is also one of the older standards. It supports much higher speeds, up to 54 Mbps, but over the 5GHz band.

Back in those days when the 802.11a standard was first created the 5GHz band couldn't really be used in most places of the world (it wasn't yet an "unlicensed radio band") so 802.11a didn't become very popular back then. These days, however, the 5GHz band is vital to any newer Wi-Fi product, and 802.11a can be used as a fallback for newer standards if necessary.

The most important information you can glean from the WiFi Certified logo is that if the logo contains the "a" label then you know that the product can communicate over the 5GHz band, which is a huge benefit. So keep a look out for that "a" in the WiFi Certified logo!

802.11g

802.11g is the successor of 802.11b. It runs over 2.4GHz and supports up to 54 Mbps.

802.11n

802.11n is a big upgrade compared to the previously mentioned standards. It is backward compatible with 802.11b and 802.11g and must support 2.4GHz communication. But it can also offer optional support on the 5GHz band and then also has backward compatibility with 802.11a. If a product supports both 2.4GHz and 5GHz communication it is called *Dual-Band*.

802.11n can utilise SU-MIMO with *spatial streams* to combine several streams into a higher total throughput speed, together with increased channel widths. Each stream can provide up to 150Mbps max theoretical throughput.
Depending on the number of antennas that the devices have they can support 150Mbps, 300Mbps or 450Mbps throughput using SU-MIMO.

The 150Mbps option is often a sign of a low budget Wi-Fi router because that option didn't even exist when 802.11n was first created. It was added later to offer a cheaper option for consumers. Those access points that can only offer 150Mbps 802.11n often lack a bunch of other 802.11n features and should probably be avoided.

More about Dual Band

There are two types of Dual Band when it comes to 802.11n.

The cheaper low budget type is often called "*Selective Dual Band*" and means that you have to pick in the router configuration if you want it to use either only the 2.4GHz band or only the 5GHz band.

This is bad because a lot of Wi-Fi clients only have support for 2.4GHz. That, in turn, means that you are more or less forced to select to use the 2.4GHz band, which deactivates the 5GHz band communication completely.

The better type of Dual Band where both the 2.4GHz and 5GHz bands can be active simultaneously is the only one that can really be recommended. This is called "*True Dual Band*" or "*Simultaneous Dual Band*".
What this means is that the access point will actually run two wireless networks at the same time, one for each radio band.

You could then connect laptops or tablets that don't require much throughput on the 2.4GHz band while connecting your video game and your Smart TV for streaming media on the 5GHz band. The two networks will work side by side completely without interference between the two radio bands.

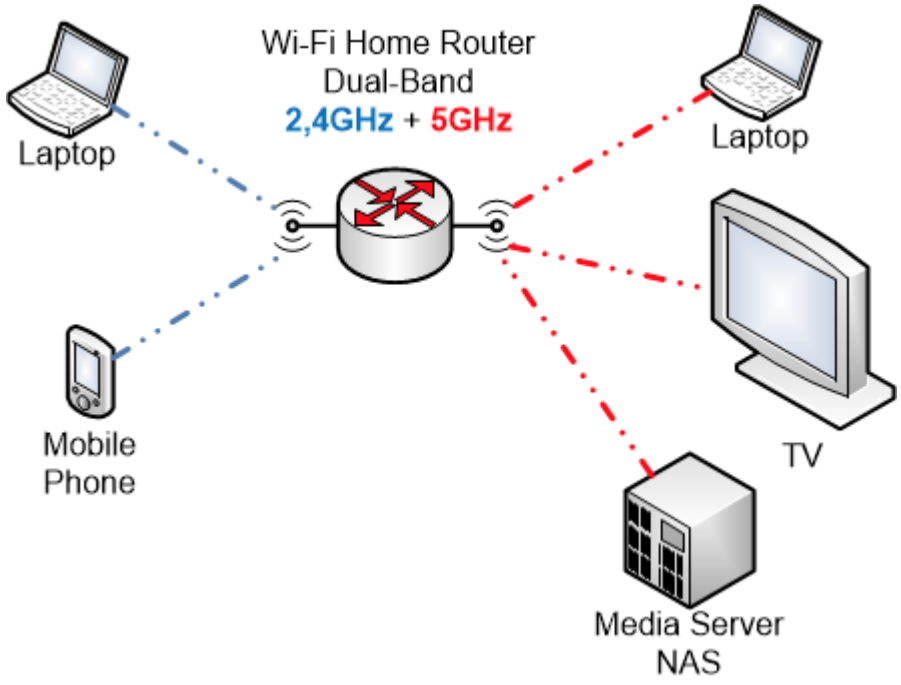

A lot of manufacturers have started using a somewhat unofficial naming convention or rating for their products to clarify what the products can do and which transfer speeds that they support. The rating contains a number showing the theoretical maximum throughput that the device supports, and will also indirectly show if the access point supports simultaneous dual band or not.

Rating	2,4GHz throughput	5GHz throughput (simultaneous Dual-Band)
N150	150 Mbps	- (no dual-band)
N300	300 Mbps	- (no dual-band)
N450	450 Mbps	- (no dual-band)
N600	300 Mbps	300 Mbps
N750	300 Mbps	450 Mbps
N900	450 Mbps	450 Mbps

If you are buying a 802.11n Wi-Fi router then I would suggest buying one that supports N600 or higher. You would then get the benefit of getting a router with simultaneous Dual Band where you can use the 5GHz band for higher throughput. You might also look to the newer 802.11ac routers for even higher data transfer rates.

802.11ac

The latest Wi-Fi standard (as of 2015-2016) is called 802.11ac. All 802.11ac devices are backward compatible with 802.11n.

802.11ac is very similar to 802.11n. 802.11ac is also built on transmitting one or multiple data streams that can be combined to increase the throughput. But whereas 802.11n can sometimes be limited to 2.4GHz communication only, 802.11ac requires support for simultaneous Dual Band. So buying an 802.11ac device is always a safe bet. The new increased transfer speeds of 802.11ac are only used on the 5GHz band.

The data streams of 802.11n can transmit up to 150Mbps each. Devices compatible with 802.11ac can use streams on the 5GHz band that are up to 433Mbps or even 866Mbps in future iterations of the standard. That is why you can reach much higher speeds with 802.11ac. If one stream is used you can reach 433Mbps. Two streams equal 867 Mbps and three streams 1300Mbps.

To reach those speeds, both MIMO and wider channels are used. The maximum channel width for 802.11ac is 160MHz.

All 802.11ac devices must be SU-MIMO compatible which requires multiple antennas. Newer 802.11ac standards can also support MU-MIMO which lets them act more like switches, communicating simultaneously with multiple clients on the same radio channel. If you are buying new equipment today and are looking into getting the very best equipment, then you should look for 802.11ac access points with MU-MIMO support.

An access point that supports 802.11ac can run simultaneous Dual Band mode in two different ways depending on the needs of the connecting devices:

- 802.11ac 5GHz + 802.11n 2,4GHz
- 802.11n 5GHz + 802.11n 2,4GHz

This is what the Rating table for 802.11ac transfer speeds could look like:

Rating	5GHz 802.11ac	2,4GHz 802.11n	5GHz 802.11n
AC583 (also known as AC600)	433 Mbps	150 Mbps	150 Mbps
AC1167 (also known as AC1200)	867 Mbps	300 Mbps	300 Mbps
AC1750	1300 Mbps	450 Mbps	450 Mbps
AC1900	1300 Mbps	600 Mbps	600 Mbps
AC2350	1750 Mbps	600 Mbps	600 Mbps
AC3200	2600 Mbps	600 Mbps	600 Mbps

Some of the higher numbered ratings in the table aren't actually fully completed yet as of 2015, so if you buy equipment (as of 2015) with a higher rating than AC2350 then you are moving into unchartered territory. Most probably you are then buying equipment that is built according to best guesses of the manufacturer as to how the final standard will eventually work, but since the standard might not yet have been tested within the Wi-Fi alliance against equipment from other manufacturers it cannot be guaranteed to work across devices from different manufacturers.

Often when it comes to such new technology you have to buy all devices from the same manufacturer to be sure that it works as announced, and you might encounter limitations.

802.11ac has existed for a while now, and there are a lot of good equipment out there. The 802.11ac standard will also continue to improve for at least another year or two with new equipment and improved transfer speeds.

Building blocks of Wi-Fi networks

This section brings up different building blocks that are used to build wireless networks.

In a Wi-Fi home network, you often have one or more wireless laptops, computers, Smart TV's, tablets, phones and other wireless devices. What all of those have in common is that they are commonly referred to as "clients". Clients are devices that want to utilise the wireless service that a Wi-Fi router or access point is providing.

So the key to the perfect wireless home network is the Wi-Fi router or access point that all other clients connect to. Most home users probably just have a single home router with Wi-Fi. What that means is simply that the home router has a built-in Wi-Fi access point with one or more antennas.

A wireless network is not limited to having a single access point. It could consist of up to hundreds of access points in an enterprise network. Such a network is out of scope for this guide, but the components and building blocks are the same no matter the size of the wireless network.

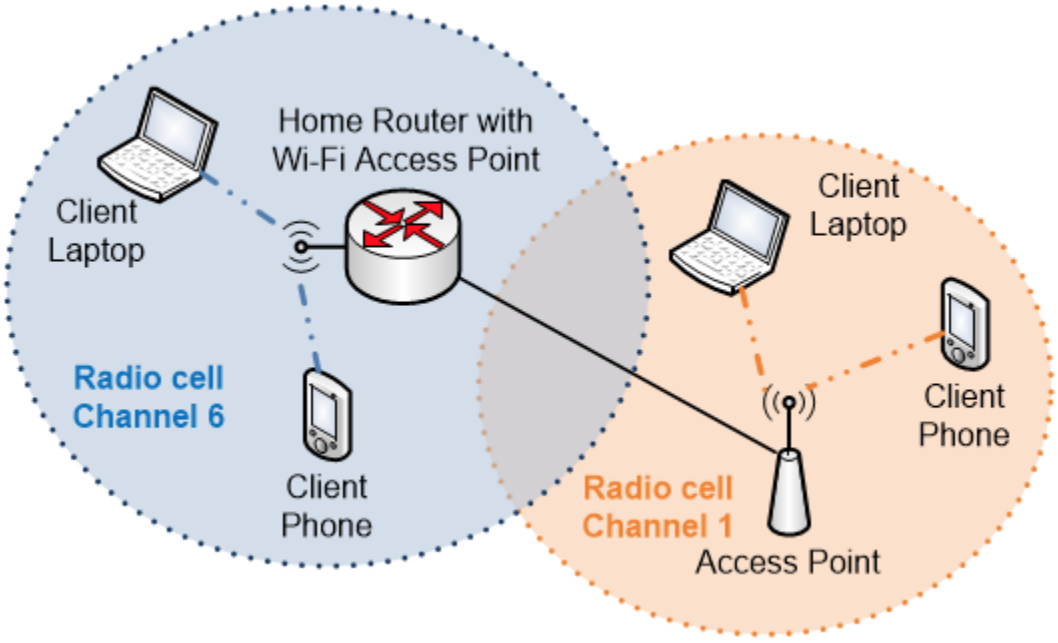

Product Types

There are a number of different types of products that have built-in Wi-Fi in one way or another. Here we will go through the most common.

Home router with built-in Wi-Fi

Many routers for home use have built-in Wi-Fi today. Just like the router has a built-in switch where you can connect computers using network cables the router also has a built-in Wi-Fi access point. Both the switch and the access point belong to the inside LAN of the router.

These devices are of very varying quality and there are many low-cost, low-quality Wi-Fi routers out there. If you are buying or getting a new router today and you have a smaller apartment or house then you should probably be okay with just getting a home router with built-in Wi-Fi. Opt for a router that can handle at least 802.11n with N300 rating. Much better is if you could get a router that is 802.11n capable with simultaneous Dual Band and communication over both the 2.4GHz and 5GHz bands, which means you will get a router with N600 or better rating.

However, if you have the extra money to spare you should absolutely choose a router with 802.11ac Wi-Fi with AC1300 or higher rating. If the device also has MU-MIMO support then you would potentially improve your wireless home network greatly.

Access Point

An access point is a Wi-Fi device that acts as a hub device for your wireless network. An access point can announce a wireless network that clients can connect to, just like a wireless home router does. Then the access point will let wireless clients communicate with each other via the access point, and it will also forward traffic between its wireless antenna and its wired LAN connection to the rest of the network.

However, the access point does only that and nothing else. It doesn't do any of the other things that a router can do.

These access points are very common at enterprises and bigger corporations because a corporation doesn't want to litter the place with routers. Instead, they want to use simpler access point devices for their wireless network, and they want to be able to spread those access points out all over the company so that no matter where you are you will be within reach of a wireless access point. The access points will then be mounted in the ceiling or on the walls.

So an access point just does one single thing, and that is to act as a connection point for a wireless network.

If you would use multiple home routers to accomplish the same thing then you would run into problems because one of the things that each home router will do is to create a separate LAN network on the "inside" of the router, which the Wi-Fi network belongs to. So unless you take quite a few steps to avoid it then if you were to use home routers with built-in Wi-Fi your network would become segmented into multiple smaller LAN's, making it impossible for the computers to talk with each other.

By instead using access points then each access point just announces a wireless network, but all of them will belong to the same inside LAN that your home router has created. So all computers and other devices will be connected to the same LAN still, even if they have connected to different access points.

Most access points are developed for the enterprise market segments, so they tend to be more expensive to buy compared to just buying a home router with built-in Wi-Fi. As long as you can cover your whole living space with a single access point it is usually both simpler and cheaper to just buy a single home router with built-in Wi-Fi. But if you want to build a bigger and better home network then the best way of doing so is by using access points.

(Wireless) Client

Clients are all wireless devices that connect to wireless access points. They can be laptops, mobile phones, tablets, computers with wireless network cards, Smart TV's with built-in Wi-Fi, NAS connected storage devices, wireless printers and so on. Anything that can connect to a wireless network to communicate can be called a "client" within wireless.

Building Blocks

The most basic type of wireless network is something called a *Basic Service Set* or "BSS". It consists of a wireless access point that announces a wireless network, also known as an *SSID* or *Service Set Identifier,* which is the name of the network paired with the settings of the network. The clients are the devices that connect to that wireless network.
The BSS forms a radio cell which reaches as far as the radio waves can travel from the single access point.

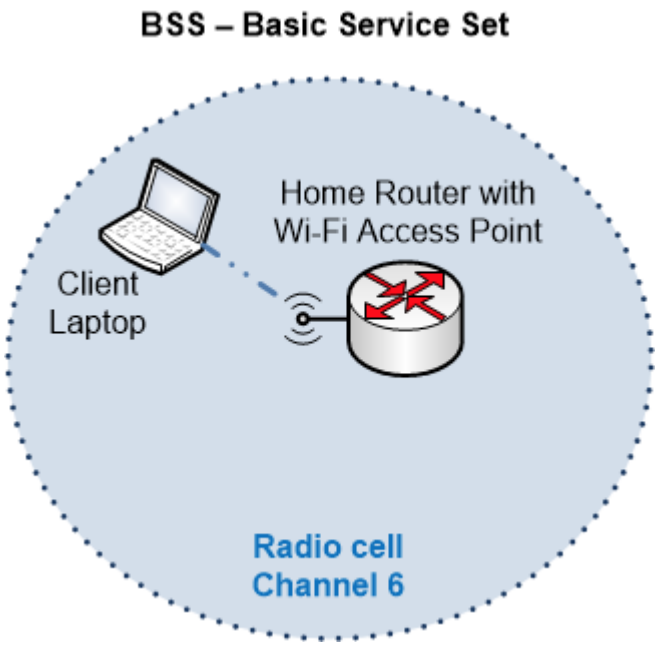

A common problem that a lot of people encounter is that your villa, house or apartment is simply too big, so the radio signals from a single Wi-Fi router can't reach every part of your home. Or perhaps the signals do reach everywhere you intend to use the wireless network, but the signal strength or signal quality is too poor so using the wireless network becomes almost impossible.

What you could do then is to add another access point that you connect to the same LAN network. If you configure the exact same settings for the wireless network on both access points then you can get something called an ESS or *Extended Service Set*.

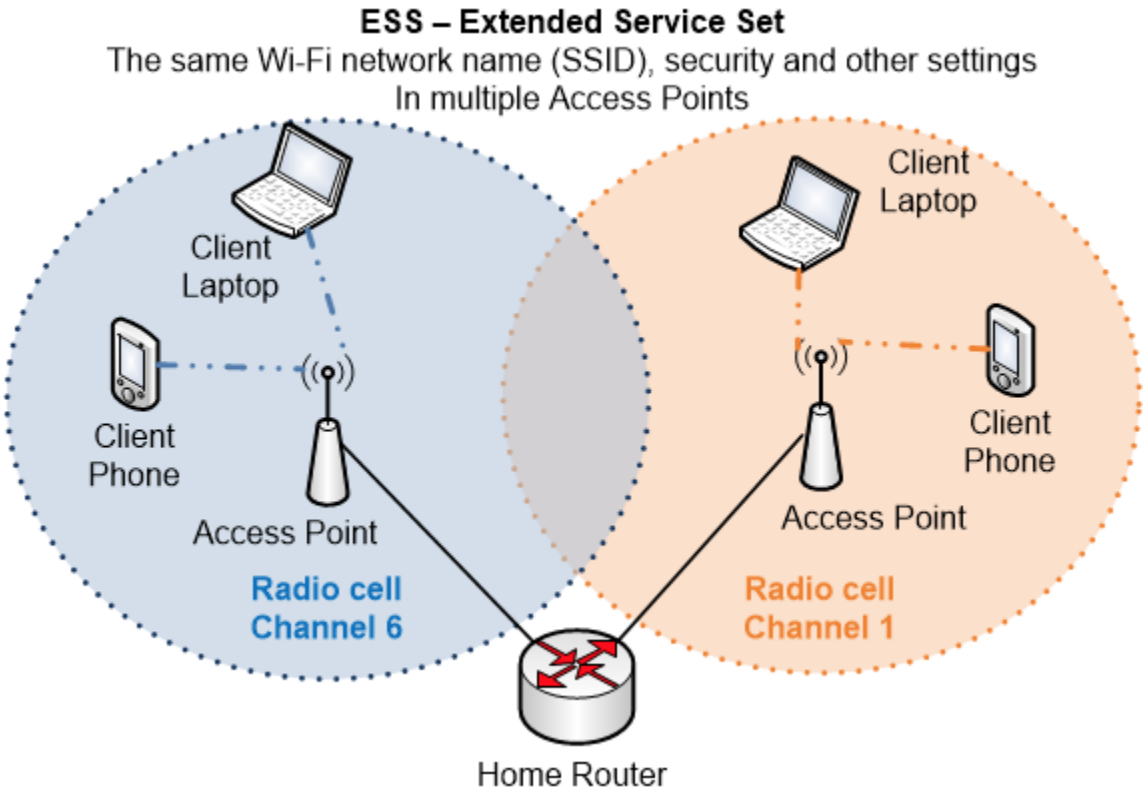

But there are multiple caveats with this type of solution that most people run into when they try to extend their wireless network. The different caveats can lead to different type of problems with the wireless network.

Later in the wireless section we will go through what you can do to extend your network by showing you both good ways of doing so as well as other more sub-optimal but still common ways of extending your network..

SSID: The Wi-Fi network name

The SSID is the name of the wireless network as the Wi-Fi Access Point announces it. This is the name that you will see when you scan for available networks within reach on your client, for example when you want to connect to a wireless network on your phone or your computer.

Home routers with built-in Wi-Fi often come pre-configured with an SSID that can be used. But you can also change the SSID to another name of your own choosing.

A lot of routers will let you choose to either show or hide the SSID. A hidden SSID will not show up if you scan for wireless networks, but you can still connect to that SSID if you know the details and enter them manually in your client.

Some home routers list this function as a "security feature". Do not let this fool you however into believing that hiding an SSID has anything to do with security or increasing security for your wireless network.

A hidden SSID will not protect you against hackers or prevent people from accessing your wireless network. The only thing it achieves is that the SSID will not show up if you look in the regular wireless network scan on your client. But for somebody who actually wants to find your network they will find it no matter what.

In fact, it is actually slightly more "secure" to show your SSID instead of hiding it. If the SSID is hidden, then the wireless clients that you have manually configured to connect to that SSID must continually try to send out probes to detect if your SSID is available nearby. This would then have some potential impact on your privacy since your client will be acting as a beacon in your pocket. Not perhaps a big issue for most people, but at least remember that hiding your SSID will not increase your wireless security.

It is easier for both yourself, your family and your friends to connect to your SSID if it is visible. With a hidden SSID you have to manually configure any devices that you want to connect to the wireless network.

Security

You should always enable encryption for your wireless networks. Encryption is a built-in function that is easy to activate in both your access point or home router and your clients. Most home routers these days come pre-configured with encryption enabled for the wireless network. Remember that a wireless network sends all communication via radio through the air. Anybody within reach of the signal can read the radio messages. Unless you have encrypted the communication anything you send over the wireless network will be accessible to anybody within reach.

You also have to choose which type of encryption that you want to enable on your wireless network. Actually, it is not so much of a choice since all but one of the options are more or less obsolete. If your router comes pre-configured with encryption enabled (which it should) then you should check which encryption that is implemented and change it if necessary. Here are the available encryption options:

WEP: extremely bad security

WEP encryption was useless even when it was first released in 1997. To use WEP is almost just as bad as sending your secret messages in clear text over your wireless network. It is completely broken. Do not use WEP encryption on your wireless network.

WPA: Okay security

WPA has got a few lesser security flaws but has not yet been completely broken. But why would you use anything that is less than perfect? There is no reason today why you would choose WPA security on your home network in favor of WPA2. Do not use WPA unless you have to for some reason.

WPA2: Good security

WPA2 is the only encryption type that you should consider using for your wireless network. But you still need to use a secure password for your wireless network.
You will configure the password on your Wi-Fi router, and then again on the clients the first time you connect them to your network.

MAC address filtering

MAC address filtering is based on building up a list of accepted MAC addresses in your home router or access point. Only clients with the listed MAC addresses are allowed to connect to your wireless network.
This initially sounds like a proper security feature and it makes people feel more secure. MAC addresses are normally unique, so surely if you make a list of allowed MAC addresses then nobody else will be able to connect to your wireless network?

Unfortunately, MAC address filtering is completely unreliable and should not be used as a security feature on your wireless network.

Actually implementing MAC address filtering will make it more complicated to connect to your wireless network. Every time you want to connect a new device you have to check what MAC address that device has on its wireless network card, and then add that MAC address to the MAC address filtering function.

But in reality, the function doesn't add any security. All MAC addresses are completely visible in the wireless communication even if you use encryption. So anybody who is listening on your radio communication can see which MAC addresses that are in use and that are allowed. Then they can simply change their own MAC address to one of the allowed MAC addresses to gain access to the wireless network.

Instead just simply rely on the WPA2 encryption which you should have enabled anyway. The encryption with the secret key will protect your network.

Common Wi-Fi network solutions

So how do you actually build a wireless Wi-Fi network? It depends a little bit on the scenario, but the basics are always the same. First let's discuss what most people might do, and which problems that might lead to.

Most people will probably either receive a home router with built-in Wi-Fi from their Internet Service Provider (ISP) as part of the service, or they will buy a home router with built-in Wi-Fi on their own.

Then they place the router somewhere close to where the Internet Connection is delivered. It might be in the living room by the cable TV jack, by the phone jack in the kitchen, or by the fiber converter box in the hallway. But since the Wi-Fi router is often in the way and people think it is ugly and doesn't fit the rest of the interior design they choose to hide their Wi-Fi router. Maybe underneath a TV table, behind a cupboard or lying down on top of a shelf.

Often they then get bad Wi-Fi coverage in general. Especially in other rooms further away in the house. If you have a second floor then the coverage might be even worse up there. But since you just have a single router which might be difficult to move away from where it is currently installed people don't know how to solve the issue.

Some might have googled for a solution to extend their wireless network coverage area. They might have stumbled upon solutions including buying yet another Wi-Fi router that they connect to the first router using a network cable. They place the second router on the second floor to improve the coverage up there. Or they might buy a Wi-Fi repeater that they install.

What they don't know is that there are a whole lot of reasons for why those are less than optimal solutions and that the solutions themselves are creating other types of related problems. The end result is often that people are dissatisfied with their wireless network and think that it works bad in general.

Here we look at a few common solutions for building Wi-Fi networks.

Adding a repeater

A *repeater* is a special type of wireless device which is sort of a mix between an access point and a client. If you have bad Wi-Fi coverage in a part of your apartment because the signal can't reach that far, then you could put a repeater halfway there. The repeater will connect to the wireless network (acting as a client) and will extend the signal by repeating the signal to the clients further away (acting more like an access point).

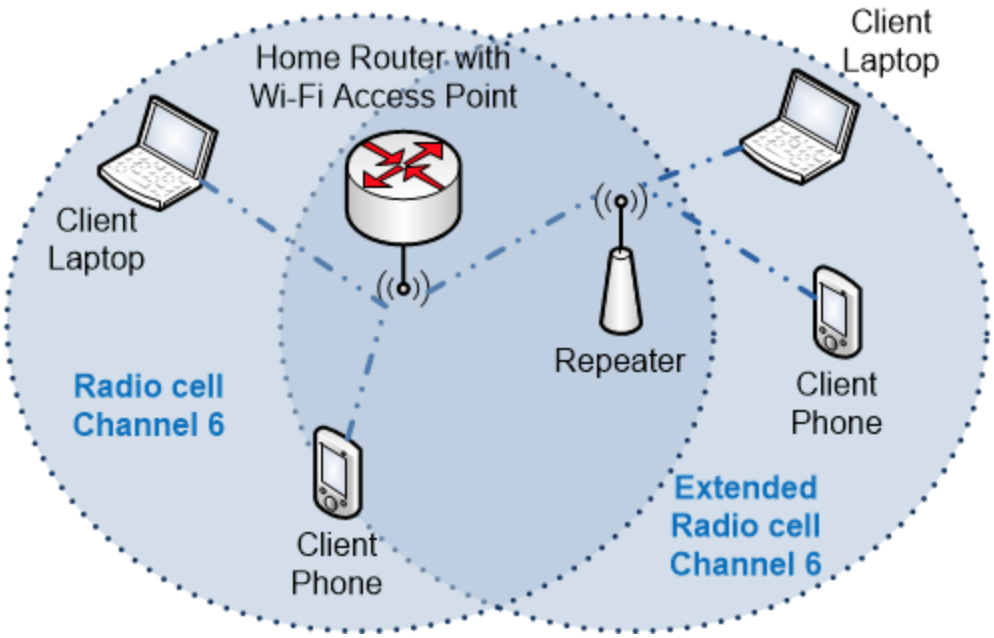

The first thing you should know is that a repeater is avoided by anybody who works professionally with wireless networks. It is a makeshift solution that might look good on paper if you don't know how it works, but if you do understand how it works then you also understand why it is a bad solution.

First of all you are going to extend the signal with the help of the repeater. The repeater doesn't have any better reach than the Wi-Fi router. So to extend the signal from the access point by 10 meters you have to put the repeater something like 10 meters away from the access point. If you put it too far away then the signal from the access point gets too weak. If you put the repeater too close then the signal from the access point isn't really going to be extended by much.
But the further away you put the repeater the worse signal it will receive from the access point. This also means that the throughput drops which limits the maximum connection speed of anybody who connects to the repeater (remember, the speed decreases the further away you get from the access point).

Wireless networks are Half-Duplex. Only one device can talk at a time. So when the repeater is re-transmitting the messages that it receives then everybody else must be silent. Each message going back and forth between the original access point and the client connected to the repeater will in effect be transmitted twice, which takes time during which all other clients must be silent.

Since both the repeater and the client laptop are probably further away from each other with worse signal quality they have to lower their transmission speeds, causing the whole network to have to wait even longer for them to communicate.

Anything that the client says must be repeated by the repeater. So in effect any client that is connected to the repeater will automatically have their throughput cut in half.

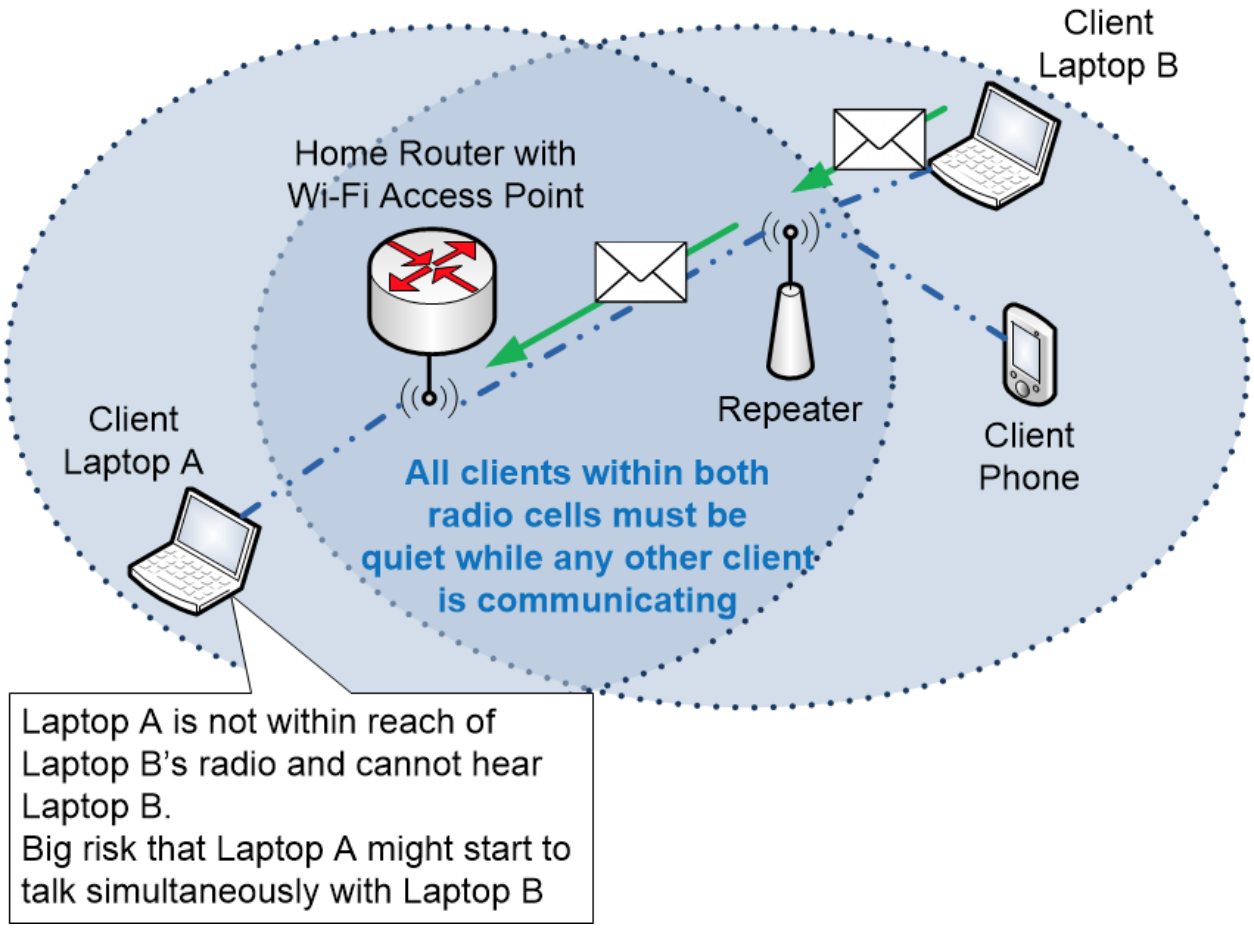

The repeater will extend the wireless network all the way to the client that previously had bad coverage. But clients that are connected to the repeater in the extended area might not be able to hear if a client that is connected to the Wi-Fi router far away in the other side of the house is talking. This causes the mechanisms for how wireless clients try to avoid collisions to become less effective, which causes collisions to occur more often, which causes more communication downtime and lower throughput for all of the connected clients.

So even if a repeater solves the immediate problem of trying to extend your wireless network to an area further away the solution has lots of problems.

The only real benefit of using a repeater is that you do not have to run a network cable from your home router to the repeater. It can talk wirelessly to the Wi-Fi router. The only real benefit of using a repeater is that you do not have to run a network cable from your home router to the repeater. It can talk wirelessly to the Wi-Fi router. If running a cable between your router and an additional access point is your main concern then you could look at solutions based on

HomePlug (https://en.wikipedia.org/wiki/HomePlug) to see if that could be a viable option. HomePlug is a system for communicating over the existing electrical grid that is already installed in your home.

The bottom line is that you shouldn't use a repeater unless you really have to for some reason. Your wireless network will be limited by your repeater based solution.

Adding a second home router with built-in Wi-Fi

Some people who notice that they have Wi-Fi coverage issues at home solve the problem by adding another home router with built in Wi-Fi to their network. Remember how a home router is built up? It has a WAN port (outside) and one or more LAN ports (inside). The built in access point counts as if it belongs to the LAN.

So if you connect the WAN port of a second home router to one of the LAN ports of the first router, then you have created a new small LAN segment which is separate from the first one. This means that clients that are connected to the different home routers will have trouble talking with each other.

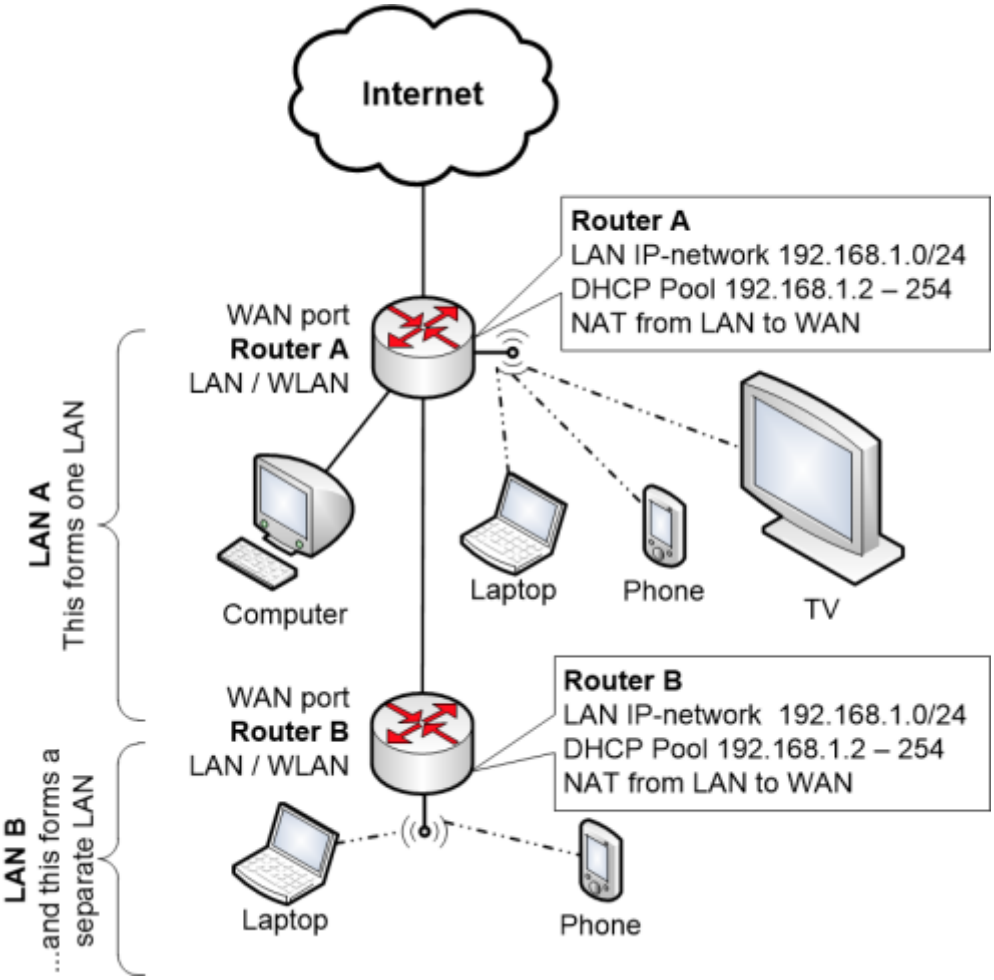

In the picture above, devices that are connected to Home Router A will go through a NAT address translations when they talk to the Internet.

Devices connected to Router B however must first pass through NAT address translation on Router B, and then they must pass through yet another address translation in Router A to get to the Internet.

Since Router B is using NAT, then devices that are connected on the outside of Router B can't communicate to devices that are connected to the inside LAN of Router B. To be able to do that you would have to add Port Forwards on Router B, which quickly gets really complex to work with.

One problem that this causes is that a lot of solutions today that stream music and video require that devices are connected to the same Wireless Network. Apple TV, AirPlay and Google Chromecast all require that the devices are connected to exactly the same wireless home network. Otherwise they cannot connect to each other.

In the solution above where you add a second home router, you actually have two separate wireless networks on two separate routers. So things like AirPlay, Chromecast, Steam In-Home Streaming and similar might not work as expected, or at all.

You can also have issues if a wireless client connects to the "wrong" wireless network. If you have an Apple TV connected to the Wi-Fi router in the living room, but your mobile phone is connected to the Wi-Fi router on the second floor, then you can't stream music or film from your mobile to the Apple TV. The only way to solve this would be to try to make your mobile jump over to the "right" wireless network. Maybe by disabling the Wi-Fi in the mobile phone and then enabling it again to try to make it connect to the closest Wi-Fi router with the best signal.

So having two home routers with separate LAN networks and two separate wireless networks will cause all sorts of problems with your home network.

Improving upon the "second home router" solution

There is actually something you could do to improve the above situation considerably. The solution would still not be perfect, but far better than above.

If you disable both NAT and DHCP in the second home router, and then connect a network cable from a LAN port of Router A directly to one of the LAN ports on on Router B, then you form a shared LAN segment between the two routers that all clients will belong to. You also have to change the IP address of router B's LAN port to avoid having IP address conflicts, and you must make sure that their IP addresses are on the same IP network.

This would at least prevent the two routers from belonging to two separate LAN networks. Instead they will have a common shared inside LAN IP network between them.

Since only Router A is handing out IP addresses via DHCP then all clients will have Router A as their default gateway.

Hopefully Router B that you use in this solution will let wireless clients that connect to Router B obtain an IP address from the DHCP pool on Router A. Otherwise you will have to keep a DHCP pool active on Router B as well, but you must make sure that both of the DHCP pools (in Router A and Router B) hand out IP addresses within the shared LAN segment between the routers, and you must also make sure that the addresses in the DHCP pools do not overlap.

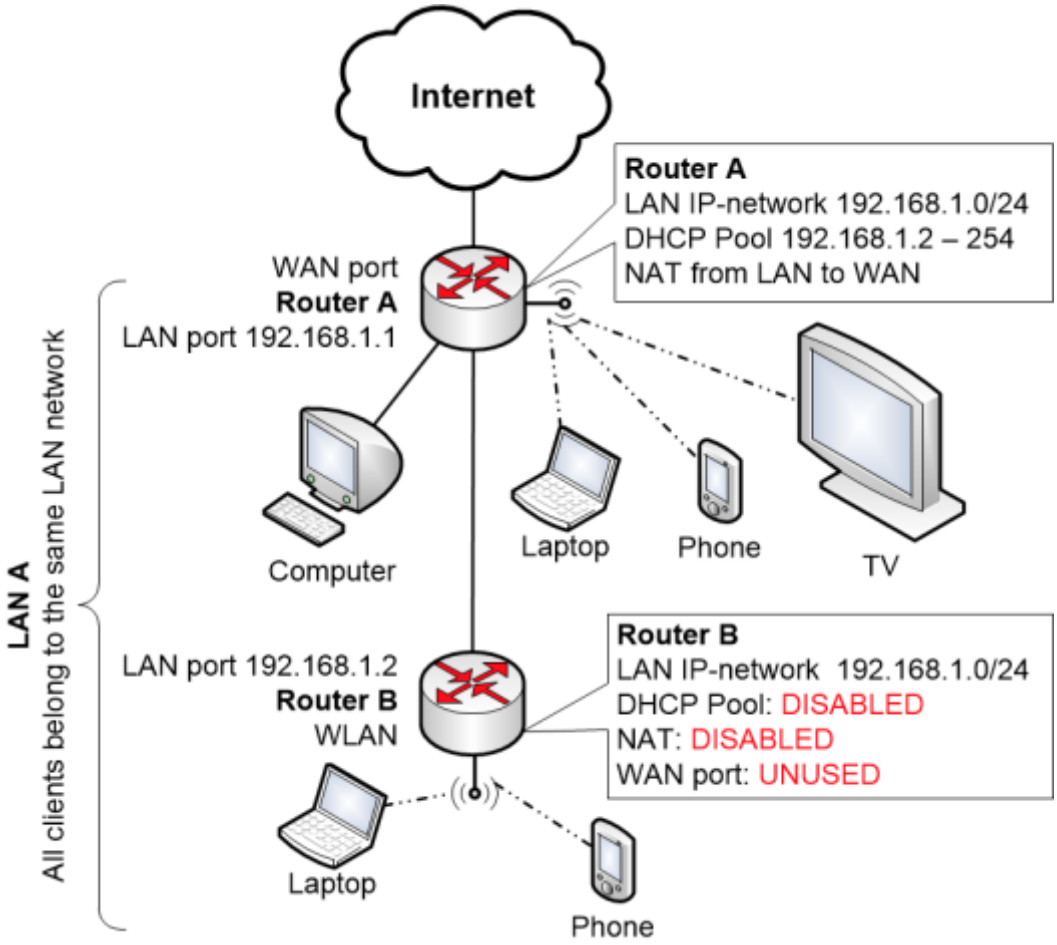

As you can tell even if this solution would improve the situation there is a lot of work involved, and the end result is still not perfect. But it is a cheap and effective solution to a common problem!

Building a better Wi-Fi network

First we can look at a few quick general tips that could improve your wireless network dramatically:

Quick tips for a better Wi-Fi network

- Place the Wi-Fi router in a central location in the area where you want good wireless coverage. Wi-Fi routers radiate radio signals in all directions or omnidirectionally. You can increase both the coverage and the overall quality of the wireless network by moving it to an open area.

- Wi-Fi signals are blocked by most everything. Walls, glass doors, furniture, floors, cloth and so on. So it is absolutely best if you have a completely free line of sight between the antenna of the router and your wireless client. Many people wish to hide away their ugly Wi-Fi routers, putting it behind things, underneath stuff or in cupboards. This is generally a bad strategy for obtaining a good wireless signal.

- It is simpler if you imagine radio waves behaving much like light. If you can see the light source (you have a clear line of sight) the light is brighter. You will still see the light indirectly even if you don't have a clear line of sight because the light bounces around the room. But the best signal quality is achieved if your computer can "see" straight to the antenna of the Wi-Fi router.

- You can use apps on your phone to measure how good signal strength you get from your Wi-Fi router in different rooms. WiFi Analyzer is such an app for Android. Then you can measure and customise your Wi-Fi installation to work in the best way possible. Try moving the Wi-Fi router around to see what effect it has on your wireless network.

- You can also use speed tests on the Internet (google for "speed test") to measure your download and upload speeds at different locations. First run a test using a network cable between your access point and your laptop to see which results you should expect under perfect conditions. Then disconnect the network cable and run the test again over the wireless network. Start off nearby the access point with a clear line of sight between your laptop and the access point, and then move around to various locations where you might want to use your wireless network

- Sometimes, moving the access point just up to a few meters can drastically change how the radio signals bounce around in your home. Experiment to see if you can improve the signal quality further away in your home by moving your access point.

- Wi-Fi radio waves dislike closed doors. Closing the door to a room could be the tipping point that changes the wireless signal strength in the room from acceptable levels to being more or less unusable. While testing your wireless network during the implementation you have to be aware of this, so that you perform the tests under similar conditions to how you will be using the network later on.

Setting up a wireless network

Here are the general steps you should take when setting up a wireless network if you want to follow in the footsteps of more professional Wi-Fi engineers. We encourage taking notes and sketching down floor plans to visualise your network environment while you are working with your plan.

The steps described below can also give you clues to the type of things that you need to consider. For example, if you right away come to the conclusion that some family members will be streaming films to the TV while others will be using their gaming computers in rooms on the second floor, then you have already identified a potential problem that you have to take into consideration.

1. Investigate your current situation

a. Where is the Internet connection delivered at your house?
   i. Through a phone jack perhaps?
   ii. In a fiber box in your basement?
   iii. In your living room through a cable connection box?

b. What type of Internet connection is it?

c. Based on the above, is it possible to move the placement of your home router?
   i. This is important to determine if you will be able to choose where your Wireless home router will be located.
   ii. Can you extend the cable to move the home router if required?
   iii. Perhaps you can run a network cable from your router to a switch further away and connect devices to the switch?

d. How big is your home?
   i. How many square meters and how many floors do you have?

e. What type of walls does your home have and where are they located?
   i. Concrete?
   ii. Wood?
   iii. Plaster?

f. Where will your users be?
   i. Draw hot spots on your floor plan to mark areas where your users will be using their wireless devices

g. What type of throughput or speed does your applications or programs require?
   i. Video streaming (youtube, Smart TV, Media Server) is among the most demanding use cases

h. How many people or devices will be using the Wi-Fi at once?
    i. Are they all in the same place or spread out over multiple rooms, floors, areas?

2. Produce a plan based on the prerequisites

a. Make a guess as to how big coverage area you can get from the access points you would like to install. Usually the supplier or manufacturer can at least list some highly optimistic numbers for “indoor use” that you probably need to cut at least in half to get a more realistic overview.

b. Place the access point(s) where you would want them on your floor plan.
    i. Start with a single access point.
    ii. Use line of sight from the perspective of the access point on the floor plan. Are there a lot of walls in between the access point and other areas you want it to cover? Are those concrete walls? The more walls and the thicker and more solid they are, the worse the signal generally gets.
    iii. Where would you potentially have to add extra access points?

c. Use this floor plan to try to make guesses as to how you should place your access point(s)

3. Test your plan

a. You need at least one access point that you temporarily connect and set up in the way you imagined in step two. Perhaps you could borrow a similar access point or router from a friend, or your ISP could have provided you with a home router with built in Wi-Fi as part of the service.

b. Take a laptop and walk around the house to test in different spots what type of connection you actually get. Make sure to perform the tests in a way that mirrors how you would later be using the wireless network at that location.

c. Run speed tests to see what internet connection speeds you get at the different spots. Concentrate on the areas where you imagine that you will need to use your wireless network the most, and where you want the highest throughput, but also check other areas where you know that you will definitely need or want wireless coverage.

4. Adjust the plan

a. Based on your results you might move the access point to another area to improve the coverage of one or more of your important usage spots.

b. You might also conclude after trying multiple placements of the Wi-Fi router or access point that a single access point is simply not going to be enough for the whole house, and you might have to get another access point.

c. Go back to main step 2 and produce a new adjusted plan using another access point if necessary, then test the new revised plan again

Example: Small apartment

Most people start off with the assumption that you probably only need a single wireless access point and that it will be enough with the built-in Wi-Fi that you have in your home router. This is usually a good strategy for smaller apartments. Up to 80-90 square meters is rarely a big problem, depending on the layout and building materials. You might end up with a somewhat worse wireless connection in some rooms, but you will probably have good enough wireless network quality no matter which room you are in.

Always plan to put your Wi-Fi router in the room where you will require the best wireless connection. For a lot of people, this would be the living room where you might end up sitting with both mobile phones, laptops, Smart TV's, Apple TV's, streaming boxes and other media devices. The further away and the more obstructions between the client and the Wi-Fi router, the worse signal quality you get, which leads to lower throughput and a Wi-Fi connection with worse quality.

So you really want to put the Wi-Fi router in the room where you have the biggest need for a fast and undisrupted wireless network.

Once you picked a room where you would like to install the Wi-Fi router you have to investigate if you can place the router there. Sometimes the Internet connection is handed off in a cable jack, phone jack or fiber box somewhere else in your house.

Depending on what type of Internet connection you have it can be easy or more difficult to extend the cable to the location where you want to place your router. A phone line can often be extended from the phone wall jack quite far away before the signal gets too bad. The same goes for a fiber box, which typically lets you run a network cable from the fiber box to your router. The network cable can be run for a length of up to about 70-100 meters before the signal becomes too bad, so you shouldn't have any issues with extending such a connection to where you want to place your Wi-Fi router.

If you have a cable provider then maybe you can put a separate cable modem by the cable wall jack, and then run a UTP network cable between the cable modem and your router. This is unless the cable modem is integrated into the router that the cable service provider has given you. Then you would have to buy a separate cable modem and a separate router.

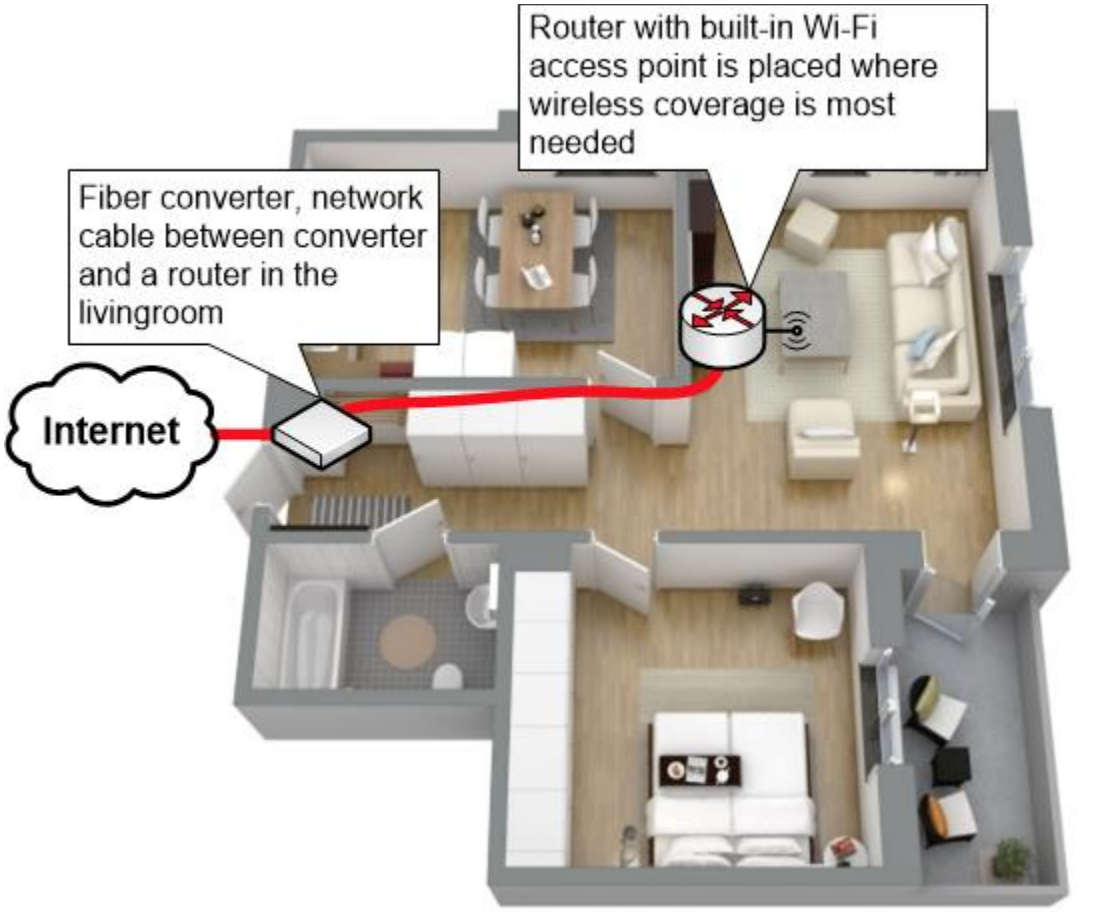

If you can extend the connection to your Wi-Fi home router location of choice, then do so and place the router there. You can then test your wireless network to see if it works as well as you want. Concentrate on the more important rooms, but try the wireless network everywhere where you think you might want to use it.
During your tests, use multiple devices including your phones and your computers that you are actually going to be using later.

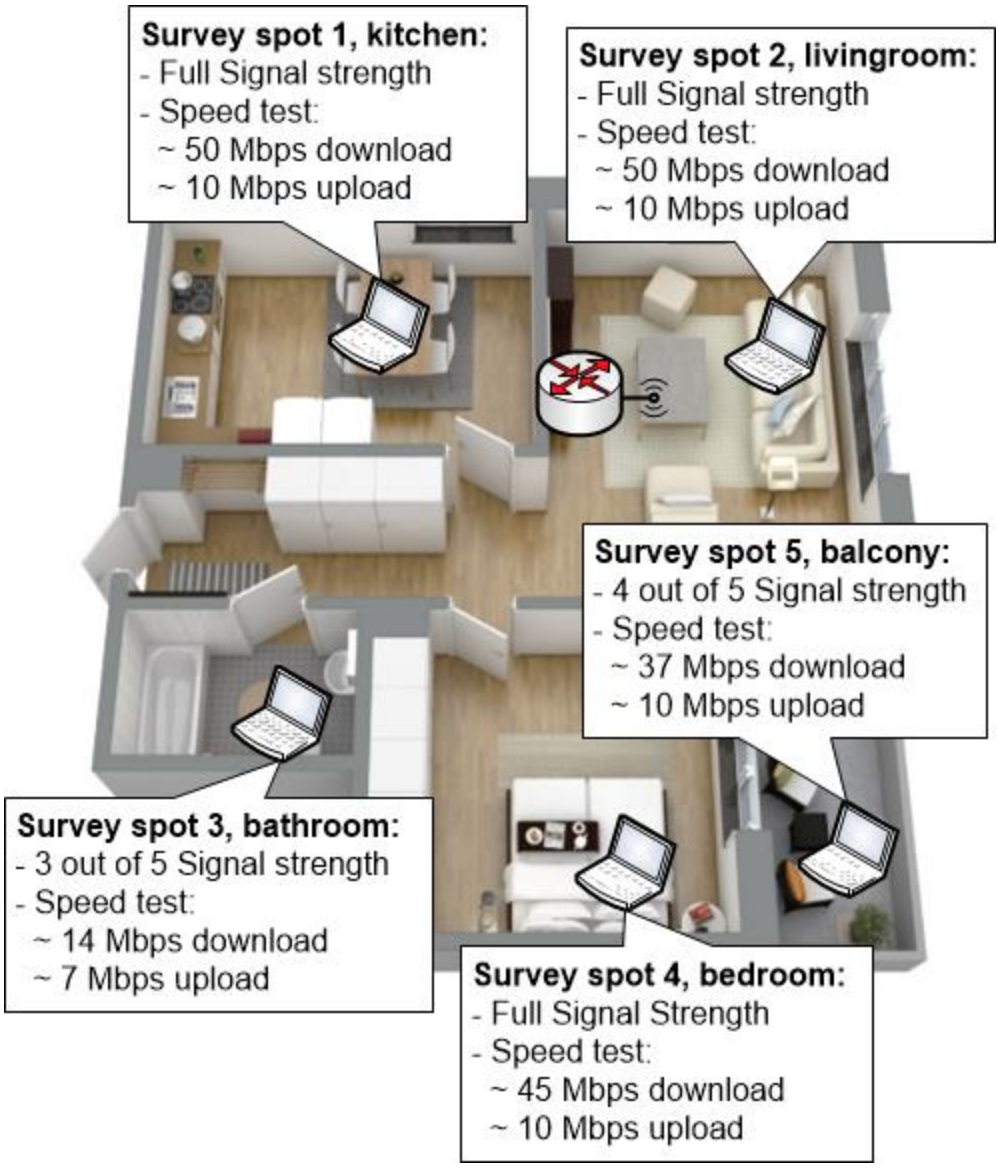

If you get bad results anywhere then you can try moving the access point up to a couple of meters in different directions to see if it makes any difference.

Example: Bigger apartment or house

When would you need more than one single access point? When is a single access point is not enough?

The answer to that question is simply *when it is no longer enough with a single access point to get adequate wireless network coverage in all of the areas that you want to have good network coverage in.*

Often this would be something that you would notice after the fact, once you have already installed your wireless network. This is why we recommend that you should never install your equipment permanently until you have made a few tests, to see if you obtain results that you are satisfied with.

Professionals that install wireless networks for a living often fall back on their experience, looking at floor plans with their knowledge about wall materials, furniture etc to try to estimate how many access points that are needed for a specific building. But it also depends on exactly which access points that are used, which clients that will connect to the network and how they work, which requirements that the applications, programs, clients and users put on the wireless network and so on.

A real professional will never just trust experience, because there are simply too many unknowns. They would also use a predictive program that can calculate the estimated number of required access points and their placements based on a number of input factors. Then once they have that estimation they will use it to set up the network and do a so-called Site Survey, where they simply test the network layout in the building to see if they get the results they wanted.

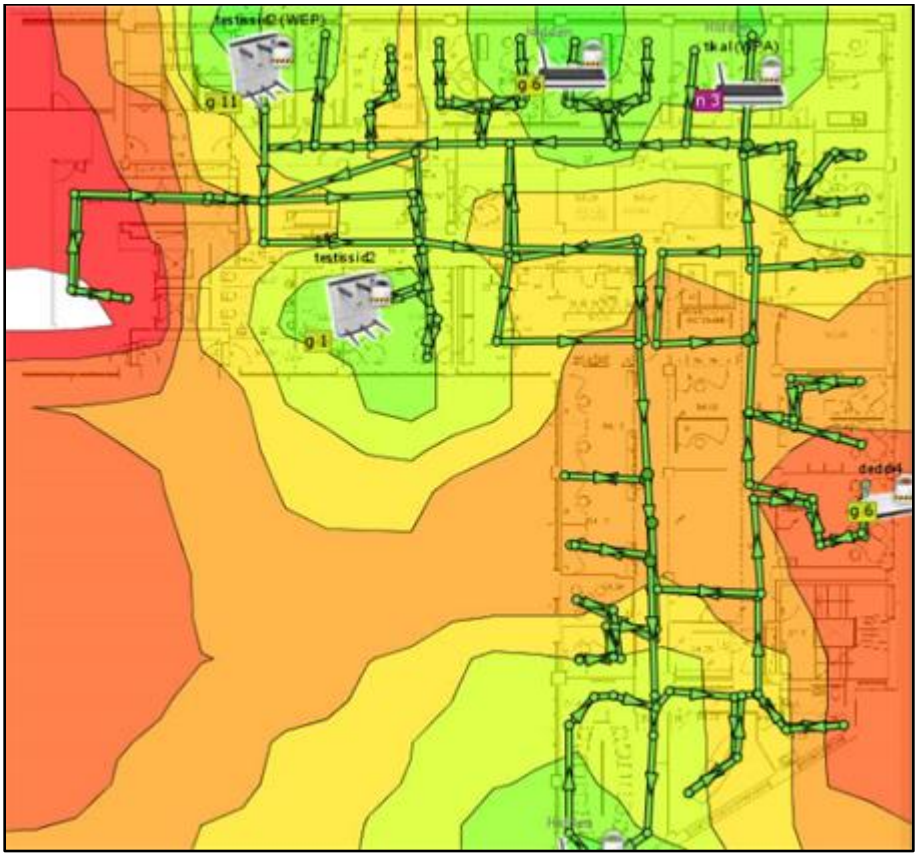

Example of a site survey application that can try to estimate the number of required access points based on a number of input factors

During the site survey the network engineer will often find multiple issues that have to be corrected before the network is delivered and handed off to the customer. As soon as the installation requires more than a few access points it actually gets really difficult to plan a wireless network, at least if the applications have high requirements.

Luckily, in a home network there is rarely a big need for clients to have an undisrupted Wi-Fi connection as they are moving around in the home. Rather there are a few hot spots where users will be using their wireless devices. This makes it a lot easier to setup the wireless network.

There are also a few shortcut questions that would normally indicate if you will require more than one access point:

- Are there multiple floors where you need good wireless coverage?
- Does each floor have more than maybe 80 square meters?
- Do you have multiple concrete walls, stone walls or other dense building materials that cut through the house?
- Do you need to cover any additional guest houses, outdoor areas or gardens?

If the answer to at least one of the above questions is yes then you might need more than one access point. Often a good basic estimate is one access point per floor, unless each floor is also really big. Basements where all surrounding floors, walls and ceilings are made out of concrete or stone materials are especially tricky radio environments and will require an access point of its own for good wireless signal quality.

Another reason for why you might want to have more than one access point per floor is if you want especially good radio signal quality and throughput in more than one place. You might have a living room where you need good throughput, but you also have a recreation room with entertainment systems that also need the best possible Wi-Fi connection. You might then have to put an extra Wi-Fi access point in the recreation room to achieve your goal.

When you build such a network you can do it in two different ways. The cheaper typical do-it-yourself way, or the more optimal but a bit more expensive way.

The cheaper way consists of using your home router with built-in Wi-Fi as the first access point and place it in the first location. Then you add a second pure access point that does nothing but act as a wireless access point and place it at the second location. The second access point will be connected with a network cable to your Wi-Fi home router.

This setup is very similar to the last setup in the previous section, where the LAN port of a second Wi-Fi router was connected to a LAN port of the first Wi-Fi router. The main difference is that using an access point is a cleaner solution using equipment that is meant for the task, and you can also more easily add a third or a fourth access point to the solution.

Each access point must use different radio channels that do not overlap. On the 2.4GHz band you should pick from channels 1, 6 and 11 and let each access point use a unique channel out of those three.

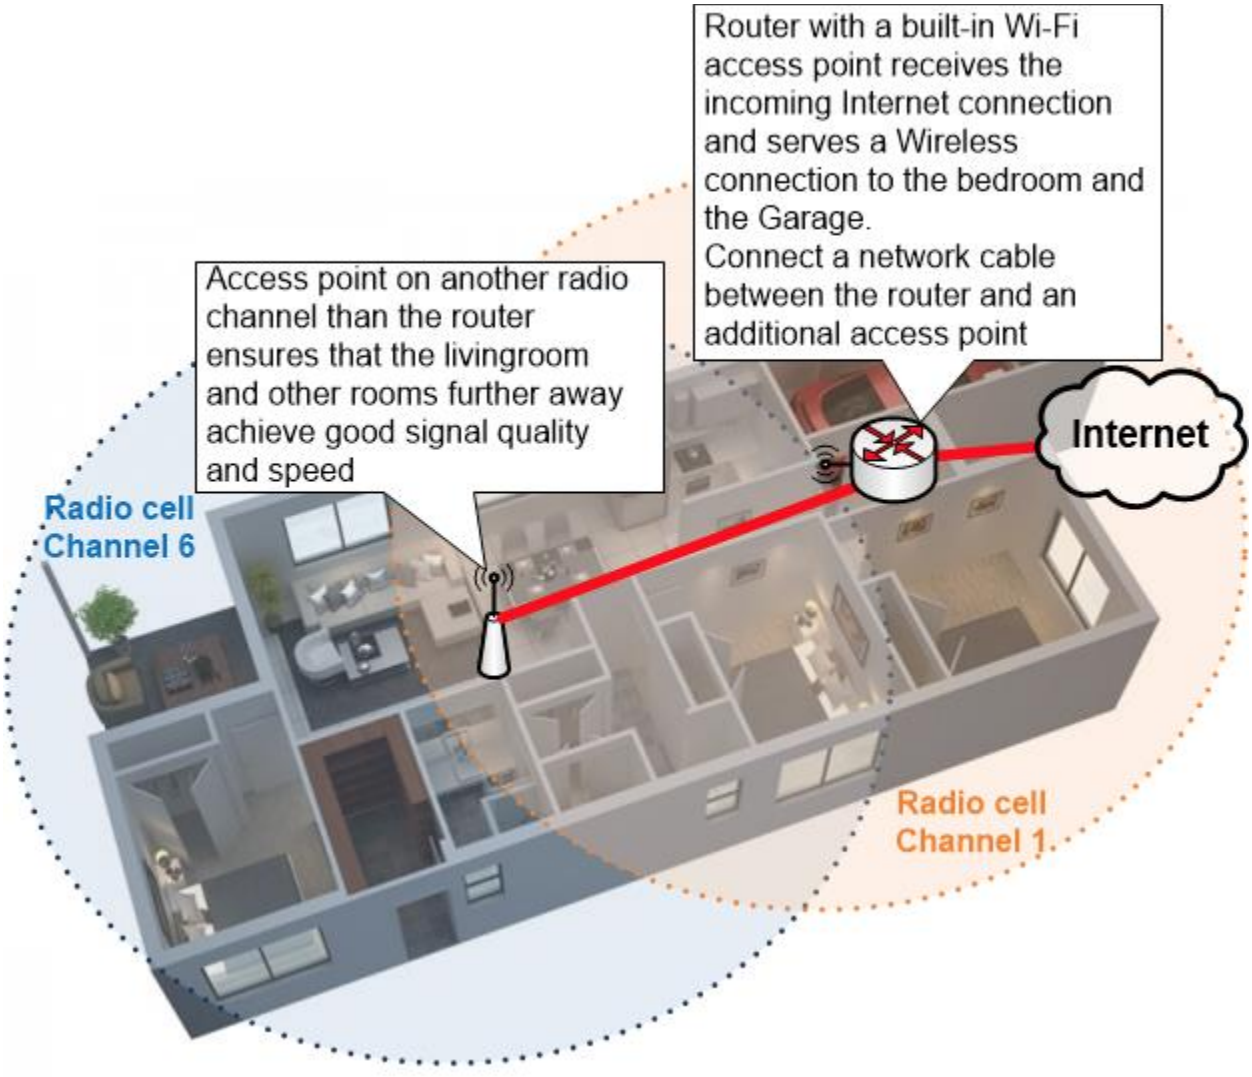

The only downside of this solution is that you still do not have a completely uniform wireless network. Even if the access point and the home router are configured with the same wireless settings they will still not cooperate in any particular way. They are just two separate devices that happen to announce the exact same wireless network. Clients that connect to this wireless network will perceive it as if there are two access points on the same wireless network, but it will be up to the clients to determine when and how they want to roam or switch over between the two.

The more expensive but also much better way of doing it is to buy into a wireless controller based solution that is discussed below.

Wireless LAN Controller based solution

A *Wireless LAN Controller* (or *WLC* for short) is a device that controls access points which are registered in the Controller and makes the setup of the wireless network more automatic. This includes which channels that are used, which effect the access points operate on etc. The WLC will optimise and adapt the radio environment of your access points automatically and can also keep optimising it continuously based on changes in the radio environment, for example if your neighbours set up new access points of their own.

A WLC based solution consists of the following parts:

- Wireless access points
- Wireless LAN Controller

The WLC is what you will configure, often via a graphical interface in the WLC that you connect to in your web browser. The WLC then tells the access points what they should do, which channels they should operate on, what SSID network name they should announce and so on.

The WLC can be either a physical device that looks something like your home router, or it can be just a program that you can install on your computer, which turns your computer into a WLC by the use of a program or service that is running in the background.

Ubiquiti UniFi (https://www.ubnt.com/unifi/unifi-ap/) is an example of a popular and not too expensive Wireless LAN Controller based solution for both home and corporate use. In their case the WLC consists of an application that you install on a computer in your network. Their WLC application is free-of-charge but can only be used together with the Ubiquiti access points. The price point of the solution is therefore decided by which access point model that you are interested in and the number of access points that your home network requires.

When you install a WLC based solution you would set up your home router as usual and place it wherever you like. You would then log into your home router and disable its built-in Wi-Fi. Once that is done you would install the Access Points in your home, connect them to the router or a switch using network cables, and then start the Wireless LAN Controller program on your computer to configure the access points.

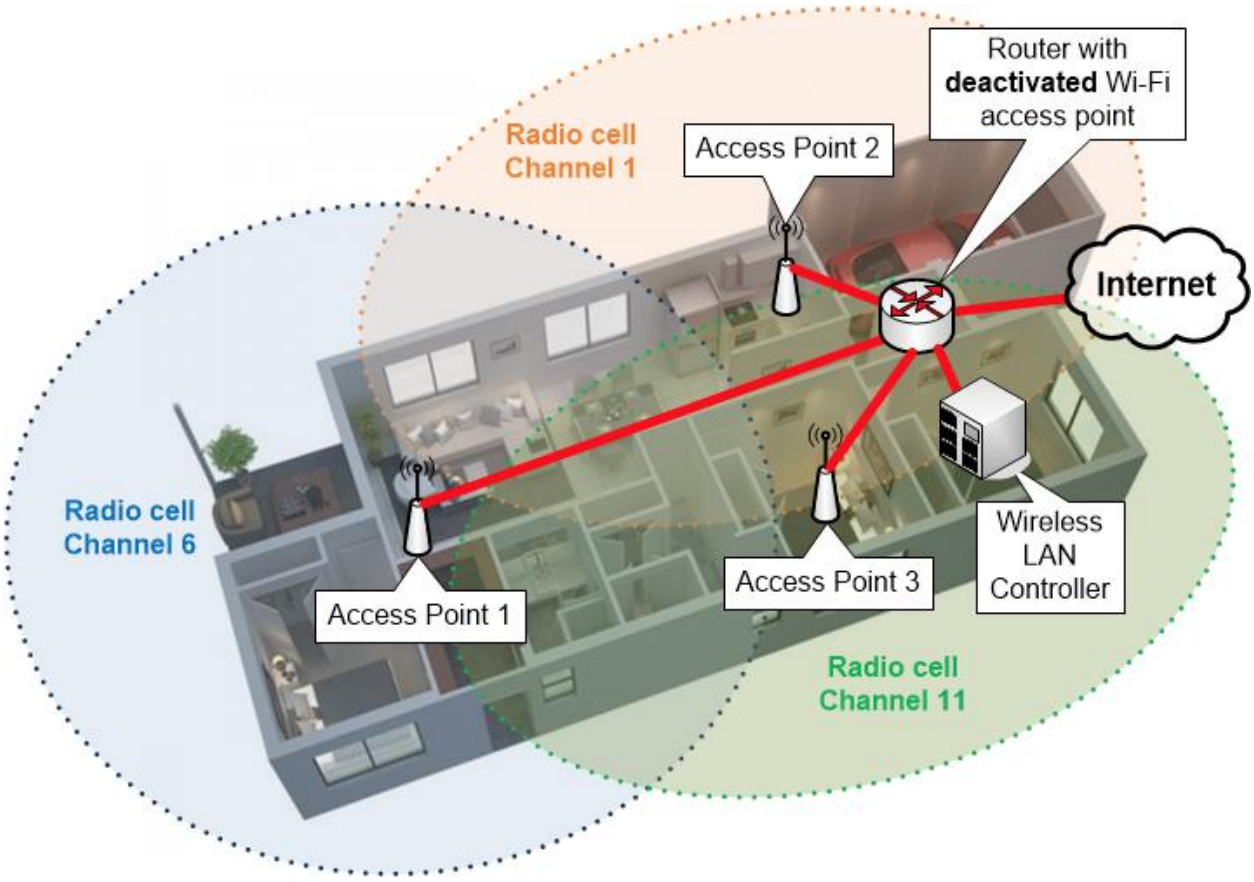

By doing this you would end up with a network that costs a bit more, but you would also have a wireless network of professional quality. You can easily add additional access points later on to increase the total coverage area or to increase the throughput of specific hotspots.

Most Wireless LAN Controllers will support a mixed environment with different access points. You could install an access point with 802.11ac support in the livingroom, while your guest house gets a cheaper 802.11n capable access point.

If you want to make a configuration change you will only have to log into your WLC instead of having to connect to each single access point individually to perform the same change multiple times. The WLC will push any configuration changes automatically to your access points. Wireless networks also benefit a lot from being updated with more recent software versions in the access points, to fix issues and to add support for new features which will improve your network. Software updates can also be installed on your access points from the WLC.

The Wireless LAN Controller also takes care a lot of other things for your wireless network, making it function much better. One such thing is Roaming, which is when a wireless client such as a mobile phone is moving from one access point closer to another access point and wants to swap over between the two. This is usually handled somewhat poorly by the wireless client if it has to do it on its own. But by using a Wireless LAN Controller based network the Wireless LAN

itself will help the client in different ways to roam between access points so that it can become completely seamless to the client and to the user.

The WLC can also detect disturbances and changes in the radio environment and make automatic adjustments to your access points, making them function better over time.

All in all, if you can afford buying into a solution with a Wireless LAN Controller based Wi-Fi network, then it will be the biggest long term improvement you can possibly make for your wireless network.

Using cables

*"But I want to build a **Wireless** network!"*

Many people want to move completely over to a fully wireless network, but never end up with perfect results no matter how hard they try.

The truth is that wireless networks aren't the perfect solution to any given problem. Even though the wireless network standards improve all the time there are lots of problems that are still difficult to solve with a wireless network.

One of the most important tips we can give you is therefore to reconsider and actually connect your more important devices using network cables. Consider those things that never move around and which require the best quality network connection, such as your gaming computer, the Smart TV, your media player and the NAS device where you store those films that you want to stream over the network.

All of the mentioned devices place high demands upon the network communication and are typically sensitive to packet loss, latency and other types of disturbances. When such disturbances do occur you also notice them very easily through decreased or varying video quality, lagging or stuttering streaming and buffering. Many of the devices are also often physically placed permanently close to each other, nearby your TV set.

In truth, the best way to solve the network connection for devices like these is to place a small switch close to the devices and then connect them to the switch using network cables. No matter what happens on your wireless network they will still have a consistently performing network connection.

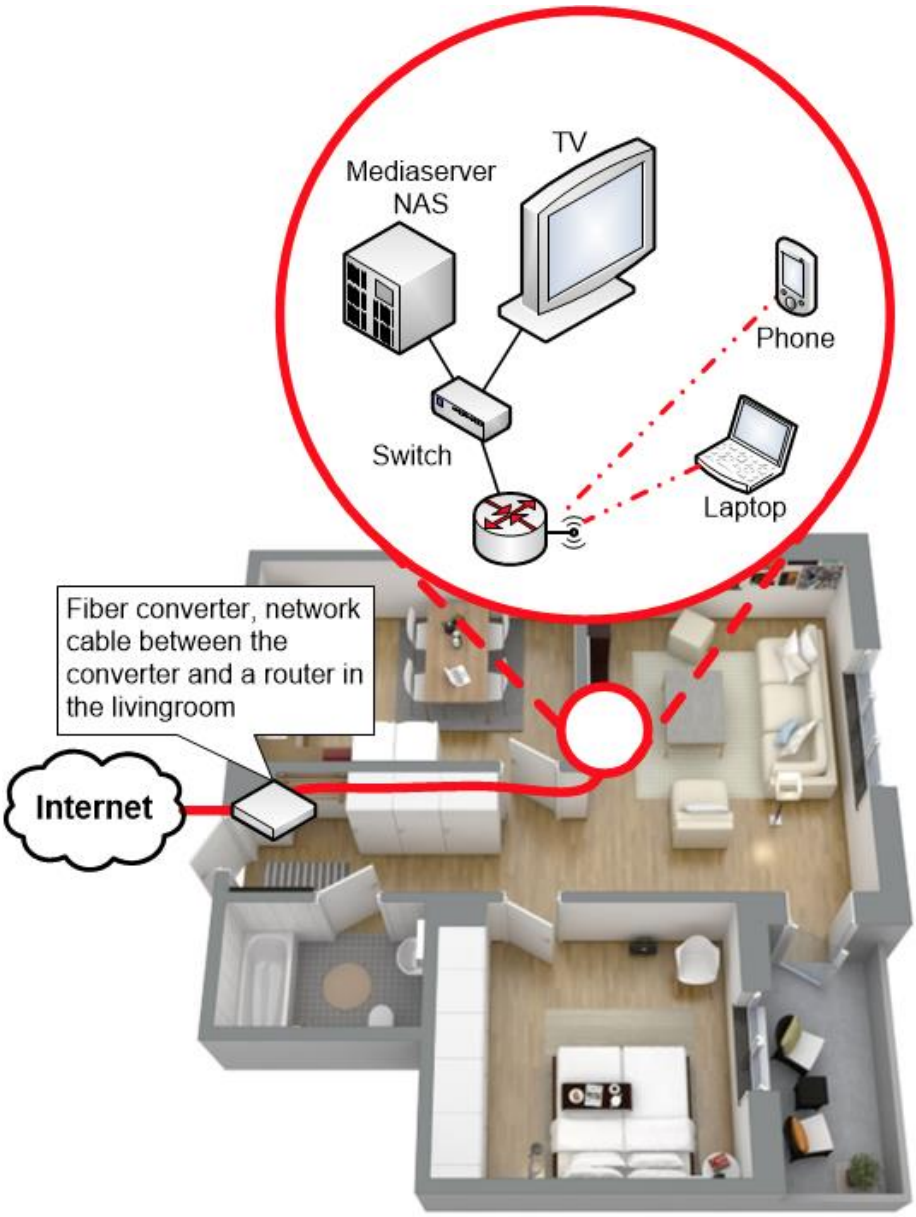

With a wireless network there are no guarantees that you could ever hope to avoid disturbances completely. Even with the best available equipment you could be unlucky. Civilian or military Radar can disturb 5GHz equipment. Microwave ovens will have negative impact on 2.4GHz. Or maybe your neighbour buys some other type of solution that operates on the 5GHz band and which interferes with your wireless network.

All we are saying is that the one way to guarantee a good network connection is to use network cables to connect your most demanding devices.

Final words

This material has covered all different areas of theory and basic network knowledge from many of the areas that home networks span over.

Of course there are also many areas that are difficult to cover in a material such as this. Some standards and areas such as wireless networks simply change so quickly that some information might be outdated if you go too deep into detail.
Other subjects might be difficult to describe in detail since each explanation leads to another even more complex subject that also needs an explanation of its own in a downward spiral of more and more complex subjects.

We hope that the level of theory and basic knowledge you have gotten from the material has given you a good basic understanding of the communication that takes place in home networks, and that you can find use for this knowledge both in your own home network as well as in other areas of general computing!

With regards from the team behind www.homenethowto.com