


Variable Length Subnet Masking (VLSM)

Conventional Subnet masking replaces the two-level IP addressing scheme with a more flexible three-level method. Since it lets network administrators assign IP addresses to hosts based on how they are connected in physical networks, subnetting is a real breakthrough for those maintaining large IP networks. It has its own weaknesses though, and still has room for improvement. The main weakness of conventional subnetting is in fact that the subnet ID represents only one additional hierarchical level in how IP addresses are interpreted and used for routing.

The Problem with Single-Level Subnetting

It may seem “greedy” to look at subnetting and say “what, only one additional level”? J However, in large networks, the need to divide our entire network into only one level of subnetworks doesn't represent the best use of our IP address block. Furthermore, we have already seen that since the subnet ID is the same length throughout the network, we can have problems if we have subnetworks with very different numbers of hosts on them—the subnet ID must be chosen based on whichever subnet has the greatest number of hosts, even if most of subnets have far fewer. This is inefficient even in small networks, and can result in the need to use extra addressing blocks while wasting many of the addresses in each block.

For example, consider a relatively small company with a Class C network, 201.45.222.0/24. They have six subnetworks in their network. The first four subnets (S1, S2, S3 and S4) are relatively small, containing only 10 hosts each. However, one of them (S5) is for their production floor and has 50 hosts, and the last (S6) is their development and engineering group, which has 100 hosts.



Class C (/24) Network (254 Hosts)

Figure 70: Class C (/24) Network Split Into Eight Conventional Subnets

The total number of hosts needed is thus 196. Without subnetting, we have enough hosts in our Class C network to handle them all. However, when we try to subnet, we have a big problem. In order to have six subnets we need to use 3 bits for the subnet ID. This leaves only 5 bits for the host ID, which means every subnet has the identical capacity of 30 hosts, as shown in Figure 70. This is enough for the smaller subnets but not enough for the larger ones. The only solution with conventional subnetting, other than shuffling the physical subnets, is to get another Class C block for the two big subnets and use the original for the four small ones. But this is expensive, and means wasting hundreds of IP addresses!

With traditional subnetting, all subnets must be the same size, which creates problems when there are some subnets that are much larger than others. Contrast to Figure 71.

The Solution: Variable Length Subnet Masking

The solution to this situation is an enhancement to the basic subnet addressing scheme called Variable Length Subnet Masking (VLSM). VLSM seems complicated at first, but is easy to comprehend if you understand basic subnetting. The idea is that you subnet the network, and then subnet the subnets just the way you originally subnetted the network. In fact, you can do this multiple times, creating subnets of subnets of subnets, as many times as you need (subject to how many bits you have in the host ID of your address block). It is possible to choose to apply this multiple-level splitting to only some of the subnets, allowing you to selectively cut the "IP address pie" so that some of the slices are bigger than others. This means that our example company could create six subnets to match the needs of its networks, as shown in Figure 71.




Figure 71: Class C (/24) Network Split Using Variable Length Subnet Masking (VLSM)

Using VLSM, an organization can divide its IP network multiple times, to create subnets that much better match the size requirements of its physical networks. Contrast to Figure 70.

An Example: Multiple-Level Subnetting Using VLSM

VLSM subnetting is done the same way as regular subnetting; it is just more complex because of the extra levels of subnetting hierarchy. You do an initial subnetting of the network into large subnets, and then further break down one or more of the subnets as required. You add bits to the subnet mask for each of the "sub-subnets" and "sub-sub-subnets" to reflect their smaller size. In VLSM, the slash notation of classless addressing is commonly used instead of binary subnet masks—VLSM is very much like CIDR in how it works—so that's what I will use.

Let's take our example above again and see how we can make everything fit using VLSM. We start with our Class C network, 201.45.222.0/24. We then do three subnettings as follows (see Figure 72 for an illustration of the process):

1. We first do initial subnetting by using one bit for the subnet ID, leaving us 7 bits for the host ID. This gives us two subnets: 201.45.222.0/25 and 201.45.222.128/25. Each of these can have a maximum of 126 hosts. We set aside the first of these for subnet S6 and its 100 hosts.
2. We take the second subnet, 201.45.222.128/25, and subnet it further into two sub-subnets. We do this by taking one bit from the 7 bits left in the host ID. This gives us the sub-subnets 201.45.222.128/26 and 201.45.222.192/26, each of which can have 62 hosts. We set aside the first of these for subnet S5 and its 50 hosts.

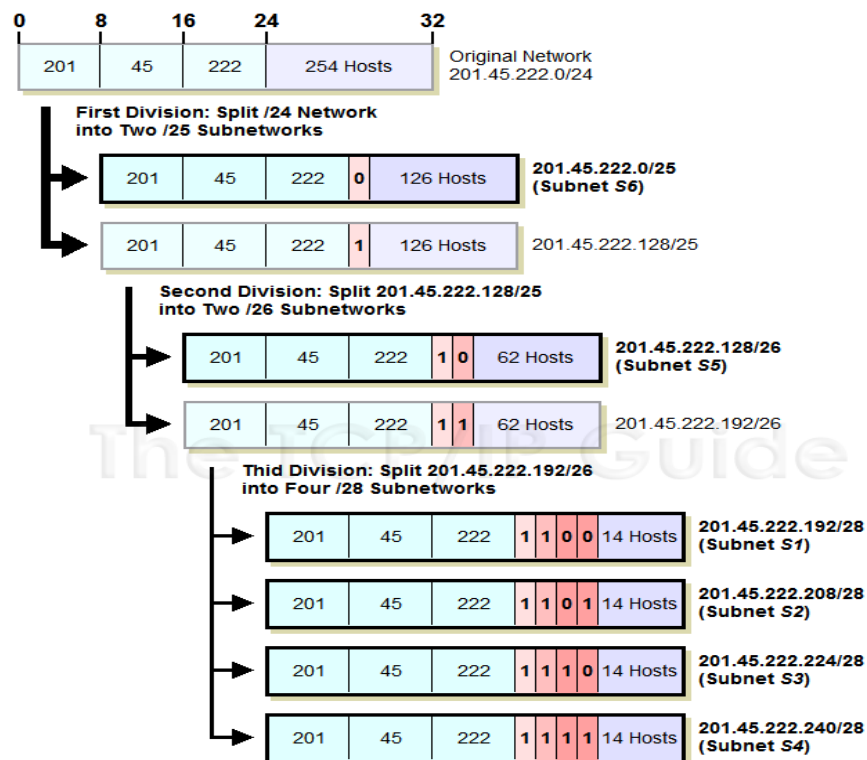


Figure 72: Variable Length Subnet Masking (VLSM) Example

This diagram illustrates the example described in the text, of a Class C (/24) network divided using three hierarchical levels. It is first divided into two subnets; one subnet is divided into two *sub-subnets*; and one sub-subnet is divided into four *sub-sub-subnets*. The resulting six subnets are shown with thick black borders, and have a maximum capacity of 126, 62, 14, 14, 14 and 14 hosts.

3. We take the second sub-subnet, 201.45.222.192/26, and subnet it further into four sub-sub-subnets. We take 2 bits from the 6 that are left in the host ID. This gives us four sub-sub-subnets that each can have a maximum of 14 hosts. These are used for S1, S2, S3 and S4.

Okay, I *did* get to pick the numbers in this example so that they work out just perfectly, but you get the picture. VLSM greatly improves both the flexibility and the efficiency of subnetting. In order to use it, routers that support VLSM-capable routing protocols must be employed. VLSM also requires more care in how routing tables are constructed to ensure that there is no ambiguity in how to interpret an address in the network.

As I said before, VLSM is similar in concept to the way classless addressing and routing (CIDR) is performed. The difference between VLSM and CIDR is primarily one of focus. VLSM deals with subnets of a single network in a private organization. CIDR takes the concept we just saw in VLSM to the Internet as a whole, by changing how organizational networks are allocated by replacing the single-level “classful” hierarchy with a multiple-layer hierarchy.

Variable Length Subnet Masking (VLSM)

Variable Length Subnet Masking (VLSM) is a way of further subnetting a subnet. Using Variable Length Subnet Masking (VLSM) we can allocate IP addresses to the subnets by the exact need. Variable Length Subnet Masking (VLSM) allows us to use more than one subnet mask within the same network address space. If we recollect from the previous lessons, we can divide a network only into subnets with equal number of IP addresses. Variable Length Subnet Masking (VLSM) allows creating subnets from a single network with unequal number of IP addresses.

Example: We want to divide 192.168.10.0, which is a Class C network, into four networks, each with unequal number of IP address requirements as shown below.

Subnet	A:	126 IP	Addresses.
Subnet	B:	62 IP	Addresses.
Subnet	C:	30 IP	Addresses.
Subnet	D:	30 IP	Addresses.

This type of division is not possible as described in previous lessons, since it divide the network equally, but is possible with Variable Length Subnet Masking (VLSM).

Original Network (Network to be sub-netted) – 192.168.10.0/24

Variable Length Subnet Masking (VLSM) - First Division

Divide the two networks equally with 128 IP Addresses (126 usable IP addresses) in each network using 255.255.255.128 subnet mask (192.168.10.0/25).

We will get two subnets each with 128 IP Addresses (126 usable IP addresses).

1) 192.168.10.0/25, which can be represented in binaries as below.

```
11000000.10101000.00001010.0|00000000
11111111.11111111.11111111.1 | 00000000
```

2) 192.168.10.128/25, which can be represented in binaries as below.

```
11000000.10101000.00001010.1|00000000
11111111.11111111.11111111.1 | 00000000
```

Variable Length Subnet Masking (VLSM)- Second Division

Divide second subnet (192.168.10.128/25) we got from the first division again into two Networks, each with 64 IP Addresses (62 usable IP Addresses) using 255.255.255.192

subnet mask.

We will get two subnets each with 64 IP Addresses (62 usable IP Addresses).

1) 192.168.10.128/26, which can be represented in binaries as below.

```
11000000.10101000.00001010.1|0|000000  
11111111.11111111.11111111.1|1|000000
```

2) 192.168.10.192/26

```
11000000.10101000.00001010.1|1|000000  
11111111.11111111.11111111.1|1|000000
```

Variable Length Subnet Masking (VLSM) - Third Division

Divide 192.168.10.192/26 Network again into two Networks, each with 32 IP Addresses (30 usable IP addresses) using 255.255.255.224 subnet mask

We will get two subnets each with 32 IP Addressee (30 usable IP addresses).

1) 192.168.10.192/27, which can be represented in binaries as below.

```
11000000.10101000.00001010.11|0|00000  
11111111.11111111.11111111.11|1|00000
```

2) 192.168.10.224/27, which can be represented in binaries as below.

```
11000000.10101000.00001010.11|1|00000  
11111111.11111111.11111111.11|1|00000
```

Now we have split the 192.168.10.0/24 network into four subnets using Variable Length Subnet Masking (VLSM), with unequal number of IP addresses as shown below. Also note that when you divide a network using Variable Length Subnet Masking (VLSM), the subnet masks are also different.

- | | | | | | | | |
|----|----------------|---|-----------------|------|---------|-----------|------------|
| 1) | 192.168.10.0 | - | 255.255.255.128 | (126 | (128-2) | usable IP | Addresses) |
| 2) | 192.168.10.128 | - | 255.255.255.192 | (62 | (64-2) | usable IP | Addresses) |
| 3) | 192.168.10.192 | - | 255.255.255.224 | (30 | (32-2) | usable IP | Addresses) |
| 4) | 192.168.10.224 | - | 255.255.255.224 | (30 | (32-2) | usable IP | Addresses) |

You have learned the term Variable Length Subnet Masking (VLSM) and how we can use Variable Length Subnet Masking (VLSM) for subnetting a subnet. Click "Next" to continue.

