

eSubnet Fragment Article

Four Golden Rules for IP Address Deployment

Internet Protocol Version 4 (IPv4) addressing, more specifically network subnets, is an often misunderstood technology. Assigning IP address ranges usefully within a network can be a daunting task. To design an effective IP addressing plan for an organization, keep in mind what I call the Four Golden Rules for IP Address Deployment:

1. All IP subnets have natural boundaries.
2. Smaller adjacent subnets fit into larger subnets, most of the time.
3. Smaller routing tables require less management and resources.
4. Grouping like objects into “natural” subnets makes them easier to control.

This is the background for a solid foundation for handling IP address assigning which makes it easier to identify IP traffic by source or destination address, have control over network objects in firewalls and routers, and use tighter router and firewall configurations.

I am not going to go into how to calculate IP subnets. There are extensive resources online to learn about subnet calculation. For the purposes of writing this article, I used an online tool, available here: <http://www.esubnet.net/subnet-calc.html>

As well, for all examples I use the three IP address ranges reserved for internal networks equipped with Network Address Translation (NAT). The Internet Assigned Numbers Authority (IANA) has reserved the following three blocks of the IP address space for private networks:

Table 1: IP address space for private networks

10.0.0.0	10.255.255.255 (10/8 prefix)
172.16.0.0	172.31.255.255 (172.16/12 prefix)
192.168.0.0	192.168.255.255 (192.168/16 prefix)

For more information on these IP address ranges please refer to RFC 1918: Address Allocation for Private Internets

The Golden Rules

1. IP address natural boundaries

In the dawn of IPv4 usage, three different subnet sizes were used to assign network addresses. They were labeled by their size differences. The largest, Class A, could accommodate 16,777,214 IP addresses per subnet (or, $2^{24} - 1$). The next size down, Class B, accommodated 65,534 IP addresses per subnet (or, $2^{16} - 1$). The smallest, Class C, accommodated only 254 IP addresses (or, $2^8 - 1$).

In calculating the number of available IP addresses, don't forget to subtract one for the router. Without a router, the traffic bound for other networks including the public Internet will never leave the local network.

The simple class-full networks have subnet boundaries based on the size of the network as defined by their subnet mask. The per class subnets are shown below:

Table 2: Per class subnets

CLASS	Decimal	Binary Notation or Bits
Class-A	255.0.0.0	11111111.00000000.00000000.00000000
Class-B	255.255.0.0	11111111.11111111.00000000.00000000
Class-C	255.255.255.0	11111111.11111111.11111111.00000000

In the case of class-full network subnets, there is a large amount of waste with small networks such as point-to-point links. To resolve this, a technique known as Variable Length Subnet Masking (VLSM) was developed, which allows for the class-based networks to be broken in to small chunks or to be bundled in to large chunks. VLSM is also known as Classless Internet Domain Routing (CIDR) and is usually represented by a / followed by a number indicating the size of the subnet.

Table 3: Variable length subnet masking

CLASS	Decimal	Binary Notation or Bits	CIDR Notation
Class-A	255.0.0.0	11111111.00000000.00000000.00000000	/8
Class-B	255.255.0.0	11111111.11111111.00000000.00000000	/16
Class-C	255.255.255.0	11111111.11111111.11111111.00000000	/24

The number after the / represents the number of masked bits. In IPv4 addressing the masked bits are that portion of the IP address which represent the network and the unmasked bits represent the possible unique IP addresses of any host. As the masked bit size increases the subnet gets smaller. The table below demonstrates this.

Table 4: Masked bit size vs. Network size

/24	/25	/26	/27	
172.16.1.0 – 172.16.1.255	172.16.1.0 – 172.16.1.127	172.16.1.0-63	172.16.1.0-31	
			172.16.1.32-63	
		172.16.1.64-127	172.16.1.64-95	
			172.16.1.96-127	
	172.16.1.128 – 172.16.1.255	172.16.1.128-191		172.16.1.128-159
				172.16.1.160-191
		172.16.1.192-255		172.16.1.192-223
				172.16.1.224-255

Table 4 above shows the effect of making subnets. The area of each box in the table is one half of the area of box to its left; the horizontal lines in the table represent the natural boundaries for the subnets.

2. Fitting smaller ranges into large ranges: super netting

In Table 4, we progressed from a /24 through to a /27, thereby halving the size of the network with each increase in mask size as we moved from left to right, this is subnetting. This same table may be read right to left to achieve super netting. Two /25 subnets fit into a /24, and two /26 subnets fit into a /25. This means that four /26 subnets fit into a /24. When you wish to combine four /27 subnets into a single /25 subnet ensure you pick the right four.

Example of super netting that does NOT work

The following four subnets; 172.16.1.64-95, 172.16.1.96-127, 172.16.1.128-159, 172.16.1.160-191 overlap between the two /25 subnets in the table and therefore can not be super netted. Natural boundaries are not respected.

Example that works

The following four subnets; 172.16.1.0-31, 172.16.1.32-63, 172.16.1.64-95, 172.16.1.96-127 can be super netted by the range 172.16.1.0-127 which as our table shows is indeed a /25. All natural boundaries respected.

The proper application of super netting allows for few routes in your routing table this is referred to as route summarization.

3. Routing table, the smaller the better

By following the rules and practices for super netting and route summarization you can change three routing statements in to one. The following scenario shows the benefit.

Table 5: Example of shortening routing statements

Subnet	Building #1	Building #2	Building #3	Building #4	Building #5
Building Net	172.16.1.0/24	172.16.2.0/24	172.16.3.0/24	172.16.4.0/24	172.16.5.0/24
Network gear and Printers	172.16.1.0/26	172.16.2.0/26	172.16.3.0/26	172.16.4.0/26	172.16.5.0/26
Servers	172.16.1.64/26	172.16.2.64/26	172.16.3.64/26	172.16.4.64/26	172.16.5.64/26
Workstations	172.16.1.128/25	172.16.2.128/25	172.16.3.128/25	172.16.4.128/25	172.16.5.128/25

There is a campus area network which has 5 buildings each containing 3 subnets. Each subnet is to be provided with its own VLAN and subsequent IP range. The IP addressing could look as shown above in Table 5.

The entire campus can be summarized as 172.16.0.0/21 which covers the 7 buildings, and 172.16.0.0/24 which covers the network between buildings. When you know that there will never be more than 7 buildings the next campus can be summarized with the following address 172.16.8.0/21. The campus after that would be 172.16.16.0/21. Leaving the geo-political boundary is simple, using the 172.16.0.0 - 172.31.255.255 range we can increment the second number (also known as an octet) to define areas such as city, province or country

Route summarization provides one major advantage, speed. The one campus network described above would require, without summarization, 15 lines in the routing table of the routers connecting the buildings, and this increases in a linear fashion with each campus added. We have only 5 routing statements with summarization for the campus, due to this the other campuses only require an additional line; as do groups of campuses in other regions.

By keeping the routing table small, the hardware works less. A smaller, less expensive router will do the job. Additionally, if dynamic routing protocols are used, there is less routing data transmitted between the routers, leaving more bandwidth to handle the user data. The links will be more stable and the user data will flow better. And therefore cheaper WAN links can be provisioned.

By deploying proper route summarization significant cost savings can be had by lower bandwidth costs and reduced system requirements resulting in smaller therefore less expensive devices.

4. Grouping for easier control

By grouping the devices on your network along the natural boundaries for subnets, even if all of the IP addresses are within a larger subnet, you

can easily divide the larger subnet into smaller ones on your firewalls, or within your routers, for the purpose of defining an Access Control List (ACL).

Returning to the earlier scenario we reexamine Building #1:

Table 6: Example of grouping along natural boundaries

Subnet	CIDR Network
Network gear and Printers	172.16.1.0/26
Servers	172.16.1.64/26
Workstations	172.16.1.128/25

The network gear and printers never need to access the public Internet. And so a single ACL can be used to deny that traffic. Since the servers should never be accessed from the public Internet, a single ACL can be used. Without this sort of grouping, each of the servers needs to be defined as an object, and then placed in a group with an ACL applied. Adding a new server and forgetting to place it in the group means that server falls outside of the company security policy.

A note on the use of numbering systems We're used to thinking in base-10 or decimal number systems: our number system, the metric system, and our currency are all using base-10.

IPv4 addressing uses base-2 or binary numbering, instead. So keep this in mind. When you are grouping hosts into ranges of IP addresses, remember to look to the natural boundaries for subnets in base-2. As an example the range 0-64 (/26) works much better than 10-75, as the later range is really 0-128 (/27).

My favourite .sig file quote for this topic: *"There are 10 kinds of people in the world, those who understand binary numbers, and those who don't."*

Conclusion

Remember to respect the Four Golden Rules of IP Address deployment. Respect the natural boundaries of your IP address ranges. Try to fit smaller adjacent subnets into large subnets. Summarize your routing tables for ease of management, and ultimately cost savings and, group like objects into "natural" IP address ranges to allow for easier control

Originally published July, 2009

About the Author

Richard Danielli is the founder and President of eSubnet Enterprises. He has broad expertise in the fields of networking and data security. Based in Toronto, eSubnet provides superior customized solutions for networking and data security. Additional articles are available at <http://www.esubnet.com/fragment.html>.

Reprints

1. Commercial use: Commercial newspapers, magazines and websites that wish to reprint an article should contact eSubnet for permission.
2. Non-Commercial use: Any reproduction of content (excerpt or full text) by nonprofit organizations or companies wishing to share information with their members or clients requires that you notify us, whether for print or online usage.

Regarding reprints, contact Mr. Danielli at rdanielli@esubnet.com or at (416) 203-1223.