


Subnetwork

A **subnetwork** or **subnet** is a logical subdivision of an IP network.^{[1]:1,16} The practice of dividing a network into two or more networks is called **subnetting**.

Computers that belong to the same subnet are addressed with an identical most-significant bit-group in their IP addresses. This results in the logical division of an IP address into two fields: the *network number* or *routing prefix* and the *rest field* or *host identifier*. The *rest field* is an identifier for a specific host or network interface.



Creating a subnet by dividing the host identifier

The *routing prefix* may be expressed in Classless Inter-Domain Routing (CIDR) notation written as the first address of a network, followed by a slash character (/), and ending with the bit-length of the prefix. For example, *198.51.100.0/24* is the prefix of the Internet Protocol version 4 network starting at the given address, having 24 bits allocated for the network prefix, and the remaining 8 bits reserved for host addressing. Addresses in the range *198.51.100.0* to *198.51.100.255* belong to this network. The IPv6 address specification *2001:db8::/32* is a large address block with 2⁹⁶ addresses, having a 32-bit routing prefix.

For IPv4, a network may also be characterized by its **subnet mask** or **netmask**, which is the bitmask that when applied by a bitwise AND operation to any IP address in the network, yields the routing prefix. Subnet masks are also expressed in dot-decimal notation like an address. For example, *255.255.255.0* is the subnet mask for the prefix *198.51.100.0/24*.

Traffic is exchanged between subnetworks through routers when the routing prefixes of the source address and the destination address differ. A router serves as a logical or physical boundary between the subnets.

The benefits of subnetting an existing network vary with each deployment scenario. In the address allocation architecture of the Internet using CIDR and in large organizations, it is necessary to allocate address space efficiently. Subnetting may also enhance routing efficiency, or have advantages in network management when subnetworks are administratively controlled by different entities in a larger organization. Subnets may be arranged logically in a hierarchical architecture, partitioning an organization's network address space into a tree-like routing structure, or other structures such as meshes.

Contents

Network addressing and routing

Internet Protocol version 4

Determining the network prefix

Subnetting

Special addresses and subnets

Subnet host count

Internet Protocol version 6

See also

References

[Further reading](#)

[External links](#)

Network addressing and routing

Computers participating in a network such as the [Internet](#) each have at least one [network address](#). Usually, this address is unique to each device and can either be configured automatically with the [Dynamic Host Configuration Protocol \(DHCP\)](#) by a network server, manually by an administrator, or automatically by [stateless address autoconfiguration](#).


An address fulfills the functions of identifying the host and locating it on the network. The most common network addressing architecture is [Internet Protocol version 4 \(IPv4\)](#), but its successor, [IPv6](#), has been increasingly [deployed](#) since approximately 2006. An IPv4 address consists of 32 bits, for readability written in a form consisting of four [decimal octets](#) separated by dots, called [dot-decimal notation](#). An IPv6 address consists of 128 bits written in [hexadecimal notation](#) and groupings of 16 bits, called [hextets](#), separated by colons. An IP address is divided into two logical parts, the *network prefix* and the *host identifier*. All hosts on a subnetwork have the same network prefix. This prefix occupies the most-significant bits of the address. The number of bits allocated within a network to the prefix may vary between subnets, depending on the network architecture. The host identifier is a unique local identification and is either a host number on the local network or an interface identifier.

This addressing structure permits the selective [routing](#) of IP packets across multiple networks via special gateway computers, called [routers](#), to a destination host if the network prefixes of origination and destination hosts differ, or sent directly to a target host on the local network if they are the same. Routers constitute logical or physical borders between the subnets, and manage traffic between them. Each subnet is served by a designated default router but may consist internally of multiple physical [Ethernet](#) segments interconnected by [network switches](#).

The routing prefix of an address is identified by the *subnet mask*, written in the same form used for IP addresses. For example, the subnet mask for a routing prefix that is composed of the most-significant 24 bits of an IPv4 address is written as `255.255.255.0`.

The modern standard form of specification of the network prefix is [CIDR notation](#), used for both IPv4 and IPv6. It counts the number of bits in the prefix and appends that number to the address after a *slash (/)* character separator. This notation was introduced with [Classless Inter-Domain Routing \(CIDR\)](#).^[2] In IPv6 this is the only standards-based form to denote network or routing prefixes.

For example, the IPv4 network `192.0.2.0` with the subnet mask `255.255.255.0` is written as `192.0.2.0/24`, and the IPv6 notation `2001:db8::/32` designates the address `2001:db8::` and its network prefix consisting of the most significant 32 bits.



The concept of subnetting the IPv4 address space `200.100.10.0/24`, which contains 256 addresses, into two smaller address spaces, namely `200.100.10.0/25` and `200.100.10.128/25` with 128 addresses each.

In classful networking in IPv4, before the introduction of CIDR, the network prefix could be directly obtained from the IP address, based on its highest order bit sequence. This determined the class (A, B, C) of the address and therefore the subnet mask. Since the introduction of CIDR, however, the assignment of an IP address to a network interface requires two parameters, the address and a subnet mask.

Given an IPv4 source address, its associated subnet mask, and the destination address, a router can determine whether the destination is on a locally connected network or a remote network. The subnet mask of the destination is not needed, and is generally not known to a router.^[3] For IPv6, however, on-link determination is different in detail and requires the Neighbor Discovery Protocol (NDP).^{[4][5]} IPv6 address assignment to an interface carries no requirement of a matching on-link prefix and vice versa, with the exception of link-local addresses.

Since each locally connected subnet must be represented by a separate entry in the routing tables of each connected router, subnetting increases routing complexity. However, by careful design of the network, routes to collections of more distant subnets within the branches of a tree hierarchy can be aggregated into a supernet and represented by single routes.

Internet Protocol version 4

Determining the network prefix

An IPv4 subnet mask consists of 32 bits; it is a sequence of ones (1) followed by a block of zeros (0). The ones indicate bits in the address used for the network prefix and the trailing block of zeros designates that part as being the host identifier.

The following example shows the separation of the network prefix and the host identifier from an address (192.0.2.130) and its associated /24 subnet mask (255.255.255.0). The operation is visualized in a table using binary address formats.

	Binary form	Dot-decimal notation
IP address	11000000.00000000.00000010.10000010	192.0.2.130
Subnet mask	11111111.11111111.11111111.00000000	255.255.255.0
Network prefix	11000000.00000000.00000010.00000000	192.0.2.0
Host identifier	00000000.00000000.00000000.10000010	0.0.0.130

The result of the bitwise AND operation of IP address and the subnet mask is the network prefix 192.0.2.0. The host part, which is 130, is derived by the bitwise AND operation of the address and the one's complement of the subnet mask.

Subnetting

Subnetting is the process of designating some high-order bits from the host part as part of the network prefix and adjusting the subnet mask appropriately. This divides a network into smaller subnets. The following diagram modifies the above example by moving 2 bits from the host part to the network prefix to form four smaller subnets each one quarter of the previous size.

	Binary form	Dot-decimal notation
IP address	11000000.00000000.00000010.10000010	192.0.2.130
Subnet mask	11111111.11111111.11111111.11000000	255.255.255.192
Network prefix	11000000.00000000.00000010.10000000	192.0.2.128
Host part	00000000.00000000.00000000.00000010	0.0.0.2

Special addresses and subnets

IPv4 uses specially designated address formats to facilitate recognition of special address functionality. The first and the last subnets obtained by subnetting a larger network have traditionally had a special designation and, early on, special usage implications.^[6] In addition, IPv4 uses the *all ones* host address, i.e. the last address within a network, for broadcast transmission to all hosts on the link.

The first subnet obtained from subnetting a larger network has all bits in the subnet bit group set to zero (0). It is therefore called *subnet zero*.^[7] The last subnet obtained from subnetting a larger network has all bits in the subnet bit group set to one (1). It is therefore called the *all-ones subnet*.^[8]

The IETF originally discouraged the production use of these two subnets. When the prefix length is not available, the larger network and the first subnet have the same address, which may lead to confusion. Similar confusion is possible broadcast address at the end of the last subnet. Therefore, reserving the subnet values consisting of all zeros and all ones on the public Internet was recommended,^[9] reducing the number of available subnets by two for each subnetting. This inefficiency was removed, and the practice was declared obsolete in 1995 and is only relevant when dealing with legacy equipment.^[10]

Although the all-zeros and the all-ones host values are reserved for the network address of the subnet and its broadcast address, respectively, in systems using CIDR all subnets are available in a subdivided network. For example, a /24 network can be divided into sixteen usable /28 networks. Each broadcast address, i.e. *.15, *.31, ..., *.255, reduces only the host count in each subnetwork.

Subnet host count

The number of subnetworks available and the number of possible hosts in a network may be readily calculated. For instance, the 192.168.5.0/24 network may be subdivided into the following four /26 subnets. The highlighted two address bits become part of the network number in this process.

Network	Network (binary)	Broadcast address
192.168.5.0/26	11000000.10101000.00000101.00000000	192.168.5.63
192.168.5.64/26	11000000.10101000.00000101.01000000	192.168.5.127
192.168.5.128/26	11000000.10101000.00000101.10000000	192.168.5.191
192.168.5.192/26	11000000.10101000.00000101.11000000	192.168.5.255

The remaining bits after the subnet bits are used for addressing hosts within the subnet. In the above example, the subnet mask consists of 26 bits, making it 255.255.255.192, leaving 6 bits for the host identifier. This allows for 62 host combinations (2^6-2).

In general, the number of available hosts on a subnet is $2^h - 2$, where h is the number of bits used for the host portion of the address. The number of available subnets is 2^n , where n is the number of bits used for the network portion of the address.

There is an exception to this rule for 31-bit subnet masks,^[11] which means the host identifier is only one bit long for two permissible addresses. In such networks, usually point-to-point links, only two hosts (the end points) may be connected and a specification of network and broadcast addresses is not necessary.

Internet Protocol version 6

The design of the IPv6 address space differs significantly from IPv4. The primary reason for subnetting in IPv4 is to improve efficiency in the utilization of the relatively small address space available, particularly to enterprises. No such limitations exist in IPv6, as the large address space available, even to end-users, is not a limiting factor.

As in IPv4, subnetting in IPv6 is based on the concepts of variable-length subnet masking (VLSM) and the Classless Inter-Domain Routing methodology. It is used to route traffic between the global allocation spaces and within customer networks between subnets and the Internet at large.

A compliant IPv6 subnet always uses addresses with 64 bits in the host identifier.^[12] Given the address size of 128 bits, it therefore has a /64 routing prefix. Although it is technically possible to use smaller subnets,^[13] they are impractical for local area networks based on Ethernet technology, because 64 bits are required for stateless address autoconfiguration.^[14] The Internet Engineering Task Force recommends the use of /127 subnets for point-to-point links, which have only two hosts.^{[15][16]}

IPv6 does not implement special address formats for broadcast traffic or network numbers,^[17] and thus all addresses in a subnet are acceptable for host addressing. The all-zeroes address is reserved as the subnet-router anycast address.^[18]

In the past, the recommended allocation for an IPv6 customer site was an address space with a 48-bit (/48) prefix.^[19] However, this recommendation was revised to encourage smaller blocks, for example using 56-bit prefixes.^[20] Another common allocation size for residential customer networks has a 64-bit prefix.

See also

- Autonomous system (Internet)

References

1. Jeffrey Mogul; Jon Postel (August 1985). *Internet Standard Subnetting Procedure* (<https://tools.ietf.org/html/rfc950>). IETF. doi:10.17487/RFC0950 (<https://doi.org/10.17487%2FRFC0950>). RFC 950 (<https://tools.ietf.org/html/rfc950>). Updated by RFC 6918.
2. V. Fuller; T. Li (August 2006). *Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan* (<https://tools.ietf.org/html/rfc4632>). Network Working Group. doi:10.17487/RFC4632 (<https://doi.org/10.17487%2FRFC4632>). RFC 4632 (<https://tools.ietf.org/html/rfc4632>).
3. R. Braden, ed. (October 1989). *Requirements for Internet Hosts -- Communication Layers* (<https://tools.ietf.org/html/rfc1122#section-3.3.1>). Network Working Group IETF. sec. 3.3.1. doi:10.17487/RFC1122 (<https://doi.org/10.17487%2FRFC1122>). RFC 1122 (<https://tools.ietf.org/html/rfc1122>). Updated by RFC 1349, RFC 4379, RFC 5884, RFC 6093, RFC 6298, RFC 6633, RFC 6864, RFC 8029.

4. T. Narten; E. Nordmark; W. Simpson; H. Soliman (September 2007). *Neighbor Discovery for IP version 6 (IPv6)* (<https://tools.ietf.org/html/rfc4861>). Network Working Group. doi:10.17487/RFC4861 (<https://doi.org/10.17487%2FRFC4861>). RFC 4861 (<https://tools.ietf.org/html/rfc4861>).
5. H. Singh; W. Beebe; E. Nordmark (July 2010). *IPv6 Subnet Model: The Relationship between Links and Subnet Prefixes* (<https://tools.ietf.org/html/rfc5942>). IETF. doi:10.17487/RFC5942 (<https://doi.org/10.17487%2FRFC5942>). RFC 5942 (<https://tools.ietf.org/html/rfc5942>).
6. "Document ID 13711 - Subnet Zero and the All-Ones Subnet" (http://www.cisco.com/en/US/tech/tk648/tk361/technologies_tech_note09186a0080093f18.shtml). Cisco Systems. 2005-08-10. Retrieved 2010-04-25. "Traditionally, it was strongly recommended that subnet zero and the all-ones subnet not be used for addressing. [...] Today, the use of subnet zero and the all-ones subnet is generally accepted and most vendors support their use."
7. "Document ID 13711 - Subnet Zero and the All-Ones Subnet" (http://www.cisco.com/en/US/tech/tk648/tk361/technologies_tech_note09186a0080093f18.shtml). Cisco Systems. 2005-08-10. Retrieved 2010-04-23. "the first [...] subnet[...], known as subnet zero"
8. "Document ID 13711 - Subnet Zero and the All-Ones Subnet" (http://www.cisco.com/en/US/tech/tk648/tk361/technologies_tech_note09186a0080093f18.shtml). Cisco Systems. 2005-08-10. Retrieved 2010-04-23. "[...] the last subnet[...], known as [...] the all-ones subnet"
9. Jeffrey Mogul; Jon Postel (August 1985). *Internet Standard Subnetting Procedure* (<https://tools.ietf.org/html/rfc950#page-6>). IETF. p. 6. doi:10.17487/RFC0950 (<https://doi.org/10.17487%2FRFC0950>). RFC 950 (<https://tools.ietf.org/html/rfc950>). "It is useful to preserve and extend the interpretation of these special addresses in subnetted networks. This means the values of all zeros and all ones in the subnet field should not be assigned to actual (physical) subnets."
10. Troy Pummill; Bill Manning (December 1995). *Variable Length Subnet Table For IPv4* (<https://tools.ietf.org/html/rfc1878>). IETF. doi:10.17487/RFC1878 (<https://doi.org/10.17487%2FRFC1878>). RFC 1878 (<https://tools.ietf.org/html/rfc1878>). "This practice is obsolete! Modern software will be able to utilize all definable networks." (Informational RFC, demoted to category *Historic*)
11. A. Retana; R. White; V. Fuller; D. McPherson (December 2000). *Using 31-Bit Prefixes on IPv4 Point-to-Point Links* (<https://tools.ietf.org/html/rfc3021>). doi:10.17487/RFC3021 (<https://doi.org/10.17487%2FRFC3021>). RFC 3021 (<https://tools.ietf.org/html/rfc3021>).
12. R. Hinden; S. Deering (February 2006). *IP Version 6 Addressing Architecture - section 2.5.1. Interface Identifiers* (<https://tools.ietf.org/html/rfc4291#section-2.5.1>). IETF. sec. 2.5.1. doi:10.17487/RFC4291 (<https://doi.org/10.17487%2FRFC4291>). RFC 4291 (<https://tools.ietf.org/html/rfc4291>). "For all unicast addresses, except those that start with the binary value 000, Interface IDs are required to be 64 bits long and to be constructed in Modified EUI-64 format." (Updated by RFC 5952, RFC 6052, RFC 7136, RFC 7346, RFC 7371, RFC 8064.)
13. S. Thomson; T. Narten; T. Jinmei (September 2007). *IPv6 Stateless Address Autoconfiguration - section 5.5.3.(d) Router Advertisement Processing* (<https://tools.ietf.org/html/rfc4862#section-5.5.3>). IETF. sec. 5.5.3. doi:10.17487/RFC4862 (<https://doi.org/10.17487%2FRFC4862>). RFC 4862 (<https://tools.ietf.org/html/rfc4862>). "It is the responsibility of the system administrator to ensure that the lengths of prefixes contained in Router Advertisements are consistent with the length of interface identifiers for that link type. [...] an implementation should not assume a particular constant. Rather, it should expect any lengths of interface identifiers." (Updated by RFC 7527.)
14. M. Crawford (December 1998). *Transmission of IPv6 Packets over Ethernet Networks - section 4 Stateless Autoconfiguration* (<https://tools.ietf.org/html/rfc2464#section-4>). IETF. sec. 4. doi:10.17487/RFC2464 (<https://doi.org/10.17487%2FRFC2464>). RFC 2464 (<https://tools.ietf.org/html/rfc2464>). "The Interface Identifier [AARCH] for an Ethernet interface is based on the EUI-64 identifier [EUI64] derived from the interface's built-in 48-bit IEEE 802 address. [...] An IPv6 address prefix used for stateless autoconfiguration [ACONF] of an Ethernet interface must have a length of 64 bits." (Updated by RFC 6085, RFC 8064.)

15. M. Kohno; B. Nitzan; R. Bush; Y. Matsuzaki; L. Colitti; T. Narten (April 2011). *Using 127-Bit IPv6 Prefixes on Inter-Router Links* (<https://tools.ietf.org/html/rfc6164>). IETF. doi:10.17487/RFC6164 (<https://doi.org/10.17487%2FRFC6164>). RFC 6164 (<https://tools.ietf.org/html/rfc6164>). "On inter-router point-to-point links, it is useful, for security and other reasons, to use 127-bit IPv6 prefixes."
16. W. George (February 2012). *RFC 3627 to Historic Status* (<https://tools.ietf.org/html/rfc6547>). IETF. doi:10.17487/RFC6547 (<https://doi.org/10.17487%2FRFC6547>). RFC 6547 (<https://tools.ietf.org/html/rfc6547>). "This document moves "Use of /127 Prefix Length Between Routers Considered Harmful" (RFC 3627) to Historic status to reflect the updated guidance contained in "Using 127-Bit IPv6 Prefixes on Inter-Router Links" (RFC 6164)."
17. R. Hinden; S. Deering (February 2006). *IP Version 6 Addressing Architecture - section 2 IPv6 Addressing* (<https://tools.ietf.org/html/rfc4291#section-2>). IETF. sec. 2. doi:10.17487/RFC4291 (<https://doi.org/10.17487%2FRFC4291>). RFC 4291 (<https://tools.ietf.org/html/rfc4291>). "There are no broadcast addresses in IPv6, their function being superseded by multicast addresses. [...] In IPv6, all zeros and all ones are legal values for any field, unless specifically excluded."
18. R. Hinden; S. Deering (February 2006). *IP Version 6 Addressing Architecture - section 2.6.1 Required Anycast Address* (<https://tools.ietf.org/html/rfc4291#section-2.6.1>). IETF. sec. 2.6.1. doi:10.17487/RFC4291 (<https://doi.org/10.17487%2FRFC4291>). RFC 4291 (<https://tools.ietf.org/html/rfc4291>). "This anycast address is syntactically the same as a unicast address for an interface on the link with the interface identifier set to zero."
19. "IPv6 Addressing Plans" (http://www.getipv6.info/index.php?title=IPv6_Addressing_Plans&oldid=2998). ARIN IPv6 Wiki. Retrieved 2010-04-25. "All customers get one /48 unless they can show that they need more than 65k subnets. [...] If you have lots of consumer customers you may want to assign /56s to private residence sites."
20. T. Narten; G. Huston; L. Roberts (March 2011). *IPv6 Address Assignment to End Sites* (<https://tools.ietf.org/html/rfc6177>). IETF. doi:10.17487/RFC6177 (<https://doi.org/10.17487%2FRFC6177>). ISSN 2070-1721 (<https://www.worldcat.org/issn/2070-1721>). BCP 157. RFC 6177 (<https://tools.ietf.org/html/rfc6177>). "APNIC, ARIN, and RIPE have revised the end site assignment policy to encourage the assignment of smaller (i.e., /56) blocks to end sites."

Further reading

- *Requirements for IPv4 Routers* (<https://tools.ietf.org/html/rfc1812>). doi:10.17487/RFC1812 (<https://doi.org/10.17487%2FRFC1812>). RFC 1812 (<https://tools.ietf.org/html/rfc1812>).
- *Utility of subnets of Internet networks* (<https://tools.ietf.org/html/rfc917>). doi:10.17487/RFC0917 (<https://doi.org/10.17487%2FRFC0917>). RFC 917 (<https://tools.ietf.org/html/rfc917>).
- *DNS Encodings of Network Names and Other Type* (<https://tools.ietf.org/html/rfc1101>). doi:10.17487/RFC1101 (<https://doi.org/10.17487%2FRFC1101>). RFC 1101 (<https://tools.ietf.org/html/rfc1101>).
- Blank, Andrew G. (2006). *TCP/IP Foundations*. Wiley. ISBN 9780782151138.
- Lammle, Todd (2005). *CCNA Cisco Certified Network Associate Study Guide 5th Edition*. San Francisco, London: Sybex.
- Groth, David; Skandier, Toby (2005). *Network + Study Guide* (4th ed.). San Francisco, London: Wiley.

External links

- [Cisco-IP Addressing and Subnetting for New Users](http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a00800a67f5.shtml) (http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a00800a67f5.shtml)
- [Subnetworking](https://curlie.org/Computers/Internet/Protocols/IP/Addressing/) (<https://curlie.org/Computers/Internet/Protocols/IP/Addressing/>) at Curlie

- [Netmask Quick Reference Chart \(http://unixwiz.net/techtips/netmask-ref.html\)](http://unixwiz.net/techtips/netmask-ref.html)
-

Retrieved from "<https://en.wikipedia.org/w/index.php?title=Subnetwork&oldid=1019946344>"

This page was last edited on 26 April 2021, at 09:05 (UTC).

Text is available under the Creative Commons Attribution-ShareAlike License; additional terms may apply. By using this site, you agree to the Terms of Use and Privacy Policy. Wikipedia® is a registered trademark of the Wikimedia Foundation, Inc., a non-profit organization.