Computer Networking

Computer Networking

Resources

Visual Subnet Calculator: https://www.davidc.net/sites/default/subnets/subnets.html

Overview

Premise

This guide is not designed to be all-encompassing nor a replacement for professional learning. This guide is purely to help describe overall concepts in TCP/IP networks and serve as a basis for further independent learning.

Traffic Types

A single packet of information on a network has three different means of communicating across that network: Unicast, Broadcast, and Multicast

Unicast

Unicast traffic is any traffic sent from one source to one destination and no other destinations receive the message.

Broadcast

Broadcast traffic is unlike unicast traffic as a single source will send a packet which is received by *every* device in the network. Broadcast packets should be used sparingly and with small payloads in order to minimize the impact that they have on the network as having too many unnecessary broadcast packets can flood a network with traffic and saturate the link speed across the network, effectively bringing the network to a crawl in terms of speed and latency. Broadcast packets are sent by setting either or both the IP address to 255.255.255.255 or MAC address to FFFF::FFFF (all F). As such, the top-most IP address of a subnet (x.x.x.255) is reserved for broadcast traffic and should not be assigned to devices.

Multicast

Multicast is a derrivative of broadcast as it allows one or more sources to send traffic to one or more destinations. Multicast is defined by the Internet Group Management Protocol or IGMP and is the most complex of the different means of sending traffic over the network and requires special configuration. The synopsis of multicast is that receiving devices inform the network that they want to listen to a specific multicast IP address, then senders can transmit data to that reserved IP address and it is received by the receivers. This requires careful configuration of the IGMP Querier and IGMP Snooping. The querier will periodically quiery devices to see what multicast traffic they are interested in receiving. IGMP snooping on switches will allow switches to read those packets to view the membership reports generated by the receivers so that they know what traffic to forward to them and what traffic to not. If the querier's interval is higher than the switches' IGMP snooping timeout interval, then the switches will stop sending traffic to receivers for a period until the querier re-establishes what the receivers are interested in listening to.

In the event IGMP snooping is disabled, multicast traffic will be forwarded as broadcast traffic by switches, which in video over IP or similarly high bandwidth workflows can cripple networks.

Multicast traffic is organized by desination IP address utilizing the set-aside range of 224.0.0.0 to 239.255.255.255 IPv4 addresses, however some addresses are reserved for IGMP management of the streams so referencing documentation prior to assigning transmit and receive IP addresses on devices is important.

OSI / TCP/IP Models

Overview

The foundation of modern computer networks is derrived from the OSI and TCP/IP models. These models were developed to help describe how computer networks operate on a logical level where changes can be made within any level of the models without affecting other levels. For example, say a new standard for how to transmit data over a new medium becomes available, it would only necessitate adaptions in layer 1 without requiring changes at other layers (usually).

The OSI model is generally described as 7 layers enumerated 1 through 7 while the TCP/IP model is generally either 4 or 5 layers. Both models are applicable for different applications, however when it comes to network design there is a seemingly standardized understanding of the layers and how they relate to a network as described below:

Layer	Name / Description	Examples
4	Transport Layer	TCP, UDP, Ports, Firewalls
3	Network Layer	IP Addresses, Routing, Subnets
2	Data Link Layer	MAC Addresses, Switching, VLANs
1	Physical Layer	Fiber, Twisted Pair, Coax, Encoding schemes

This descibed terminology is often used by manufacturers for marketing. A layer 3 switch for example is a standard ethernet switch which operates on layer 2, but contains layer 3 features such as routing between networks / subnets.

Layer 1 - Physical

The physical layer is responsible for how the data gets from one location to another. This means that the mediums themselves that the data is transported by and the encoding schemes used on those mediums are included within this layer. Generally speaking, modern high-speed networks rely primarily on using twisted pair cable, such as CAT-6, CAT-6A, etc. and fiber optics (either multimode or singlemode) for the core of the network. Historically, coaxial cable (for example DSL lines) has been used and is included in this layer, alongside other more specialized or experimental mediums such as laser communications between satellites.

Wireless devices such as those implementing Wi-Fi and wireless point-to-point links *can* technically be included in this layer, however those devices tend to spread into layer 2, so it is not entirely accurate to claim that they are encompassed entirely within the physical layer.

Most items within this layer are commonly-known and are easily researchable so this guide will not go in-depth into the physical layer, but the key take away is to know the limitations of each common modern medium that is used, such as distance limitations with twisted pair cable, or cost limitations with fiber optics including transceiver modules.

Layer 2 - Data Link

The data link layer is where Ethernet packets are formed and live. Each unit of infromation transmitted over the network is separated into one or more packets and then sent. Each packet has a header which has information on how to get the packet from its source to its destination, followed by the payload which is the raw data that is being sent.

Each endpoint device has a MAC address assigned to its interface, either manually or more typically directly from the manufacturer of the interface. Each packet then contains a source and destination MAC address within its header which helps to inform the network on where to send the packet, and tells the receiving device where the packet came from. Switches will read the header of a packet, then reference a table it has stored in its memory of what MAC addresses are on which port of the switch, then will send the packet out the port that corresponds to the destination MAC address. In the event that the switch does not have an entry in its MAC address table, the swtich will send the packet out of all ports in a process known as "flooding." In this way, the packet becomes broadcast rather than multicast.

Address Resolution Protocol (ARP) is commonly used on the layer 2 network in order for endpoint devices to find out the MAC address of a destination they want to reach. The simplified way it works is that a sender will send a broadcast packet out to the network asking which device has a certain IP address. Devices that receive the packet which do not have that IP address simply forget about the packet and do not respond. In the event that a device does correspond to that IP address, it will send a unicast response to the sender of the original packet stating its own MAC address. The sender will then cache the MAC address into its own MAC address table, so that whenever it needs to send data to that same IP address it can reference that table without sending another ARP packet. The entire conversation is also being listened two by the switches which take notice on the source addresses within the packets and record them alongside the port they arrived on in its own MAC address table, so that further communication to those devices can be done over unicast.

VLANs, or Virtual Local Area Networks, also live on layer 2. These are specified by the IEEE 802.1Q standard and as such are also referenced as dot1q. VLANs work as a means of separating traffic by effectively creating separate network topologies per VLAN whilst using the same hardware. For example, a single switch could have half of its ports on one VLAN while the other half are on a second VLAN. When a broadcast packet arrives at the switch, it will only be forwarded to the ports which are also assigned to the VLAN. This allows multicast and broadcast traffic to be separated between different devices. In the video engineering world, it is a good practice to separate different kinds of devices onto different VLANs. This way if a manufacturer uses an odd protocol or incorrectly implements the Ethernet protocol which results in significant broadcast traffic being sent, it minimizes which devices are forced to listen to those packets, freeing up bandwidth for those devices and potentially fixing other problems.

Layer 3 - Network

Layer 3 generally deals with IP addresses, subnets, and routing between different subnets. Like every device has a MAC address, they generally also have an IP address. Unlike with MAC addresses, IP addresses are typically assigned over the network by a DHCP server, however they can also be assigned manually directly on the device.

Subnets are used to separate traffic logically by dividing IP addresses into ranges of IP addresses. A subnet is also called a network as traffic cannot pass from one subnet to another without the need of a router to interconnect the two networks, however this can get confusing as the term network is

also used to refer to overarching communications systems involving multiple subnets. A subnet is defined in two parts: the address and the mask. An IPv4 address involves 4 separate bytes which represent the subnet. The mask is similarly also 4 bytes, however each logical `1` in the mask means that the corresponding digit in the address is part of the prefix defining the range of the subnet. The mask is organized with logical `1`s on the left and logical `0`s on the right. There are various ways of describing a subnet in more human-friendly non-binary. First, the subnet and mask can be separate in decimal such as an address of 10.10.123.0 and a mask of 255.255.255.0. This can however be abreviated by using CIDR notation which involves stating the address in decimal as normal, followed by a forward slash and then in decimal the number of leading `1`s in the subnet mask. The same subnet could therefore be defined in CIDR notation as 10.10.123.0/24. Regardless of notation, this subnet would define all IP addresses between 10.10.123.0 and 10.10.123.255.

Routers work by having multiple interfaces onto different subnets, and have IP addresses themselves unlike switches. The interfaces can be either physical ports, or virtual by using VLANs. When a packet arrives at the router, the router will check the desination IP against a table of connected subnets known as the routing table. If the router does not know how to get the packet to its desination subnet, unlike switches which flood, the router will delete the packet in a process referred to as "dropping" it. The routing table can include subnets not directly attached, but are directly attached to a router on a subnet it is directly attached to. Furthermore, some routers offer the ability to set a default gateway themselves, so instead of dropping a packet, the router can forward that packet to its own default gateway to process.

When a device on one subnet wants to communicate to a device on a different subnet, one or more routers needs to interconnect them. The sending device would recognize that it is sending data to an IP address outside of its subnet, reference its IP configuration for the default gateway, then send the packet over the network by substituting the MAC address of the default gateway for the receiver's MAC address. The default gateway is a router which then would reference its own configuration and re-send the packet to the receiving device over the correct network, substituting the MAC addresses again. This process can be repeated for however many routers the packet needs to go through for it to finally reach the subnet that the destination IP address corresponds with.

As many subnets can exist within a network, for example the entire internet, it can be tedious to manually program the entire routing table for each router. As such, route sharing protocols exist which allow engineers to just program the routing table for a router's directly attached subnets. Those route sharing protocols can then be used for the router to discover other routers on the network and their own directly attached subnets. Common protocols include eBGP and iBGP, OSPF, RIP, IS-IS, IGRP and EGRIP.

As a final note, you can view subnets and VLANs as being nearly logically equivilant but on different levels. It is generally best practice for every subnet to have its own VLAN, and for every VLAN to have its own subnet. This often causes some confusion as some people will refer to subnets as VLANs, or VLANs as subnets, so understanding that they work hand-in-hand will help in pushing past the confusion.

Layer 4 - Transport

Common Networking Protocols

TCP UDP Websocket SRT

Revision #2
Created 2026-07-30 20:24:57 UTC by iantech
Updated 2026-07-30 20:25:31 UTC by iantech