Contents

Multicast overview

Introduction to multicast

As a technique that coexists with unicast and broadcast, the multicast technique effectively addresses the issue of point-to-multipoint data transmission. By enabling high-efficiency point-to-multipoint data transmission over a network, multicast greatly saves network bandwidth and reduces network load.

By using multicast technology, a network operator can easily provide bandwidth-critical and time-critical information services. These services include live webcasting, Web TV, distance learning, telemedicine, Web radio, and real-time video conferencing.

Information transmission techniques

The information transmission techniques include unicast, broadcast, and multicast.

Unicast

In unicast transmission, the information source must send a separate copy of information to each host that needs the information.

Figure 1 Unicast transmission

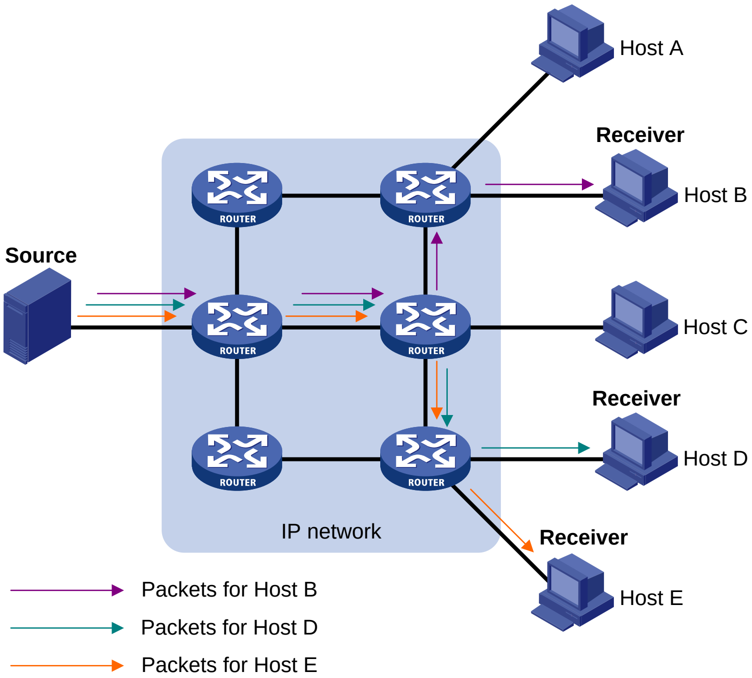

In Figure 1, Host B, Host D, and Host E need the information. A separate transmission channel must be established from the information source to each of these hosts.

In unicast transmission, the traffic transmitted over the network is proportional to the number of hosts that need the information. If a large number of hosts need the information, the information source must send a separate copy of the same information to each of these hosts. Sending many copies can place a tremendous pressure on the information source and the network bandwidth.

Unicast is not suitable for batch transmission of information.

Broadcast

In broadcast transmission, the information source sends information to all hosts on the subnet, even if some hosts do not need the information.

Figure 2 Broadcast transmission

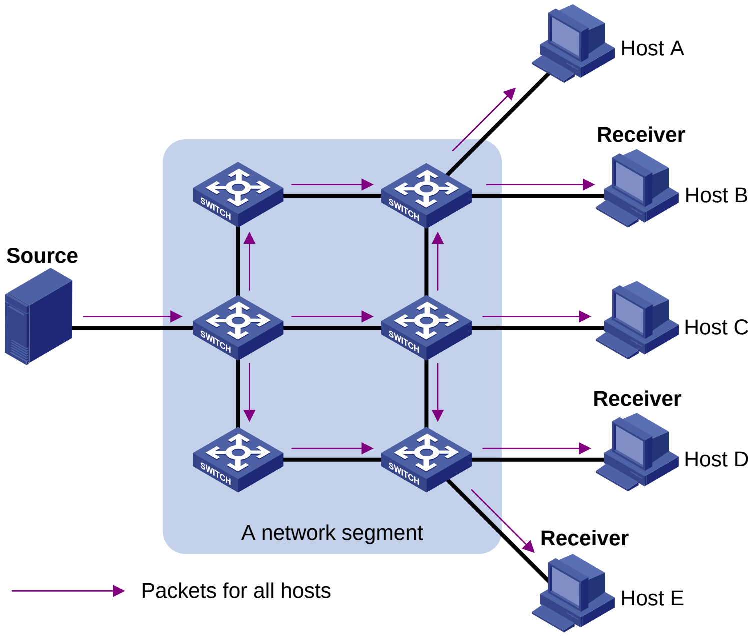

In Figure 2, only Host B, Host D, and Host E need the information. If the information is broadcast to the subnet, Host A and Host C also receive it. In addition to information security issues, broadcasting to hosts that do not need the information also causes traffic flooding on the same subnet.

Broadcast is disadvantageous in transmitting data to specific hosts. Moreover, broadcast transmission is a significant waste of network resources.

Multicast

Multicast provides point-to-multipoint data transmissions with the minimum network consumption. When some hosts on the network need multicast information, the information sender, or multicast source, sends only one copy of the information. Multicast distribution trees are built through multicast routing protocols, and the packets are replicated only on nodes where the trees branch.

Figure 3 Multicast transmission

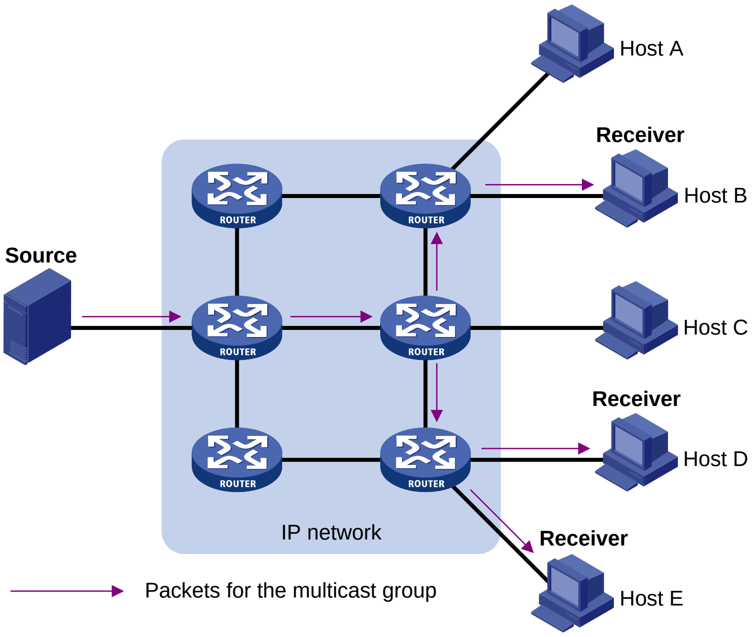

In Figure 3, the multicast source sends only one copy of the information to a multicast group. Host B, Host D, and Host E, which are information receivers, must join the multicast group. The routers on the network duplicate and forward the information based on the distribution of the group members. Finally, the information is correctly delivered to Host B, Host D, and Host E.

To summarize, multicast has the following advantages:

- **Advantages over unicast**—Multicast data is replicated and distributed until it flows to the farthest-possible node from the source. The increase of receiver hosts will not remarkably increase the load of the source or the usage of network resources.
- **Advantages over broadcast**—Multicast data is sent only to the receivers that need it. This saves network bandwidth and enhances network security. In addition, multicast data is not confined to the same subnet.

Multicast features

- A multicast group is a multicast receiver set identified by an IP multicast address. Hosts must join a multicast group to become members of the multicast group before they receive the multicast data addressed to that multicast group. Typically, a multicast source does not need to join a multicast group.
- A multicast source is an information sender. It can send data to multiple multicast groups at the same time. Multiple multicast sources can send data to the same multicast group at the same time.
- The group memberships are dynamic. Hosts can join or leave multicast groups at any time. Multicast groups are not subject to geographic restrictions.
- Multicast routers or Layer 3 multicast devices are routers or Layer 3 switches that support Layer 3 multicast. They provide multicast routing and manage multicast group memberships on stub subnets with attached group members. A multicast router itself can be a multicast group member.

For a better understanding of the multicast concept, you can compare multicast transmission to the transmission of TV programs.

Table 1 Comparing TV program transmission and multicast transmission

TV program transmission	Multicast transmission
A TV station transmits a TV program through a channel.	A multicast source sends multicast data to a multicast group.
A user tunes the TV set to the channel.	A receiver joins the multicast group.
The user starts to watch the TV program transmitted by the TV station on the channel.	The receiver starts to receive the multicast data sent by the source to the multicast group.
The user turns off the TV set or tunes to another channel.	The receiver leaves the multicast group or joins another group.

Common notations in multicast

The following notations are commonly used in multicast transmission:

- **(*, G)**—Rendezvous point tree (RPT), or a multicast packet that any multicast source sends to multicast group G. The asterisk (*) represents any multicast source, and "G" represents a specific multicast group.
- **(S, G)**—Shortest path tree (SPT), or a multicast packet that multicast source "S" sends to multicast group "G." "S" represents a specific multicast source, and "G" represents a specific multicast group.

For more information about the concepts RPT and SPT, see "Configuring PIM" and "Configuring IPv6 PIM."

Multicast benefits and applications

Multicast benefits

- **Enhanced efficiency**—Reduces the processor load of information source servers and network devices.
- **Optimal performance**—Reduces redundant traffic.
- **Distributed application**—Enables point-to-multipoint applications at the price of minimum network resources.

Multicast applications

- Multimedia and streaming applications, such as Web TV, Web radio, and real-time video/audio conferencing
- Communication for training and cooperative operations, such as distance learning and telemedicine
- Data warehouse and financial applications (stock quotes)
- Any other point-to-multipoint application for data distribution

Multicast models

Based on how the receivers treat the multicast sources, the multicast models include any-source multicast (ASM), source-filtered multicast (SFM), and source-specific multicast (SSM).

ASM model

In the ASM model, any multicast sources can send information to a multicast group. Receivers can join a multicast group and get multicast information addressed to that multicast group from any multicast sources. In this model, receivers do not know the positions of the multicast sources in advance.

SFM model

The SFM model is derived from the ASM model. To a multicast source, the two models appear to have the same multicast membership architecture.

The SFM model functionally extends the ASM model. The upper-layer software checks the source address of received multicast packets and permits or denies multicast traffic from specific sources. The receivers obtain the multicast data from only part of the multicast sources. To a receiver, multicast sources are not all valid, but are filtered.

SSM model

The SSM model provides a transmission service that enables multicast receivers to specify the multicast sources in which they are interested.

In the SSM model, receivers have already determined the locations of the multicast sources. This is the main difference between the SSM model and the ASM model. In addition, the SSM model uses a different multicast address range than the ASM/SFM model. Dedicated multicast forwarding paths are established between receivers and the specified multicast sources.

IP multicast architecture

IP multicast addresses the following issues:

- Where should the multicast source transmit information to? (Multicast addressing.)
- What receivers exist on the network? (Host registration.)
- Where is the multicast source that will provide data to the receivers? (Multicast source discovery.)
- How is the information transmitted to the receivers? (Multicast routing.)

IP multicast is an end-to-end service. The multicast architecture involves the following parts:

- **Addressing mechanism**—A multicast source sends information to a group of receivers through a multicast address.
- **Host registration**—Receiver hosts can join and leave multicast groups dynamically. This mechanism is the basis for management of group memberships.
- **Multicast routing**—A multicast distribution tree (a forwarding path tree for multicast data on the network) is constructed for delivering multicast data from a multicast source to receivers.
- **Multicast applications**—A software system that supports multicast applications, such as video conferencing, must be installed on multicast sources and receiver hosts. The TCP/IP stack must support reception and transmission of multicast data.

Multicast addresses

IP multicast addresses

- IPv4 multicast addresses:

  IANA assigned the Class D address block (224.0.0.0 to 239.255.255.255) to IPv4 multicast.

  Table 2 Class D IP address blocks and description

Address block	Description
224.0.0.0 to 224.0.0.255	Reserved permanent group addresses. The IP address 224.0.0.0 is reserved. Other IP addresses can be used by routing protocols and for topology searching, protocol maintenance, and so on. Table 3 lists common permanent

Address block	Description
	group addresses. A packet destined for an address in this block will not be forwarded beyond the local subnet regardless of the TTL value in the IP header.
224.0.1.0 to 238.255.255.255	Globally scoped group addresses. This block includes the following types of designated group addresses: • **232.0.0.0/8**—SSM group addresses. • **233.0.0.0/8**—Glop group addresses.
239.0.0.0 to 239.255.255.255	Administratively scoped multicast addresses. These addresses are considered locally unique rather than globally unique. You can reuse them in domains administered by different organizations without causing conflicts. For more information, see RFC 2365.

NOTE:

Glop is a mechanism for assigning multicast addresses between different ASs. By filling an AS number into the middle two bytes of 233.0.0.0, you get 255 multicast addresses for that AS. For more information, see RFC 2770.

Table 3 Common permanent multicast group addresses

Address	Description
224.0.0.1	All systems on this subnet, including hosts and routers.
224.0.0.2	All multicast routers on this subnet.
224.0.0.3	Unassigned.
224.0.0.4	DVMRP routers.
224.0.0.5	OSPF routers.
224.0.0.6	OSPF designated routers and backup designated routers.
224.0.0.7	Shared Tree (ST) routers.
224.0.0.8	ST hosts.
224.0.0.9	RIPv2 routers.
224.0.0.11	Mobile agents.
224.0.0.12	DHCP server/relay agent.
224.0.0.13	All Protocol Independent Multicast (PIM) routers.
224.0.0.14	RSVP encapsulation.
224.0.0.15	All Core-Based Tree (CBT) routers.
224.0.0.16	Designated SBM.
224.0.0.17	All SBMs.
224.0.0.18	VRRP.

- IPv6 multicast addresses:

Figure 4 IPv6 multicast format

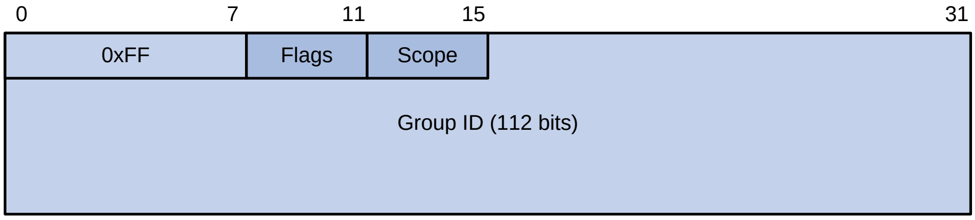

The following describes the fields of an IPv6 multicast address:

- **0xFF**—The most significant eight bits are 11111111.
- **Flags**—The Flags field contains four bits.

  Figure 5 Flags field format

  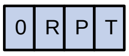

  Table 4 Flags field description

Bit	Description
0	Reserved, set to 0.
R	• When set to 0, this address is an IPv6 multicast address without an embedded RP address. • When set to 1, this address is an IPv6 multicast address with an embedded RP address. (The P and T bits must also be set to 1.)
P	• When set to 0, this address is an IPv6 multicast address not based on a unicast prefix. • When set to 1, this address is an IPv6 multicast address based on a unicast prefix. (The T bit must also be set to 1.)
T	• When set to 0, this address is an IPv6 multicast address permanently-assigned by IANA. • When set to 1, this address is a transient or dynamically assigned IPv6 multicast address.

- **Scope**—The Scope field contains four bits, which represent the scope of the IPv6 internetwork for which the multicast traffic is intended.

  Table 5 Values of the Scope field

Value	Meaning
0, F	Reserved.
1	Interface-local scope.
2	Link-local scope.
3	Subnet-local scope.
4	Admin-local scope.
5	Site-local scope.
6, 7, 9 through D	Unassigned.
8	Organization-local scope.

Value	Meaning
E	Global scope.

- **Group ID**—The Group ID field contains 112 bits. It uniquely identifies an IPv6 multicast group in the scope that the Scope field defines.

Ethernet multicast MAC addresses

- IPv4 multicast MAC addresses:

  As defined by IANA, the most significant 24 bits of an IPv4 multicast MAC address are 0x01005E. Bit 25 is 0, and the other 23 bits are the least significant 23 bits of an IPv4 multicast address.

  Figure 6 IPv4-to-MAC address mapping

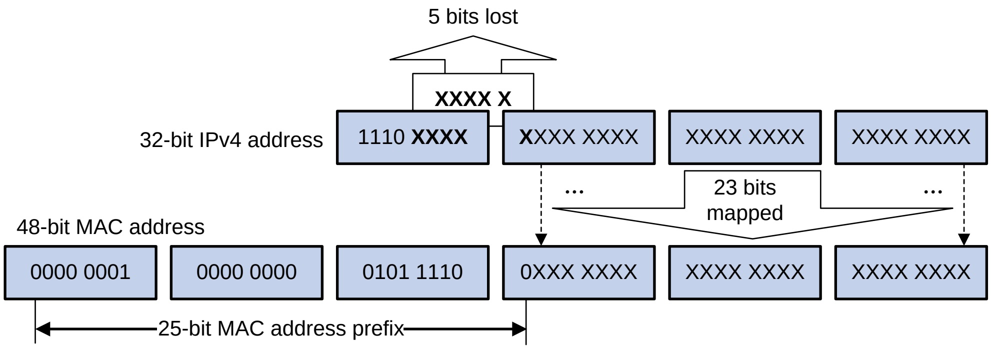

  The most significant four bits of an IPv4 multicast address are fixed at 1110. In an IPv4-to-MAC address mapping, five bits of the IPv4 multicast address are lost. As a result, 32 IPv4 multicast addresses are mapped to the same IPv4 multicast MAC address. A device might receive unwanted multicast data at Layer 2 processing, which needs to be filtered by the upper layer.

- IPv6 multicast MAC addresses:

  As defined by IANA, the most significant 16 bits of an IPv6 multicast MAC address are 0x3333. The least significant 32 bits are mapped from the least significant 32 bits of an IPv6 multicast address. Therefore, the problem of duplicate IPv6-to-MAC address mapping also arises like IPv4-to-MAC address mapping.

  Figure 7 IPv6-to-MAC address mapping

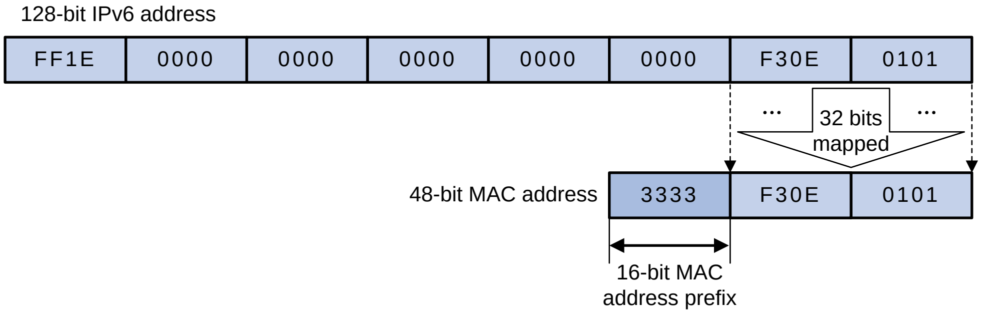

IMPORTANT:

Because of the duplicate mapping from multicast IP address to multicast MAC address, the device might inadvertently send multicast protocol packets as multicast data in Layer 2 forwarding. To avoid this, do not use the IP multicast addresses that are mapped to multicast MAC addresses 0100-5E00-00xx and 3333-0000-00xx (where "x" represents any hexadecimal number from 0 to F).

Multicast protocols

Multicast protocols include the following categories:

- Layer 3 and Layer 2 multicast protocols:
  - Layer 3 multicast refers to IP multicast operating at the network layer.

    Layer 3 multicast protocols—IGMP, MLD, PIM, IPv6 PIM, MSDP, MBGP, and IPv6 MBGP.
  - Layer 2 multicast refers to IP multicast operating at the data link layer.

    Layer 2 multicast protocols—IGMP snooping, MLD snooping, PIM snooping, IPv6 PIM snooping, multicast VLAN, and IPv6 multicast VLAN.
- IPv4 and IPv6 multicast protocols:
  - **For IPv4 networks**—IGMP snooping, PIM snooping, multicast VLAN, IGMP, PIM, MSDP, and MBGP.
  - **For IPv6 networks**—MLD snooping, IPv6 PIM snooping, IPv6 multicast VLAN, MLD, IPv6 PIM, and IPv6 MBGP.

This section provides only general descriptions about applications and functions of the Layer 2 and Layer 3 multicast protocols in a network. For more information about these protocols, see the related chapters.

Layer 3 multicast protocols

In Figure 8, Layer 3 multicast protocols include multicast group management protocols and multicast routing protocols.

Figure 8 Positions of Layer 3 multicast protocols

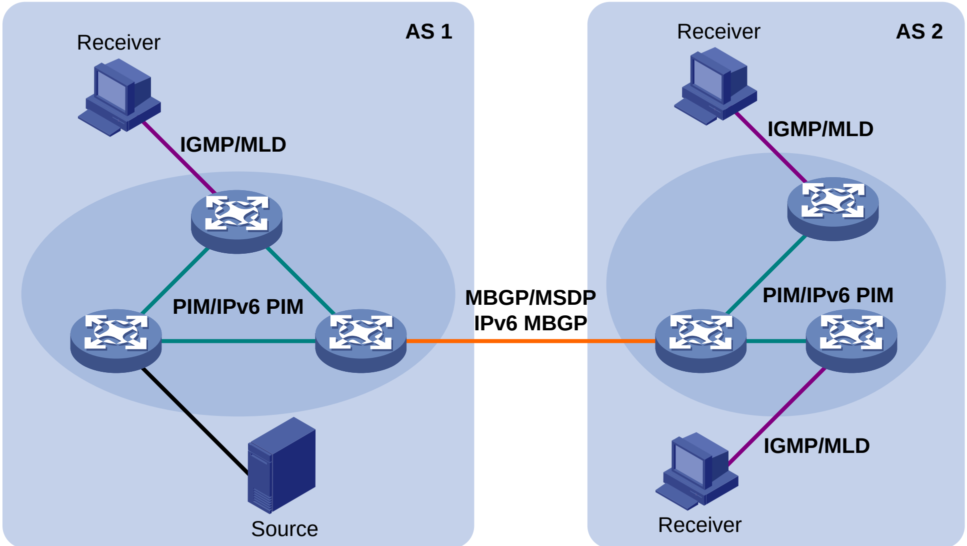

- Multicast group management protocols:

  Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) protocol are multicast group management protocols. Typically, they run between hosts and Layer 3 multicast devices that directly connect to the hosts to establish and maintain multicast group memberships.
- Multicast routing protocols:

  A multicast routing protocol runs on Layer 3 multicast devices to establish and maintain multicast routes and correctly and efficiently forward multicast packets. Multicast routes constitute loop-free data transmission paths (also known as multicast distribution trees) from a data source to multiple receivers.

In the ASM model, multicast routes include intra-domain routes and inter-domain routes.

- An intra-domain multicast routing protocol discovers multicast sources and builds multicast distribution trees within an AS to deliver multicast data to receivers. Among a variety of mature intra-domain multicast routing protocols, PIM is most widely used. Based on the forwarding mechanism, PIM has dense mode (often referred to as PIM-DM) and sparse mode (often referred to as PIM-SM).
- An inter-domain multicast routing protocol is used for delivering multicast information between two ASs. So far, mature solutions include Multicast Source Discovery Protocol (MSDP) and MBGP. MSDP propagates multicast source information among different ASs. MBGP is an extension of the MP-BGP for exchanging multicast routing information among different ASs.

For the SSM model, multicast routes are not divided into intra-domain routes and inter-domain routes. Because receivers know the positions of the multicast sources, channels established through PIM-SM are sufficient for the transport of multicast information.

Layer 2 multicast protocols

In Figure 9, Layer 2 multicast protocols include IGMP snooping, MLD snooping, PIM snooping, IPv6 PIM snooping, multicast VLAN, and IPv6 multicast VLAN.

Figure 9 Positions of Layer 2 multicast protocols

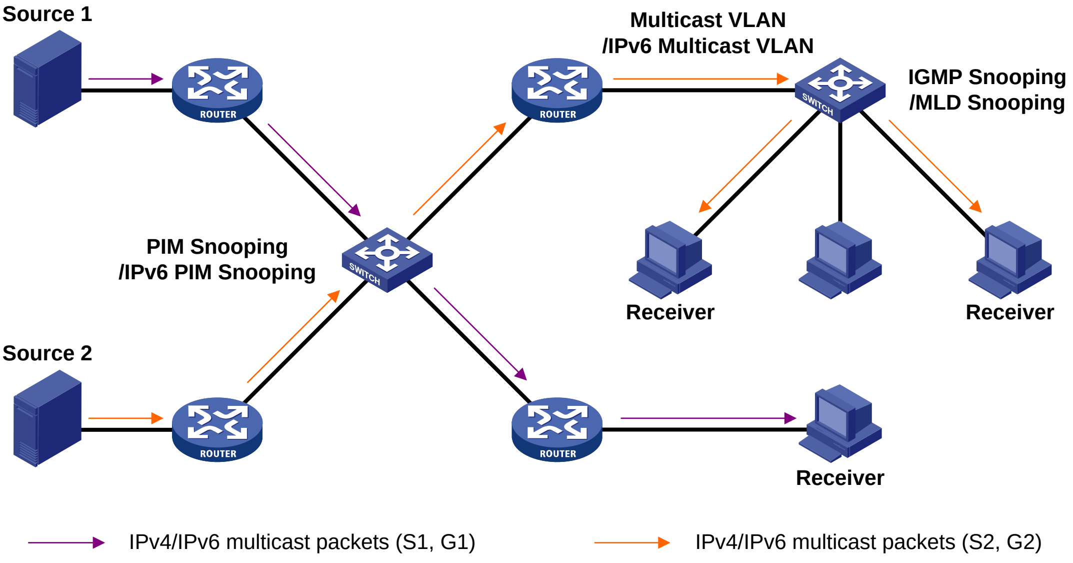

- IGMP snooping and MLD snooping:

  IGMP snooping and MLD snooping are multicast constraining mechanisms that run on Layer 2 devices. They manage and control multicast groups by monitoring and analyzing IGMP or MLD messages exchanged between the hosts and Layer 3 multicast devices. This effectively controls the flooding of multicast data in Layer 2 networks.
- PIM snooping and IPv6 PIM snooping:

  PIM snooping and IPv6 PIM snooping run on Layer 2 devices. They work with IGMP snooping or MLD snooping to analyze received PIM messages. Then, they add the ports that are interested in specific multicast data to a PIM snooping routing entry or IPv6 PIM snooping routing entry. In this way, multicast data can be forwarded to only the ports that are interested in the data.
- Multicast VLAN and IPv6 multicast VLAN:

  Multicast VLAN or IPv6 multicast VLAN runs on a Layer 2 device in a multicast network where multicast receivers for the same group exist in different VLANs. With these protocols, the Layer 3 multicast device sends only one copy of multicast to the multicast VLAN or IPv6 multicast

VLAN on the Layer 2 device. This method avoids waste of network bandwidth and extra burden on the Layer 3 device.

Multicast packet forwarding mechanism

In a multicast model, receiver hosts of a multicast group are usually located at different areas on the network. They are identified by the same multicast group address. To deliver multicast packets to these receivers, a multicast source encapsulates the multicast data in an IP packet with the multicast group address as the destination address. Multicast routers on the forwarding paths forward multicast packets that an incoming interface receives through multiple outgoing interfaces. Compared to a unicast model, a multicast model is more complex in the following aspects:

- To ensure multicast packet transmission on the network, different routing tables are used to guide multicast forwarding. These routing tables include unicast routing tables, routing tables for multicast (for example, the MBGP routing table), and static multicast routing tables.
- To process the same multicast information from different peers received on different interfaces, the multicast device performs an RPF check on each multicast packet. The RPF check result determines whether the packet will be forwarded or discarded. The RPF check mechanism is the basis for most multicast routing protocols to implement multicast forwarding.

  For more information about the RPF mechanism, see "Configuring multicast routing and forwarding" and "Configuring IPv6 multicast routing and forwarding."

Multicast support for VPNs

Multicast support for VPNs refers to multicast applied in VPNs.

Introduction to VPN instances

VPNs are isolated from one another and from the public network. As shown in Figure 10, VPN A and VPN B separately access the public network through PE devices.

Figure 10 VPN networking diagram

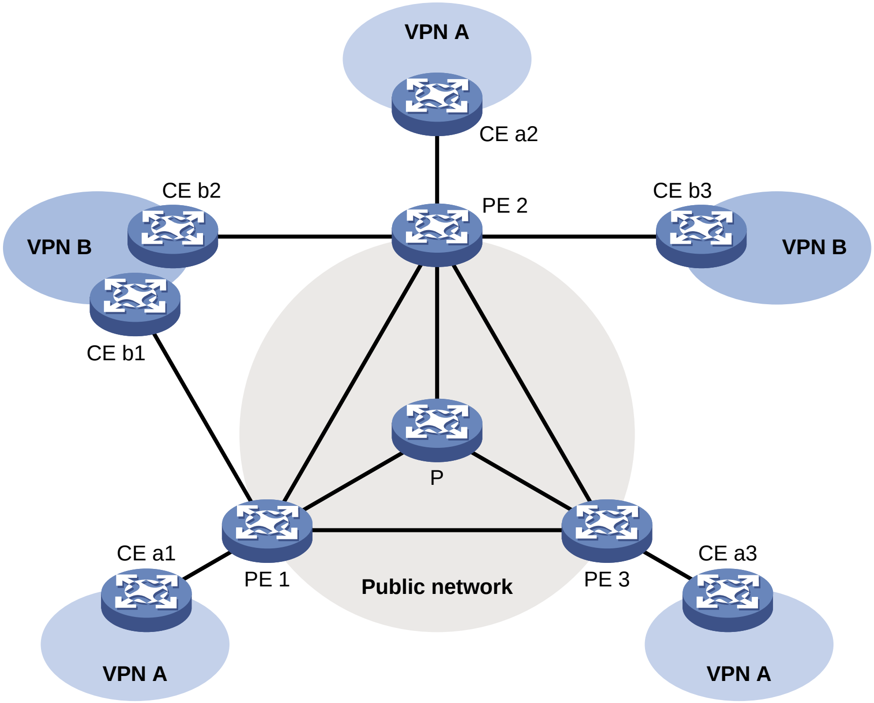

- The P device belongs to the public network. The CE devices belong to their respective VPNs. Each CE device serves its own VPN and maintains only one set of forwarding mechanisms.
- The PE devices connect to the public network and the VPNs. Each PE device must strictly distinguish the information for different networks, and maintain a separate forwarding mechanism for each network. On a PE device, a set of software and hardware that serve the same network forms an instance. Multiple instances can exist on the same PE device, and an instance can reside on different PE devices. On a PE device, the instance for the public network is called the public network instance, and those for VPNs are called VPN instances.

Multicast application in VPNs

A PE device that supports multicast for VPNs performs the following operations:

- Maintains an independent set of multicast forwarding mechanisms for each VPN, including the multicast protocols, PIM neighbor information, and multicast routing table. In a VPN, the device forwards multicast data based on the forwarding table or routing table for that VPN.
- Implements the isolation between different VPNs.
- Implements information exchange and data conversion between the public network and VPN instances.

For example, as shown in Figure 10, a multicast source in VPN A sends multicast data to a multicast group. Only receivers that belong to both the multicast group and VPN A can receive the multicast data. The multicast data is multicast both in VPN A and on the public network.