Multicast Technology White Paper

Contents

Overview

Technical background

Traditional IP communication includes unicast and broadcast. Unicast implements point-to-point communication between the source host and the destination host, and broadcast implements point-to-multipoint communication between the source host and all other hosts in the same network segment. With broadcast, the information is delivered to all hosts, rather than some specific hosts that need the information, resulting in waste of network bandwidth. In addition, broadcasts are confined only to the local subnet. With unicast, as a separate copy of information is sent to each host, the duplicate IP packets not only use a tremendous amount of network bandwidth but also add to the burden of the source host. Therefore, the conventional unicast and broadcast technologies cannot effectively address the issue of point-to-multipoint data transmission.

Multicast provides a best-effort service to deliver data packets to a specific set of receiver hosts, known as a multicast group, on the network. With multicast, the source host, known as a multicast source, sends only one copy of data packets destined for a multicast group address, and each receiver host of the multicast group can receive the data packets. Only the hosts that have joined the multicast group can receive the traffic addressed to the multicast group, while hosts not in the multicast group cannot receive the traffic.

Benefits

Compared with unicast and broadcast, the multicast technique effectively addresses the issue of point-to-multipoint data transmission. By allowing high-efficiency point-to-multipoint data transmission over an IP network, multicast greatly saves network bandwidth and reduces network load. More importantly, multicast allows convenient deployments of new value-added services in Internet-based information service areas, such as live Webcasting, Web TV, distance learning, telemedicine, Web radio, and real-time videoconferencing.

Multicast implementation

Multicast implementation needs to resolve the following issues:

- Multicast addressing mechanism: As a multicast source sends information to a certain group of receivers, a multicast addressing mechanism is needed to identify multicast groups.
- Group membership management: As a receiver host needs to join a multicast group before receiving the traffic destined for that group, a group membership management mechanism is needed to allow receiver hosts to join or leave a multicast group dynamically.
- Multicast packet forwarding: The process a multicast stream is forwarded and delivered to the receiver hosts over the network.
- Multicast routing protocol: Multicast routing protocols for constituting multicast forwarding trees.

Multicast addressing mechanism

IP multicast addresses

An IP multicast address identifies a specific IP multicast group. IANA has assigned the Class D address space (224.0.0.0 to 239.255.255.255) for multicast.

Figure 1 IP multicast address format

As shown in Figure 1, the high-order four bits of a multicast address are 1110. Figure 2 shows the specific address blocks and usages.

Figure 2 IP multicast address blocks

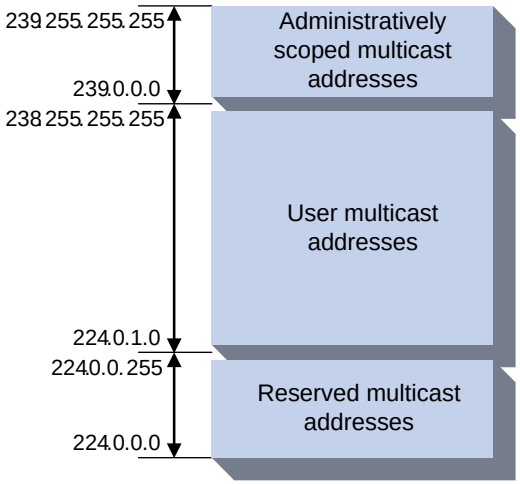

- **224.0.0.0 to 224.0.0.255**—Reserved permanent group addresses. The IP address 224.0.0.0 is reserved and other IP address can be used by routing protocols and for topology searching, protocol maintenance, and so on. Addresses in this range identify local subnets, namely, a packet destined for an address in this block will not be forwarded beyond the local subnet regardless of the Time to Live (TTL) value.
- **224.0.1.0 to 238.255.255.255**—User-available, globally scoped group addresses, among which 232.0.0.0/8 are SSM group addresses while the others are ASM group addresses. For more information about ASM and SSM, see "Multicast models."
- **239.0.0.0 to 239.255.255.255**—Administratively scoped multicast addresses, which are considered to be locally rather than globally unique ASM group addresses. In other words, these addresses can be reused in domains administered by different organizations without causing conflicts.

NOTE:

Some addresses in the 224.0.1.0/24 segment have also been reserved by IANA for particular multicast applications. For example, 224.0.1.1 has been reserved for the Network Time Protocol (NTP).

Ethernet multicast MAC addresses

An Ethernet multicast MAC address identifies receivers that belong to the same multicast group at the data link layer.

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 3 IPv4-to-MAC address mapping

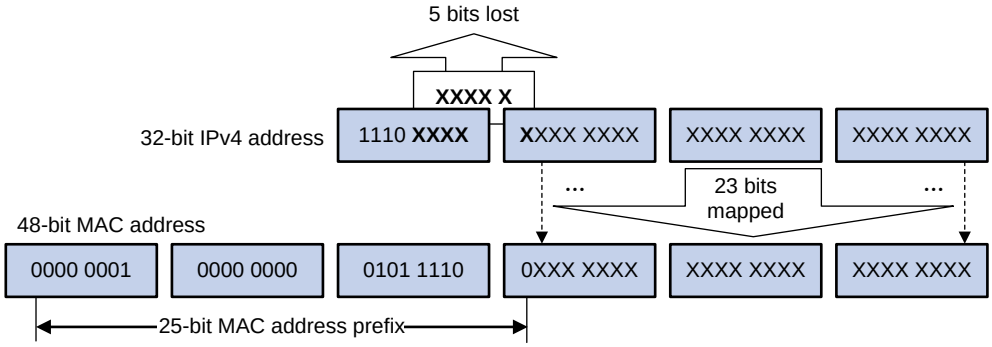

The most significant four bits of an IPv4 multicast address are fixed at 1110. In an IPv4-to-MAC address mapping, five bits of the IPv4 multicast address are lost. As a result, 32 IPv4 multicast addresses are mapped to the same IPv4 multicast MAC address. A device might receive unwanted multicast data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

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.

Group membership management

Group membership management means the establishment and maintenance of multicast group memberships on a multicast router or switch for the subnets directly connected to it, namely, the management of multicast group members attached to the interfaces or ports of the multicast device.

The Internet Group Management Protocol (IGMP) runs between IP hosts and the immediately connected router. IGMP functions in both directions:

- A host informs the router of its interest in specific multicast groups through IGMP.
- A multicast router uses IGMP to periodically query group memberships on the local subnet to collect and maintain the group memberships.

Through IGMP, a router learns whether a specific multicast group has active members on the local subnet, rather than the interrelationships between a multicast group and a host.

So far, there are three IGMP versions:

- IGMPv1 (documented in RFC 1112) defines the group member query and report processes;
- IGMPv2 (documented in RFC 2236) adds a leave-group mechanism based on IGMPv1;
- A major enhancement in IGMPv3 (documented in RFC 3376) is that it allows hosts to specify a list of sources they expect or do not expect multicast data from. The Source-Specific Multicast (SSM) model needs the support of IGMPv3.

This section mainly describes how IGMPv2 works.

Figure 4 How IGMPv2 works

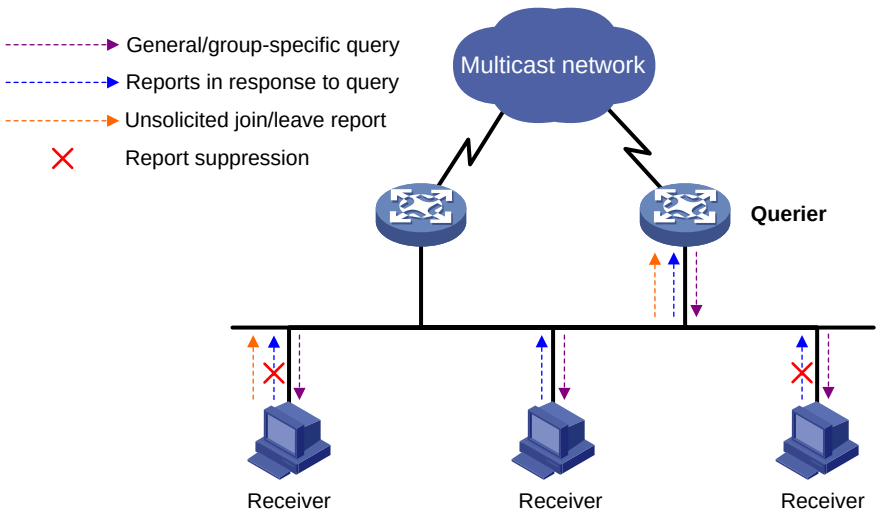

As shown in Figure 4, when multiple IGMP routers are connected to a subnet, a unique querier needs to be chosen from these routers through the querier election mechanism provided by IGMPv2. A querier periodically sends general query messages (often referred to as general queries) to learn the group memberships and a host responds with an IGMP report. A router also responds to the queries as a receiver host if it has joined a multicast group.

A host sends an unsolicited IGMP report when it needs to join a multicast group, without having to wait for an IGMP query. When a host leaves a multicast group, it sends a leave group message (often referred to as leave message). Upon receiving the leave message, the router sends group-

specific queries to that group to determine whether the group still has any active members on the local subnet.

After a member in a multicast group sends an IGMP report to join the group, the IGMP device already knows that the group has a minimum of one member. Other members do not need to report their memberships. This mechanism, known as the host IGMP report suppression, helps reduce traffic on the local subnet.

Through the mechanism discussed above, a router creates a table that records the multicast group members on the subnet attached to each of its interfaces. When receiving multicast traffic destined for multicast group G, the router forwards the traffic only to the interfaces with members of multicast group G. The forwarding of multicast traffic between multicast routers is not implemented by IGMP but by the multicast routing protocols.

Multicast packet forwarding

Multicast forwarding tree

Multicast packets travel along tree-shaped forwarding paths known as multicast forwarding trees over the network to the receivers. Multicast forwarding trees fall into two types: source tree and shared tree.

Source tree

Rooted at the multicast source, a source tree is a forwarding tree with the shortest path from the multicast source to the receivers; therefore, it is also called a shortest path tree (SPT). An SPT needs to be constructed per source per group.

As the shortest forwarding path between a multicast source and the receivers, the source tree minimizes the end-to-end transmission latency. However, this does come at a price. As the router must maintain the routing information for each multicast source, a great deal of system resource is used and the routing table is very large.

Shared tree

Rooted at a router called a rendezvous point (RP), a shared tree is a forwarding tree with the shortest path from the RP to each receiver. It is also called a rendezvous point tree (RPT). There is only one RPT per multicast group on the network. All multicast sources and receivers use the RPT tree for multicast data transmission and reception. The multicast sources send data to the RP and the RP forwards the data down the RPT to all the receivers.

The main advantage of an RPT is that it allows a router to maintain a small number of routing entries. However, as the multicast traffic from a multicast source must pass through the RP before it reaches the receivers, this forwarding tree is not the shortest path from the source to the receivers, and the RP must be highly reliable and powerful.

Multicast packet forwarding mechanism

Upon receiving a multicast packet, a router searches its multicast forwarding table according to the destination address and then forwards the packet accordingly. Forwarding a multicast packet is more complex than forwarding a unicast packet. In unicast, a router does not care about the source address; it cares only about the destination address of the packet, based on which the router determines the interface to forward the packet to. In multicast, multicast traffic is destined for to a group of receivers identified by a logical address known as multicast address. Upon receiving a multicast packet, a router checks whether the packet has arrived to the correct incoming interface, namely whether the incoming interface leads to the multicast source, based on the source address before forwarding the packet out the outgoing interface. This process is known as the reverse path forwarding (RPF) check.

The RPF check process is as follows:

1. The router searches the unicast routing table for the RPF interface.
   - If an SPT is used, the RPF interface is the outgoing interface of the unicast route to the multicast source.
   - If an RPT is used, the RPF interface is the outgoing interface of the unicast route to the RP.
2. If the packet is received on the RPF interface, it passes the RPF check and then forwarded to downstream node; otherwise, the packet is discarded.

The RPF check not only ensures multicast data forwarding along the correct forwarding path but also helps avoid loops.

Multicast routing protocols

Similar to unicast protocols, multicast routing protocols fall into intra-domain and inter-domain protocols:

- Based on the group memberships maintained by IGMP, an intra-domain multicast routing protocol builds multicast distribution trees according to certain multicast routing algorithms and creates multicast routing entries on multicast routers, which forward multicast traffic as per the routing entries.
- Based on the inter-domain multicast routing policy configured in the network, an inter-domain multicast routing protocol propagates multicast source information and exchanges multicast routing information among autonomous systems (ASs), thus ensuing multicast forwarding among different domains.

Intra-domain multicast routing protocols

Among a variety of intra-domain multicast routing protocols, Protocol Independent Multicast (PIM) is a popular one. Based on the forwarding mechanism, PIM includes the following modes:

Based on the implementation mechanism, PIM includes the following modes:

- Protocol Independent Multicast–Dense Mode (PIM-DM).
- Protocol Independent Multicast–Sparse Mode (PIM-SM).
- Bidirectional Protocol Independent Multicast (BIDIR-PIM).
- Protocol Independent Multicast Source-Specific Multicast (PIM-SSM).

PIM-DM

In a PIM-DM domain, each PIM interface on a device periodically multicasts PIM hello messages to all other PIM devices (identified by the address 224.0.0.13) on the local subnet. Through the exchanging of hello messages, all PIM devices on the subnet determine their PIM neighbors, maintain PIM neighboring relationship with other devices, and build and maintain SPTs. PIM devices also exchange hello messages to elect the DR.

NOTE:

PIM-DM does not require a DR. However, if IGMPv1 runs on any shared-media LAN in a PIM-DM domain, a DR must be elected to act as the IGMPv1 querier for the LAN.

As a dense mode multicast routing protocol, PIM-DM uses the push mode for multicast forwarding, and is suitable for small-sized networks with densely distributed multicast members. PIM-DM works as follows:

- PIM-DM assumes that at least one multicast group member exists on each subnet of a network, and therefore multicast data is flooded to all nodes on the network. Then, branches without receivers downstream are pruned from the forwarding tree, leaving only those branches with receivers. This flood and prune process takes place periodically.

- When a host attached to a pruned node joins the multicast group, the node sends a graft message toward the upstream node. Then the node resumes multicast traffic forwarding.

PIM-SM

In a PIM-SM domain, each PIM interface on a device periodically multicasts PIM hello messages to all other PIM devices (identified by the address 224.0.0.13) on the local subnet. Through the exchange of hello messages, all PIM devices determines their PIM neighbors and elect the DR. On a multi-access network, the DR sends join/prune messages toward the root of the multicast forwarding tree for the receiver host attached to it, or forwards multicast traffic from the directly connected multicast source onto the multicast distribution tree.

As a sparse mode multicast routing protocol, PIM-SM uses the pull mode for multicast forwarding, and is suitable for large- and medium-sized networks with sparsely and widely distributed multicast members. The basic implementation of PIM-SM is as follows:

- PIM-SM assumes that no hosts need to receive multicast data. In the PIM-SM mode, and delivers multicast data only to those hosts that have explicitly requested for the data. The core task for PIM-SM in multicast forwarding is to build and maintain rendezvous point trees (RPTs). An RPT is rooted at a router in the PIM domain as the common node, referred to as the rendezvous point (RP), through which the multicast data travels down the RPT to the receivers.
- When a receiver is interested in the multicast data addressed to a specific multicast group, the router connected to this receiver sends a join message to the RP for that multicast group. The path along which the message goes hop by hop to the RP forms a branch of the RPT.
- When a multicast source sends multicast data to a multicast group, the source-side DR first registers the multicast source with the RP by sending register messages to the RP by unicast. The arrival of a register message at the RP triggers the establishment of an SPT. Then, the multicast source sends subsequent multicast packets along the SPT to the RP. Upon reaching the RP, the multicast packet is duplicated and delivered to the receivers down the RPT.

PIM-SSM

The SSM model provides a solution for source-specific multicast.

If receivers join a multicast group in the SSM group range (232.0.0.0/8 reserved by IANA for SSM), multicast devices will use the PIM-SSM mode to build the multicast forwarding tree.

PIM-SSM mechanisms include neighbor discovery and DR election, which are the same as those in PIM-SM.

PIM-SSM also uses the pull mode for multicast forwarding. However, because receivers have located a multicast source, no RP or RPT is required. Multicast sources do not register with the RP, and the MSDP is not needed for discovering multicast sources in other PIM domains. The basic implementation of PIM-SSM is as follows:

- A receiver sends IGMPv3 reports to the DR for expecting to receive traffic of multicast source and group (S, G).
- The multicast device connected to the receiver discovers that the multicast group is in the SSM group range. Then, it sends a join message toward the multicast source. All multicast devices on the path to the multicast source create the (S, G) entry, which builds an SPT.
- The multicast source sends the multicast traffic of the multicast group. The multicast traffic is forwarded along the SPT and reaches the receiver.

BIDIR-PIM

BIDIR-PIM is suitable for many-to-many applications, where multiple receivers of a multicast group might be interested in the multicast data from multiple multicast sources. With PIM-DM or PIM-SM, each device along the SPT must create an (S, G) entry for each multicast source, consuming a lot of system resources. Derived from PIM-SM, BIDIR-PIM builds and maintains a bidirectional RPT, which is rooted at the RP and connects the multicast sources and the receivers. Along the bidirectional RPT, the multicast sources send multicast data to the RP, and the RP forwards the

data to the receivers. Each device along the bidirectional RPT needs to maintain only one (*, G) entry, saving system resources.

BIDIR-PIM mechanisms include neighbor discovery, RP discovery, DF election, and bidirectional RPT building. Neighbor discovery and RP discovery are the same as those of PIM-SM.

BIDIR-PIM uses a designated forwarder (DF) election mechanism to elect a unique DF for each RP on a subnet. Only the DFs can forward multicast data to the RP.

The process of building the bidirectional RPT is as follows:

1. When a multicast source sends multicast packets to the multicast group G, the DF in each subnet unconditionally forwards the packets to the RP.
2. The devices along the path from the source's directly connected device to the RP constitute an RPT branch. Each device on this branch adds to its forwarding table a (*, G) entry.

After a bidirectional RPT is built, the multicast sources send multicast traffic to the RP along the source-side RPT. Then, the RP forwards the traffic to the receivers along the receiver-side RPT.

Inter-domain multicast routing protocols

Inter-domain multicast routing protocols are used to propagate multicast information among different ASs. So far, mature solutions include:

- Multicast Border Gateway Protocol (MBGP) is used for exchanging multicast routing information between ASs.
- Multicast Source Discovery Protocol (MSDP) is used for advertising multicast source information between ASs.

MBGP

The exchange of routing information (reachability information) between different ASs is the first-line issue to address for inter-domain communications. Because different ASs may be operated by different service providers, unlike intra-domain routing information, inter-domain routing information needs to contain not only the distance information but also the service provider policies.

The multicast topology may be different from the unicast topology due to both physical and policy-related reasons. Some routers on the network may support only unicast, while some others, though multicast capable, may be configure not forward multicast packets. In order to construct inter-domain multicast forwarding trees, in addition to unicast routing information, the multicast topology information is also needed. In short, an inter-domain multicast routing protocol needs to meet the following requirements:

- Able to differentiate the unicast topology and the multicast topology.
- Having a set of stable methods for peering and policy control.

As the most popular inter-domain unicast routing protocol so far, the Border Gateway Protocol version 4 (BGP-4) already satisfies the latter requirement and is proven to be effective and stable. Therefore, a reasonable solution for inter-domain propagation of multicast routing information is to enhance and extend BGP-4 rather than to devise a set of entirely new protocols. Multiprotocol extensions for BGP are defined in RFC 2858. The extended BGP (known as MP-BGP or BGP-4+) can carry not only IPv4 unicast routing information but also the routing information for other network layer protocols (such as multicast and IPv6). The capability of carrying multicast routing information is only one of the functions of these extensions. The multicast extension of BGP is referred to as multicast BGP (MBGP).

With MBGP, both unicast routing information and multicast routing information can be exchanged in the same process, but are stored in different routing tables. MBGP is an enhanced version of BGP-4; therefore, all the common policies and configuration methods supported by BGP-4 can be applied to multicast.

MSDP

In the basic PIM-SM mode, a multicast source registers only with the RP in the local PIM-SM domain, and the multicast source information of a domain is confined within the domain. As a result, the RP is aware of the source information only within the local domain and multicast distribution trees are built only within the local domain to deliver multicast data from a local multicast source to local receivers. A service provider does not want to rely on other ASs' RP routers to forward multicast traffic to its own customers, but it does want to be able to get information from multicast sources, wherever the RPs for these sources are, and send the information to its customers.

The Multicast Source Discovery Protocol (MSDP) is an inter-domain multicast solution developed to discover multicast sources in other PIM-SM domains thus to allow interconnection among these domains. MSDP establish peer relationships between appropriate devices to interconnect RPs in PIM-SM domains. The MSDP peers send source active (SA) messages to exchange multicast source information.

In addition to inter-domain propagation of multicast source information, MSDP has a special application for PIM-SM: Anycast RP. Anycast RP refers to such an application that implements load balancing and redundancy backup between two or more RPs within a PIM-SM domain by configuring the same IP address for, and establishing MSDP peering relationships between, these RPs.

Layer 2 multicast protocols

As protocols designed for IP multicast at the network layer, IGMP and PIM maintain only the relationships between Layer 3 interfaces and IP multicast addresses. In most situations, however, multicast traffic inevitably goes through some Layer 2 switches. Without a mechanism to establish mappings between Layer 2 ports and multicast MAC addresses, multicast traffic will be flood to all ports of a switch, and this wastes a great deal of system resource. Layer 2 multicast protocols can resolve this issue.

IGMP snooping

IGMP snooping runs on Layer 2 devices. When an IGMP snooping switch snoops an IGMP report message a host sent to the IGMP querier, the switch creates a mapping between the port connected to the host and the multicast MAC address corresponding to the reported multicast group. Then, upon receiving multicast data for that group, the switch forwards the multicast data just to that port based on the created mapping.

PIM snooping

PIM snooping runs on Layer 2 devices. It works with IGMP snooping to analyze received PIM messages, and adds the ports that are interested in specific multicast data to a PIM snooping routing entry. In this way, the multicast data can be forwarded to only the ports that are interested in the data.

When the Layer 2 switch runs only IGMP snooping, it performs the following actions:

1. Maintains the router ports according to the received PIM hello messages that PIM routers send.
2. Floods all other types of received PIM messages except PIM hello messages in the VLAN.
3. Forwards all multicast data to all router ports in the VLAN.

   Each PIM router in the VLAN, whether interested in the multicast data or not, can receive all multicast data and all PIM messages except PIM hello messages.

When the Layer 2 switch runs both IGMP snooping and PIM snooping, it performs the following actions:

1. Examines whether a PIM router is interested in the multicast data addressed to a multicast group according to the received PIM messages that the router sends.
2. Adds only the ports that connect to the router and are interested in the data to a PIM snooping routing entry.
3. Forwards PIM messages and multicast data to only the routers that are interested in the data, which saves network bandwidth.

Multicast VLAN

Multicast VLAN is suitable for the multicast network where receivers of the same multicast groups scatter in multiple VLANs.

As shown in Figure 5, Host A, Host B, and Host C are in three different VLANs and the same multicast group. When Switch A (Layer 3 device) receives multicast data for that group, it sends three copies of the multicast data to Switch B (Layer 2 device). This occupies a large amount of bandwidth and increases the burden on the Layer 3 device.

Figure 5 Multicast transmission without the multicast VLAN feature

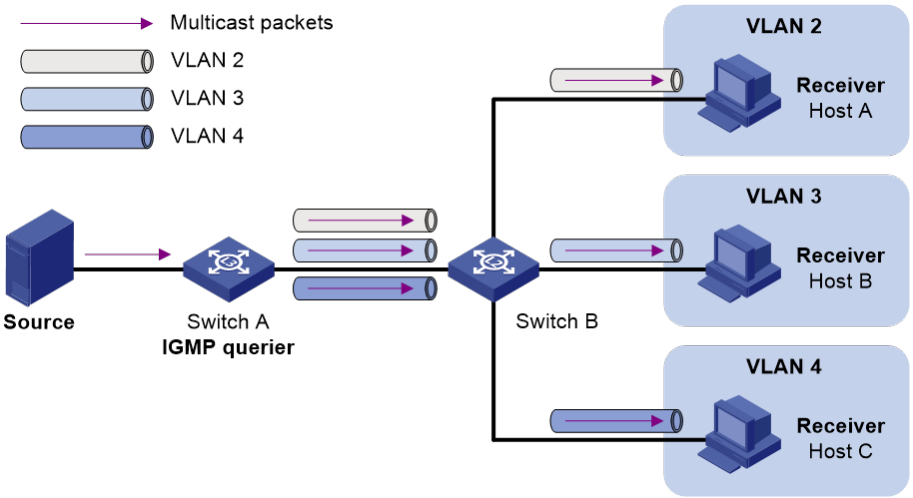

After a multicast VLAN is configured on Switch B, Switch A sends only one copy of the multicast data to the multicast VLAN on Switch B. This method saves network bandwidth and lessens the burden on the Layer 3 device.

Multicast VLANs include sub-VLAN-based multicast VLANs and port-based multicast VLANs.

Application scenarios

Intra-domain multicast

Currently, PIM-SM is the recognized standard protocol for implementing intra-domain multicast. In a single-AS network or a PIM-SM domain, multicast can be implemented with PIM-SM alone running in the network. In such a network, multiple RPs can be configured to implement anycast RP to enhance the RP reliability and achieve load-balancing and redundancy backup.

Figure 6 shows the network diagram for intra-domain multicast through PIM-SM.

Figure 6 Network diagram for intra-domain multicast

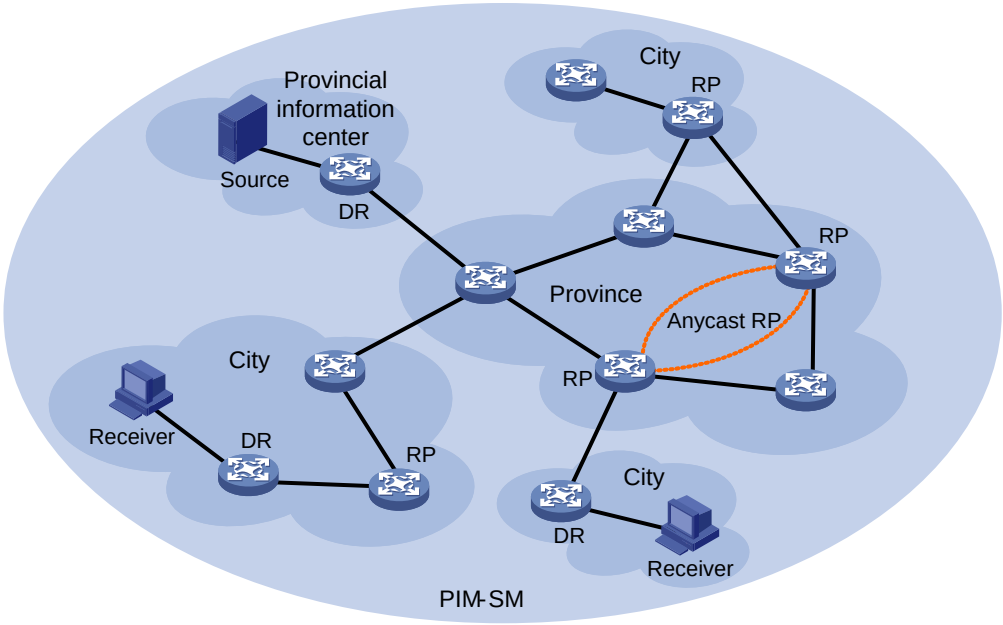

Inter-domain multicast

PIM-SM+MBGP+MSDP solution

Currently, a mature solution for implementing inter-domain multicast is the combined PIM-SM+MBGP+MSDP solution, which requires all ASs to support PIM-SM, MBGP, and MSDP simultaneously. As shown in Figure 7, PIM-SM runs in all ASs, and MBGP and MSDP run among these PIM-SM domains.

Figure 7 PIM-SM+MBGP+MSDP solution

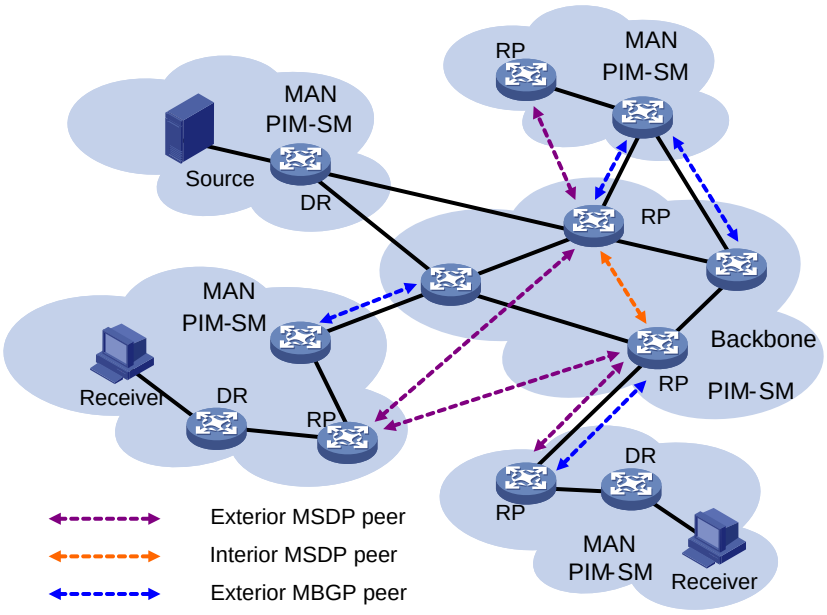

This solution is an extension of PIM-SM in an inter-domain environment. If the entire PIM-SM+MBGP+MSDP solution is considered as PIM-SM, the RPs in all the domains put together can be considered as the RP in the PIM-SM domain with two more processes:

1. Flooding of multicast source information within the RP set so that the multicast sources and receivers meet at the “RP”.
2. Exchange of multicast routing information between different domains to ensure the normal forwarding of multicast traffic between the domains. In these processes, the multicast topology information propagated by MBGP is leveraged to establish the reverse forwarding paths from the RPs in one AS to the peers in another AS.

In the PIM-SM+MBGP+MSDP solution, it is required that exterior MBGP peering relationships to be established between the routers at AS borders, exterior MSDP peering relationships to be established between RPs of different ASs, interior MBGP peering relationships to be established between routers in the same AS, and interior MSDP peering relationships to be established between RPs in the same AS, anycast RP to be configured, and PIM-SM to be enabled in all ASs. PIM-SM is responsible for collecting multicast and routing and multicast source information within each PIM-SM domain, MBGP for exchanging multicast topology information, and MSDP for exchanging multicast source information among the domains.

PIM-SM+tunneling (MBGP&MSDP) solution

As shown in Figure 8, to provide multicast service while the routers on the backbone network is not multicast capable or not multicast enabled, run PIM-SM on each metropolitan area network (MAN), set up a virtual network by establishing tunnels between the RPs on different MANs, and run PIM-SM, MBGP, and MSDP on this virtual network. The advantage of this solution is that the backbone network does not need to support PIM-SM, MBGP, or MSDP. In other words, multicast traffic is transparent to the backbone network. This avoids the influence on the device performance during multicast traffic transmission. However, this solution requires that PIM-SM, MBGP and MSDP tunnels must be simultaneously supported between RPs. The configuration and management are complex and the devices must be highly powerful.

Figure 8 PIM-SM+tunneling (MBGP&MSDP) solution

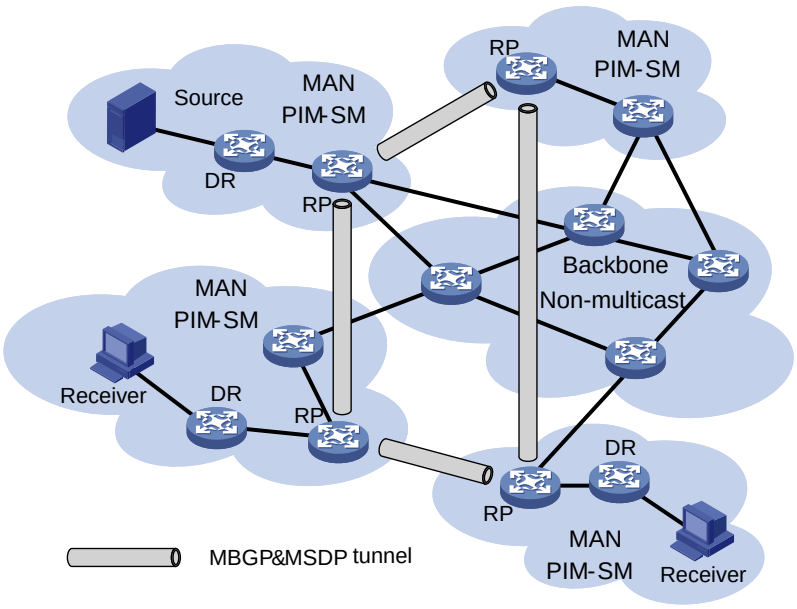

Traversal of multicast traffic across firewalls

As shown in Figure 9, if multicast-capable networks are separated by a multicast-incapable network (such as the Internet) or across firewalls configured with NAT or IPSec VPN, the routers across the Internet or firewalls cannot establish PIM neighboring relationships with each other and thus cannot implement multicast routing. In this scenario, a GRE tunnel can be configured to interconnect these multicast networks, thus allowing multicast implementation.

Figure 9 Network diagram for multicast traversal across firewalls