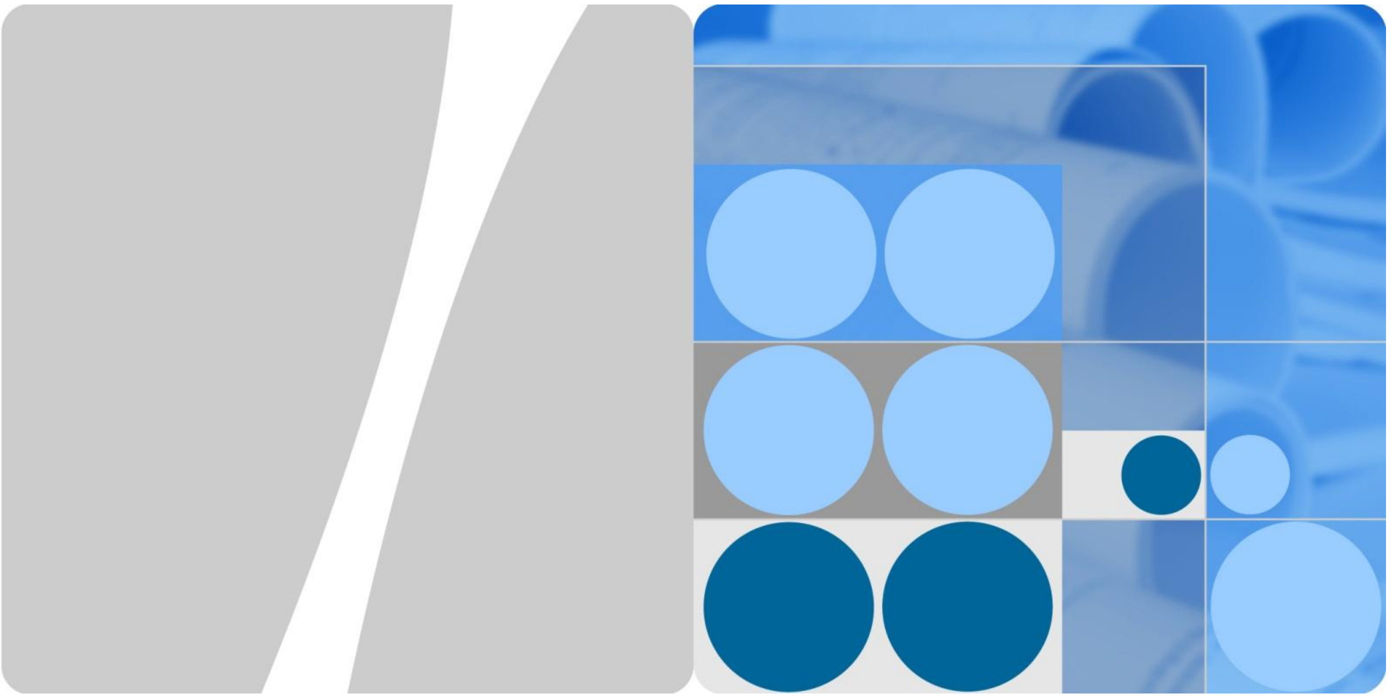

Multicast Technology White Paper

Issue	**2.0**
Date	**2011-10-31**

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions

Notice

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Huawei Technologies Co., Ltd.

Address: Huawei Industrial Base
Bantian, Longgang
Shenzhen 518129
People's Republic of China

Website: http://enterprise.huawei.com

About This Document

Keywords

Multicast, operations management, controllable multicast, IGMP, PIM-SM, PIM-DM, MBGP, MSDP

Abstract

IP multicast technology transmits data efficiently from one point to multiple points over an IP network. Multicast can effectively conserve network bandwidth and reduce network load. It is widely used in various applications, including real-time data transmission, multimedia conferencing, data copy, games, and simulation. This document describes the basic concepts of multicast, common multicast protocols, and multicast networking solutions. To meet multicast service requirements and solve the problems that occur in operations, this document puts forward an operable, manageable controllable multicast solution that covers source management, user management, and multicast security control.

List of Acronyms

Acronym	Full Name
IGMP	Internet Group Management Protocol
MBGP	Multiprotocol Border Gateway Protocol
MSDP	Multicast Source Discovery Protocol
PIM-DM	Protocol Independent Multicast-Dense Mode
PIM-SM	Protocol Independent Multicast-Sparse Mode

Contents

Figures

1 Multicast Overview

1.1 Background

In traditional IP communication, two modes of data transmission are used: unicast and broadcast. In unicast mode, a source IP host communicates with a single destination IP host. In broadcast mode, a source IP host communicates with all other IP hosts on a network. To send a message to multiple hosts on a network, the source host can either broadcast the message to all hosts or send the message to the hosts one by one in unicast mode. When broadcast is used, messages are sent to the hosts that do not need the messages, which wastes bandwidth. In addition, a broadcast storm may occur because of routing loops. When unicast is used, repeated transmission of IP packets wastes much bandwidth and burdens the server. Traditional unicast and broadcast modes cannot effectively address the problem of sending messages from one point to multi points.

IP multicast sends data packets to a subset of nodes on an IP network with best-effort delivery. The subset is called a multicast group. The basic concept of IP multicast is that the source host sends only one copy of data to the destination hosts over a network. The destination of the data is the address of a multicast group. All receivers in the multicast group receive the same data copy. In this way, all destination hosts in the multicast group can receive the data and the other hosts on the network cannot receive the data. Multicast groups are identified with Class D IP addresses in the range from 224.0.0.0 to 239.255.255.255.

1.2 Applications

IP multicast addresses the problem of sending packets from one point to multiple points. This technology efficiently transmits data from a single source to multiple points over an IP network, which saves substantial network bandwidth and reduces the network load. As a communications mode compared to unicast and broadcast, multicast has more benefits. Multicast can be used to provide new value-added services conveniently over networks, including live online broadcast, online TV, distance learning, remote medical care, online radio, real-time video conferencing, and other information services on the Internet.

First put forward in 1988, multicast has been developed over more than 20 years. Several international organizations have devoted a large amount of work to the technical research and deployment of multicast. With the rapid development of the Internet and continuous emergence of new services, multicast communication has matured. At present, some operators or ISPs have deployed end-to-end full-mesh multicast services. Major ISPs have run inter-domain multicast routing protocols to exchange multicast routes so that multicast peer relationships can be established. In the context of ever increasing multimedia services on IP networks, multicast has huge market potential.

2 Implementation

Multicast involves a variety of technologies from address allocation and membership management to multicast packet forwarding, route setup, and reliability. This chapter describes the overall structure of the multicast protocol system and key protocols and mechanisms from the perspectives of multicast address, membership management, multicast packet forwarding, and intra-domain and inter-domain multicast routing.

2.1 Structure of the Multicast Protocol System

Depending on the application scope, multicast protocols are classified as those between hosts and routers or those between routers. The protocols between hosts and routers are multicast membership management protocols that include the Internet Group Management Protocol (IGMP). Protocols between routers are further classified into intra-domain and inter-domain multicast routing protocols. Intra-domain multicast routing protocols include PIM-SM and PIM-DM. Inter-domain routing protocols include MBGP and MSDP. In addition, Layer 2 multicast functions such as IGMP snooping effectively suppress the broadcasting of multicast packets on a Layer 2 network.

Through IGMP and Layer 2 multicast protocols, memberships on directly connected network segments are set up on routers and switches. The memberships pertain to multicast group member devices and interfaces connected to the member devices. Based on the multicast group memberships maintained by IGMP, intra-domain routing protocols use certain multicast routing algorithms to construct a multicast distribution tree. A router maintains multicast routes and then forwards multicast packets based on these routes. Based on the inter-domain multicast routing policies configured on the network, inter-domain multicast routing protocols advertise information about multicast routes and sources among Autonomous Systems (ASs). As a result, multicast packets can be forwarded between different ASs.

2.2 Mechanism of Multicast Addresses

2.2.1 Multicast IP Address

A multicast IP address identifies an IP multicast group. The Internet Assigned Numbers Authority (IANA) assigns Class D IP addresses in the range 224.0.0.0 through 239.255.255.255 to multicast communication. As shown in Figure 2-1 (in binary), the first 4 bits of an IP address are always 1110.

Figure 2-1 Format of a multicast IP address

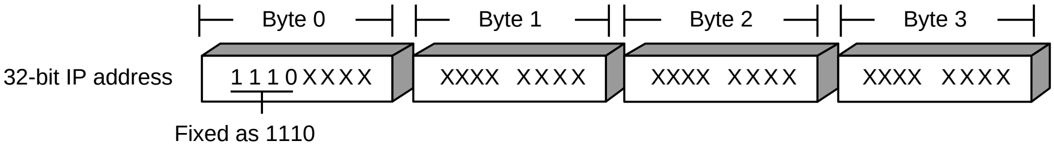

2.2.2 Classification of Multicast Addresses

Figure 2-2 shows how IP multicast addresses are classified.

Figure 2-2 Classification of multicast addresses

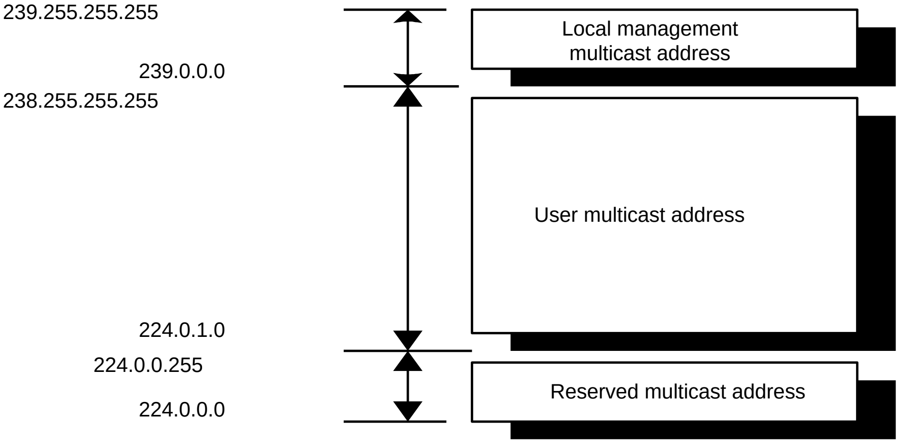

Addresses from 224.0.0.0 to 224.0.0.255 are reserved by the IANA. Address 224.0.0.0 is not allocated and the other addresses are used by routing protocols, topology discovery, and protocol maintenance. These addresses are locally valid and routers cannot forward any multicast packet with a destination address in this range, regardless of the time to live (TTL).

Addresses from 224.0.1.0 to 238.255.255.255 are used as user multicast addresses that are globally valid. 233/8 is a GLOP address. GLOP is a mechanism to assign multicast addresses between ASs. The AS number is directly inserted into the middle of a multicast address. Each AS is assigned 255 multicast addresses.

Addresses from 239.0.0.0 to 239.255.255.255 are administratively scoped addresses that remain valid only in certain locations.

When a multicast data packet is received by the IP layer, a router uses the destination address of the packet to search the multicast forwarding table and then forwards the packet.

2.2.3 Mapping from an IP Multicast Address to a MAC Address

The IANA assigns MAC addresses in the range 01:00:5E:00:00:00 through 01:00:5E:7F:FF:FF to multicast communication. The 28-bit multicast IP address needs to be mapped to the 23-bit MAC address. As shown in Figure 2-3, the low order 23 bits of a multicast address are mapped to the low order 23 bits of the MAC address.

Figure 2-3 Mapping from an IP multicast address to a MAC address

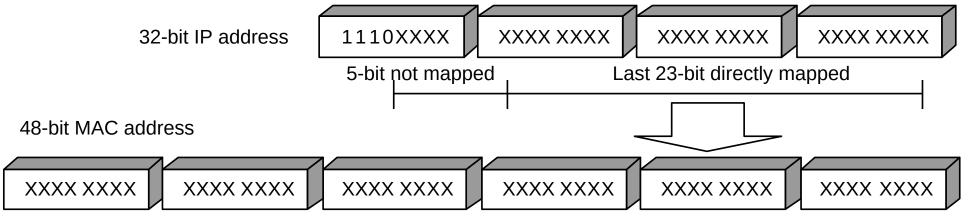

Of the last 28 bits of an IP multicast address, 23 bits are mapped to a MAC address. The 23 bits of the IP address and the 5 remaining bits plus the 4 high order bits, 1110, allow you to map 32 multicast IP addresses to one MAC address.

2.3 Layer 2 Multicast

IP multicast addresses are needed on the network layer to enable communication between a source and its receivers across the Internet. When the multicast source and the multicast group are deployed on the local physical network, link layer multicast must be enabled to guarantee normal transmission of multicast information. When the link layer runs the Ethernet protocol, hardware-based multicast uses multicast MAC addresses, and IP multicast addresses must be mapped to multicast MAC addresses. Generally, link layer multicast is called Layer 2 multicast.

2.3.1 Static Layer 2 Multicast

When multicast packets are transmitted on an Ethernet network, the packet destination is a group of unspecified members instead of a specific receiver. Multicast forwarding entries cannot be generated when multicast packets are forwarded to the link layer from the IP layer. As a result, the multicast packets are broadcast on the link layer. This wastes bandwidth, makes the accounting of user services inefficient, and poses a threat to information security.

Interfaces can be added to a multicast group statically or dynamically. Adding interfaces to the multicast group statically has the following advantages:

- Prevents attacks from protocol packets.
- Reduces the network delay by directly searching the multicast forwarding table to transmit multicast packets.
- Improves information security and implements the provision of paid services by preventing unregistered hosts from receiving multicast traffic.

When a switch is deployed between a router and user hosts and needs to have Layer 2 forwarding enabled, you can configure static Layer 2 multicast, that is, configure forwarding entries manually. In this manner, multicast data is always forwarded to the users who need the multicast data.

You can set the forwarding mode of multicast data so that the multicast flows can be forwarded based on IP addresses or MAC addresses. When multicast IP addresses are mapped to MAC addresses, a maximum of 32 IP addresses can be mapped to one MAC address. Therefore, it is recommended that multicast data be forwarded based on IP addresses. Otherwise, unregistered users may receive the multicast data.

2.3.2 Multicast Replication in VLANs

In traditional multicast forwarding mode, if hosts in different VLANs need to receive the multicast data from the same source through a router, the router needs to replicate and send multicast packets for each Virtual Local Area Network (VLAN). This wastes bandwidth and burdens routers.

After multicast VLAN replication is configured, you can configure different VLANs as user VLANs of one multicast VLAN, if hosts in different VLANs need to receive the multicast data from the same source. In this manner, routers only send multicast packets to the multicast VLAN, and do not need to replicate multicast packets for each VLAN.

2.3.3 IGMP Snooping/MLD Snooping

Internet Group Management Protocol (IGMP) snooping and Multicast Listener Discovery (MLD) snooping monitor the multicast protocol packets exchanged between a router and host to maintain information about the outbound interface of multicast packets. In this way, IGMP snooping manages and controls forwarding of multicast data packets to implement Layer 2 multicast.

As specified by the IGMP/MLD protocol, a host needs to send an IGMP/MLD Report message to a switch before joining a multicast group, and then the switch can send multicast packets to the host. IGMP/MLD messages, encapsulated in IP packets or IPv6 packets, are Layer 3 packets. A Layer 2 device, however, cannot process Layer 3 packets. Therefore, it does not detect that a host sent an IGMP/MLD message to join a multicast group. In addition, the switch cannot learn multicast MAC addresses by learning the source MAC addresses of the multicast packets (the source MAC addresses are not the multicast MAC addresses). In this case, when a Layer 2 switch receives a data frame with the destination MAC address that is a multicast MAC address, the switch cannot find a matching entry in its existing MAC address table. The switch then broadcasts the received multicast packets. This wastes bandwidth and degrades network security.

IGMP snooping/MLD snooping implements Layer 2 multicast on switches. After a switch enabled with IGMP snooping/MLD snooping receives IGMP messages transmitted from hosts, the IGMP snooping/MLD snooping module analyzes the messages and sets up and then maintains a multicast forwarding table. The multicast forwarding table contains VLAN IDs, multicast addresses, router interfaces (upstream interfaces), and member interfaces (downstream interfaces). When the switch receives a multicast packet, it searches for a matching forwarding entry according to the VLAN ID and destination address (multicast group address) of the packet. If a matching entry is found, the switch forwards the packet to all member interfaces in the multicast group. If no matching entry is found, the switch discards the multicast packet or broadcasts it in the VLAN, depending on the configuration.

IGMP snooping/MLD snooping also provides functions such as interface fast leave, multicast group policy, and limiting the number of multicast groups.

2.3.4 IGMP Proxy

IGMP proxy is deployed on a Layer 2 device between a router and host so that the Layer 2 device can function as a proxy server (IGMP proxy agent). The IGMP proxy agent sends IGMP query packets to the host and processes the IGMP response packets sent from the host. In addition, the IGMP proxy agent responds to query packets sent from the router, summarizes the messages that the host sends to join or leave the multicast group, and notifies the router of the host activity. For the host, the IGMP proxy agent functions as a router; for the router, the IGMP proxy agent functions as a host. Forwarding entries are created to implement Layer 2 multicast.

2.3.5 Controllable Multicast

Traditional multicast services are uncontrollable. That is, any user can send IGMP Report messages to join a multicast group, and receive multicast packets from the group. Controllable multicast was developed to control the authorization by which a user can join a certain multicast group. When a user requests to join a multicast group, the switch must authenticate the request, and reject unauthorized requests.

The switch provides the controllable multicast mechanism based on VLANs. Through multicast profiles, the switch provides hierarchical mechanisms to control the users' authorization to join corresponding multicast groups, as shown in Figure 2-4.

Figure 2-4 Hierarchical control mechanisms of controllable multicast

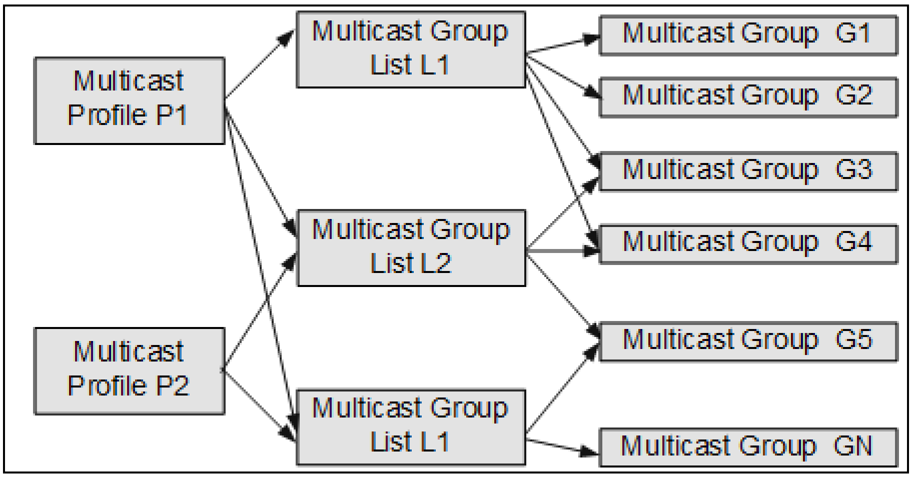

A multicast group has a multicast address, and a multicast group list is a set of multicast groups. A multicast profile is a set of multicast group lists, and defines the framework of users' rights to join multicast groups.

The switch on which controllable multicast is applied can control the generation of Layer 2 multicast forwarding entries by intercepting IGMP Report messages. When a switch receives an IGMP Report message from a user, the switch obtains the multicast profile based on the VLAN to which the message belongs. If the group is not in the list of the profile, the user cannot join the group. The switch intercepts the IGMP Report message and does not generate the related forwarding entry. The user cannot receive data flows of this group. If the multicast group is in the list of the profile, the switch checks the mode in which the list is added to the profile. If the list is added in watch mode, the switch allows the IGMP Report message to pass through. If the list is added in preview mode, the switch allows the IGMP Report message to pass through and starts a timer. When the preview period expires, the switch deletes the forwarding entry of the group and intercepts subsequent IGMP Report messages of the group.

2.3.6 Layer 2 Multicast CAC

CAC is short for Call Admission Control. Layer 2 multicast CAC is a technology to control user access. It accurately controls multicast services and ensures service quality for most users. To some extent, it reduces multicast attacks.

Layer 2 multicast CAC is implemented by controlling the establishment of Layer 2 multicast forwarding entries. The establishment and deletion of Layer 2 multicast forwarding entries depend on IGMP Report messages and IGMP Leave messages, respectively. Therefore, in the scenario where IGMP or IGMP snooping is deployed for multicast services, Layer 2 multicast CAC can be used. Layer 2 multicast CAC restrains the generation of multicast forwarding entries. When the upper threshold for the number of multicast forwarding entries is reached, no more forwarding entries are generated. This ensures the processing capacity of the device and controls link bandwidth.

Layer 2 multicast CAC can restrict the following items:

- Number of multicast group members
- Bandwidth of multicast groups
- Number of multicast groups on a channel
- Bandwidth of member multicast groups on a channel

Based on control dimensions, Layer 2 multicast CAC is classified into:

- Global Layer 2 multicast CAC
- Layer 2 multicast CAC in a VLAN
- Layer 2 multicast CAC on an interface
- Layer 2 multicast CAC on an interface in a specified VLAN
- Layer 2 multicast CAC in a VSI
- Layer 2 multicast CAC on a sub-interface
- Layer 2 multicast CAC on a PW

You can configure multicast CAC based on network types and service requirements. You can configure the upper threshold for the number of multicast group members to control the generated multicast forwarding entries. If the number of members in some multicast groups does not need to be controlled, you can define access control list (ACL) rules. Moreover, Layer 2 multicast CAC supports the definition of channel-based rules to limit the number or bandwidth of channels, to meet service requirements.

2.4 Multicast Membership Management

2.4.1 IGMP

Internet Group Management Protocol (IGMP) runs on hosts and multicast routers that are directly connected to the hosts. Hosts send messages to join a certain multicast group and receive messages from the multicast group to a local router through IGMP. The router periodically searches a LAN for the status of known members in a multicast group through IGMP. That is, the router checks whether members of a multicast group exist. The router then collects and maintains information about the multicast group members connected to itself. Through IGMP, the router checks whether a member in a multicast group exists. The router does not record the mapping between multicast groups and hosts.

IGMP has three versions. IGMPv1 (RFC 1112) defines the basic process of membership query and report. IGMPv2, which is widely used now, is defined by RFC 2236. Compared with IGMPv1, IGMPv2 adds the member fast leave mechanism. In IGMPv3, members can be configured to receive or not receive packets from a certain multicast source.

Figure 2-5 and the discussion that follows describe the working principle of IGMPv2.

Figure 2-5 Working principle of IGMPv2

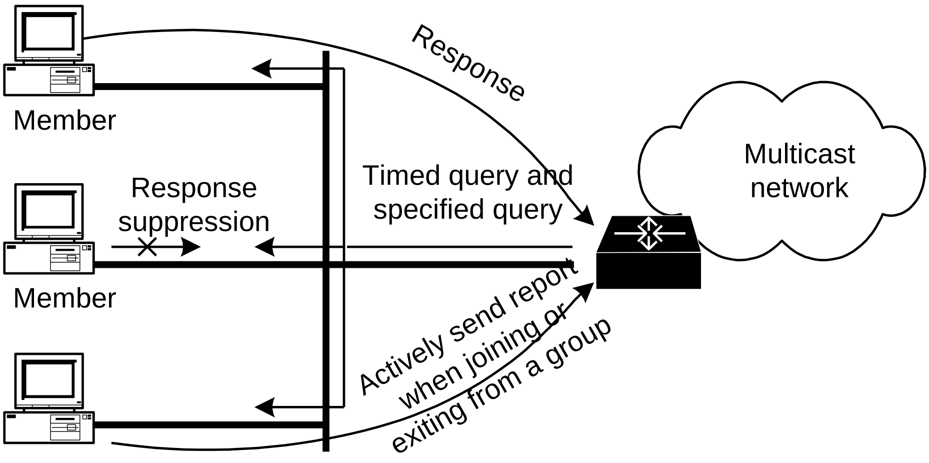
Response
Member
Response suppression
Timed query and specified query
Multicast network
Member
Actively send report when joining or exiting from a group

When several multicast routers exist on the same network segment, IGMPv2 elects only one multicast router to act as a querier through the querier election mechanism. The querier periodically sends General Query messages to search for memberships. To respond to such queries, hosts send Membership Report messages. The time hosts send Membership Report messages, however, is random. When finding that other members on the same network segment are sending the same messages, a host suppresses its own Membership Report messages. If a new host joins a multicast group, the host sends a Membership Report message voluntarily without waiting for the query message from the querier. To leave a multicast group, the host sends a Leave message voluntarily. After receiving a Leave message, the querier sends Group-Specific Query messages to check whether members of the multicast group still exist.

If a router is a member of a certain multicast group, the router functions the same as a common host to respond to the query of other routers. In IGMP, an interface indicates a main interface connected to a network. If a router has multiple interfaces connected to the same network, you need to enable IGMP on only one interface. The interfaces on hosts, however, need to be enabled with IGMP if they are members of a multicast group.

Through the querier election mechanism, a table is set up on the multicast router to record the multicast group members on the subnets to which interfaces on the router belong. When receiving the data packets of multicast group G for example, the router forwards the data packets only to the interfaces connected to members of multicast group G. Forwarding of data packets between routers is determined by routing protocols rather than IGMP.

Huawei switches support the following basic IGMP functions:

- IGMPv1, IGMPv2, and IGMPv3 and configurable versions
- Static IGMP
- Configurable multicast group range that an interface can join
- Suppression for IGMP Membership Report and Leave messages sent continuously by hosts in a VLAN or VSI and setting the suppression time

2.4.2 MLD

On IPv6 networks, after Multicast Listener Discovery (MLD) is configured on receiver hosts and the multicast router to which the hosts are directly connected, the hosts can dynamically join related multicast groups, and the multicast router can manage members on the local network.

MLDv1 is defined in RFC 2710 and MLDv2 is defined in RFC 3810. Both MLD versions support the Any-Source Multicast (ASM) model. MLDv2 supports the Source-Specific Multicast (SSM) model, whereas MLDv1 supports the SSM model only with the help of SSM mapping.

MLD functions the same as the Internet Group Management Protocol (IGMP) for IPv6. Implementation of MLD and is similar to that of IGMP. For example, MLDv1 is similar to IGMPv2, and MLDv2 is similar to IGMPv3. Features unique to MLD include the principles of MLDv1 and MLDv2, MLD querier election mechanism, and MLD group compatibility.

Figure 2-6 shows the networking of MLD.

Figure 2-6 MLD networking

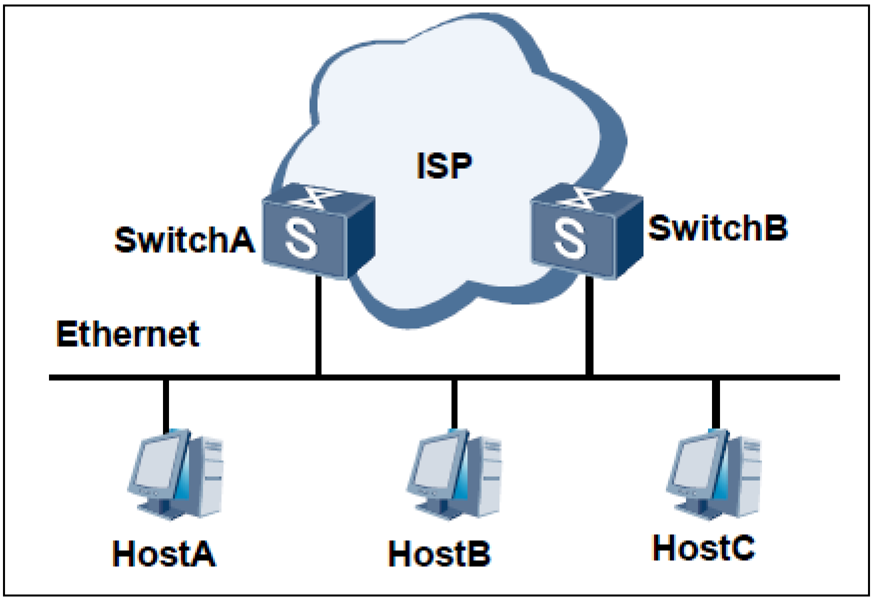

By sending Multicast Listener Query messages to hosts and receiving Multicast Listener Report messages and Multicast Listener Done messages from hosts, the switch detects which multicast group contains receivers on the relevant network segment. If receivers exist on the network segment, the switch forwards the corresponding multicast data to the network segment. If no receivers exist on the network segment, the switch forwards no multicast data. With MLD, hosts can determine whether to join or leave a multicast group.

The multicast device enabled with MLD assumes one of two roles on the network segment: querier or non-querier. There is normally only one querier on a network segment. Therefore, a querier has to be elected from among the multicast devices. The querier sends Multicast Listener Query messages to hosts and receives Multicast Listener Report messages and Multicast Listener Done messages from hosts. In this manner, the querier detects which multicast group contains receivers on the relevant network segment. The non-querier receives only Multicast Listener Report messages from hosts. According to the querier behavior, the non-querier detects which multicast group contains receivers on the relevant network segment and identifies hosts that leave the multicast group.

2.4.3 Membership Management on a Layer 2 Network

IGMP is designed for Layer 3 networks. To control the forwarding of multicast packets on a Layer 3 network, a router only needs configurations on its interfaces to check the time to live (TTL) value of the packets. In many cases, particularly in a LAN, multicast packets have to pass through certain Layer 2 switching devices. If no configuration is performed on a Layer 2 device, the multicast packets are sent to all interfaces on the device, which wastes vast system resources. IGMP snooping addresses the flooding problem. This section describes the working principles of IGMP snooping.

When the IGMP Membership Report message sent by a host to a router passes through a switch, the switch detects and records the message to form mappings between multicast group members and interfaces. Therefore, when the switch receives multicast data packets, it forwards the packets to only the interfaces that are connected to the members according to the mappings.

IGMP snooping prevents flooding of multicast packets on a Layer 2 network but requires switches to extract Layer 3 information from the packets. In addition, IGMP snooping requires switches to detect and interpret all multicast packets, leading to inefficiency in processing. The detection and interpretation of multicast packets also take an inordinate amount of CPU processing time.

For improved user experience, you need to enable fast leave on switch interfaces, so the switch does not wait for an interface to age after the interface receives an IGMP Leave message. The interface becomes invalid immediately.

Huawei S series switches support IGMPv1 snooping, IGMPv2 snooping, and IGMPv3 snooping, as well as fast leave of interfaces. You can configure multicast group policies and source-specific multicast (SSM) group policies on the switch so that interfaces in a VLAN can be added only to the multicast group that matches the specified ACL rules.

2.5 Broadcast Packet Forwarding

Compared with unicast packet forwarding, multicast packet forwarding is more complex. Unlike unicast routes, multicast routes are on a routing tree from one point to multiple points, which requires a different forwarding procedure for multicast packets.

2.5.1 Classification of Multicast Routes

Multicast routes are classified as source tree routes and shared tree routes. A source tree takes the multicast source as the root and constructs a forwarding tree by connecting the shortest path from the source to each receiver. The source tree is therefore also called the shortest path tree (SPT). For a multicast group, an SPT is set up from the multicast source that sends packets to the group. A shared tree takes a router as the root and constructs a forwarding tree based on the shortest paths from the router to all receivers. The router is referred to as the Rendezvous Point (RP). When a shared tree is used, one multicast group corresponds to only one forwarding tree on the network. All multicast sources and receivers use the tree to receive and send packets. A multicast source first sends data packets to the RP, and the RP forwards the packets to all receivers.

Figure 2-7 shows examples of source tree and shared tree routing.

Figure 2-7 Examples of a source tree and shared tree

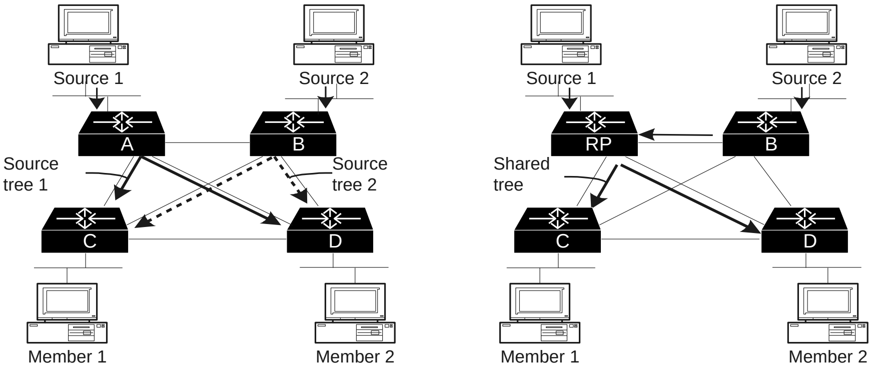

Using the shortest path from the source to each receiver, the source tree minimizes end-to-end delay. The limitation, however, is that routers need to store routing information for each multicast source. This consumes vast system resources and leads to a large routing table.

On the other hand, the shared tree minimizes the number of multicast group entries stored on routers. The paths of packets sent by a multicast source, however, may not be the shortest because packets are sent to the RP before being forwarded to the receivers. In addition, the RP needs to provide high reliability and processing capability.

2.5.2 Multicast Packet Forwarding

In unicast packet forwarding, a router does not need to search for the address of the multicast source. Instead, the router forwards packets to interfaces based on the destination address of the packets. In multicast, packets are sent to a group of receivers that are identified by a logical address. After receiving packets, a router must determine the upstream (multicast source) and downstream directions based on the source and destination addresses of the packets. The router then forwards the packets against the path to the multicast source. This process is called Reverse Path Forwarding (RPF).

In the RPF process, the original unicast routing table is used to determine the upstream and downstream neighboring nodes. Packets are forwarded to a downstream node only when they are received through the RPF interface that corresponds to a neighboring upstream node. In addition to correctly forwarding packets based on the configured multicast routes, RPF also prevents loops. Eliminating loops is of utmost significance in multicast routing, and the RPF check is the main function of RPF. After receiving a multicast packet, a router performs the RPF check and forwards the packet that passes the check. The packet that does not pass the RPF check is discarded. The process of the RPF check is as follows:

- The router searches the unicast routing table for the RPF interface that corresponds to the multicast source or the RP.
  - If the source tree is in use, the router searches for the RPF interface that corresponds to the multicast source.
  - If the shared tree is in use, the router searches for the RPF interface that corresponds to the RP.

  The RPF interface to which an address corresponds is the outbound interface for packets sent by the router to the address.
- If the multicast packet is received by the RPF interface, the RPF check succeeds and the packet is sent to a downstream interface. Otherwise, the packet is discarded.

Figure 2-8 shows the RPF check process in a source tree scenario.

Figure 2-8 RPF check

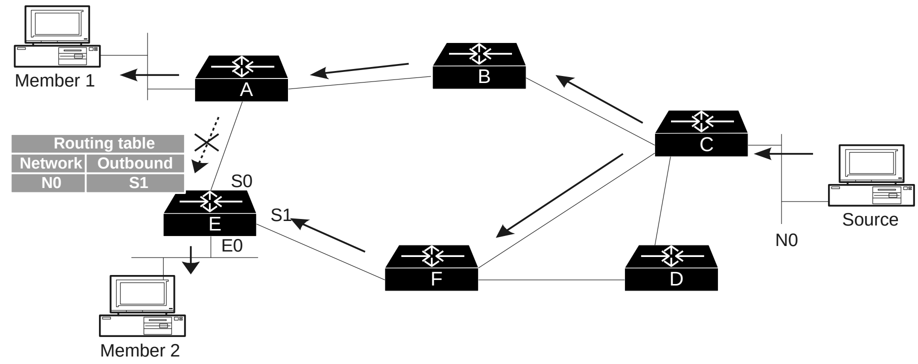

Routing table	
Network	Outbound
N0	S1

Interface S0 on router E receives a multicast packet whose source address is on network segment N0. Router E searches the routing table and finds that the outbound interface for packets destined to N0 is S1. Therefore, router E discards the packet. If a multicast packet is received from S1, router E forwards the packet.

RPF uses the interface on the shortest path from the router to multicast source or RP.

2.6 Intra-Domain Multicast Routing Protocols

Similar to unicast routes, multicast routes are classified as intra-domain or inter-domain. Intra-domain multicast routing protocols have become quite sophisticated. Among the many intra-domain routing protocols, Protocol Independent Multicast - Dense Mode (PIM-DM) and Protocol Independent Multicast - Sparse Mode (PIM-SM) are the most widely used.

2.6.1 PIM-DM

In a PIM-DM domain, routers enabled with PIM-DM periodically send Hello messages to discover adjacent PIM routers and determine leaf networks and leaf routers. PIM-DM routers are also responsible for electing a designated router (DR) on a multi-access network.

PIM-DM is developed under the assumption that when a multicast source sends data, all network nodes in the domain need to receive the data. Packets are forwarded in the flood and prune method. When the source sends data, all interfaces on the router forward the data, except the RPF interface that corresponds to the source. As a result, all network nodes in the PIM-DM domain receive the multicast packets.

To forward the multicast packets, intermediate routers need to create multicast entries (S, G) for multicast group G and multicast source S. Multicast entries (S, G) include multicast source address, group address, inbound interface, list of outbound interfaces, timer, and flag. If no multicast group member exists in a certain area, routers in the area send Prune messages to temporarily prune away or suspend the interface that sends multicast packets to the area. When the interface enters the pruned state, a timeout timer starts. When the timer expires, the state of the interface changes from pruned to forwarding again. In addition, the Prune message contains information about the multicast source and group. If a new multicast group member is detected in the pruned area, the downstream device does not wait for the pruned state to expire. The device voluntarily sends a Graft message to the upstream device to change the pruned state to the forwarding state and reduce response time.

Figure 2-9 shows the flood and prune process in a network scenario.

Figure 2-9 Flood and prune process in PIM-DM

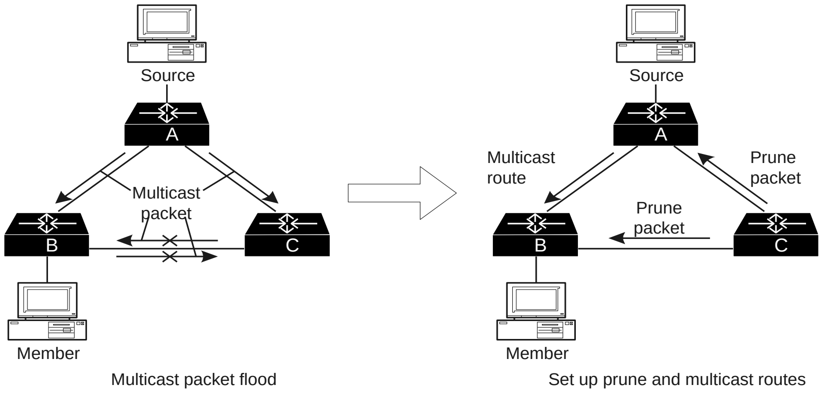

At first, the data sent by the multicast source is flooded across the network. Routers perform the RPF check when forwarding packets. As a result, the flooded packets sent by routers B and C to each other are rejected because the RPF check failed. Because the area where router C resides has no multicast group member, router C sends a Prune message to routers A and B. As a result, routers A and B set corresponding interfaces to the pruned state, and the multicast packets are sent to all group members along the correct paths.

In addition to DR election, the following mechanisms related to pruning are available on a PIM-DM multi-access network:

- An assert mechanism is used to elect the only forwarder to prevent the repeated transmission of multicast packets to the same network segment.
- A Join/Prune suppression mechanism is available to reduce the number of Join and Prune messages.
- The prune override mechanism is used to veto unwanted prune actions.

2.6.2 PIM-SM

In a PIM-SM domain, routers enabled with PIM-SM periodically send Hello messages to discover adjacent PIM routers. These PIM-SM routers are also responsible for electing a DR on a multi-access network. The DR is responsible for sending Join and Prune messages to the multicast tree root on behalf of the multicast group members directly connected to the DR. The DR also sends the data from the directly connected multicast source to the multicast distribution tree.

PIM-SM forwards multicast data packets by building up a multicast distribution tree. The multicast distribution tree is one of two types: the shared tree that takes the RP of multicast group G as a root or the shortest path tree that takes the multicast source as a root. PIM-SM sets up and maintains the multicast distribution tree by an explicit join and prune mechanism. Figure 2-10 and Figure 2-11 show the setup process of PIM-SM routes.

As shown in Figure 2-10, when an active member of group G exists on the network directly connected to the DR, the DR sends Join messages to join the shared tree (1) hop by hop to the RP of group G. When the Join messages are sent upstream on the tree, intermediate routers record the multicast forwarding status (2), that is, the routers create routing entries. Routing entries contain fields including source address, group address, inbound interface of multicast data packets, and list of outbound interfaces of multicast data packets, timer, and flag bits. The data enables routers that receive the multicast packets to forward them along the distribution tree. If the multicast packets are no longer required, the DR sends Prune messages hop by hop to the RP of group G to prune the shared tree. When pruning is performed upstream along the tree, intermediate routers update their routing entries, such as deleting an outbound interface. The routers on the distribution tree periodically send Join and Prune messages to the RP to maintain the status of the multicast distribution tree.

When the source host sends multicast data to the multicast group, the data is encapsulated into Register messages and then sent to the RP by the DR (3). The RP decapsulates the data from the Register messages and forwards the data to all group members along the shared tree. Then the RP sends Join or Prune messages (4) for a specific source to join the shortest path tree of the multicast source. As a result, data packets from the source are directly sent to the RP without encapsulation along the shortest path tree (5). When the multicast data packets are received along the shortest path, the RP sends Register-Stop messages to the DR of the source (5), and the DR then stops the register encapsulation process. After that, multicast data from the source is sent to the RP along the shortest path tree (B→A→RP) without register encapsulation. The RP forwards the data to the shared tree and sends the data to all group members along the shared tree (RP→D→C).

Figure 2-10 Working process of PIM-SM (a)

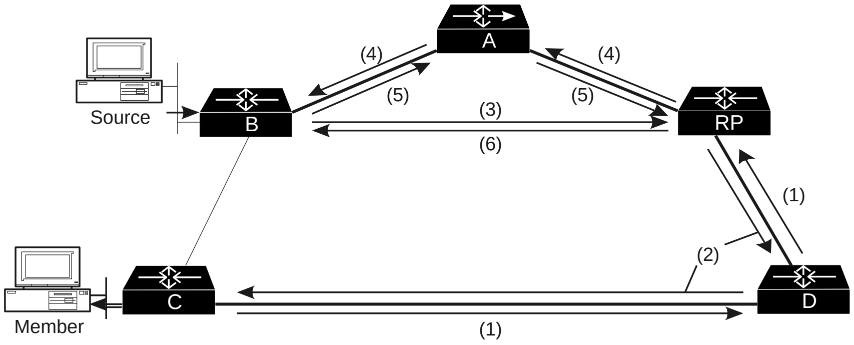

Figure 2-11 Working process of PIM-SM (b)

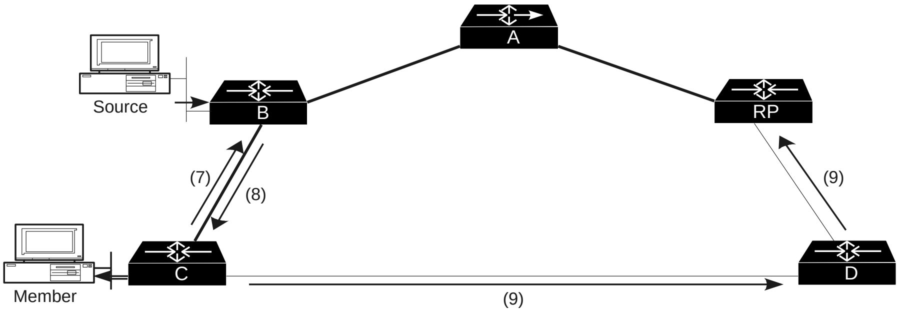

As shown in Figure 2-11, if a certain data transmission rate is reached, the DR sends explicit Join messages to the shortest path tree of the source (7) and the multicast packets are then forwarded along the shortest path tree (8). The DR updates the shared tree and deletes the corresponding shared forwarding route (9).

RP election also applies on a PIM-SM network. One or multiple candidate-BSRs are configured in the PIM-SM domain. A BSR is then elected according to certain rules. In the PIM-SM domain, candidate-RPs are also configured. These candidate-RPs send packets that include information about their addresses and available multicast groups to the BSR in unicast mode. The BSR then generates Bootstrap messages about the candidate-RPs and their corresponding group addresses at the specified time. The Bootstrap messages are sent hop by hop across the entire domain. Routers receive and store Bootstrap messages.

After the DR receives an IGMP Membership Report message from a directly connected host to join a multicast group, the DR uses the hash algorithm to map the group address to a candidate-RP, if no routing entry is available for the multicast group. The DR then sends Join or Prune messages hop by hop to the RP. If the DR receives multicast data packets from a directly connected host, but no routing entry is available for the multicast group, the DR also uses the hash algorithm to map the group address to a candidate-RP. The DR then encapsulates the multicast packets to Register messages and forwards the Register messages to the RP in unicast mode.

A PIM-SM multi-access network also includes the following mechanisms:

- An assert mechanism is used to elect the only forwarder to prevent the repeated transmission of multicast packets to the same network segment.
- A Join/Prune suppression mechanism is available to reduce the number of Join and Prune messages.
- The prune override mechanism is used to veto unwanted prune actions.

2.7 Inter-Domain Multicast Routing Protocols

Inter-domain multicast routing protocols are being studied and tested. At present, the mature solution is the combination of the following protocols:

- Multiprotocol Border Gateway Protocol (MBGP), used to exchange multicast routing information between ASs
- Multicast Source Discovery Protocol (MSDP), used to exchange multicast source information between ISPs

- PIM-SM, used as a multicast routing protocol in an AS

The following discussion focuses on the working process of MBGP and MSDP and the combined solutions of PIM-SM, MBGP, and MSDP.

2.7.1 MBGP

The primary problem of inter-domain routes is how to transmit routing information (or reachability information) between ASs that may belong to different operators or ISPs. In addition to the distance, inter-domain routing information must include the policies of operators or ISPs, unlike intra-domain routing information. At present, the most widely used inter-domain unicast routing protocol is BGP-4.

Multicast network topology differs from unicast both physically and strategically. Some routers on the network may support only unicast. Other routers may not forward multicast packets according to configured policies. In addition to unicast routing information, the administrator must know which routers on the network support multicast, that is, the multicast network topology, to construct an inter-domain multicast routing tree. To transmit multicast route information between different domains, BGP must be modified. In brief, inter-domain multicast routing protocols must meet the following requirements:

- Differentiate unicast and multicast topologies.
- Have peer and policy control methods.

BGP-4 meets the second criterion and is an effective and steady unicast inter-domain routing protocol. Therefore, the proper solution is an improved extension to BGP-4 rather than a new set of protocols. RFC 2858 defines the multi-protocol extensions for BGP. The extended BGP protocol, Multiprotocol Border Gateway Protocol (MBGP) or BGP4+, carries IPv4 unicast routing information and routing information about other network protocols, such as multicast and IPv6.

With MBGP, unicast and multicast routes can be exchanged in the same process and stored in different routing tables. Common policies and configurations supported by BGP-4 can be applied to multicast.

Compared with BGP-4, MBGP adds two attributes in Update messages: MP_REACH_NLRI (multi-protocol reachable NLRI) and MP_UNREACH_NLRI (multi-protocol unreachable NLRI). The attributes are optional non-transitive. Routers that do not support MBGP ignore these attributes and do not forward packets based on them.

MP_REACH_NLRI

MP_REACH_NLRI is identified by the combination of AFI, NHI, and NLRI.

The Address Family Information (AFI) consists of an address family identifier (AFI) and subsequent address family identifier (SAFI). The address family identifier indicates the communication protocol in use and the address type in the Next Hop Identifier (NHI). The protocol identification methods are defined in RFC 1700. The value 1 means an IPv4 address and the value 2 means an IPv6 address. SAFI is a supplement to the Network Layer Reachability Information (NLRI). As defined in RFC 2858, SAFI has three identifier values. The value 1 means the NLRI is of unicast mode. The value 2 means the NLRI is of multicast mode. The value 3 means the NLRI can be unicast or multicast mode.

The NHI contains information about the next hop, including the next-hop network address. Packets can be sent to all destinations specified in the NLRI, by the next-hop network address.

The NLRI contains a destination address that can be reached by the next-hop address in the NHI.

MP_UNREACH_NLRI

Used to withdraw one or more existing routes, MP_UNREACH_NLRI is identified by the combination of AFI and WR.

- The meaning and structure of AFI are the same as those of the AFI in MP_REACH_NLRI.
- The withdrawn route (WR) contains one or more NLRIs that contain unreachable destination addresses.

By introducing SAFI in Address Family Information, MBGP can support both unicast and multicast. MBGP also supports different policies in different unicast or multicast topologies constructed over the network. Therefore, the MBGP-generated inter-domain unicast route and multicast forwarding route in a routing policy may be different.

2.7.2 MSDP

ISPs do not want to forward multicast traffic to the competitors' RPs, but they need to receive multicast traffic from the multicast source no matter where the source's RP is located. The Multicast Source Discovery Protocol (MSDP) can meet the ISPs' requirements. In MSDP, an inter-domain source tree is used, rather than a common tree, and the intra-domain multicast routing protocol must be PIM-SM.

In MSDP, an RP in a domain sets up the MSDP peer relationship with an RP in another domain through the TCP connection. Based on the peer relationship, the two RPs exchange multicast source information. If a local receiver intends to receive packets from a source in another domain, a multicast source tree is built according to the methods in PIM-SM. Figure 2-12 shows the working process of the combined solution of PIM-SM, MSDP, and MBGP.

Figure 2-12 Working process of PIM-SM, MSDP, and MBGP combined solution

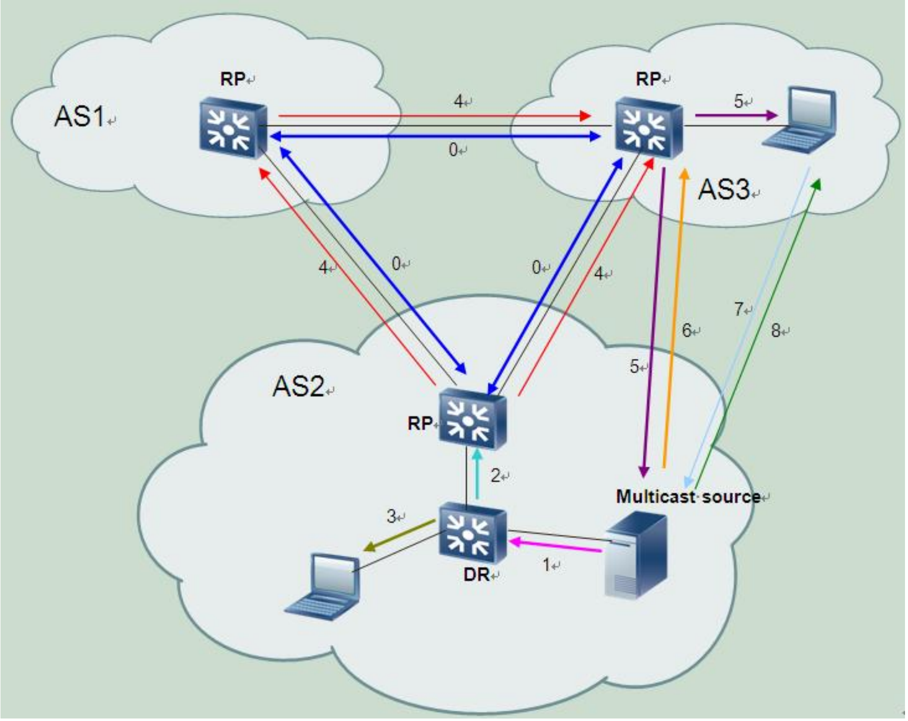

Figure 2-12 shows the following steps in the process:

0. MBGP peer relationships are established between ASs and MSDP peer relationships are established between RPs.

1. The multicast source in domain AS2 starts to send a packet.

2. The designated PIM-SM router of the source (generally the local PIM-SM router closest to the multicast source) encapsulates the packet into a Register message and then sends the message to the RP in the domain.

3. After receiving the packet, the RP decapsulates the packet and forwards it to all members in the domain along the shared tree. The members in the domain can choose whether to switch to the source tree. Steps 1 through 3 are typical PIM-SM working processes.

4. The RP also generates a Source Active (SA) message and sends it to the MSDP peers (the RPs in AS1 and AS3). Each SA message contains the IP address of the multicast source, address of the multicast group, and address of the RP that generated the message. The RP sends SA messages periodically as long as the source continues to send packets.

5. If a receiver exists in the domain where the peer RP resides (as in AS3), the peer RP sends a PIM-SM Join message to the source when forwarding packets to the receive end. Note that the peer RP does not send the Join message to the RP in AS2.

6. The reverse forwarding path is then set up, and the packets from the multicast source are sent directly to the RP in AS3. The AS3 RP starts to forward packets.

7 and 8. Multicast group members in AS3 can choose whether to switch to the shortest path tree (SPT). At this time, the group members know the IP address of the multicast source. This step is the same as that in PIM-SM.

The combined solution of PIM-SM, MBGP, and MSDP is an extension of PIM-SM in the inter-domain scenario. If the combined solution as a whole is considered as a PIM-SM domain, the set of RPs in all domains is equivalent to the single RP in the PIM-SM domain. In the combined solution, the following two steps are added:

1. Flooding of multicast source information in the set of RPs accomplishes the convergence of multicast source and group members at the RP in PIM-SM.
2. Transmission of inter-domain multicast routes ensures the successful forwarding of multicast packets in the domains. As shown in the preceding process, the multicast topology information transmitted by MBGP is applied in the setup of reverse paths from the RP and receivers in AS3 to the peer RP in AS2.

3 Multicast Networking

3.1 Multicast Networking in a Single Domain

PIM-SM is a recognized intra-domain multicast routing standard. If a network consists of only one AS or if multicast is required in only one domain, PIM-SM is sufficient for the deployment of multicast services. To strengthen the reliability of the RP and balance the multicast traffic on the network, you can select several RP nodes to configure anycast RP. This achieves the purposes of redundancy backup and load balancing.

The anycast RP mechanism works as follows: Multiple RPs are configured with the same anycast RP address that is the address of an RP interface. The address of a logical interface on the RP, such as a loopback interface, is generally used. The RP uses the address to advertise mappings from multicast groups to the RP. Because the anycast RP address is used, group members send Join messages to join a group to the closest RP in the topology. In addition, MSDP connections are set up between RPs using their different addresses. MSDP is then applied to synchronize multicast source information among all the RPs. Anycast RP is actually a special application of MSDP in a domain.

Figure 3-1 shows multicast networking in a single PIM-SM domain.

Figure 3-1 Multicast networking in a single PIM-SM domain

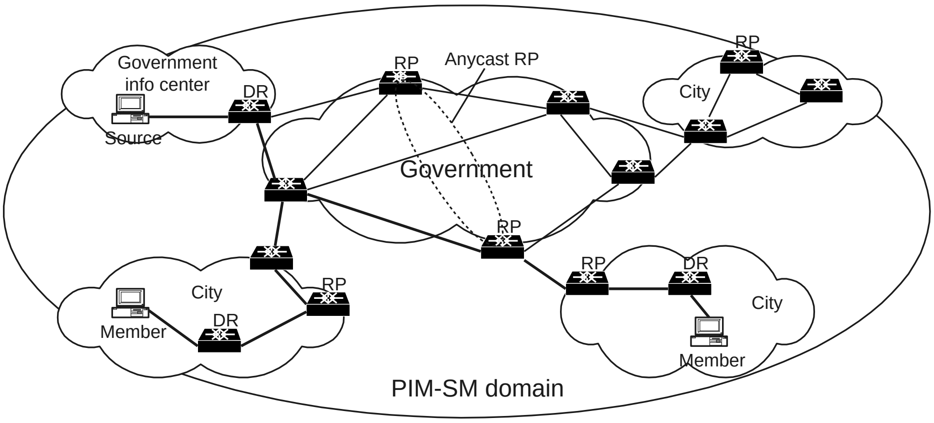

3.2 Inter-AS Multicast Networking

In accordance with different multicast capabilities that networks support, inter-AS multicast networking can be deployed in the following ways.

3.2.1 Full-Mesh Multicast Networking – Combination of PIM-SM, MBGP, and MSDP

In Figure 3-2, PIM-SM runs on the entire network. MBGP and MSDP are enabled between domains. That is, collection of multicast routing information in a domain and multicast sources is completed by PIM-SM. Between domains, MBGP transmits information about multicast topologies and MSDP transmits information about multicast sources. This scheme requires all ASs to support PIM-SM, MBGP, and MSDP. The combined solution of PIM-SM, MBGP, and MSDP is a scheme in inter-domain multicast networking and is deployed by UUNet and Sprint in inter-AS multicast networking schemes.

Figure 3-2 Full-mesh multicast networking

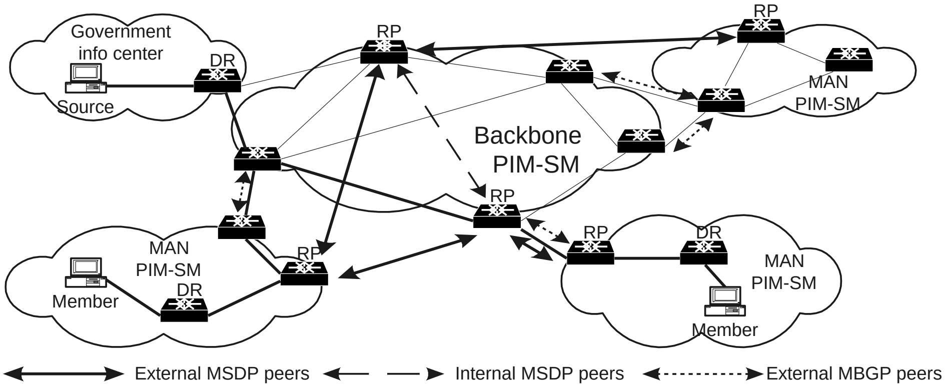

In the combined solution of PIM-SM, MBGP, and MSDP, MBGP peer relationships are configured between Autonomous System Border Routers (ASBRs). MSDP peer relationships are configured between RPs. Inside an AS, routers can be configured as internal MBGP peers. RPs are configured as internal MSDP peers to run anycast RP. PIM-SM runs in all ASs.

3.2.2 Multicast Networking Not Supported by the Backbone – PIM-SM and Tunnels (MBGP and MSDP)

In Figure 3-3, either the backbone network does not support multicast or multicast does not run on the backbone network. PIM-SM runs inside each MAN. The RPs on all MANs constitute a virtual network through tunnels. PIM-SM, MBGP, and MSDP run on the virtual network. The advantage of this scheme is that the backbone network is not required to support PIM-SM, MBGP, and MSDP. Multicast traffic is transparent to the backbone network. This prevents the adverse impact of multicast packet forwarding on device performance. The weakness is that RPs have to support PIM-SM, MBGP and MSDP tunnels. Configuration and management of RPs are complex, and the requirements for RP devices are high.

Figure 3-3 PIM-SM and tunnel (MBGP and MSDP) solution

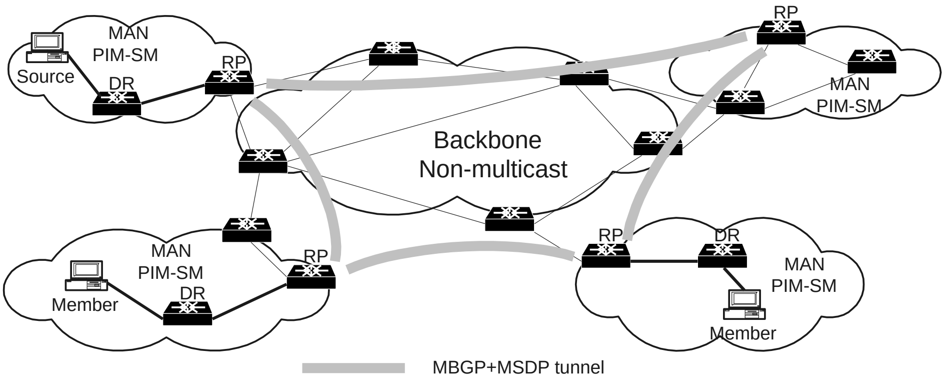

3.2.3 PIM-SM and Tunnel (PIM-DM)

In Figure 3-4, PIM-SM runs inside each MAN. The RPs on all MANs constitute a virtual network through tunnels. PIM-DM runs on the virtual network. The advantage of this scheme is that the backbone network is not required to support PIM-SM, MBGP, and MSDP. Multicast traffic is transparent to the backbone network. Therefore, a large amount of multicast routing information does not need to be stored on the backbone network. The weakness is that PIM-DM runs among RP nodes. The timed flooding of multicast packets may waste bandwidth resources on the backbone network.

Figure 3-4 PIM-SM and tunnel (PIM-DM) solution

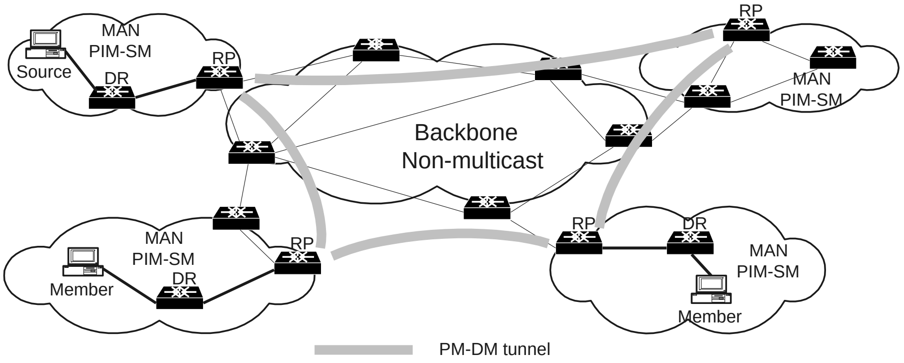

3.2.4 Typical IPTV Networking

In Figure 3-5, Layer 2 multicast CAC is preferred in the IPTV solution to limit the number of IPTV channels to ensure high service quality for existing users. With the development of IPTV, the number of channels increases rapidly. If the number of channels offered to users keeps increasing, it results in overload on the devices at the convergence layer and degrades user satisfaction. Moreover, if multicast-based network attacks occur, devices on the network may be tied up in processing a high volume of attack packets and cannot respond to normal processing requests.

Figure 3-5 Typical IPTV Networking

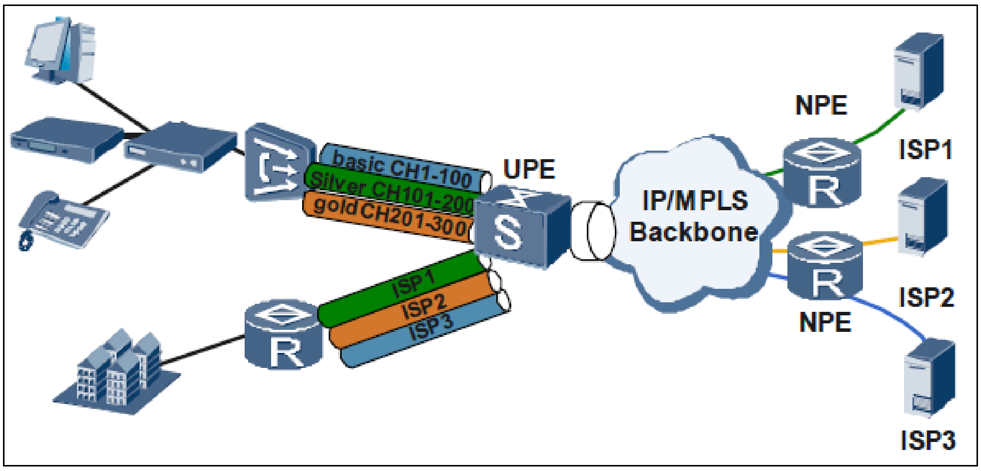

As described in section 2.3.6 "Layer 2 Multicast CAC", Layer 2 multicast CAC is a technology to control user access from different dimensions. It accurately controls multicast services and ensures service quality for most users. To some extent, it reduces multicast attacks.

4 Networking Service Management – Controllable Multicast

IP multicast technology plays an important role in the development of new multimedia services. However, problems of user management and service management still exist in the operation of multicast services.

No multicast protocol provides user authentication. A user can join a multicast group and leave the group at will. The multicast source does not know when a user joins or leaves a multicast group, so the number of users receiving multicast traffic on a network during a certain period is unknown. The multicast source also lacks effective measures to control the direction and scope of transmitting multicast information over networks.

In addition, multicast protocols provide no reliable security guarantee, because they lack effective control over multicast sources. Any user on the network can act as a multicast source to send multicast packets. There is also no effective control over receivers. On a multicast network, multicast programs may clash with each other and illegal multicast sources may exist.

Nevertheless, multicast technologies are being improved to bring out many benefits in deploying new services. Deployment of multicast services, however, is still faced with problems of user authentication, source security, and security of multicast packet flooding. Considering existing network characteristics, multicast technology, and practical applications, Huawei has put forward the controllable multicast technology that fully complies with the standard multicast protocols. Controllable multicast consists of multicast source management, user management, and service management, to address problems encountered in the deployment of multicast services.

4.1 Multicast Source Management

Before a multicast flow enters the backbone network, a multicast service controller is responsible for distinguishing legal media servers from illegal ones. In this process of multicast source management, the legal multicast flow is transmitted and the illegal flow is blocked.

On a large-scale network, configuring source management information manually is quite complex and slows network development. To address this configuration problem, Huawei uses the Quidway service management platform to implement multicast source management. The Quidway platform, shown in Figure 4-1, facilitates configuration of source management by adding or deleting access control entries and ensures consistency of entries on the entire network. With multicast source management, unauthorized media servers are prevented from using network bandwidth, which might otherwise occur due to the limitations of multicast protocols. The platform ensures the security and stability of the backbone network.

Figure 4-1 Multicast source management through the Quidway platform

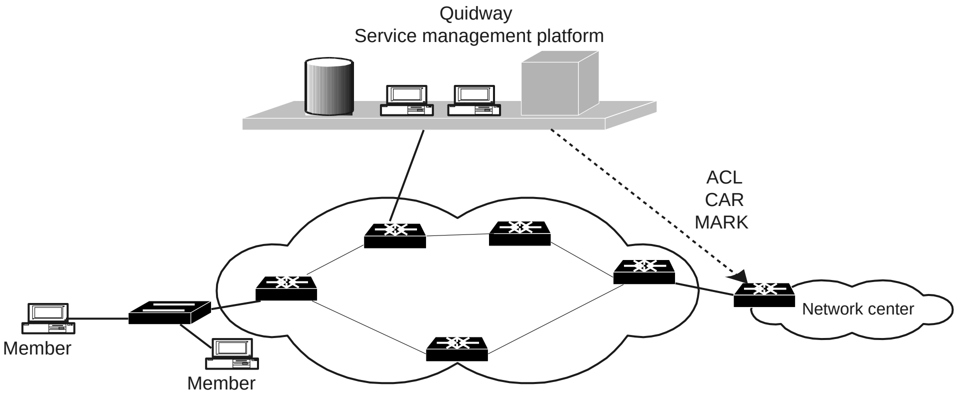

4.2 Multicast User Management

Standard multicast protocols fail to consider user management. User management is not properly addressed in multicast services operations for many existing applications, including deployment sites being operated or tested locally and abroad. As a value-added service, user control management is an indispensable component of multicast services.

Broadly speaking, user management of multicast services is of two types. In the first type, a user already receives the multicast packets but can read the contents only after authentication. In the second type, a user cannot receive or read multicast packets until the user passes authentication.

The first type of user management has low requirements for intermediate devices. The intermediate devices only need to support standard multicast protocols and ensure that users can receive multicast packets, and other operations are performed by the client software. This management mode is easy to implement and can control specific users. However, it consumes network bandwidth because unauthorized users also receive the multicast packets. In addition, this mode does not facilitate unified multicast management.

In the second type of user management, interfaces or VLANs on a Layer 2 switch are used to control users according to the 802.1X protocol. The Layer 2 device first authenticates the multicast rights of a user based on 802.1X. If authentication succeeds, the Layer 2 device receives IGMP Membership Report or Leave messages from the user, creates forwarding entries, and allows the user to receive multicast packets. Otherwise, the Layer 2 device discards the IGMP messages from the user and prevents the user from receiving multicast packets.

To further strengthen user access control, the Quidway service management platform can create multicast access rules for the user after authentication. The user can access only authorized multicast services based on the access rules. When a user joins a multicast group, the Layer 2 device first authenticates the user on the Quidway platform. If authentication succeeds, the Layer 2 device generates a multicast path to the user. Otherwise, the Layer 2 device prevents the user from joining the multicast group.

By combining Layer 2 devices and the Quidway service management system, you can authenticate and authorize multicast users on an entire network as shown in Figure 4-2.

Figure 4-2 Network authentication and authorization of multicast users

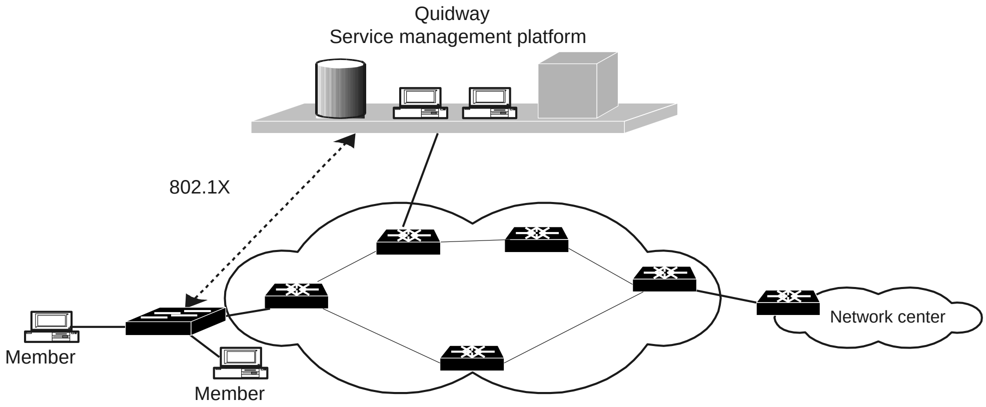

4.3 Multicast Security Control

In standard multicast, a receiver can join any multicast group and the branches of the multicast tree cannot be controlled. The multicast source does not know the scope and direction of the multicast tree, which results in low security. Huawei static multicast tree scheme meets the security requirement to control the scope of certain important information. A static multicast tree is pre-configured to control the scope and direction of the multicast tree. The static multicast tree does not allow joining of other dynamic multicast group members. As a result, packets from the multicast source are flooded in only the specified area, as shown in Figure 4-3. By configuring a static multicast tree, you can meet the security demand of high-value users.

Figure 4-3 Security control through a static multicast tree

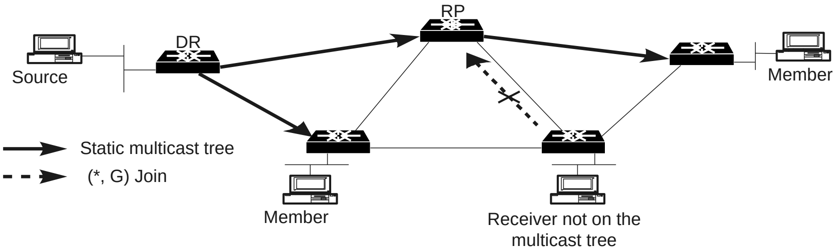

5 Conclusion

Broadband is the key to construction of high-speed information network architectures. Many MANs use broadband from the access layer to core layer and build the IP-based non-blocking switching platform.

The application of broadband on networks allows people to communicate more freely on the information highway. At the same time, broadband networks are required to support increasing multimedia applications. Multicast provides transmission technologies in the deployment of multimedia services.

Multicast technology involves address allocation and membership management, as well as routes and security. The assignment mode of multicast addresses, inter-domain multicast routing, and multicast security still receive wide attention. In the existing applications:

- Multicast packet replication across VLANs is adopted in Layer 2 multicast.
- IGMPv2 is generally used for membership management.
- PIM-SM is the first choice in intra-domain multicast routing protocol because of good scalability and convenient switching from the shared tree to the source tree.
- The combined solution of PIM-SM, MBGP, and MSDP is widely used in inter-domain deployment scenarios.

To meet different networking requirements, Huawei has put forward other schemes described in this white paper. Multicast technology provides a variety of broadband value-added services including stream media and video conferencing. The successful deployment of multicast services still relies on effective service management, monitoring, and security control. Based on understanding and experience in service operations, Huawei provides controllable multicast solutions that can be operated, managed, and continuously improved. As a major organization involved in setting multicast international standards, Huawei is dedicated to the development of multicast technologies, the popularity of multicast services, and the perfecting of multicast functions.

6 References

S.E. Deering. "Host extensions for IP multicasting," RFC1112, August 1,1989.

W. Fenner. "Internet Group Management Protocol, Version 2," RFC2236, November 1997.

D. Waitzman, C. Partridge, and S.E. Deering. "Distance Vector Multicast Routing Protocol." RFC1075. November 1,1988.

J. Moy. "Multicast Extensions to OSPF." RFC1584, March 1994.

Brad Cain, Ajit Thyagarajan, and Steve Deering. "Internet Group Management Protocol, Version 3," <draft-cain- igmp-00.txt>, Expires March 8, 1996.

D. Estrin, D. Farinacci, A. Helmy, D. Thaler, S. Deering, M. Handley, V. Jacobson, C. Liu, P. Sharma, and L. Wei. "Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification." RFC2362, June 1998.

D. Estrin, D. Farinacci, V. Jacobson, C. Liu, L. Wei,P. Sharma, and A. Helmy. "Protocol-Independent Multicast (PIM), Dense-Mode Protocol Specification," <draft-ietf-idmr-PIM-DM-spec-01.ps>, January 17, 1996.

Andrew Adams, Jonathan Nicholas, and William Siadak. "Protocol-Independent Multicast (PIM), Dense-Mode Protocol Specification (Revised)," <draft-ietf-pim-dm-new-v2-01.txt>, February 15, 2002.

S. Deering, D. Estrin, D. Farinacci, V. Jacobson, C. Liu, and L. Wei. "Protocol-Independent Multicast (PIM): Motivation and Architecture," <draft-ietf-idmr-pim-arch-01.ps> January 11, 1995.

Y. Rekhter and T. Li. "A border gateway protocol 4 (BGP-4)." Internet Engineering Task Force (IETF), RFC 1771, March 1995.

T. Bates, R. Chandra, D. Katz, and Y. Rekhter. "Multiprotocol extensions for BGP-4." Internet Engineering Task Force (IETF), RFC2283, February 1998.

D. Farinacci and Bill Fenner, "Multicast source discovery protocol (MSDP)." Internet Engineering Task Force (IETF), <draft-farinacci-msdp-13.txt>, November 2001.

Dorian Kim, David Meyer, Henry Kilmer, and Dino Farinacci. "Anycast RP mechanism using PIM and MSDP", <draft-ietf-mboned-anycast-rp-08.txt>, May, 2001.

J.M.Pullen, M. Mytak, and C. Bouwens. "Limitations of Internet Protocol Suite for Distributed Simulation in the Large Multicast Environment", RFC 2502, February 1999.

P. Bagnall, R. Briscoe, and A. Poppitt. "Taxonomy of Communication Requirements, for Large-scale Multicast Applications," <draft-ietf-lsma-requirements-03.txt>, May 1999.

K. Obraczka. "Multicast Transport Mechanisms: A Survey and Taxonomy", IEEE Communications Magazine, Vol. 36 No. 1, January 1998.

Kevin C. Almeroth, Santa Barbara. "The Evolution of Multicast: From the MBone to Inter-Domain Multicast to Internet2 Deployment," 1999.

A Acronyms and Abbreviations

AFI	Address-Family Identifier
AS	Autonomous System
BGP	Border Gateway Protocol
BSR	BootStrap Router
DR	Designated Router
IANA	Internet Assigned Numbers Authority
IETF	Internet Engineering Task Force
IGMP	Internet Group Management Protocol
IP	Internet Protocol
ISP	Internet Service Provider
MAC	Media Access Control
MBGP	Multiprotocol Border Gateway Protocol
MBONE	Multicast backBONE
MSDP	Multicast Source Discovery Protocol
NLRI	Network Layer Reachable Information
PIM	Protocol Independent Multicast
PIM-DM	Protocol Independent Multicast-Dense Mode
PIM-SM	Protocol Independent Multicast-Sparse Mode

RIP	Routing Information Protocol
RP	Rendezvous Point
RPF	Reverse Path Forwarding
RPM	Reverse Path Multicast
SA	Source Active
SAFI	Subsequent AFI
SPT	Shortest Path Tree
TTL	Time to Live