
Networkforyou
Welcome
To
Network for you
Multicast
Networkforyou

There are three types of traffic that we can choose from for our networks:

Unicast

Broadcast

Multicast

- If you want to send a message from one source to one destination, we use unicast.
- If you want to send a message from one source to everyone, we use broadcast.
- What if we want to send a message from one source to a group of receivers? That's when we use multicast.
- Multicast communication is a technology that optimizes network bandwidth utilization.
- Multicast communication is a technology optimizes network conserves system resources.
- **IGMP (Internet Group Management Protocol) snooping is a technique used by Layer 2 switches to track which hosts on their LANs are interested in receiving multicast traffic. It relies on IGMP messages to learn which hosts are members of multicast groups.**
- **PIM is a multicast routing protocol that allows routers to learn about multicast groups and forward multicast packets to the correct destinations.**
- Multicast has many advantages; main advantage is scalability compared to unicast traffic.

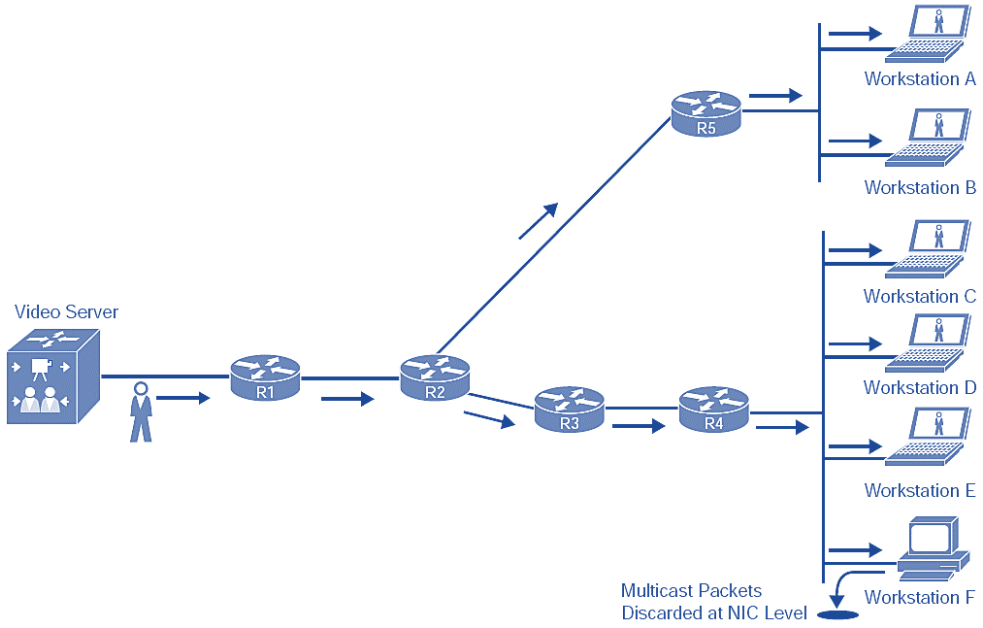

Multicast IP Addresses:

- Unicast IP addresses represent single device, multicast IP addresses represent a group.
- Internet Assigned Numbers Authority has reserved the class D range to use for multicast.
- Means we have from **224.0.0.0 through 239.255.255.255** range for IP multicast addresses.
- Some of the addresses are reserved and we cannot use them for our own applications.
- Make sure you don't use the **224.0.0.0 /24 and 224.0.1.0 /24** range and you will be safe.
- Like private and public IP addresses for unicast, IANA has reserved a range of IP addresses.
- We can use for the multicast on our local networks and this is the 239.0.0.0 /8 range.
- Everything between **239.0.0.0 - 239.255.255.255** is safe to use on your own networks.

Local Network Control Block:

- The 224.0.0.0 – 224.0.0.255 range has been reserved by IANA to use for network protocols.
- All multicast IP packets in this range are not forwarded by Cisco routers between subnets.
- (224.0.0/24) Addresses in local network control block are used for protocol control traffic.
- Internetwork control block (224.0.1.0/24) - Addresses in the internetwork control block.
- They are used for protocol control traffic that may be forwarded through the Internet.

IP Address	Usage
224.0.0.1	All Hosts
224.0.0.2	All Multicast Routers
224.0.0.3	Unassigned
224.0.0.4	DVMRP Routers
224.0.0.5	OSPF Routers
224.0.0.6	OSPF DR/BDR Router
224.0.0.7	ST Routers
224.0.0.8	ST Hosts
224.0.0.9	RIPv2 Routers
224.0.0.10	EIGRP Routers
224.0.0.11	Mobile Agents
224.0.0.12	DHCP Server / Relay
224.0.0.13	All PIM Routers
224.0.0.14	RSVP Encapsulation
224.0.0.15	All CBT Routers
224.0.0.16	Designated SBM
224.0.0.17	All SBMS
224.0.0.18	VRRP
224.0.0.19 – 255	Unassigned

Multicast Components:

Multicast is efficient but it doesn't work "out of the box". There are a number of components that we require:

- **First of all we use a designated range of IP address that is exclusively used for multicast traffic. We use the class D range for this: 224.0.0.0 to 239.255.255.255. These addresses are only used as destination addresses, not as source addresses. The source IP address will be the device that is sending the multicast traffic, for example the video server.**
- We also require **applications that support multicast**. A simple example is the **VLC** media player, it can be used to stream and receive a video on the network.
- When a **router receives multicast traffic, somehow it has to know if anyone is interested in receiving the multicast traffic**. Have look below picture.

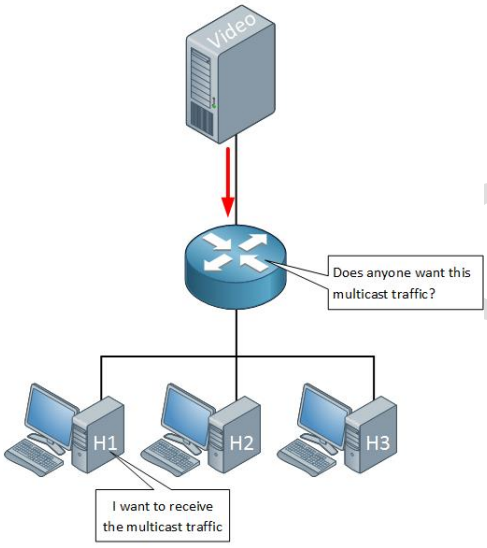

- Above you can see the router is receiving the multicast traffic from the video server. It doesn't know where and if it should forward this multicast traffic.
- We need some mechanism on our hosts that tell the router when they want to receive multicast traffic. We use the **IGMP (Internet Group Management Protocol)** for this.
- Hosts that want to receive multicast traffic will use the **IGMP protocol to tell the router which multicast traffic they want to receive.**
- **IGMP helps the router to figure out on what interfaces** it should forward multicast traffic but problem how about switches? Have a look below picture.

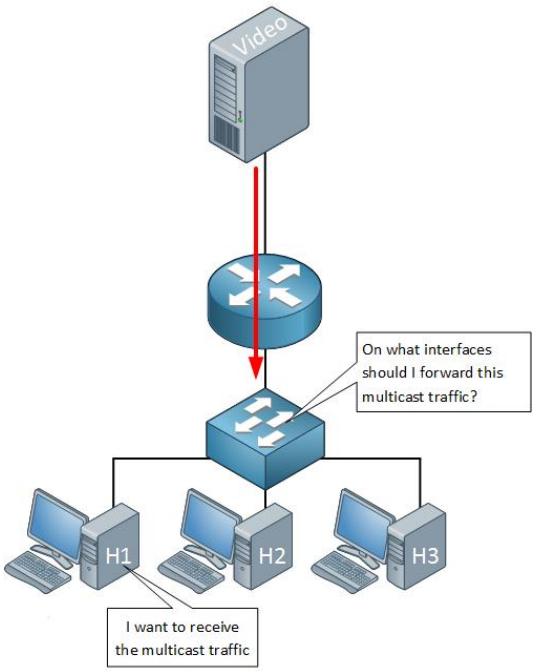

- Router knows that it has to forward the multicast traffic since a host used IGMP to tell the router it is interested.
- Once the multicast traffic arrives at the switch, we have another problem.
- Switches learn MAC addresses by looking at the source address of an Ethernet frame.
- Since we use multicast addresses only for the destination, how is the switch supposed to learn where to forward multicast traffic to?
- To help the switch figure out where to forward multicast traffic, we can use **IGMP snooping**.
- **The switch will "listen" to IGMP messages between the host(s) and router to figure out where it should forward multicast traffic to.**
- There's also a Cisco proprietary protocol called **CGMP (Cisco Group Management Protocol)** that can be used between switches and routers.
- The router will then be able to inform the switch where to forward multicast traffic. Unlike IGMP snooping, CGMP isn't used much.
- Above we have our video server that is forwarding multicast traffic to R1. On the bottom there's H1 who is interested in receiving it.
- **With unicast routing, each router advertises its directly connected interfaces in a routing protocol. Routers who receive unicast packets only care about the destination address.**
- They check their routing tables, find the outgoing interface and forward the packets towards the destination.

- With **multicast routing, things are not that simple…the destination is a multicast group address and the multicast packets have to be forwarded to multiple receivers throughout the network**.

To accomplish this, we use a multicast routing protocol:

- DVMRP (Distance Vector Multicast Routing Protocol)
- MOSPF (Multicast Open Shortest Path First)
- **PIM (Protocol Independent Multicast)**

IGMP (Internet Group Management Protocol):

- **Internet Group Management Protocol (IGMP) is the protocol that receivers use to join multicast groups.**
- When a receiver wants to receive a specific multicast feed, it sends an IGMP join using the multicast IP group address for that feed.
- The receiver reprograms its interface to accept the multicast MAC group address that correlates to the group address.
- For example, a PC could send a join to 239.255.1.1 and would reprogram its NIC to receive 01:00:5E:7F:01:01.
- IGMP must be supported by receivers and the router interfaces facing the receivers.
- Three versions of IGMP exist.
- IGMPv1, which is old and rarely used.
- IGMPv2, which is common in most multicast networks.
- IGMPv3, which is used by SSM (source-specific multicast).

IGMP (Internet Group Management Protocol) Version1:

Version 1 is the first version that hosts can use to announce to a router that they want to receive multicast traffic from a specific group. It's a simple protocol that uses only two messages:

Membership report
Membership query

When a **host wants to join a multicast group, it will send a membership report** to the group address that it wants to receive.

When the multicast-enabled router receives this message, it will start forwarding the requested multicast traffic on the interface where it received the IGMP membership report on.

The **router will periodically send a membership query to destination 224.0.0.**1 (all hosts multicast group address).

Hosts that receive this message will **respond with a membership report** to tell the router that they are still interested in receiving the multicast traffic.

When the router receives the **membership report, it's expiry timer will be refreshed**. When no hosts respond, the router knows that nobody is interested anymore in the multicast traffic and it will then remove the entry once the timer exceeds. Disadvantage is router will keep forwarding multicast traffic even no one interested.

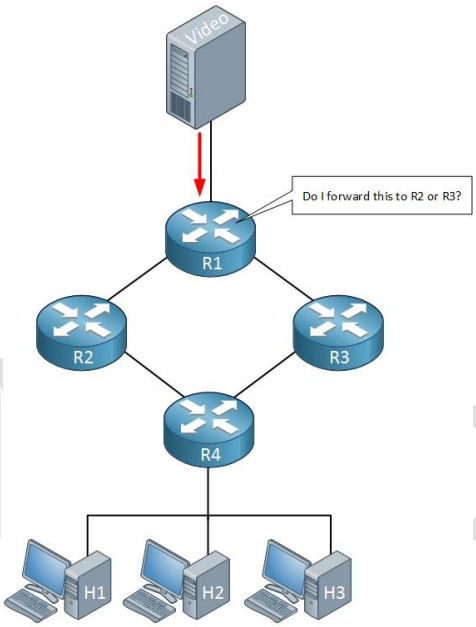

IGMP Version 2:

- In multicasting the IGMP version 2 is the enhanced version of the IGMP version 1.
- IGMP V1, hosts stop listening to multicast group address but never report to router.
- IGMP Version 2 is very similar to Version 1 but due to new features it's more efficient.
- IGMP Version 2 because of new features it leaving groups has become much faster.
- IGMP Version 2 is an enhancement of IGMP Version 1.
- IGMP Version 2 It was introduced in RFC 2236 in 1997. IGMP Version 2 adds the following features:

Leave group messages: When a host no longer wants to listen to a multicast group address, it can send a leave group message to the router. This allows the router to quickly update its multicast routing table and stop forwarding multicast packets to the host.

Group-specific membership queries: The router can send a group-specific membership query to a specific multicast group address. This allows the router to quickly determine if any hosts are still members of the group.

Maximum Response Time (MRT) field: The MRT field in the membership query message specifies the maximum amount of time that a host has to respond to the query. This allows the router to quickly determine if there are any hosts that are not responding to the query.

Querier election process: A new querier election process is used to determine which router is the IGMP querier for the network. This process helps to ensure that there is only one router sending membership queries on the network.

IGMP Version 2 is backward compatible with IGMP Version 1. This means that IGMP Version 1 hosts can coexist with IGMP Version 2 hosts on the same network.

IGMP Version 2 is the most widely used version of IGMP. It is supported by most operating systems and routers.

Feature	IGMP Version 1	IGMP Version 2
Leave group messages	No	Yes
Group-specific membership queries	No	Yes
Maximum Response Time (MRT) field	No	Yes
Querier election process	Simple	More complex
Backward compatibility	No	Yes
Widespread support	No	Yes

IGMP Version 3:

- Internet Group Management Protocol Version 3 adds support for “Source Altering”.
- IGMP v1 & v2 allow hosts to join multicast groups but don’t check source of the traffic.
- In IGMPv2, when a receiver sends a membership report to join the multicast group.
- IGMP V2 does not specify which source it would like to receive multicast traffic from.
- Internet Group Management Protocol V3 adds support for multicast source filtering.
- Giving receivers the capability to pick source they wish to accept multicast traffic from.
- It supports all IGMPv2’s IGMP message types & is backward compatible with IGMPv2.

IGMP snooping:

IGMP snooping allows us to constrain our multicast traffic. As the name implies, this is done by listening to IGMP traffic between the router and hosts: IGMP snooping, is the most widely used method and works by examining IGMP joins sent by receivers and maintaining a table of interfaces to IGMP joins. When the switch receives a multicast, frame destined for a multicast group, it forwards the packet only out the ports where IGMP joins were received for that specific multicast group.

Switches listen to IGMP messages and learn on which interfaces they have to forward multicast traffic. Without IGMP snooping, switches will flood multicast traffic everywhere, treating like broadcast traffic.

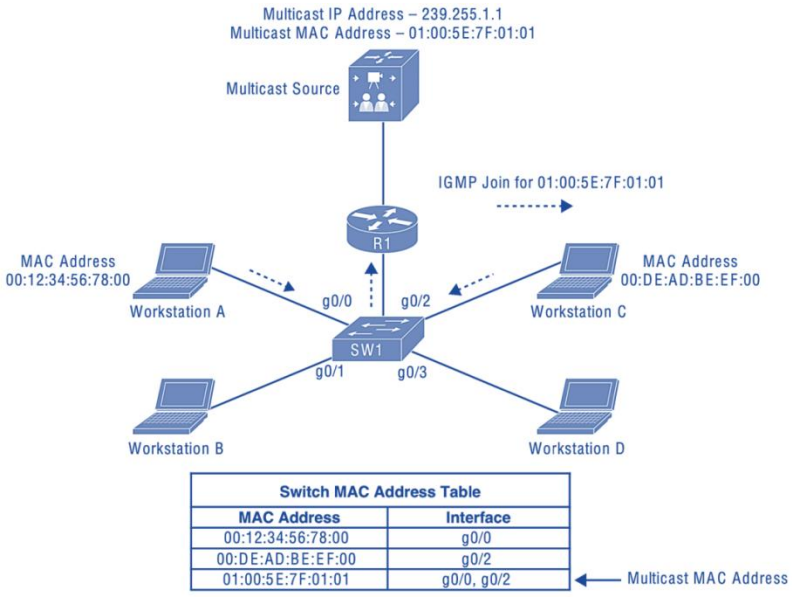

Multicast Routing Protocols:

To route multicast traffic, we need to use a multicast routing protocol. A multicast routing protocol is necessary to route the multicast traffic throughout the network so that routers can locate and request multicast streams from other router.

PIM (Protocol Independent Multicast):

A multicast routing protocol is necessary to route the multicast traffic throughout the network so that routers can locate and request multicast streams from other routers. The only multicast routing protocol that is fully supported on Cisco IOS devices is PIM (Protocol Independent Multicast).

Protocol Independent Multicast (PIM) is used between routers so that they can track which multicast packets to forward to each other and to their directly connected LANs. **Typically, either PIM Sparse Mode or PIM Dense Mode** will be used throughout a multicast domain. PIM is a multicast routing protocol that routes multicast traffic between network segments. PIM can use any of the unicast routing protocols to identify the path between the source and receivers. Multicast routers create distribution trees that define the path that IP multicast traffic follows through the network to reach the receivers.

Dense Mode:

- Dense mode multicast routing protocols are used for networks where most subnets in your network should receive the multicast traffic. When a router receives the multicast traffic, it will flood it on all of its interfaces except the interface where it received the multicast traffic on.
- PIM (Protocol Independent Multicast) is the most popular multicast routing protocol. Dense mode floods multicast traffic until a router asks you to stop.

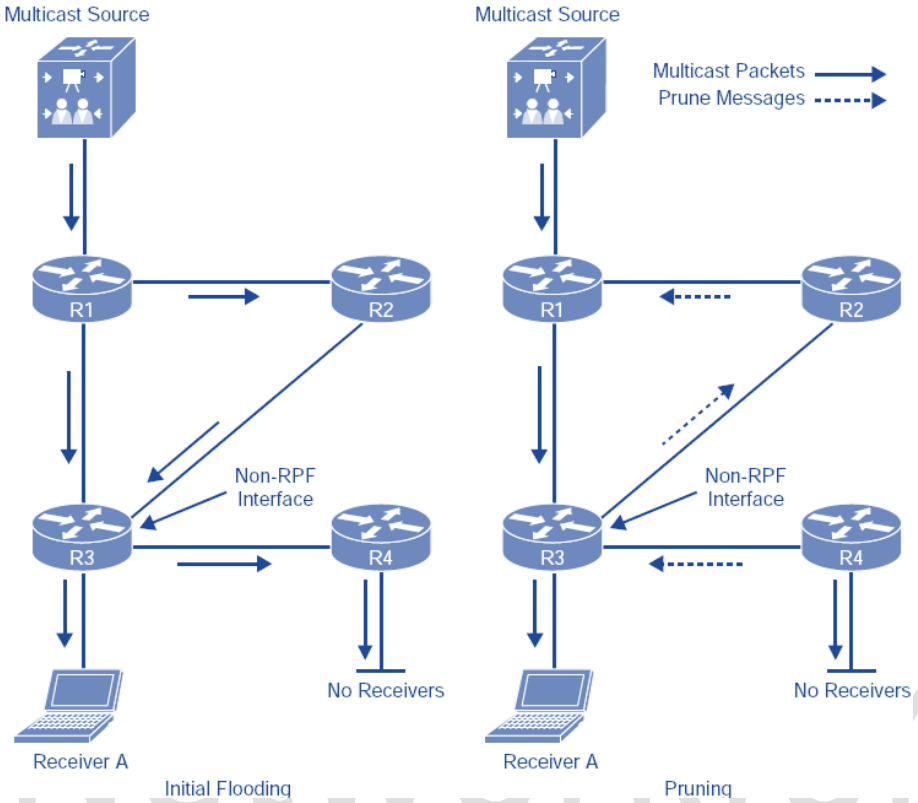

Sparse Mode:

- When you only have a few receivers on your network then yes, you will be wasting a lot of bandwidth and resources on your routers. The alternative is sparse mode which is far more efficient.
- Sparse mode multicast routing protocols only forward the multicast traffic when another router requests it.
- It's the complete opposite of dense mode. Sparse mode sends multicast traffic only when a router requests it.

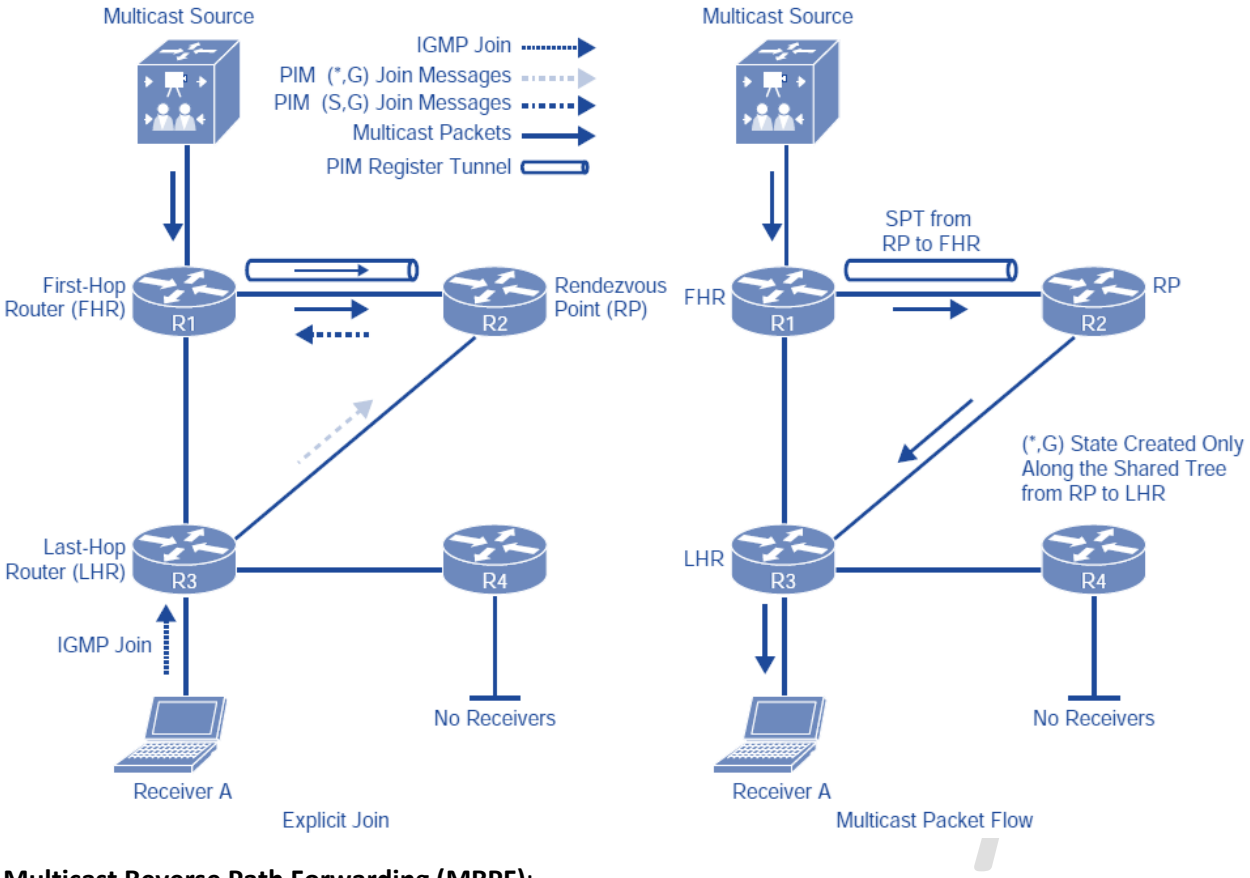

Multicast Reverse Path Forwarding (MRPF):

- It is a **technique used in multicast routing to prevent loops from occurring**.
- **RPF is important because it helps to prevent loops from occurring in multicast networks**.
- Loops can occur in multicast networks when a packet is forwarded back to the router that it came from. This can cause the packet to be forwarded around the network forever, which can consume a lot of bandwidth and resources.
- RPF (Reverse Path Forwarding) is a technique used in multicast routing to prevent routing loops.
- **It works by checking the source IP address of a multicast packet against the unicast routing table.**
- **If the source IP address is found in the routing table, the router will check the next hop and outgoing interface that is used to reach the source.**
- **If the multicast packet was received on the same interface that is used to reach the source, the RPF check succeeds.**
- **If the multicast packet was received on a different interface, the RPF check fails and the packet is discarded.**

Here is an example of how RPF works:

- A router receives a multicast packet with a source IP address of 192.168.1.1.
- The router looks up the source IP address in the unicast routing table.
- The routing table shows that the next hop for 192.168.1.1 is 192.168.2.1.
- The router also sees that the outgoing interface for 192.168.2.1 is Ethernet0/1.
- The multicast packet was received on Ethernet0/0.
- Since the multicast packet was not received on the same interface that is used to reach the source, the RPF check fails.
- The packet is discarded.
- RPF is an important security mechanism that helps to prevent multicast routing loops.
- By discarding packets that fail the RPF check, routers can prevent malicious actors from sending multicast packets that could cause damage to the network.

How Does RPF Work?

- When the router receives the multicast packet, RPF will check the routing table and check which egress interface the router would use if it were to send traffic back to the multicast sender.
- If the interface matches the interface the multicast packet has just been received on, the packet is accepted, otherwise, it is dropped.
- Below shows an example. Because R4's routing table states that the multicast sender is located via gi0/1 traffic on gi0/2 is dropped.

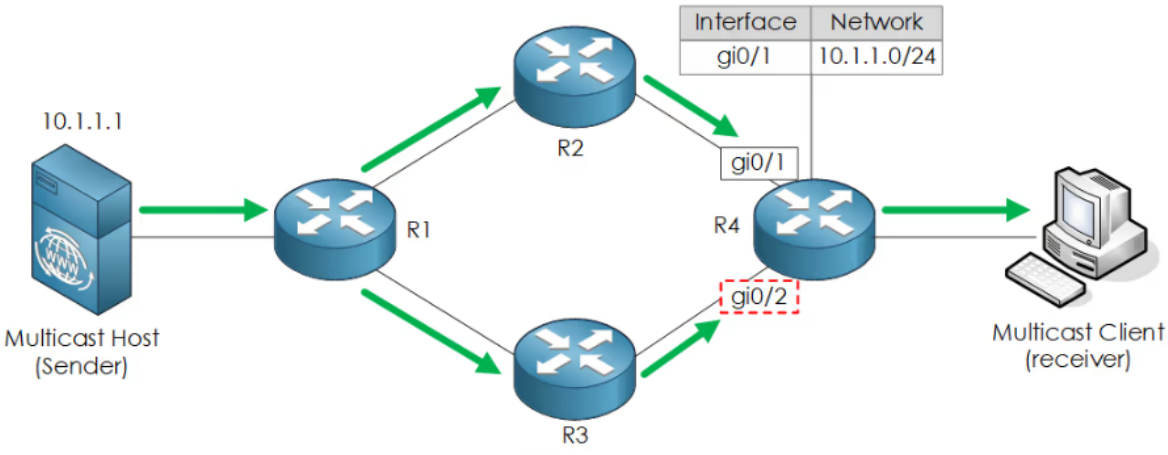

- **By default, RPF is enabled on Cisco routers and switches.** However, there are some cases where you may need to disable RPF. For example, if you are using a router as a DHCP server, you will need to disable RPF on the DHCP interface.

- This is because the DHCP server will send multicast packets to all of the interfaces on the router, and RPF will discard these packets.
- To disable RPF on a Cisco router or switch, you can use the following command:

ip pim rpfilter disable

```
Router# configure terminal
Router(config)# ip pim rpfilter disable
Router(config)# ip pim restart
Router(config)# end
```

- This command will disable RPF on all interfaces on the router or switch. If you only want to disable RPF on a specific interface, you can use the following command:

ip pim rpfilter disable interface interface_name

```
Router# configure terminal
Router(config)# ip pim rpfilter disable interface Ethernet0/0
Router(config)# ip pim restart
Router(config)# end
```


- Once you have disabled RPF, you will need to restart the multicast routing process on the router or switch. To do this, you can use the following command:

ip pim restart