



Advanced Routing

A Network Simulation & Emulation Software
By



The information contained in this document represents the current view of TETCOS LLP on the issues discussed as of the date of publication. Because TETCOS LLP must respond to changing market conditions, it should not be interpreted to be a commitment on the part of TETCOS LLP, and TETCOS LLP cannot guarantee the accuracy of any information presented after the date of publication.

This manual is for informational purposes only.

The publisher has taken care in the preparation of this document but makes no expressed or implied warranty of any kind and assumes no responsibility for errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of the use of the information contained herein.

Warning! DO NOT COPY

Copyright in the whole and every part of this manual belongs to TETCOS LLP and may not be used, sold, transferred, copied or reproduced in whole or in part in any manner or in any media to any person, without the prior written consent of TETCOS LLP. If you use this manual you do so at your own risk and on the understanding that TETCOS LLP shall not be liable for any loss or damage of any kind.

TETCOS LLP may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from TETCOS LLP, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property. Unless otherwise noted, the example companies, organizations, products, domain names, e-mail addresses, logos, people, places, and events depicted herein are fictitious, and no association with any real company, organization, product, domain name, email address, logo, person, place, or event is intended or should be inferred.

Rev 13.1 (V), Dec 2021, TETCOS LLP. All rights reserved.

All trademarks are property of their respective owner.

Contact us at

TETCOS LLP

214, 39th A Cross, 7th Main, 5th Block Jayanagar,

Bangalore - 560 041, Karnataka, INDIA.

Phone: +91 80 26630624

E-Mail: sales@tetcos.com

Visit: www.tetcos.com

Table of Contents

| | | |
|----------|---|-----------|
| 1 | Multicast Routing Protocols..... | 4 |
| 1.1 | IGMP..... | 4 |
| 1.1.1 | Introduction..... | 4 |
| 1.2 | PIM..... | 8 |
| 1.2.1 | Introduction..... | 8 |
| 2 | Access Control Lists (ACLs)..... | 12 |
| 2.1 | Introduction | 12 |
| 3 | Virtual LAN (VLAN) | 13 |
| 3.1 | Introduction | 13 |
| 3.1.1 | When do we need a VLAN?..... | 14 |
| 3.1.2 | Understanding Access and Trunk Links..... | 15 |
| 4 | Public IP Address & NAT (Network Address Translation)..... | 17 |
| 4.1 | Introduction | 17 |
| 4.1.1 | Public Address..... | 17 |
| 4.1.2 | Private Address | 17 |
| 4.1.3 | Network address translation (NAT)..... | 18 |
| 5 | Featured Examples | 20 |
| 5.1 | Multicast Routing Protocols Example | 20 |
| 5.1.1 | IGMP Example | 20 |
| 5.1.2 | Results | 24 |
| 5.2 | Access Control Lists (ACLs) Examples | 29 |
| 5.2.1 | ACL Example..... | 29 |
| 5.2.2 | Result and Observations..... | 31 |
| 6 | Advanced Routing Experiments in NetSim | 32 |
| 7 | Reference Documents | 32 |
| 8 | Latest FAQs..... | 32 |

1 Multicast Routing Protocols

Note: Multicast routing protocols can be configured and run only if licenses for component 3 (advanced routing) is available

Multicasting is one source sending a packet to multiple destinations. Group formation and management is an integral part of multicasting.

IP Multicast Group Addressing: A multicast group is identified by its multicast group address. Multicast packets are delivered to that multicast group address. Unlike unicast addresses that uniquely identify a single host, multicast IP addresses do not identify a particular host. To receive the data sent to a multicast address, a host must join the group that address identifies. The data is sent to the multicast address and received by all the hosts that have joined the group indicating that they wish to receive traffic sent to that group. The multicast group address is assigned to a group at the source.

IP Class D Addresses: IP multicast addresses have been assigned to the IPv4 Class D address space by IANA. The high-order four bits of a Class D address are 1110. Therefore, host group addresses can be in the range 224.0.0.0 to 239.255.255.255. A multicast address is chosen at the source (sender) for the receivers in a multicast group.

NetSim supports the following protocols to implement IP multicast routing:

- **IGMP** is used between hosts on a LAN and the routers on that LAN to track the multicast groups of which hosts are members.
- **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.

1.1 IGMP

1.1.1 Introduction

About Multicasting

Multicasting is a data delivery method where one sender sends data to thousands of recipients across a routed network. Multicasting is controlled broadcasting; the sender transmits data to specific recipients only.

With IP multicasting, a host sends packets to a multicast group of hosts anywhere within the IP network by using a special form of IP address called the IP multicast group address. A multicast group is made of an arbitrary number of hosts who join a group to receive packets from the source. To ensure that a host receives data, the host must join the multicast group to which the sender is sending data.

Note: You can configure and simulate multicast routing protocols such as IGMP and PIM, only if you have licenses for component 3 (advanced routing).

About IGMP

The Internet Group Management Protocol (IGMP) is a communication protocol that hosts and adjacent multicast routers on IPv4 networks use, to establish and manage the membership of hosts and routing devices in multicast groups. Hosts and multicast routers use IGMP as follows:

- The hosts use IGMP to report their multicast group memberships to neighboring multicast routers.
- The multicast routers use IGMP to know the members in multicast groups, for every physical network the multicast router is connected.

The multicast routers maintain a list of multicast group memberships for every network to which the multicast routers are connected, and a timer for each membership.

The messages that IGMP uses are encapsulated in IP datagrams, with an IP protocol number of 2. All IGMP messages are sent with an IP TTL of 1 and contain the IP Router Alert option in their IP header. All IGMP messages sent between a host and the multicast router use the following format:

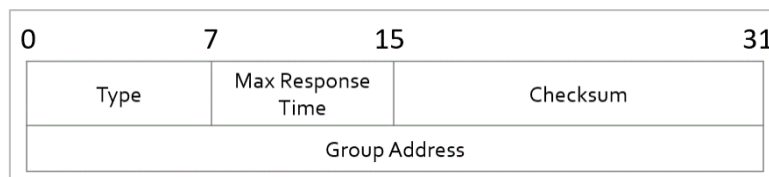


Figure 1-1: IGMP datagrams

- **Type:** There are three types of IGMP messages that hosts and multicast routers exchange, when they interact:

1. 0x11: Membership Query

There are two sub-types of Membership Query messages:

- **General Query:** The multicast router sends a General Query to all hosts to collect and update multicast group membership information for the hosts on all networks to which the multicast router is connected.
- **Group-Specific Query:** The multicast router sends Group-Specific Query to the multicast group from which it received a leave group message, to find out if other hosts in the group require multicast data.

2. 0x16: Version 2 Membership Report

Version 2 Membership Report is a message that a host sends to all other hosts in the group or all hosts on the network, in response to a General Query or a Group-Specific Query message from the multicast router.

3. 0x17: Leave Group (Not available with NetSim)

Hosts use the Leave Group message to tell the multicast router that they intend to leave the group.

- **Max Response Time:** Maximum Response Time is a random-value delay timer which a host sets, for the host to send a Version 2 Membership Report to other hosts in the group, after the host receives a Group-Specific Query message.
- **Checksum:** The Checksum is the 16-bit one's complement of the one's complement sum of the whole IGMP message (the entire IP payload).
- **Group Address:** The multicast router sets the Group Address to zero when it sends a General Query and sets to the Group Address to the address of the multicast group when it sends a Group-Specific Query.

How IGMP Works

If a router has multiple physical interfaces on a single network, IGMP runs only on one of physical interfaces. Hosts, on the other hand, need to use all interfaces that have memberships associated with them.

For every network the multicast router is connected to, the multicast router performs one of the following roles: Querier or Non-Querier. There is normally only one Querier per physical network.

At startup, all multicast routers start as a Querier on every network to which the multicast routers are connected. If a multicast router hears a Query message from another multicast router with a lower IP address, the first multicast router must perform the role of a Non-Querier on the network which has the multicast router with a lower IP address. If a multicast router does not hear a Query message from another multicast router for a time duration defined by the Other Querier Present Interval, the multicast router persists with the role of the Querier.

Now, the multicast router sends one of the two Membership Query messages:

- **General Query to all hosts to collect and update multicast group membership information:** The multicast router sends the General Query to the all-systems multicast group (224.0.0.1), with a Group Address field set to 0, and with a Max Response Time for the Query Response Interval.

When a host receives the General Query, the host sets the delay timers for every group (excluding the all-systems group) to which the host belongs, on the interface on which it received the query.

The host sets every timer to a different random value, by using the highest clock granularity available on the host, and by choosing a value between 0 and the Max Response Time. The Max Response Time is specified in the General Query packet.

- **Group-Specific Query to the multicast group from which it received a leave group message:** The multicast router sends the Group-Specific Query to the multicast group from which it received a leave group message, and with a Max Response Time for the Query Response Interval. This helps the multicast router to learn if there are other members on the group and if the group needs multicast data.

When a host receives the Group-Specific Query, the host sets the delay timers for every group to which the host belongs, on the interface on which it received the query.

The host sets every timer to a different random value by choosing a value between 0 and the Max Response Time. The Max Response Time is specified in the Group-Specific Query packet.

If the delay timer for a group has started, the host resets the delay time to a random value only if the requested Max Response Time is less than the time left in the active delay timer.

When a group's delay timer expires, the host multicasts a Version 2 Membership Report to other hosts in the group, with an IP TTL of 1. If the host receives a Version 2 Membership Report (version 1 or 2) from another host in the same group, when the host's timer is active, the host stops the timer for the group from which it received the report. The host also does send a report to other hosts, to avoid duplicate reports and conserve the bandwidth on the network.

When a multicast router receives a Version 2 Membership Report, it does the following:

- Adds the multicast group from which it received the Version 2 Membership Report, to the list of multicast group memberships on the network on which it received the Version 2 Membership Report.
- Sets the timer for the membership to the Group Membership Interval.
- Refreshes the timer, when the multicast router receives another Version 2 Membership Report from the same group.

If the multicast router does not receive any Version 2 Membership Reports from a multicast group before the Group Membership Interval timer expires, the multicast router assumes that the group has no members and that it need not forward multicast data for that group.

The multicast router may also receive an unsolicited Version 2 Membership Report from a host when the hosts intends to join a multicast group.

1.2 PIM

1.2.1 Introduction

Protocol-Independent Multicast or PIM is a group of multicast routing protocols for Internet Protocol (IP) networks. PIM distributes data in one-to-many and many-to-many multicast modes over a LAN, WAN or the Internet. PIM builds Multicast Distribution Tree (MDT) loop-free trees to enable multicast data distribution over a network.

PIM is termed protocol-independent because PIM does not include its own topology discovery mechanism; PIM uses routing information available from other routing protocols such as Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF), and static routes.

PIM also does not build its own routing tables. PIM uses the unicast routing table that IGP creates, to create a loop free MDT and uses the unicast routing table to perform the reverse path forwarding (RPF). Unlike other routing protocols, PIM does not send and receive routing updates between routers.

In a PIM-enabled network, a Rendezvous Point (RP) router is the point where other routers in the PIM protocol's domain exchange information. All routers in the PIM protocol's domain must provide a mapping to the RP router. In a PIM enabled network, only the RP router knows the active sources for the entire PIM protocol's domain. The other routers just need to know how to reach the RP router. This way, the RP router matches the receivers with the sources in the PIM protocol's domain.

The RP router is downstream from the source and forms one end of the Shortest Path Tree (SPT). The RP router is upstream from the receiver and forms one end of the Rendezvous Point Tree (RPT).

The following figure illustrates a PIM-enabled network with the routers, source node, and the destination node.

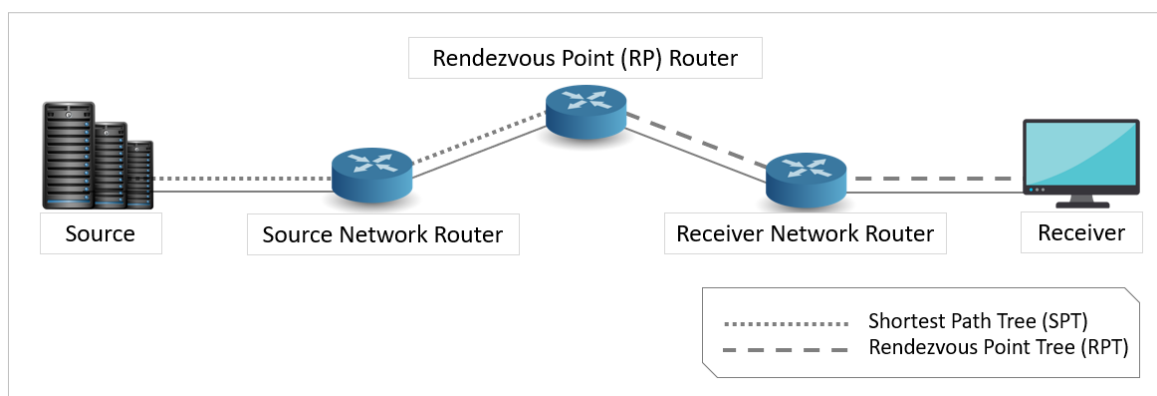


Figure 1-2: Illustrates a PIM-enabled network with the routers

To configure PIM in NetSim

Create a network as shown below Figure 1-3.

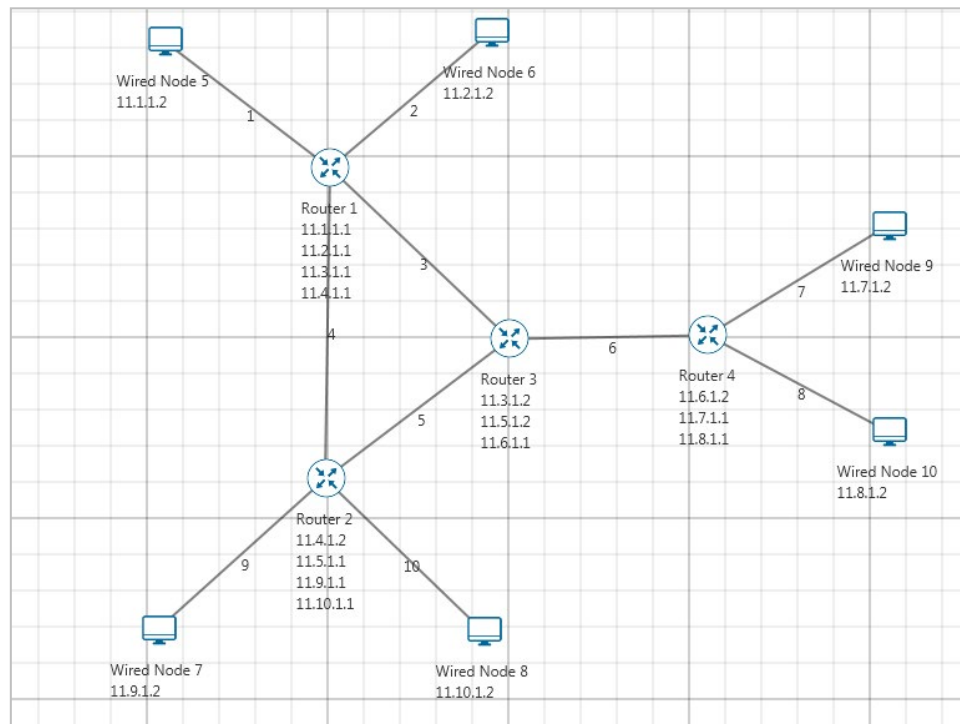


Figure 1-3: Network Topology in this experiment

Set PIM status as TRUE in all routers as shown below:

Set IGMP status as true for all devices.

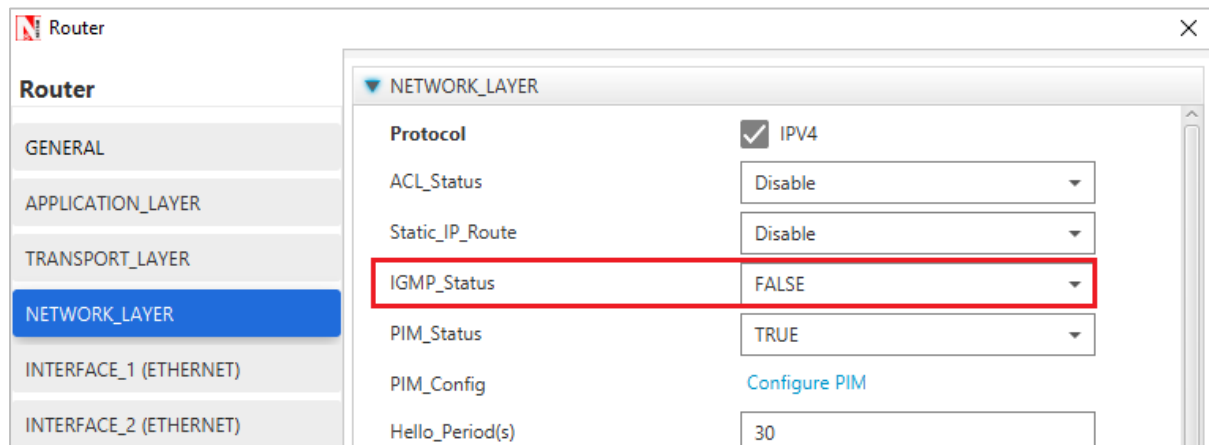


Figure 1-4: Network layer window

Configure the PIM properties as per the below screenshot and click on Add.

Figure 1-5 shows the 'Configure PIM' window. The 'Add' button is highlighted with a red box. The table below the buttons is empty.

| Multicast IP | RP ID | Router Name IF |
|---------------------|-------|----------------|
| No content in table | | |

Figure 1-5: PIM Configuration window

Then click on Accept.

Figure 1-6 shows the 'Configure PIM' window after the configuration is done. The 'Add' button is no longer highlighted. The table now contains one row with the values: Multicast IP (239.12.14.5), RP ID (11.1.1.2), and Router Name IF (Router_3_WAN1).

| Multicast IP | RP ID | Router Name IF |
|--------------|----------|----------------|
| 239.12.14.5 | 11.1.1.2 | Router_3_WAN1 |

Figure 1-6: Once PIM Configuration Done

Configure the same PIM properties for all routers in the network.

Application Properties:

Set the application properties as per the screenshot below – Multicast application with source 5 and destinations 6, 7, 8, 9, 10

Configure Application

Application + -

Application1

APPLICATION

Application_Method: MULTICAST

Application_Type: CBR

Application_ID: 1

Application_Name: App1_CBR

Source_Count: 1

Source_ID: 5

Destination_Count: 5

Destination_ID: 6,7,8,9,10

Multicast_Dest_Address: 239.12.14.5

Start_Time(s): 5

End_Time(s): 100000

Src_to_Dest: Don't show line

Encryption: NONE

Random_Startup: FALSE

OK Reset

Figure 1-7: Application properties window

Set IGMP_Status to TRUE in all wired nodes since we are running multicast application.

Enable packet Trace, Plots and run simulation for 10s. Open Packet trace and filter PACKET_ID to 1. Users can observe there is no formation of loops between source and destinations.

| PACKET_ID | SEGMENT | PACKET | CONTR | SOURCE | DESTINATION_ID | TRANSMITTE | RECEIVER | APP_LA | TRX_LA | NW_LA | MAC_L | PHY_LA | PHY_LA | PHY_LA | APP_LA | TRX |
|-----------|---------|--------|----------|--------|---------------------------------|------------|----------|---------|---------|---------|---------|---------|---------|---------|---------|------|
| 1 | 0 | CBR | APP1_CBR | NODE-5 | NODE-6; NODE-7; NODE-8; NODE-9; | NODE-5 | ROUTER-1 | 5020000 | 5020000 | 5020000 | 5020000 | 5020126 | 5020126 | 5020121 | 5020126 | 1460 |
| 1 | 0 | CBR | APP1_CBR | NODE-5 | NODE-6; NODE-7; NODE-8; NODE-9; | ROUTER-1 | ROUTER-2 | 5020000 | 5020000 | 5020126 | 5020126 | 5020126 | 5020126 | 5020245 | 5020250 | 1460 |
| 1 | 0 | CBR | APP1_CBR | NODE-5 | NODE-6; NODE-7; NODE-8; NODE-9; | ROUTER-1 | ROUTER-3 | 5020000 | 5020000 | 5020126 | 5020126 | 5020126 | 5020126 | 5020245 | 5020250 | 1460 |
| 1 | 0 | CBR | APP1_CBR | NODE-5 | NODE-6; NODE-7; NODE-8; NODE-9; | ROUTER-1 | NODE-6 | 5020000 | 5020000 | 5020126 | 5020126 | 5020126 | 5020126 | 5020247 | 5020252 | 1460 |
| 1 | 0 | CBR | APP1_CBR | NODE-5 | NODE-6; NODE-7; NODE-8; NODE-9; | ROUTER-2 | ROUTER-3 | 5020000 | 5020000 | 5020250 | 5020250 | 5020250 | 5020250 | 5020369 | 5020374 | 1460 |
| 1 | 0 | CBR | APP1_CBR | NODE-5 | NODE-6; NODE-7; NODE-8; NODE-9; | ROUTER-3 | ROUTER-2 | 5020000 | 5020000 | 5020250 | 5020250 | 5020250 | 5020250 | 5020369 | 5020374 | 1460 |
| 1 | 0 | CBR | APP1_CBR | NODE-5 | NODE-6; NODE-7; NODE-8; NODE-9; | ROUTER-3 | ROUTER-4 | 5020000 | 5020000 | 5020250 | 5020250 | 5020250 | 5020250 | 5020369 | 5020374 | 1460 |
| 1 | 0 | CBR | APP1_CBR | NODE-5 | NODE-6; NODE-7; NODE-8; NODE-9; | ROUTER-2 | NODE-7 | 5020000 | 5020000 | 5020250 | 5020250 | 5020250 | 5020250 | 5020371 | 5020376 | 1460 |
| 1 | 0 | CBR | APP1_CBR | NODE-5 | NODE-6; NODE-7; NODE-8; NODE-9; | ROUTER-2 | NODE-8 | 5020000 | 5020000 | 5020250 | 5020250 | 5020250 | 5020250 | 5020371 | 5020376 | 1460 |
| 1 | 0 | CBR | APP1_CBR | NODE-5 | NODE-6; NODE-7; NODE-8; NODE-9; | ROUTER-4 | NODE-9 | 5020000 | 5020000 | 5020374 | 5020374 | 5020374 | 5020374 | 5020495 | 5020500 | 1460 |
| 1 | 0 | CBR | APP1_CBR | NODE-5 | NODE-6; NODE-7; NODE-8; NODE-9; | ROUTER-4 | NODE-10 | 5020000 | 5020000 | 5020374 | 5020374 | 5020374 | 5020374 | 5020495 | 5020500 | 1460 |

Figure 1-8: Packet Trace

2 Access Control Lists (ACLs)

2.1 Introduction

Access Control Lists (ACLs) are filters that routers use to control which, routing updates, or packets are permitted or denied, in or out, of a network. An ACL contains a sequential list of “permit” or deny statements (rules) that apply to IP packets originating or destined to hosts, IP addresses and upper-layer IP protocols.

An ACL tells the router what types of packets to: **permit** or **deny**. The router using the ACL does the following when it finds packets inbound to or outbound from a network:

- If the router finds packets inbound or outbound categorized against the permit statements, the router forwards the packets to the next hop in the network.
- If the router finds packets inbound or outbound categorized against the deny statements, the router blocks and drops the packets at the router’s interface. The packets cannot reach the intended destination host or IP address.
- ACLs control traffic in one direction at a time, on an interface. To allow inbound and outbound traffic from a host, IP address, or for a protocol, you must create two ACLs, one for each direction, one for inbound and one for outbound traffic.
- The precedence of the ACL commands is from top to bottom.

For example, If ACL is configured in Router as follows:

```
PERMIT OUTBOUND TCP ANY ANY 0 0 3
```

```
PERMIT INBOUND TCP ANY ANY 0 0 3
```

```
DENY BOTH ANY ANY ANY 0 0 3
```

Then, the Permit statements will over-ride the deny statements. That is, Outbound TCP packets from Router through interface 3 will be permitted first, after that, the Inbound TCP packets to Router through interface 3 will be permitted. All other packets through the third interface of Router will be denied in both directions.

3 Virtual LAN (VLAN)

3.1 Introduction

VLAN is called as virtual local area network, used in Switches and it operates at layer2 and Layer3. A VLAN, is a group of hosts which communicate as if they were attached to the same broadcast domain, regardless of their physical location.

For example, all workstations and servers used by a particular workgroup team can be connected to the same VLAN, regardless of their physical connections to the network or the fact that they might be intermingled with other teams. VLANs have the same attributes as physical LANs, but you can group end stations even if they are not physically located on the same LAN segment.

A VLAN behaves just like a LAN in all respects but with additional flexibility. By using VLAN technology, it is possible to subdivide a single physical switch into several logical switches. VLANs are implemented by using the appropriate switch configuration commands to create the VLANs and assign specific switch interfaces to the desired VLAN.

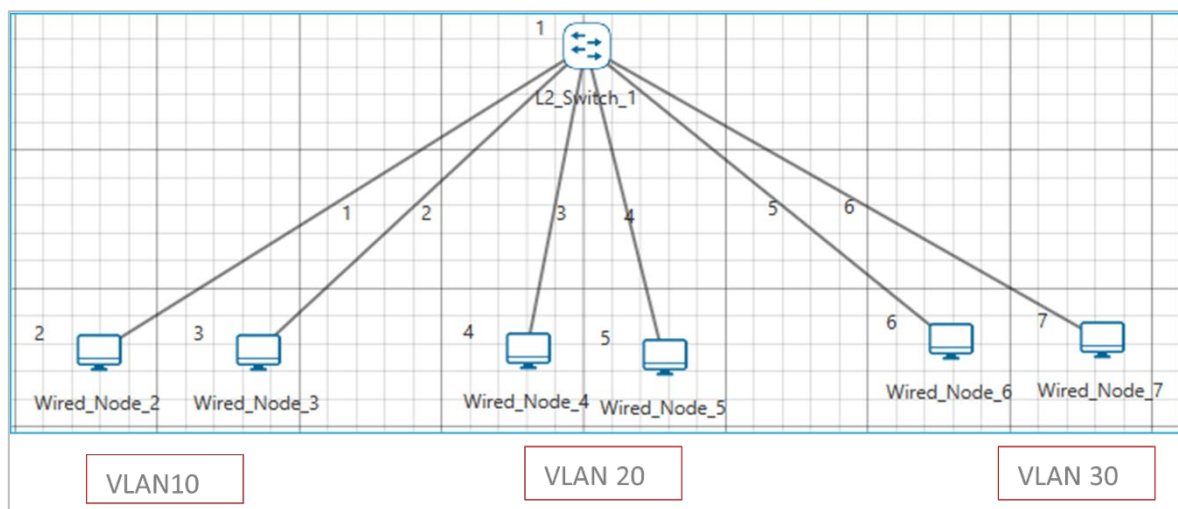


Figure 3-1: Virtual local area network (VLAN)

Switches implement VLANs by adding a VLAN tag to the Ethernet frames as they enter the switch. The VLAN tag contains the VLAN ID and other information, which is determined by the interface from which the frame enters the switch. The switch uses VLAN tags to ensure that each Ethernet frame is confined to the VLAN to which it belongs based on the VLAN ID contained in the VLAN tag. The VLAN tags are removed as the frames exit the switch on the way to their destination.

Any port can belong to a VLAN, and unicast, broadcast, and multicast packets are forwarded and flooded only to end stations in that VLAN. Each VLAN is considered a logical network.

Packets destined for stations that do not belong to the VLAN must be forwarded through a router. In the below screenshot, the stations in the development department are assigned to one VLAN, the stations in the marketing department are assigned to another VLAN, and the stations in the testing department are assigned to another VLAN.

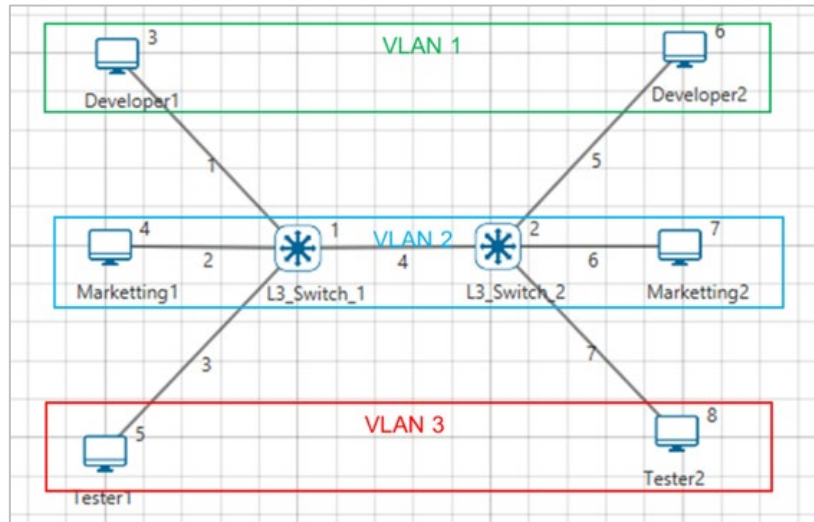


Figure 3-2: Hosts in one VLAN need to communicate with hosts in another VLAN

3.1.1 When do we need a VLAN?

You need to consider using VLAN's in any of the following situations:

- You have more than 200 devices on your LAN.
- You have a lot of broadcast traffic on your LAN.
- Groups of users need more security are being slowed down by too many broadcasts.
- Groups of users need to be on the same broadcast domain because they are running same applications or just make a single switch into multiple virtual switches.

3.1.1.1 VLAN ID

VLANs are identified by a VLAN ID (a number between 0 – 4095), with the default VLAN on any network being VLAN 2. Each port on a switch or router can be assigned to be a member of a VLAN (i.e., to allow receiving and sending traffic on that VLAN).

For example: On a switch, traffic that is sent to a port that is a member of VLAN2, may be forwarded to any other VLAN2 port on the switch, and it can also travel across a trunk port (connections between switches) to another switch and forwarded to all VLAN2 ports on that switch. Traffic will not be forwarded to ports that are on a different VLAN ID.

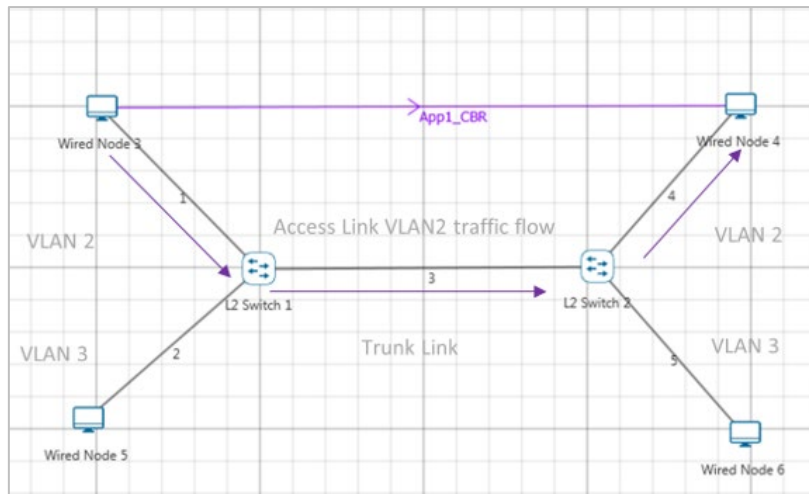


Figure 3-3: Understanding Access and Trunk Links

3.1.2 Understanding Access and Trunk Links

The links connecting the end devices are called access links. These are the links usually carrying the Data VLAN information

The link between the switches is called trunk link. It carries packets from all the VLANs.

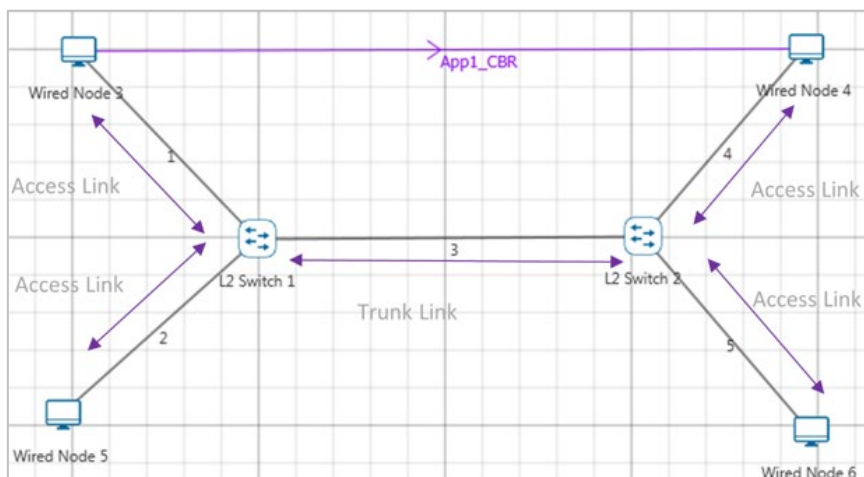


Figure 3-4: Understanding Access and Trunk Links

3.1.2.1 Access Link

Access link connection is the connection where switch port is connected with a device that has a standardized Ethernet NIC. Standard NIC only understand IEEE 802.3 or Ethernet II frames. Access link connection can only be assigned with single VLAN. That means all devices connected to this port will be in same broadcast domain.

For example, twenty users are connected to a hub, and we connect that hub with an access link port on switch, then all of these users belong to same VLAN. If we want to keep ten users in another VLAN, then we need to plug in those ten users to another hub and then connect it with another access link port on switch.

3.1.2.2 Trunk Link

Trunk link connection is the connection where switch port is connected with a device that is capable to understand multiple VLANs. Usually trunk link connection is used to connect two switches. Trunking allows us to send or receive VLAN information across the network. To support trunking, original Ethernet frame is modified to carry VLAN information.

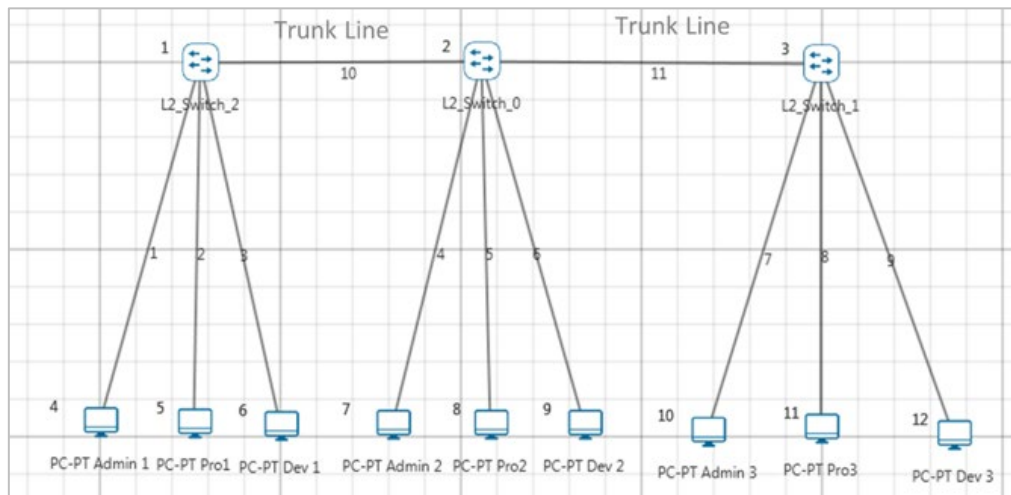


Figure 3-5: Understand multiple VLANs

4 Public IP Address & NAT (Network Address Translation)

4.1 Introduction

4.1.1 Public Address

A public IP address is assigned to every computer that connects to the Internet where each IP is unique. Hence there cannot exist two computers with the same public IP address all over the Internet. This addressing scheme makes it possible for the computers to “find each other” online and exchange information. User has no control over the IP address (public) that is assigned to the computer. The public IP address is assigned to the computer by the Internet Service Provider as soon as the computer is connected to the Internet gateway.

4.1.2 Private Address

An IP address is considered private if the IP number falls within one of the IP address ranges reserved for private networks such as a Local Area Network (LAN). The Internet Assigned Numbers Authority (IANA) has reserved the following three blocks of the IP address space for private networks (local networks):

| Class | Starting IP address | Ending IP address | No. of hosts |
|-------|---------------------|-------------------|--------------|
| A | 10.0.0.0 | 10.255.255.255 | 16,777,216 |
| B | 172.16.0.0 | 172.31.255.255 | 1,048,576 |
| C | 192.168.0.0 | 192.168.255.255 | 65,536 |

Table 4-1: Private IP address table

Private IP addresses are used for numbering the computers in a private network including home, school and business LANs in airports and hotels which makes it possible for the computers in the network to communicate with each other. For example, if a network A consists of 30 computers each of them can be given an IP starting from **192.168.0.1** to **192.168.0.30**.

Devices with private IP addresses cannot connect directly to the Internet. Likewise, computers outside the local network cannot connect directly to a device with a private IP. It is possible to interconnect two private networks with the help of a router or a similar device that supports Network Address Translation.

If the private network is connected to the Internet (through an Internet connection via ISP) then each computer will have a private IP as well as a public IP. Private IP is used for communication within the network whereas the public IP is used for communication over the Internet.

4.1.3 Network address translation (NAT)

NAT (Network Address Translation or Network Address Translator) is the virtualization of Internet Protocol (IP) addresses. NAT helps to improve security and decrease the number of IP addresses an organization needs.

A device that is configured with NAT will have at least one interface to the inside network and one to the outside network. In a typical environment, NAT is configured at the exit device between a stub domain (inside network) and the backbone. When a packet leaves the domain, NAT translates the locally significant source address into a globally unique address. When a packet enters the domain, NAT translates the globally unique destination address into a local address. If more than one exit point exists, each NAT must have the same translation table. NAT can be configured to advertise to the outside world only one address for the entire network. This ability provides additional security by effectively hiding the entire internal network behind that one address. If NAT cannot allocate an address because it has run out of addresses, it drops the packet and sends an Internet Control Message Protocol (ICMP) host unreachable packet to the destination.

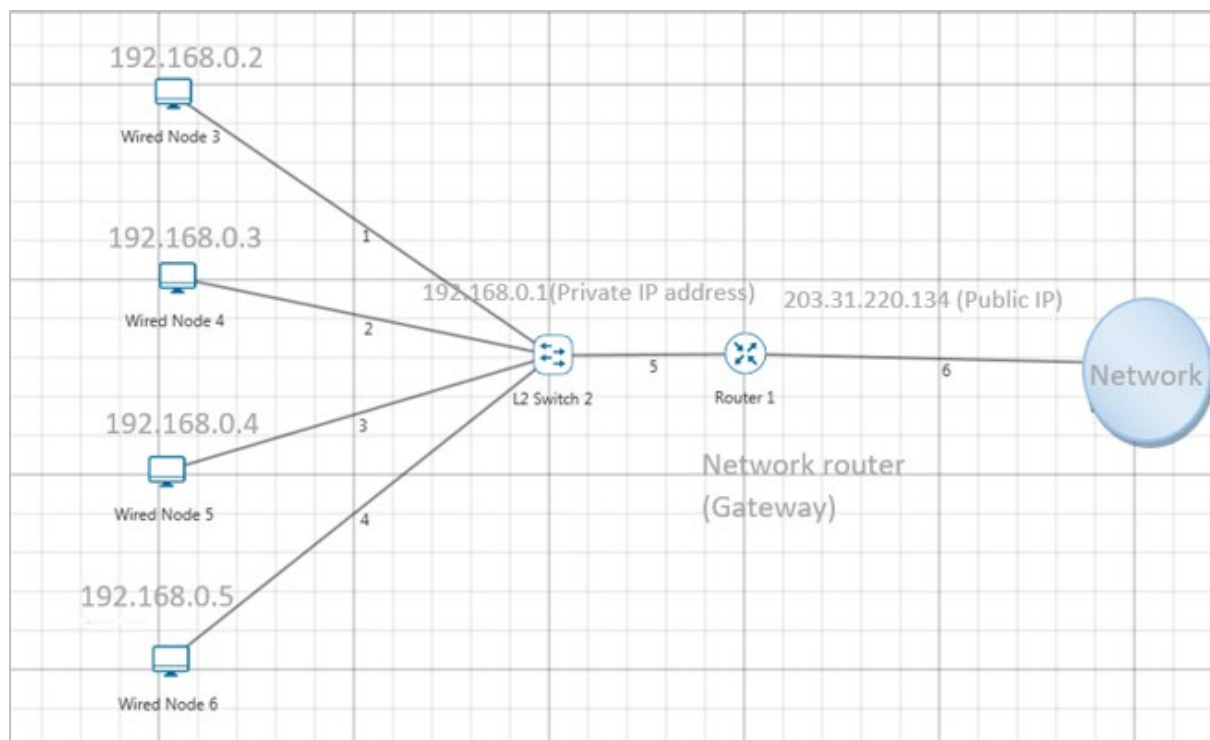


Figure 4-1: NAT implementation

NAT is secure since it hides network from the Internet. All communications from internal private network are handled by the NAT device, which will ensure all the appropriate translations are performed and provide a flawless connection between internal devices and the Internet.

In the above figure, a simple network of 4 hosts and one router that connects this network to the Internet. All hosts in the network have a private Class C IP Address, including the router's private interface (192.168.0.1), while the public interface that's connected to the Internet has a real IP Address (203.31.220.134). This is the IP address the Internet sees as all internal IP addresses are hidden.

5 Featured Examples

5.1 Multicast Routing Protocols Example

5.1.1 IGMP Example

This example explains how IGMP works to multicast data in interconnected networks.

The network modelled consists of:

- A subnet with 4 wired nodes, a multicast router, and a multicast application running on one of the wired nodes.
- IGMP is running on all the wired nodes.
- IGMP is running on the multicast router.
- Only a few nodes receive multicast traffic.

NetSim uses the following defaults for IGMP simulations:

- The multicast destination address is set to 239.12.14.5.
- The IGMP protocol starts only after 1 second into the simulation.
- The multicast application starts only after 5 seconds into the simulation.

Note that NetSim does not support the following in IGMP:

- Leave Group message
- IGMP v1 compatibility

Open NetSim, Select **Examples->Advanced Routing->Multicasting with IGMP** then click on the tile in the middle panel to load the example as shown in Figure 5-1.

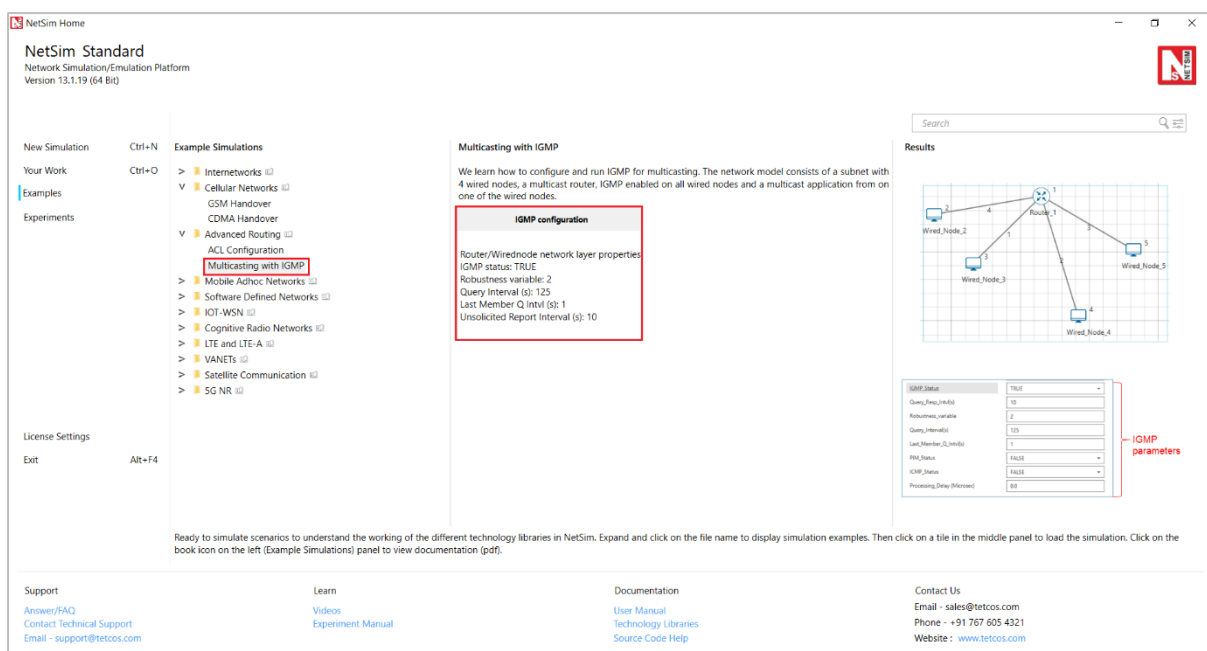


Figure 5-1: List of scenarios for the example of IGMP Configuration

The following network diagram illustrates what the NetSim UI displays when you open the example configuration file for IGMP Figure 5-2.

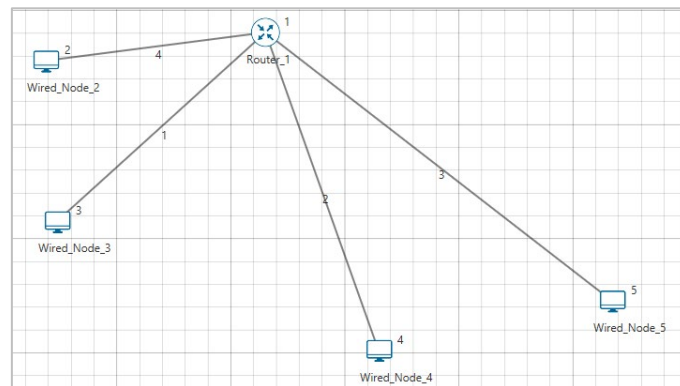


Figure 5-2: Network set up for studying the IGMP Configuration

1. See that by default, NetSim has enabled IGMP on the router, as follows:
 - a. Right-click the router and click Properties.
The Router pop-up window appears.
 - b. Click NETWORK LAYER in the left area.
 - c. IGMP_Status drop-down list is set to TRUE.
 - d. Click OK.
2. See that by default, NetSim has enabled IGMP on a node, as follows:
 - a. Right-click a wired node (say Wired_Node_2) and click Properties.
The Wired node pop-up window appears.
 - b. Click NETWORK LAYER in the left area.
 - c. IGMP_Status drop-down list is set to TRUE.

The Wired node pop-up window displays the following parameters you can configure for IGMP, on the node:

- **Robustness_variable:** The Robustness_variable parameter allows you tune your subnet to a specific number of lost packets (packet loss) in the subnet.
- **Query_Interval(s):** The Query_Interval(s) parameter allows you to specify the interval (in seconds) between two successive General Queries that a Querier multicast router sends.
- **Last_Member_Q_Intvl(s):** The Last_Member_Q_Intvl(s) parameter allows you specify the interval (in seconds) between two successive Group-Specific Query messages that a multicast router sends to hosts.
- **Unsol_Report_Intvl(s):** The Unsol_Report_Intvl(s) is the time between repetitions of a host's initial report of membership in a group.

The following image illustrates the wired node pop-up window and the parameters you can configure for IGMP, on the node.

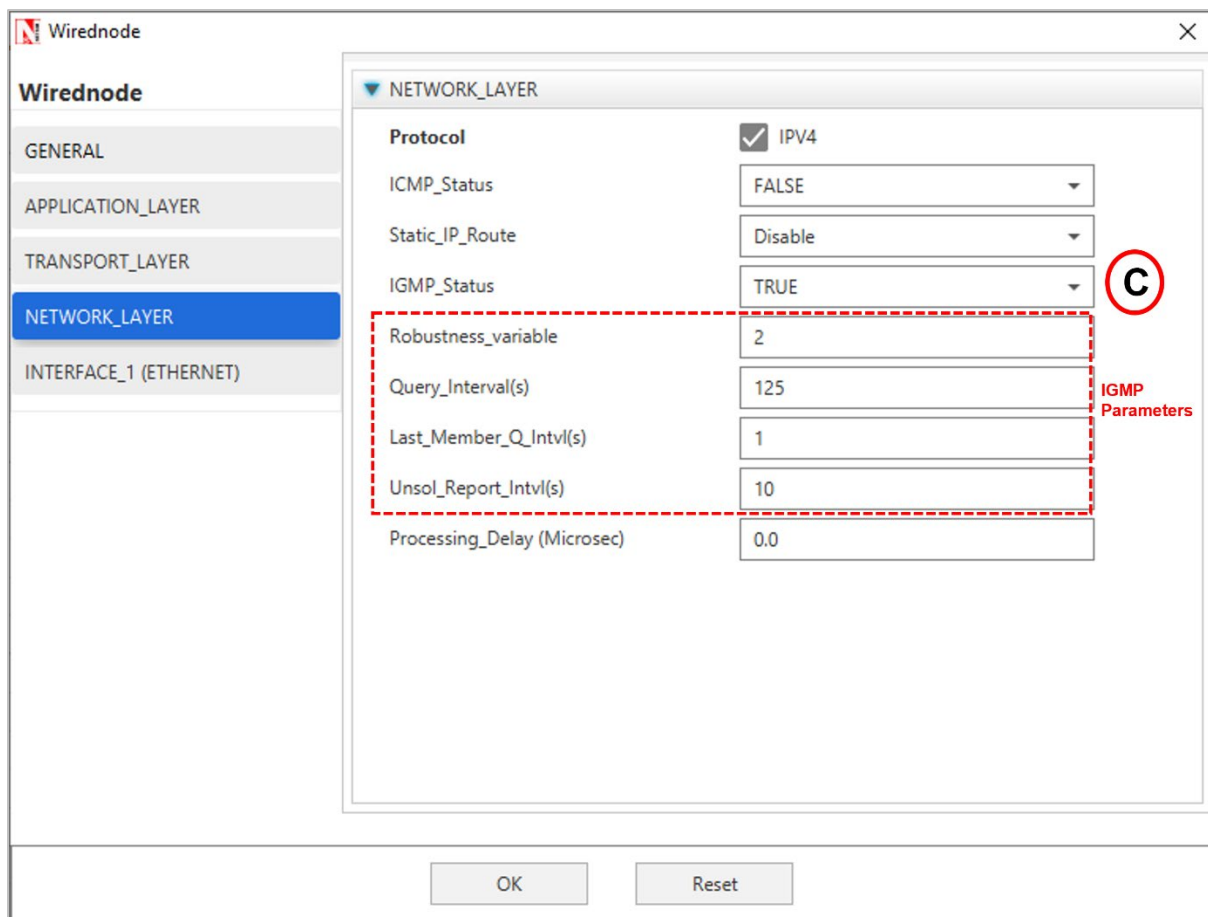


Figure 5-3: Network layer IGMP Properties

- a. Click **OK**.
3. (Optional) Do the following to modify the parameters of IGMP.
 - To modify the value of the Robustness_variable, enter a value in the **Robustness_variable** text box.
 The default value of the Robustness_variable parameter is 2.
 You can enter a value between 2 and 10.
 NetSim does not allow you to enter a value that is less than 2. If you enter a value that is less than 2, NetSim resets the value to 2.
 Increase the value of the Robustness_variable to more than 2, if you want to simulate a subnet that must lose more packets.
 By default, IGMP is robust to (Robustness Variable-1) packet losses.
 - To modify the value of the Query_Interval(s), enter a value in seconds, in the **Query_Interval(s)** text box.
 The default value of the Query_Interval(s) parameter is 125 seconds.
 You can enter a value between 1 and 3600 seconds.
 Fine-tune the Query_Interval(s) parameter to control the number of IGMP messages on the subnet.

- To modify the value of the Query_Interval(s), enter a value in seconds, in the Last_Member_Q_Intvl(s) text box.

The default value of the Last_Member_Q_Intvl(s) parameter is 1 second.

You can enter a value between 1 and 25 seconds.

Fine-tune the Last_Member_Q_Intvl(s) parameter to make your subnet less or more busy of IGMP messages.

- To modify the value of the Query_Interval(s), enter a value in seconds, in the Unsolicited_Report_Interval(s) text box.

The default value of the Last_Member_Q_Intvl(s) parameter is 10 seconds.

Fine-tune the Unsolicited_Report_Interval(s) parameter to make your subnet less or more busy of IGMP messages.

You can enter a value between 1 and 10,000 seconds.

4. Repeat steps 3 on other nodes to see that NetSim has enabled IGMP and step 4 on other nodes, if you to modify the IGMP parameters.

5. To configure a multicast application:

- a. Click the Application icon located in the toolbar.

The Application pop-up window appears.

- b. See that by default, NetSim has set the following properties for the multicast application:

- I. Application_Method = MULTICAST.
- II. **Source_ID = 2**, which means **Wired_Node_2** node is the source of the application and the multicast traffic.
- III. **Destination_Count = 2**, which means two nodes will receive multicast traffic from the multicast application.
- IV. **Destination_ID = 3, 4**, which means, **Wired_Node_3 and Wired_Node_4** nodes must receive multicast traffic from the multicast application.
- V. Set application start time to 5s.

- c. (Optional) Modify the properties except (i).

Note: You add more than one destination IDs, by separating two successive numbers by a “,” (comma). The following image illustrates the properties of the multicast application as shown Figure 5-4.

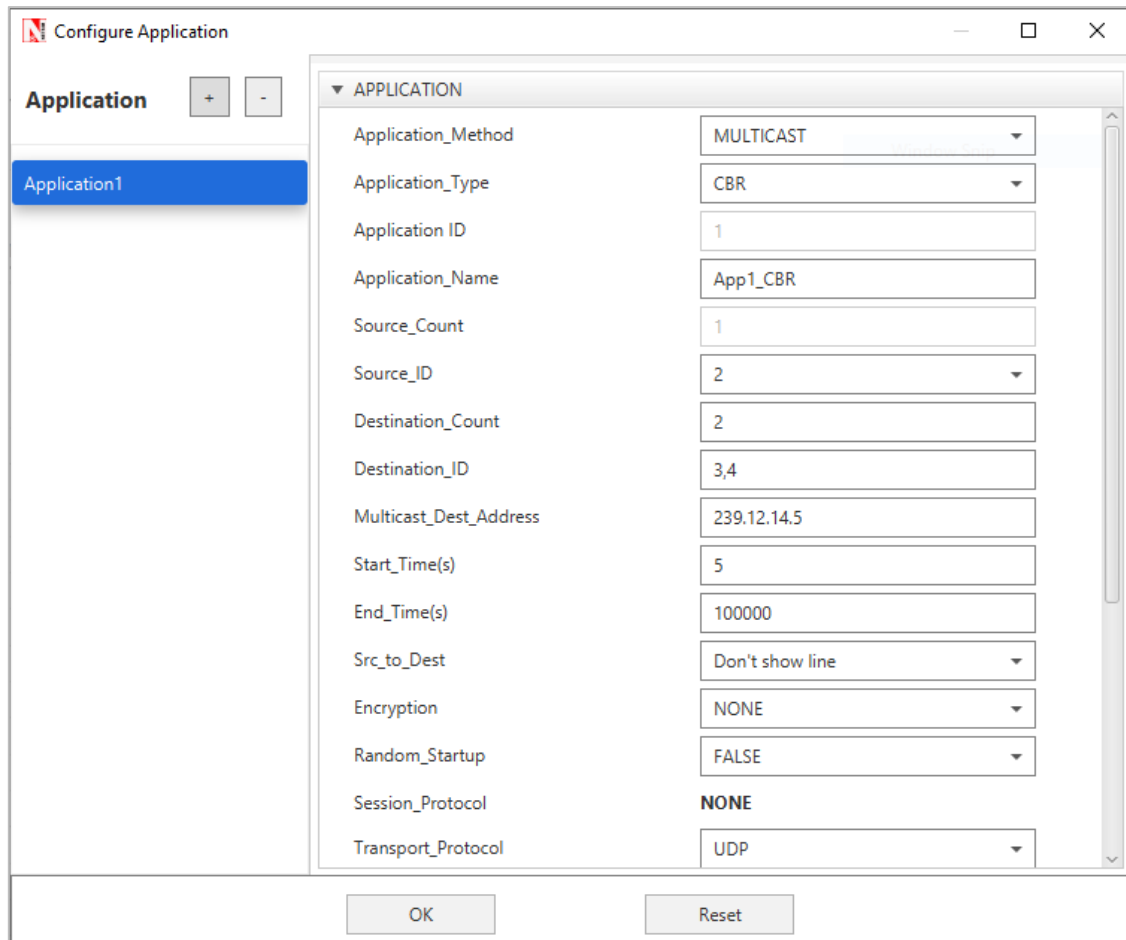


Figure 5-4: Application Properties window

- a. Click **OK**.
6. See that by default, NetSim has enabled the Packet Trace, Event Trace and Plots icons located in the toolbar.
7. To start and run the simulation:
 - a. Click the Run icon located in the toolbar.
 - b. Enter a numerical value in the Simulation Time text box, say 50s.
 - c. Click OK.

NetSim simulates IGMP for the time set.

5.1.2 Results

After NetSim simulates IGMP, a Simulation Results window appears.

You can do the following on this window:

- Print the results that NetSim displays in the Simulation Results window.
- View the packet trace details in a .CSV file and save the .CSV file to your computer.
- View the event trace details in a .CSV file and save the .CSV file to your computer.
- Export the results that NetSim displays in the Simulation Results window, in a spreadsheet.

- Close the Simulation Results window and return to your simulation.

NetSim also saves the last instance of your simulation for you to view, analyse, and download the results.

Interpreting the IGMP Simulation

Before you analyse the packet trace and event trace results, we recommend that you first interpret how IGMP worked with the parameters you specified. So, you must first view the simulation.

To view and interpret the simulation:

1. Close the **Simulation Results** window and return to your simulation.
2. Click the **View Animation** icon located on the toolbar.

The NetSim Packet Animation window appears.

3. Click the **Play** icon located on the toolbar.

You will see that the simulation runs IGMP.

The details of the packet traversing in your network appear as table located below the simulation window.

4. (Optional) To fine-tune the speed of the animation, use the **Animation Speed** slider located on the toolbar.

You will see the following happen in the animation:

- I. Initially, all nodes (Wired_Node_2, 3, 4 and 5) receive the IGMP_Memebership_Query message from Router_1.
- II. When a node receives the IGMP_Memebership_Query message, the node sends the IGMP_V2_Membership_Report to Router_1 indicating that it is interested to join the multicast group.

You can see that Wired_Node_3 sends the IGMP_V2_Membership_Report message to Router_1. Wired_Node 2, 4 and 5 also send the IGMP_V2_Membership_Report message to Router_1.

- III. Router_1 makes an entry for the membership in its routing table.

The following image illustrates IGMP at work.

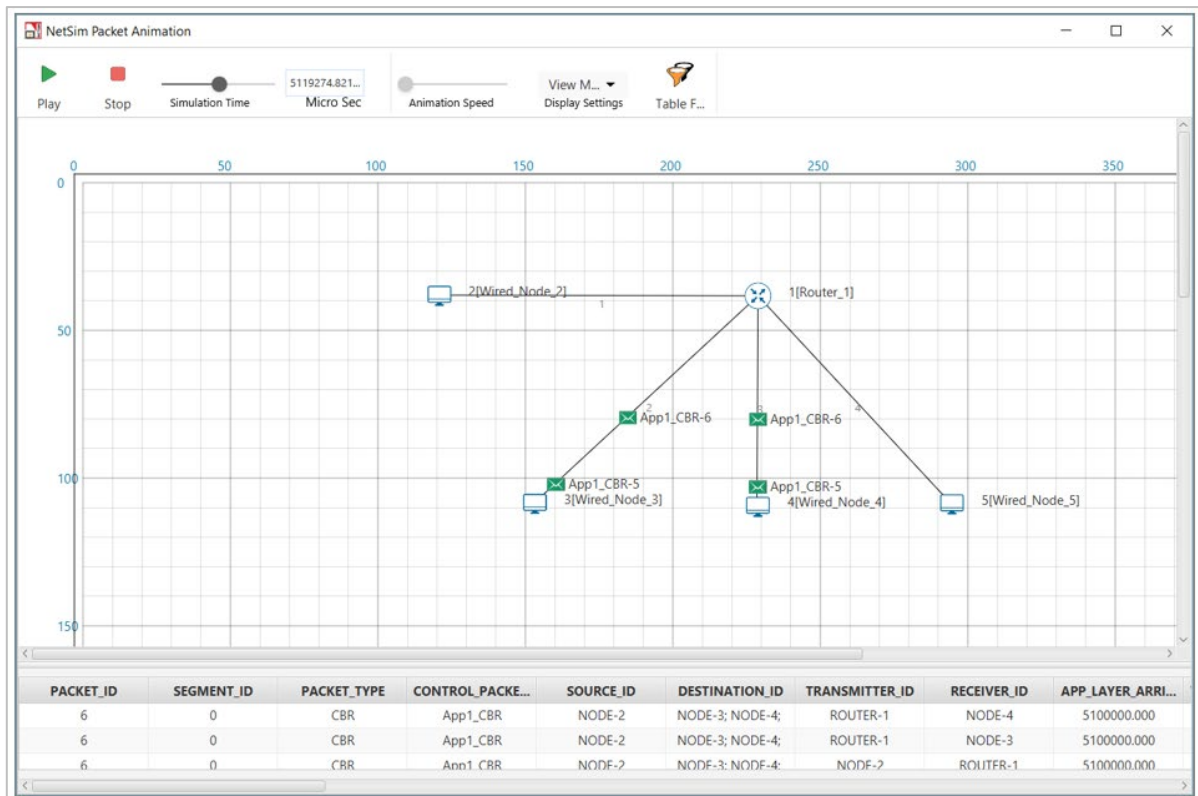


Figure 5-5: Packet animation window

When NetSim completes the simulation, the Simulation Results window appears.

Analyzing the Packet Trace Results

Now that you have seen the simulation for IGMP, we will analyze the communication between the nodes and the router.

To view and analyze the packet trace results:

1. On the Simulation Results window, click **Open Packet Trace** located in the left area.
A .CSV appears.

2. Open the .CSV file and filter the **PACKET_ID** column by **0** and **1**.

You will see the following in the .CSV file.

- I. Router_1 broadcasts the IGMP_Memebership_Query message to all the nodes.
- II. When a node receives the IGMP_Memebership_Query message, the node sends the IGMP_V2_Membership_Report message to the Router_1.
- III. The IGMP protocol starts to work only after 1 second in to the simulation.

The following image illustrates (i), (ii), and (iii).

| PACKET_ID | SEGMENT_ID | PACKET_TYPE | CONTROL_PACKET_TYPE/APP_NAME | SOURCE_ID | DESTINATION_ID | TRANSMITTER_ID | RECEIVER_ID | NW_LAYER_ARRIVAL_TIME[US] |
|-----------|------------|----------------|------------------------------|-----------|----------------|----------------|-------------|---------------------------|
| 0 | N/A | Control_Packet | IGMP_Membership_Query | ROUTER-1 | Broadcast-0 | ROUTER-1 | NODE-3 | 1000 |
| 0 | N/A | Control_Packet | IGMP_Membership_Query | ROUTER-1 | Broadcast-0 | ROUTER-1 | NODE-4 | 1000 |
| 0 | N/A | Control_Packet | IGMP_Membership_Query | ROUTER-1 | Broadcast-0 | ROUTER-1 | NODE-5 | 1000 |
| 0 | N/A | Control_Packet | IGMP_Membership_Query | ROUTER-1 | Broadcast-0 | ROUTER-1 | NODE-2 | 1000 |
| 0 | N/A | Control_Packet | IGMP_V2_Membership_Report | NODE-5 | Broadcast-0 | NODE-5 | ROUTER-1 | 252360.744 |
| 0 | N/A | Control_Packet | IGMP_Membership_Query | ROUTER-1 | Broadcast-0 | ROUTER-1 | NODE-4 | 252370.064 |
| 0 | N/A | Control_Packet | IGMP_Membership_Query | ROUTER-1 | Broadcast-0 | ROUTER-1 | NODE-5 | 252370.064 |
| 0 | N/A | Control_Packet | IGMP_Membership_Query | ROUTER-1 | Broadcast-0 | ROUTER-1 | NODE-2 | 252370.064 |
| 0 | N/A | Control_Packet | IGMP_V2_Membership_Report | NODE-3 | Broadcast-0 | NODE-3 | ROUTER-1 | 562264.331 |
| 0 | N/A | Control_Packet | IGMP_V2_Membership_Report | NODE-2 | Broadcast-0 | NODE-2 | ROUTER-1 | 595196.068 |
| 0 | N/A | Control_Packet | IGMP_V2_Membership_Report | NODE-4 | Broadcast-0 | NODE-4 | ROUTER-1 | 733877.178 |
| 0 | N/A | Control_Packet | IGMP_V2_Membership_Report | NODE-5 | Broadcast-0 | NODE-5 | ROUTER-1 | 857706.41 |
| 0 | N/A | Control_Packet | IGMP_Membership_Query | ROUTER-1 | Broadcast-0 | ROUTER-1 | NODE-3 | 3101000 |
| 0 | N/A | Control_Packet | IGMP_Membership_Query | ROUTER-1 | Broadcast-0 | ROUTER-1 | NODE-4 | 3101000 |

Figure 5-6: IGMP_Memebership_Query messages in Packet trace

IV. The multicast application Wired_Node_2 starts to send multicast traffic to Wired_Node_3 and Wired_Node_4 only after 5 seconds in to the simulation.

This is because, in NetSim, the multicast application starts after 5 seconds by default.

V. Wired_Node_2 multicasts Constant Bit Rate (CBR) packets only to Wired_Node_3 and Wired_Node_4.

The following image illustrates (iv), and (v).

| | | | | | | | |
|-----|----------------|-----------------------|----------|-----------------|----------|----------|---------|
| N/A | Control_Packet | IGMP_Membership_Query | ROUTER-1 | Broadcast-0 | ROUTER-1 | NODE-3 | N/A |
| N/A | Control_Packet | IGMP_Membership_Query | ROUTER-1 | Broadcast-0 | ROUTER-1 | NODE-4 | N/A |
| N/A | Control_Packet | IGMP_Membership_Query | ROUTER-1 | Broadcast-0 | ROUTER-1 | NODE-2 | N/A |
| N/A | Control_Packet | IGMP_Membership_Query | ROUTER-1 | Broadcast-0 | ROUTER-1 | NODE-3 | N/A |
| N/A | Control_Packet | IGMP_Membership_Query | ROUTER-1 | Broadcast-0 | ROUTER-1 | NODE-4 | N/A |
| 0 | CBR | App1_CBR | NODE-2 | NODE-3; NODE-4; | NODE-2 | ROUTER-1 | 5000000 |
| 0 | CBR | App1_CBR | NODE-2 | NODE-3; NODE-4; | ROUTER-1 | NODE-3 | 5000000 |
| 0 | CBR | App1_CBR | NODE-2 | NODE-3; NODE-4; | ROUTER-1 | NODE-4 | 5000000 |
| 0 | CBR | App1_CBR | NODE-2 | NODE-3; NODE-4; | NODE-2 | ROUTER-1 | 5020000 |
| 0 | CBR | App1_CBR | NODE-2 | NODE-3; NODE-4; | ROUTER-1 | NODE-3 | 5020000 |

Figure 5-7: Node 2 multicasting CBR packets to other Nodes

VI. Hosts send the IGMP_V2_Membership_Report to 224.0.0.1 to the multicast application sends multicast traffic to 239.12.14.5.

The following image illustrates (vi).

| | | | | | | | | | |
|---|-----|----------------|---------------------------|----------|-----------------|----------|----------|-------------|-------------|
| 0 | N/A | Control_Packet | IGMP_Membership_Query | ROUTER-1 | Broadcast-0 | ROUTER-1 | NODE-4 | 3152370.064 | 224.0.0.1 |
| 0 | N/A | Control_Packet | IGMP_Membership_Query | ROUTER-1 | Broadcast-0 | ROUTER-1 | NODE-2 | 3152370.064 | 224.0.0.1 |
| 0 | N/A | Control_Packet | IGMP_V2_Membership_Report | NODE-4 | Broadcast-0 | NODE-4 | ROUTER-1 | 3459111.886 | 224.0.0.1 |
| 0 | N/A | Control_Packet | IGMP_V2_Membership_Report | NODE-3 | Broadcast-0 | NODE-3 | ROUTER-1 | 3296053.796 | 224.0.0.1 |
| 0 | N/A | Control_Packet | IGMP_V2_Membership_Report | NODE-2 | Broadcast-0 | NODE-2 | ROUTER-1 | 3788305.986 | 224.0.0.1 |
| 0 | N/A | Control_Packet | IGMP_V2_Membership_Report | NODE-2 | Broadcast-0 | NODE-2 | ROUTER-1 | 5000000 | 239.12.14.5 |
| 0 | N/A | Control_Packet | IGMP_V2_Membership_Report | NODE-3 | Broadcast-0 | NODE-3 | ROUTER-1 | 5000000 | 239.12.14.5 |
| 0 | N/A | Control_Packet | IGMP_V2_Membership_Report | NODE-4 | Broadcast-0 | NODE-4 | ROUTER-1 | 5000000 | 239.12.14.5 |
| 0 | N/A | Control_Packet | IGMP_Membership_Query | ROUTER-1 | Broadcast-0 | ROUTER-1 | NODE-2 | 5000009.32 | 239.12.14.5 |
| 0 | N/A | Control_Packet | IGMP_Membership_Query | ROUTER-1 | Broadcast-0 | ROUTER-1 | NODE-3 | 5000009.32 | 239.12.14.5 |
| 0 | N/A | Control_Packet | IGMP_Membership_Query | ROUTER-1 | Broadcast-0 | ROUTER-1 | NODE-4 | 5000009.32 | 239.12.14.5 |
| 0 | N/A | Control_Packet | IGMP_Membership_Query | ROUTER-1 | Broadcast-0 | ROUTER-1 | NODE-2 | 5000009.32 | 239.12.14.5 |
| 0 | N/A | Control_Packet | IGMP_Membership_Query | ROUTER-1 | Broadcast-0 | ROUTER-1 | NODE-3 | 5000009.32 | 239.12.14.5 |
| 0 | N/A | Control_Packet | IGMP_Membership_Query | ROUTER-1 | Broadcast-0 | ROUTER-1 | NODE-4 | 5000009.32 | 239.12.14.5 |
| 0 | N/A | Control_Packet | IGMP_Membership_Query | ROUTER-1 | Broadcast-0 | ROUTER-1 | NODE-2 | 5000009.32 | 239.12.14.5 |
| 0 | N/A | Control_Packet | IGMP_Membership_Query | ROUTER-1 | Broadcast-0 | ROUTER-1 | NODE-3 | 5000009.32 | 239.12.14.5 |
| 0 | N/A | Control_Packet | IGMP_Membership_Query | ROUTER-1 | Broadcast-0 | ROUTER-1 | NODE-4 | 5000009.32 | 239.12.14.5 |
| 1 | 0 | CBR | App1_CBR | NODE-2 | NODE-3; NODE-4; | NODE-2 | ROUTER-1 | 5000000 | 239.12.14.5 |
| 1 | 0 | CBR | App1_CBR | NODE-2 | NODE-3; NODE-4; | ROUTER-1 | NODE-3 | 5000131.4 | 239.12.14.5 |
| 1 | 0 | CBR | App1_CBR | NODE-2 | NODE-3; NODE-4; | ROUTER-1 | NODE-4 | 5000131.4 | 239.12.14.5 |
| 2 | 0 | CBR | App1_CBR | NODE-2 | NODE-3; NODE-4; | NODE-2 | ROUTER-1 | 5020000 | 239.12.14.5 |
| 2 | 0 | CBR | App1_CBR | NODE-2 | NODE-3; NODE-4; | ROUTER-1 | NODE-3 | 5020126.12 | 239.12.14.5 |

Figure 5-8: IGMP_V2_Membership_Report to 224.0.0.1 in packet trace

IGMP Event Trace Analysis

Now that you have seen the results of packet trace, we will analyze the event trace for this IGMP simulation.

To view the event trace results:

1. On the Simulation Results window, click **Open Event Trace** located in the left area. A .CSV appears.
2. Open the .CSV file and filter the Event_Type column by **NETWORK_OUT** and **TIMER_EVENT**.

You will see the following sub-events in the **Subevent_Type** column:

- a. **IGMP_DelayTimer**: This sub-event occurs when a node sets the delay timers for every group (excluding the all-systems group) to which the node belongs, on the interface on which it received the query, after the node receives a General Query from the multicast router.
- b. **IGMP_GroupMembershipTimer**: This sub-event occurs when the multicast router refreshes the group membership interval timer, every time it receives a membership report for a multicast group. If this timer expires, the multicast router removes this group from the list of destinations for multicast traffic.
- c. **IGMP_SendQuery**: This sub-event occurs when the multicast router periodically (based on Query Interval) sends a Query message on every network to which the multicast router is connected, to solicit multicast group membership information.
- d. **IGMP_SendStartupQuery**: This subevent occurs when the multicast router sends the Startup query count to quickly and reliably determine the multicast group membership information, at startup.
- e. **IGMP_UnsolicitedReportTimer**: If the initial membership report is lost or damaged, this timer repeats once or twice after short delays, after every Unsolicited Report Interval.
- f. **JOIN_MULTICAST_GROUP**: This sub-event occurs when a node sends the join multicast group message, when the node decides to join a multicast group on an interface.

In NetSim, a node joins a multicast group only after 5 seconds into the simulation.

The following image illustrates that hosts join the multicast group after 5 seconds.

| Event_Id | Event_Type | Event_Time(US) | Device_Type | Device_Id | Interface_Id | Segment_Id | Protocol_Name | Subevent_Type |
|----------|-----------------|----------------|-------------|-----------|--------------|------------|---------------|----------------------|
| 141 | 149 NETWORK_OUT | 3459111.89 | NODE | 3 | 0 | 0 | IPV4 | 0 |
| 147 | 119 TIMER_EVENT | 3596653.74 | NODE | 2 | 0 | 0 | IPV4 | IGMP_DelayTimer |
| 148 | 155 NETWORK_OUT | 3596653.74 | NODE | 2 | 0 | 0 | IPV4 | 0 |
| 154 | 121 TIMER_EVENT | 3788305.99 | NODE | 5 | 0 | 0 | IPV4 | IGMP_DelayTimer |
| 155 | 161 NETWORK_OUT | 3788305.99 | NODE | 5 | 0 | 0 | IPV4 | 0 |
| 161 | 19 TIMER_EVENT | 5000000 | NODE | 2 | 0 | 0 | IPV4 | JOIN_MULTICAST_GROUP |
| 162 | 20 TIMER_EVENT | 5000000 | NODE | 3 | 0 | 0 | IPV4 | JOIN_MULTICAST_GROUP |
| 163 | 21 TIMER_EVENT | 5000000 | NODE | 4 | 0 | 0 | IPV4 | JOIN_MULTICAST_GROUP |
| 164 | 22 TIMER_EVENT | 5000000 | NODE | 2 | 0 | 0 | APPLICATION | |
| 165 | 167 NETWORK_OUT | 5000000 | NODE | 2 | 0 | 0 | IPV4 | 0 |
| 166 | 169 NETWORK_OUT | 5000000 | NODE | 3 | 0 | 0 | IPV4 | 0 |
| 167 | 171 NETWORK_OUT | 5000000 | NODE | 4 | 0 | 0 | IPV4 | 0 |

Figure 5-9: Node joins a multicast group

5.2 Access Control Lists (ACLs) Examples

5.2.1 ACL Example

This example models a network and simulates an ACL to understand how ACL filters inbound and outbound traffic at a router's interface.

The network modelled consists of:

- Two subnets with 2 wired nodes, 1 router each and 3 applications.
- ACLs with both permit and deny rules are defined on the interfaces of the router.

NetSim uses the following directions for ACL simulations:

- The direction of the ACL is set to both. This means the ACL applies to both inbound and outbound traffic.
- The direction of ACL is set to Inbound. This means the ACL applies to inbound traffic only.
- The direction of ACL is set to Outbound. This means the ACL applies to outbound traffic only.

Open NetSim, Select **Examples->Advanced routing->ACL Configuration** then click on the tile in the middle panel to load the example as shown below in Figure 5-10.

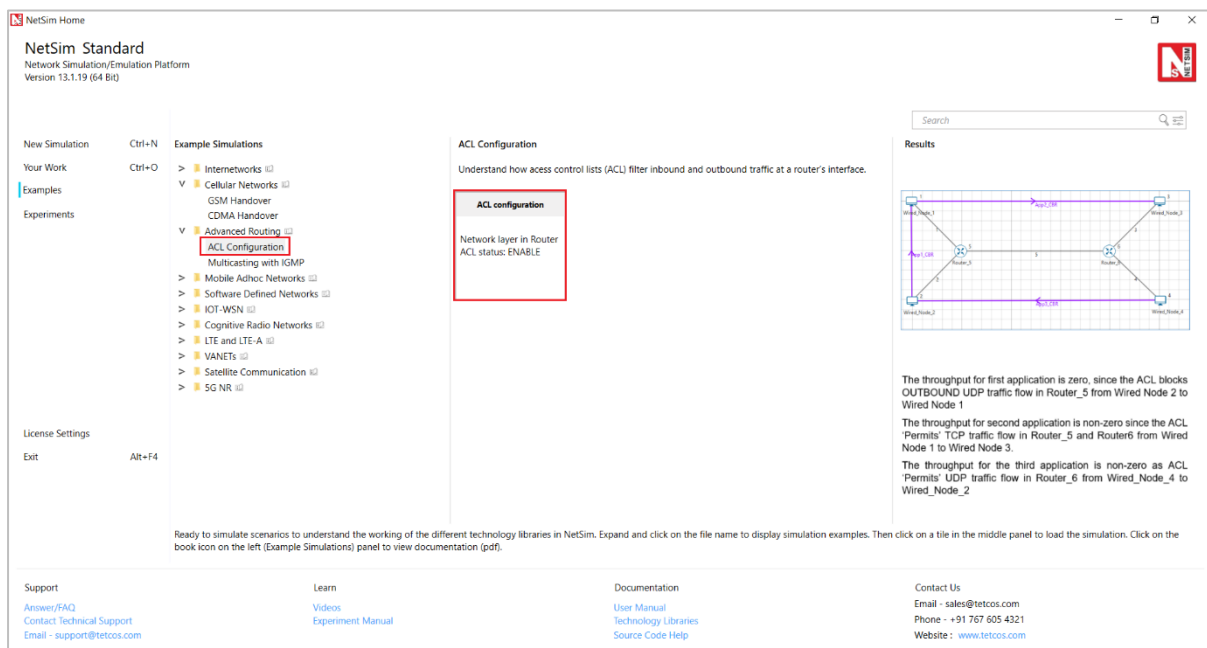


Figure 5-10: List of scenarios for the example of ACL Configuration

The following network diagram illustrates what the NetSim UI displays when you open the example configuration file for ACL as shown Figure 5-11.

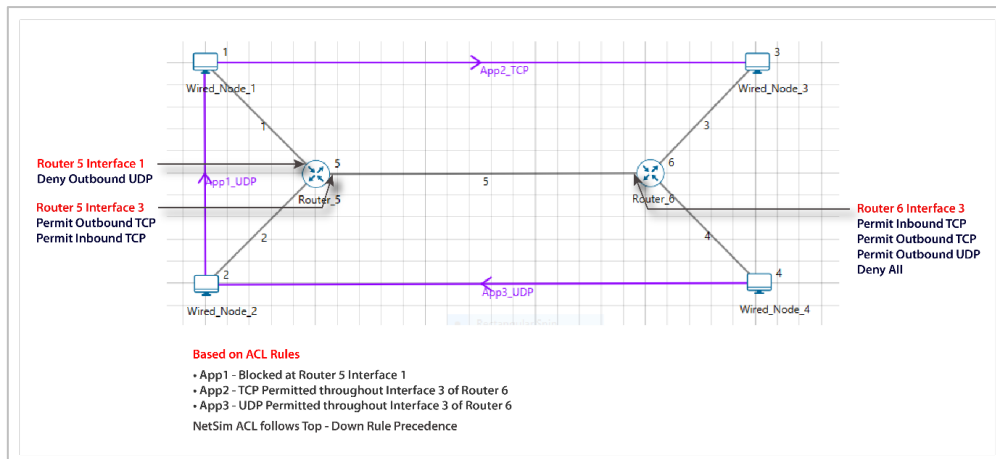


Figure 5-11: Network set up for studying the ACL Configuration

1. ACL enabled in Network Layer of Router_5 and were configured as follows as shown Figure 5-12.

Configure ACL

Action: Source Address:

Direction: Destination Address:

Protocol: Source port:

Interface ID: Destination port:

| Action | Direction | Protocol | Interface ID | Source address | Dest address | Source port | Dest port |
|--------|-----------|----------|--------------|----------------|--------------|-------------|-----------|
| PERMIT | BOTH | TCP | 3 | 11.1.1.2/31 | 11.3.1.2/31 | 0 | 0 |
| DENY | OUTBOUND | UDP | 1 | ANY | ANY | 0 | 0 |

Figure 5-12: ACL Configuration for Router 5

2. ACL enabled in Network Layer of Router_6 and were configured as follows as Figure 5-13.

Configure ACL

Action: Source Address:

Direction: Destination Address:

Protocol: Source port:

Interface ID: Destination port:

| Action | Direction | Protocol | Interface ID | Source address | Dest address | Source port | Dest port |
|--------|-----------|----------|--------------|----------------|--------------|-------------|-----------|
| PERMIT | OUTBOUND | UDP | 3 | ANY | ANY | 0 | 0 |
| PERMIT | OUTBOUND | TCP | 3 | ANY | ANY | 0 | 0 |
| PERMIT | INBOUND | TCP | 3 | ANY | ANY | 0 | 0 |
| DENY | BOTH | ANY | 3 | ANY | ANY | 0 | 0 |

Figure 5-13: ACL Configuration for Router 6

3. Transport protocol set as UDP for APP_1_CBR and APP_3_CBR.

4. Transport protocol set as TCP for APP_2_CBR.
5. Enable the plots and run Simulation for 10 seconds and observe the throughput obtained for the three applications.

5.2.2 Result and Observations

| Application_Metrics_Table | | | | | | | |
|---------------------------|---|------------------|------------------|-----------------|-------------------|-----------------|------------------|
| Application_Metrics | | | | | | | |
| Application Id | Throughput Plot | Application Name | Packet generated | Packet received | Throughput (Mbps) | Delay(microsec) | Jitter(microsec) |
| 1 | Application Throughput plot | App1_CBR | 500 | 0 | 0.000000 | 0.000000 | 0.000000 |
| 2 | Application Throughput plot | App2_CBR | 500 | 478 | 0.558304 | 26817.456362 | 4114.302997 |
| 3 | Application Throughput plot | App3_CBR | 500 | 499 | 0.582832 | 376.281924 | 0.001928 |

Figure 5-14: Application Metrics Table in result window

The throughput for first application is zero, since the ACL blocks OUTBOUND UDP traffic flow in Router_5 from Wired Node 2 to Wired Node 1

The throughput for second application is non-zero, since the ACL 'Permits' TCP traffic flow in Router_5 and Router6 from Wired Node 1 to Wired Node 3.

The throughput for the third application is non-zero as ACL 'Permits' UDP traffic flow in Router_6 from Wired_Node_4 to Wired_Node_2.

6 Advanced Routing Experiments in NetSim

Apart from examples, in-built experiments are also available in NetSim. Examples help the user understand the working of features in NetSim. Experiments are designed to help the user (usually students) learn networking concepts through simulation. The experiments contain objective, theory, set-up, results, and inference. The following experiments are available in the Experiments manual (pdf file).

1. Understanding VLAN operation in L2 and L3 Switches
2. Understanding Access and Trunk Links in VLANs
3. Understanding Public IP Address & NAT (Network Address Translation)
4. Understand the working of basic networking commands (Ping, Route Add/Delete/Print, ACL)

7 Reference Documents

1. IEEE802.1Q for Virtual LAN
2. IETF 7761 for Protocol Independent Multicast – Sparse Mode (PIM-SM)
3. RFC 2236 for Internet Group Management Protocol, Version 2

8 Latest FAQs

Up to date FAQs on NetSim's Advance Routing library is available at

<https://tetcos.freshdesk.com/support/solutions/folders/14000113123>