Networking Technologies and Applications

Rolland Vida
BME TMIT

November 3, 2016

Multicast at different layers

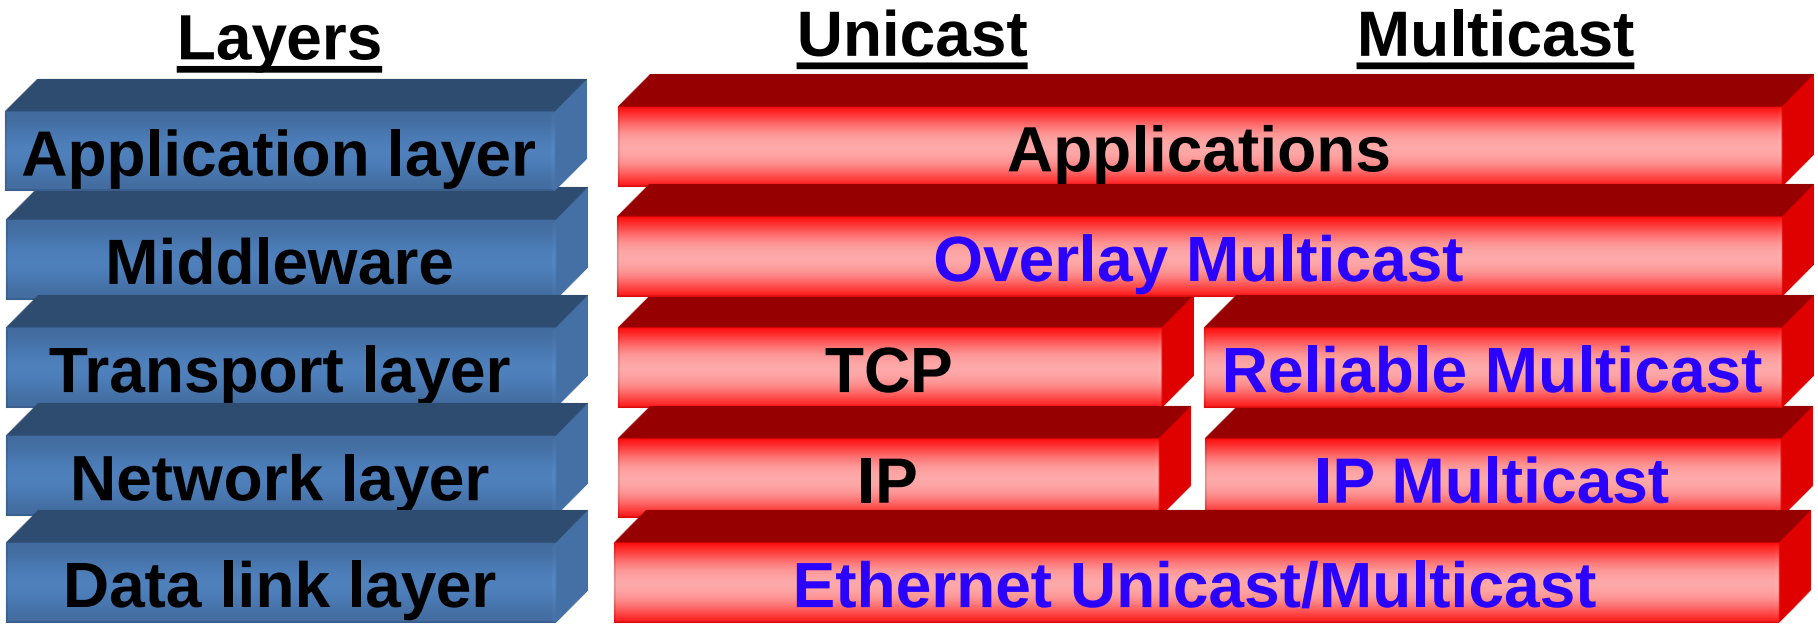

Ethernet Multicast

- Some Ethernet MAC addresses reserved for multicast
- If we want to join a G group
  - Our network interface card (NIC) normally listens only to packets sent to our unicast adress, or the network broadcast address
  - To join, it should listen also the the group address G
  - Hardware solution, efficient
- Communication in group G
  - The sender floods all the LAN segments
    - Like in case of broadcast
  - The cards that do not listen to group address G just drop the packets

Network layer multicast

- The goal is the optimisation of networ layer resources
  - One packet crosses just once a given link
- Routers build and maintain a multicast tree
  - Traffic forwarding along the tree
  - Routers duplicate packets where needed
    - Branching points on the tree

Group level unicast is not scalable

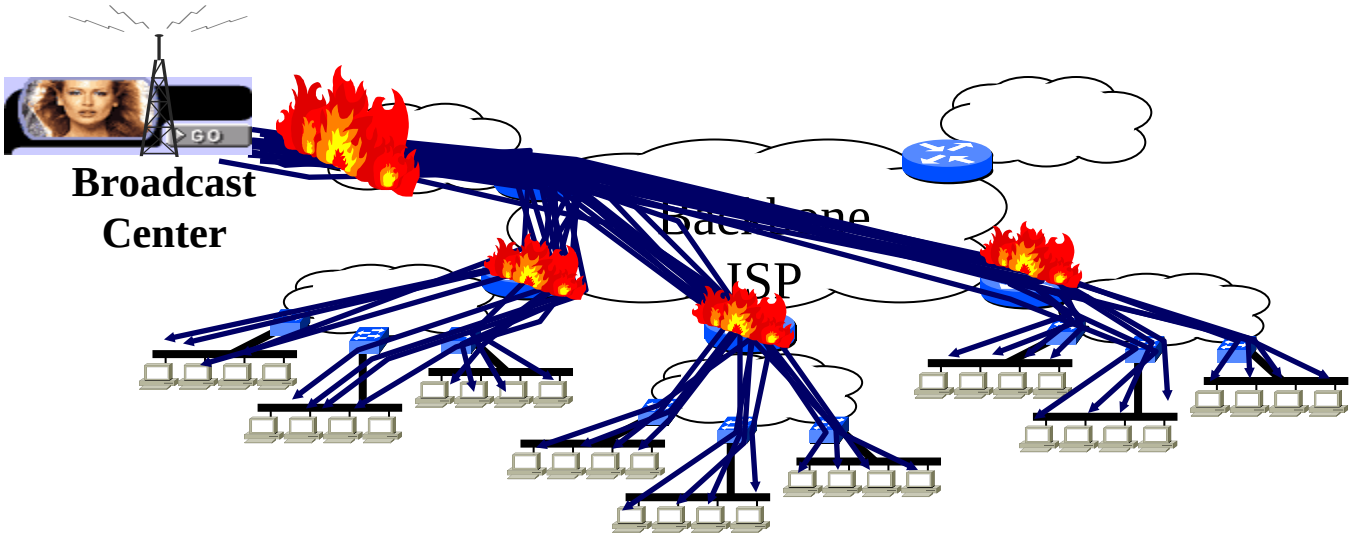

Let’s build trees instead

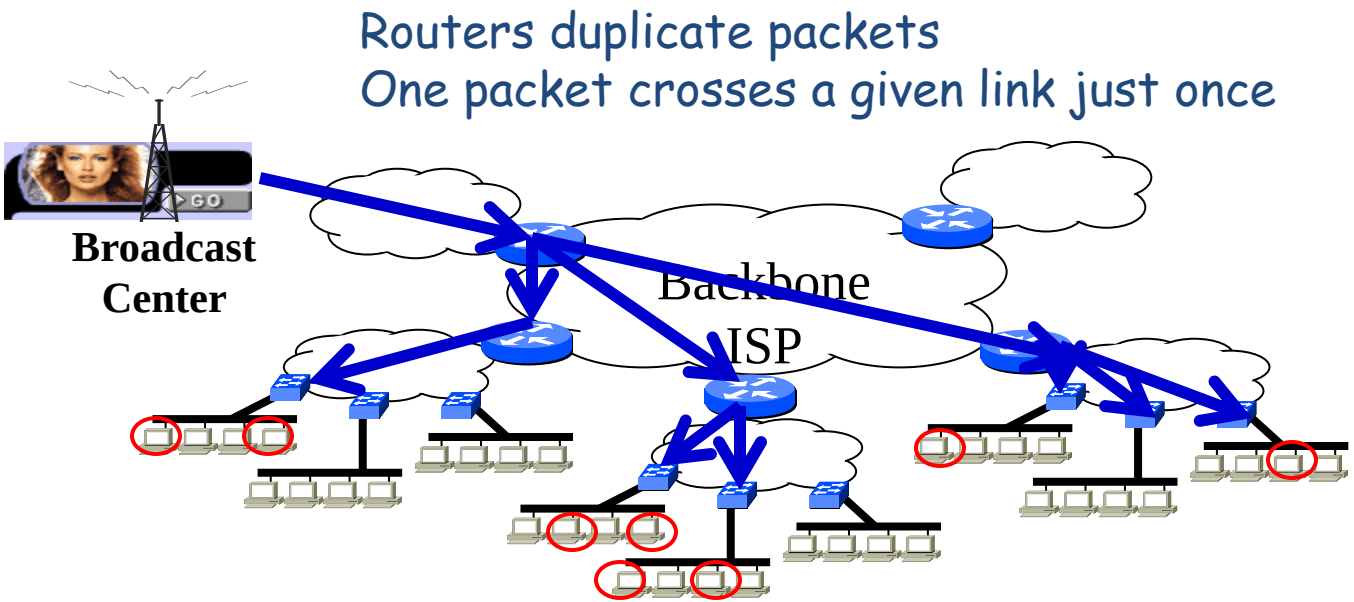

IP Multicast

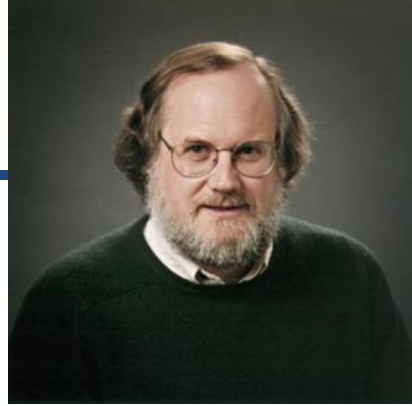

- Steve Deering PhD dissertation (1990)
  - **Any Source Multicast (ASM)**

- Open group communication model
  - Anybody can join the group, no access control
  - One user can be member of several groups in the same time
  - Anybody can send to the group, even non members
  - Group membership is dynamic
  - Nobody knows the size of the group, or its members

IP Multicast

- S. Deering, "Host Extensions for IP Multicasting", RFC 1112, 1989.

- The source sends its packets to a group address
- Anybody who joined the group is „reachable” through this address
  - Receives packets that are sent to this destination address
- A multicast group is identified by a **class D** IP address
  - **224.0.0.0 – 239.255.255.255**
  - 1110 + 28 bit group ID

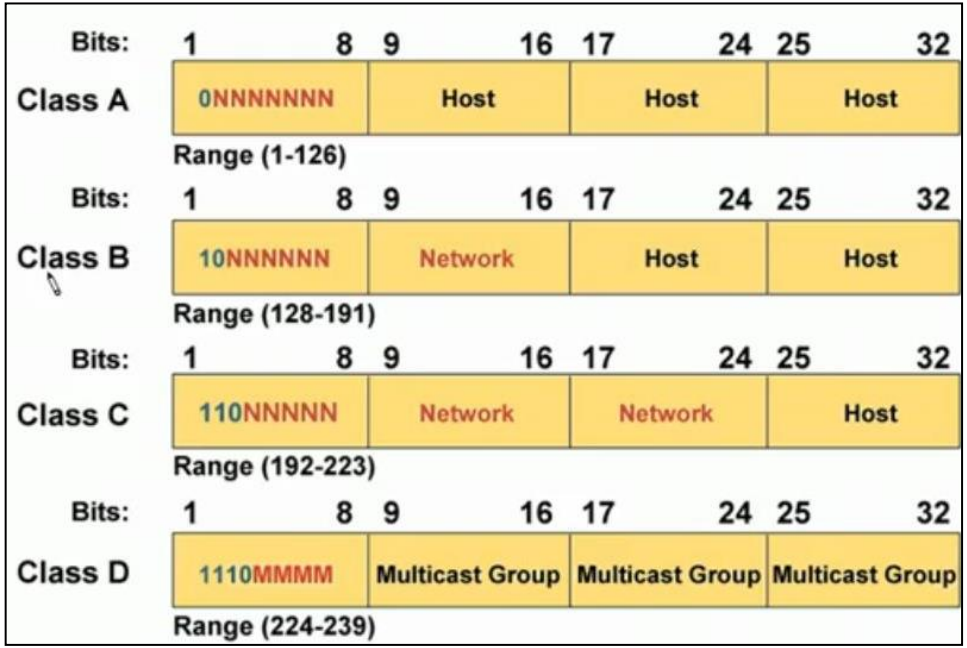

Multicast Scoping

- The scope of an IP multicast group is limited:
  - TTL based scoping
  - Administrative scoping

- TTL based scoping
  - Node-local 0
  - Link-local 1
  - Site-local < 32
  - Region-local < 64
  - Continent-local < 128
  - Global Scope < 255

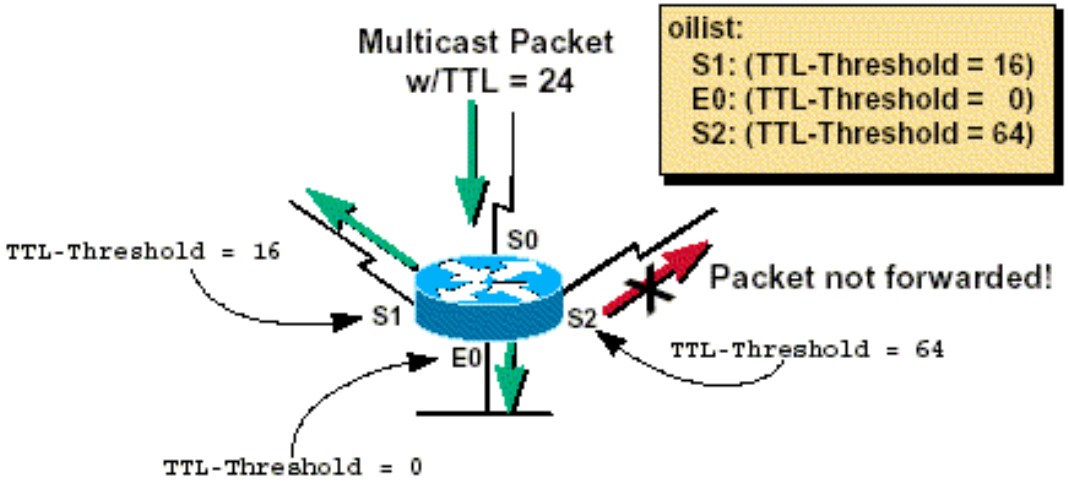

Multicast Scoping

- Administrative scoping
  - link-local scope 224.0.0.0 - 224.0.0.255
    - A router never forwards such a packet
  - global scope 224.0.1.0 - 238.255.255.255
    - Valid on the entire internet
  - administrative scope 239.0.0.0 - 239.255.255.255
    - Never forwarded outside the Intranet of a given organization

IP Multicast

- Joining a multicast tree done in two steps
  - On the local area network (LAN)
    - A user announces its local multicast routers about the groups he would like to join
    - IGMP (IPv4), MLD (IPv6)
  - Over the large Internet (WAN)
    - The local router cooperates with the other multicast routers of the network to build the tree and forward the packets along that tree
    - DVMRP, MOSPF, CBT, PIM-DM, PIM-SM, PIM-SSM

IGMP

- **Internet Group Management Protocol**
- An IPv4 protocol, running between the final users and the local multicast routers on the local network
  - Handles multicast group membership
  - Asymmetric protocol
    - User side
    - Router side
- The router learns which groups the end-users on his local network listen to
  - Not interested in how many receivers, important thing is to have at least on receiver
  - Not interested in exactly who are the receivers

IGMPv1

- S. Deering, "Host Extensions for IP Multicasting", RFC 1112, 1989.
- A multicast router sends regular **Query** messages to the multicast address of all the users (224.0.0.1)
- A user answers with a **Report** message, in which specifies the groups he listens to
  - The Report is sent to the multicast addresses of those groups
- To decrease the number of Report messages:
  - Using timers
    - A user does not answer immediately to the Query
  - Host Suppression
    - If someone else answers faster, it deletes its own Report message
- **Unsolicited Report**
  - If a user wants to listen immediately to a new group

IGMPv1 Router

- An IGMPv1 router maintains a multicast membership table
  - Which multicast groups have members on its network
  - When was the last Report message received about those groups
- **Soft-state** protocol
  - If in a given time nobody refreshes its interest in a given groups, the group will be deleted from the multicast membership table
- It forwards to the local network all packets that are sent to a multicast destination address that is contained in its membership table

IGMPv2

- W. Fenner, "Internet Group Management Protocol, Version 2", RFC 2236, November 1997.
  http://www.ietf.org/rfc/rfc2236.txt

- IPv6 version: MLD (Multicast Listener Discovery)
  - S. Deering, W. Fenner, B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", RFC 2710, November 1999.
    http://www.ietf.org/rfc/rfc2710.txt

- Introduces a **Fast Leave** mechanism
  - Do not have to wait until a timer expires to cut off a group

IGMPv2 messages

- Membership Query
  - General Query
  - Group Specific Query
- Membership Report
- **Leave Group Message**

- If a host wants to leave a group, it sends a Leave message to the multicast address of all the multicast routers (224.0.0.2)
- Before cutting off the group, the router has to ask if anybody else is still interested in that group or not
  - Group Specific Query
  - If no answer in a given limited time, the router cuts off the group from its table

- IGMPv3 – later...

Multicast Routing

- A source sends its packets to the group's multicast address
- The multicast routers in the network build and maintain a multicast tree
  - Packets are forwarded along that tree
- The local multicast router, based on its IGMP membership table, joins or leaves this tree
- A multicast routing protocol runs among the routers of the network
  - MOSPF, DVMRP, CBT, PIM

MOSPF

- **Multicast Open Shortest Path First**
  - J. Moy, „Multicast Extensions to OSPF”, RFC 1584, March 1994 http://www.ietf.org/rfc/rfc1584.txt
- Link State protocol
- Extends the OSPF unicast routing protocol
  - Multicast group membership information is also distributed among the routers
  - Each MOSPF router learns which multicast groups have listeners on which local network
  - Based on this information they build a shortest path tree for each source and each group
- Large signaling overhead
- Difficult to handle topology changes
  - All the trees have to be recalculated

DVMRP

- **Distance Vector Multicast Routing Protocol**
  - D. Waitzman, C. Partridge, S. Deering, "Distance Vector Multicast Routing Protocol", RFC 1075, November 1988

    http://www.ietf.org/rfc/rfc1075.txt

- Distance vector protocol
  - Uses the RIP unicast routing protocol

DVMRP

- Flood and prune
  - Flooding
    - Checks the incoming interfaces of the packets
    - If not over the shortest path towards the source, the packets are dropped
    - If yes, packets are flooded over all the interfaces
  - Pruning
    - If no interested receiver on the local network
    - If packet not received over the shortest path
  - An internal router learns its interfaces over which it recieved a Prune message
    - The upcoming packets are not forwarded over those interfaces anymore
    - Prune messages become obsole after a while (one minute by default)

DVMRP flooding

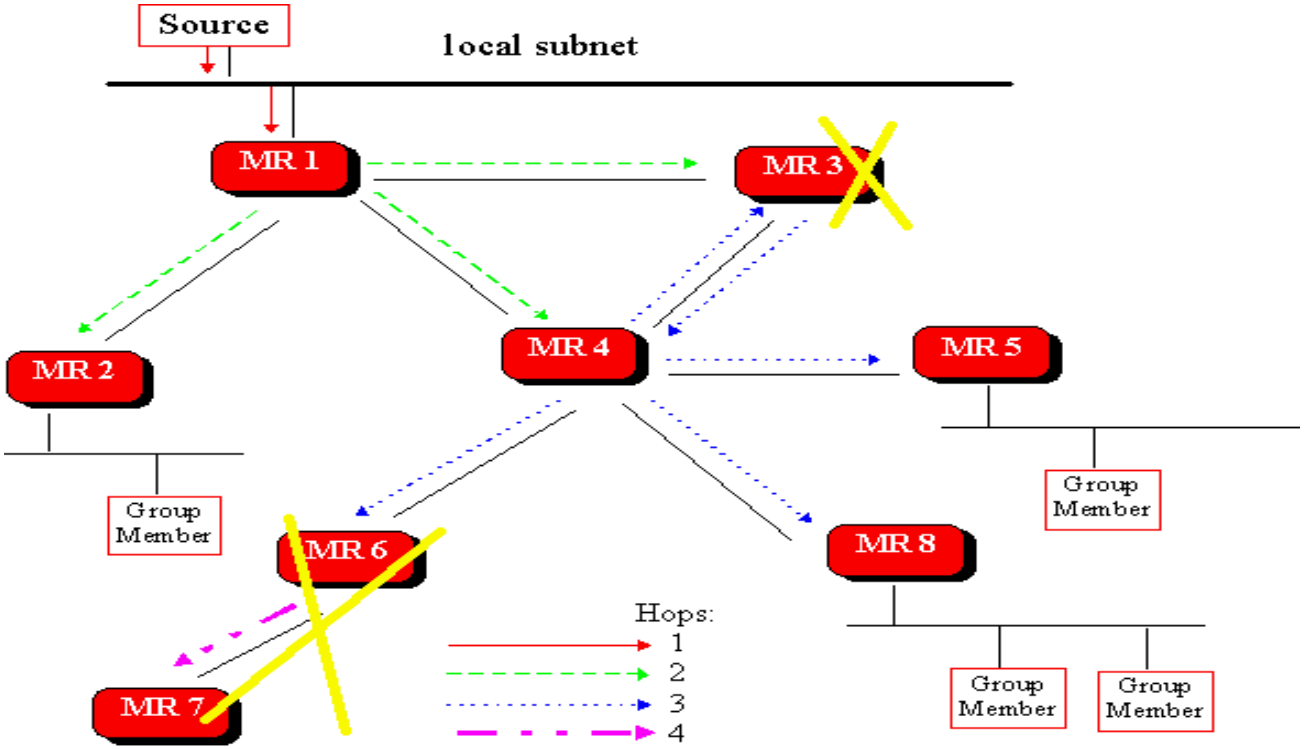

DVMRP prune

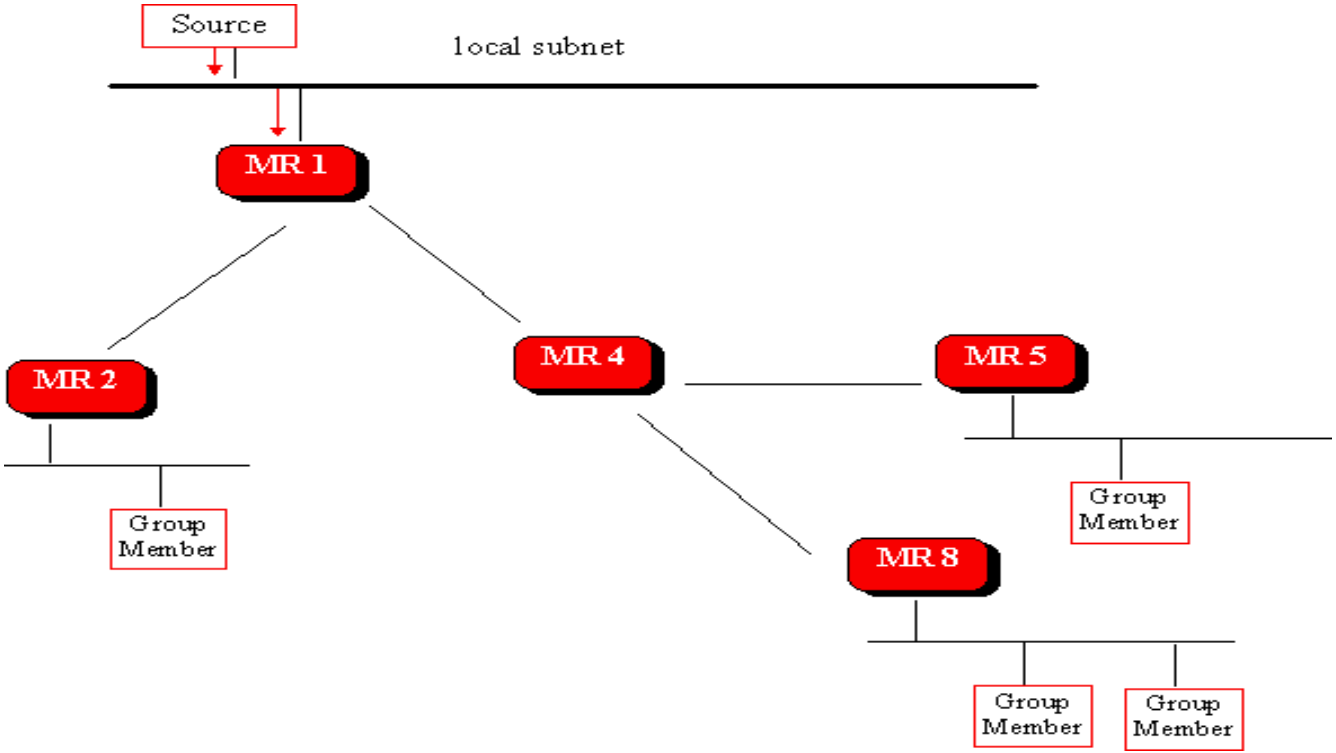
Source
local subnet
MR 1
MR 2
MR 4
MR 5
MR 8
Group Member
Group Member
Group Member
Group Member