

WHITE PAPER

MULTICAST FOR ENTERPRISE VIDEO STREAMING

Protocols and Design Guide

This document provides a network equipment neutral, technical overview of multicast protocols and a discussion of techniques and best practices for implementing multicast video on an Enterprise Network.

CONTENTS


VIDEO STREAMING ON IP NETWORKS	4
Multicast on Enterprise Networks	5
Multicast Addressing	5
Multicast on Layer 2	7
IGMP	7
IGMP Snooping	7
Example	8
Multicast on Layer 3	9
PIM	9
PIM Sparse Mode (PIM-SM)	9
Multicast Session Example	11
Other Relevant Network Configurations	13
Spanning Tree	13
Storm Control	13
IMPLEMENTING MULTICAST VIDEO STREAMING	14
Network Preparation	14
Best Practices	14
Scenario 1 Enterprise IPTV	16
Best Practices	16

Scenario 2 Room Overflow	17
Best Practices	17
Scenario 3 Non-Routed Multicast	20
Best Practices	20
APPENDIX 1, MULTICAST SECURITY	23
Blocking Multicast to Specific Networks	23
Deny Multicast Traffic	23
Administratively Scoped Multicast	24
Access Control Lists (ACL)	24
Rogue Multicast	25
Source Specific Multicast (SSM)	26
APPENDIX 2, MULTICAST OVER 802.11 (WI-FI) WIRELESS	26
Challenges to 802.11 Multicast	26
802.11 Multicast Implementation	27
Optimizing 802.11 Multicast	28
APPENDIX 3, MULTICAST OVER VPN	29
Why VPN	29
About VPN	30
Pros and Cons	30
Inter-Site Video Streaming over VPN	31
Considerations in Implementing Multicast Video Streaming Across VPN	31
Bandwidth	31
Encapsulation, Overhead, and MTU Size	31
Video Streaming Reflector	32
Considerations in Implementing Multicast Reflecting	33
Generic Routing Encapsulation (GRE) Tunnel	34
Considerations in Implementing GRE Tunnels for Multicast	35
Automatic Multicast without explicit Tunnels (AMT)	36
Considerations in Implementing AMT	37

VIDEO STREAMING ON IP NETWORKS

Most modern enterprise networks have sufficient network infrastructure to support video distribution. A gigabit network to the desktop is not required although trunks and backbone links should be reviewed to ensure they have enough available bandwidth to support the proposed application.

There are two primary ways that video is transmitted across an IP network: multicast and unicast.



UNICAST: Unicast is a one-to-one connection between the decoder and the source. Unicast uses IP delivery methods such as Transmission Control Protocol (TCP) and User Datagram Protocol (UDP), which are session-based protocols. When a decoder connects using unicast to a Vision² server, that client has a direct relationship to the server. Each unicast client that connects to the server takes up additional bandwidth. For example, if you have 10 clients all playing 1 Mbps streams, those clients as a group are taking up 10 Mbps. If you have only one client playing the 1 Mbps stream, only 1 Mbps is being used.

Unicast is used in applications like video on demand where each user is viewing the content on their own time frame, or when multicast video is reflected to an external location over non-multicast networks. Due to the increased network consumption, it is not suitable for applications where multiple viewers are receiving the same content simultaneously.

MULTICAST: Multicast is a one-to-one or more connection between multiple decoders and the source. The multicast source relies on multicast-enabled routers to forward the packets to all client subnets that have clients listening. There is no direct relationship between the decoders and the source, the decoders subscribe to a multicast group and the network ensures delivery of the stream. Each client that listens to the multicast adds no additional overhead on the server. The server sends out only one stream per source. The same load is experienced on the source whether only one client or 1,000 clients are listening.

Multicast on the Internet is not practical because the Internet is generally not multicast-enabled.
To extend Multicast streams over the Internet a Reflector is used to convert them to Unicast.

Multicast on enterprise networks

IP multicast is a mechanism for one sender sending data to multiple recipients, but only sending a single copy. It is accomplished by the sender forwarding UDP packets to a multicast IP address and port. The range of IP addresses reserved for multicast is 224.0.0.0 - 239.255.255.255. Without additional controls, such as IGMP and PIM (discussed below), multicasts are forwarded (flooded) to all ports like broadcasts. Unlike broadcasts, multicasts can be routed. Additionally while all broadcasts are processed by the network interface and passed up the stack to the host, multicasts are filtered by the NIC and only multicasts the host is subscribed to are processed.

Multicast Network Requirements

Layer 2 Services	<ul style="list-style-type: none">• Managed Switches• IGMP Snooping• IGMP Querier (One per Subnet)• Spanning Tree Protocol (STP)
Layer 3 Services (if required)	<ul style="list-style-type: none">• Multicast Routing• PIM Rendezvous point
Optional	<ul style="list-style-type: none">• Quality of Service (QoS)• Storm Control for Broadcast and Multicast

Multicast Addressing



Layer 3 Addressing

The range of IP addresses reserved for multicast is 224.0.0.0 - 239.255.255.255, however many address ranges are reserved for special purposes. Best practice for streaming is to use the range from 234.0.0.0 to 238.255.255.255, unless there is a specific reason to use other addressing.

Because of multicast Mac address overlap issue with the reserved Local Network Control Block (224.0.0.x) discussed in the next section, addresses with the format [224-239].0.x.x and [224-239].128.x.x should not be assigned.



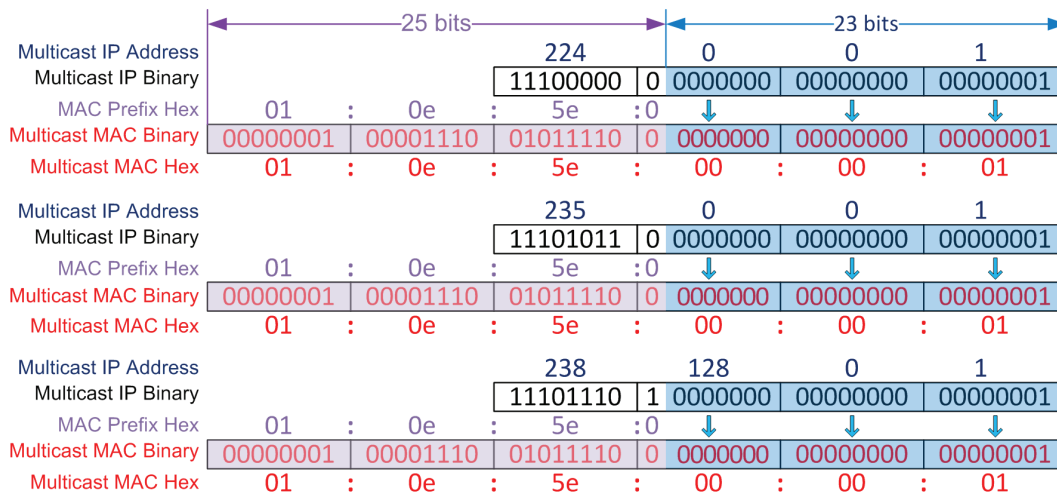
Multicast devices do not detect address conflict; many devices could transmit on the same multicast address without a failure. For example all Members of a group send IGMP Membership reports to the same multicast address. Conceivably all the streaming devices could be transmitting to the same multicast address on different ports. This is not the best practice in the case of streaming sources. If that is done all members of the group would receive all streams even if they were only viewing one.



Layer 2 Addressing

Multicast Addresses are not sent to unique physical devices. This presents a problem for Layer 2 transport which requires a destination MAC address. This has been solved by creating multicast MAC addresses. The first 25 bits of the 48 bit mac address are mapped to 0x 01-00-5E plus a 0 bit. The next 23 bits are the last 23 bits of the multicast address. There are 32 bits in an Ipv4 address, but in a multicast address the first 4 bits are always 1110. This means there are 5 bits in the IP multicast address that do not map to the MAC-layer multicast address and therefore it is possible for two different multicast groups to have the same destination MAC address.

Since there is an overlap of IP Multicast addresses to Ethernet MAC multicast addresses, care should be taken to avoid overlap by any two multicast addresses using the formula $[224-239].z.y.x$ and $[224-239].(z+128).y.x$. For example, as shown in the illustration below, 224.0.0.1 (the all hosts address), 235.0.0.1, and 238.128.0.1 all have the same multicast MAC address 01:0E:5E:00:00:01.



A good rule of thumb would be to use whatever ranges you choose, but keep the last two octets of the multicast unique within the organization within the multicast boundary...

Because the multicast MAC overlap issue is on layer 2 it is only significant within the subnet. If IGMP snooping is properly configured, it is only an issue at the host level occurring when the host is joined to a multicast group and one with and overlapping MAC is also forwarded to the same subnet. With the exception of not assigning addresses with the format $[224-239].0.x.x$ and $[224-239].128.x.x$ to avoid overlap with the reserved Local Network Control Block (224.0.0.x), the preceding is a Best Practice and not a hard rule.

Multicast on layer 2

IGMP

The Internet Group Management Protocol (**IGMP**) is a communications protocol used by hosts and their adjacent routers to allow hosts to inform the router of the desire to receive, continue receiving or stop receiving a multicast. IGMP Packets are only forwarded within a subnet and are sent with a TTL of 1.

The IGMP message is very simple. It consists of only four things, version, message type, checksum, and the group, i.e. multicast address, to be joined. There are only three message types: Membership Report, which doubles as a join message and is sent to the multicast group address; Queries, asking group members to report if they are still listening which are either sent to the multicast group address or All Hosts (224.0.0.1); and Leave Group, which is sent when a member wants to stop receiving the multicast and is sent to All Routers (224.0.0.2)

IGMP has two types of systems sending messages:

- 1 Querier:** The IGMP Querier is a process that runs on a switch or router. Its responsibility is to send out IGMP group membership queries on a timed interval, to retrieve IGMP membership reports from active members, and to allow updating of the group membership tables. There is one active Querier per subnet. If there is more than one Querier then the Queriers hold an election and the one with the lowest IP address is chosen to be active.

The Querier sends periodic Membership Queries to the All Hosts (224.0.0.1) address. It also sends out queries to a specific multicast address when it sees an IGMP leave message to check and see if there is still a listener on the network segment.

The Querier listens for Membership Reports and updates group membership tables used by the adjacent router to determine if the subnet should have a given multicast forwarded to it. The Querier removes group from the table after a timeout period if it has not seen a Membership Report during the period.

- 2 Group Member:** A group member is any client that has joined a multicast group. A Group Member joins the group by sending a membership report to the group multicast address. The network logs the Membership report in the group membership table (for the router) and the Switch IGMP cache (for IGMP Snooping) and the Member starts receiving the multicast.

A Group Member responds to IGMP Queries by sending a membership report to the group multicast address of all groups it is a member of for a general query and a membership report to the group multicast address for a group query, to keep the tables updated.

In IGMPV2 a Member sends an IGMP Leave to the to the group multicast address when it no longer wants to receive the multicast.

IGMP Snooping

The purpose of IGMP is to allow for pruning multicast forwarding at a router level. With the advent switching instead of hubs the need to prune multicast at a port level was desired. This is because a multicast group MAC address is never used as source MAC address for a packet so a 802.1 switch cannot learn them by the classical MAC address learning method and therefore, by default forwards multicast traffic to all ports. IGMP snooping was developed to deal with this issue.

In IGMP snooping the network switch listens for the IGMP messages and forwards the multicast packets from the VLAN only to the Ethernet ports that are sources of IGMP membership reports and keeps a cache, very much like the IGMP routing table to keep track of the members. Entries in the cache have a timeout function so if no Membership Reports are received the entry is removed from the cache.

This also works in switch to switch connections. The switch that supports IGMP snooping must flood all unrecognized IGMP messages to all other ports, therefore upstream switches receive new Membership reports to snoop and update their caches.

In the base IGMP protocol a Member responds to a Query after a random amount of time. If a member hears a response before the timer runs out, they do not respond. In IGMP snooping Membership Reports are suppressed to Members so all Members respond to all Queries.

There are three versions of IGMP. IGMP is fully backwards compatible, only additional features have been added in subsequent versions.

IGMPV1

DEFINED IN IETF RFC 1112

Initial Version which just handles IGMP joins. The stream is torn down by timeout.

IGMPV2

DEFINED IN IETF RFC 2236

The most common version implemented and covers most use cases. Added IGMP leave mechanism to tear down stream.

IGMPV3

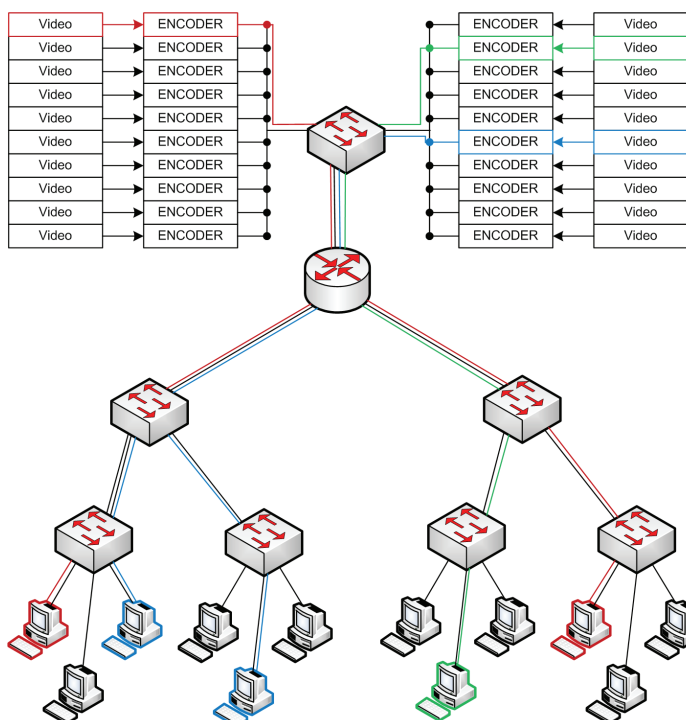
DEFINED IN IETF RFC 3376

Adds the ability for the receiving device to specify multicast source addresses as well as multicast addresses.

In IPV6, which does not use IGMP and is not covered in this document, RFC4541 defines the Multicast Listener Discovery (MLD) Snooping Switches which is based on IGMPV3.

Example

This illustration depicts multicast traffic across a properly configured layer 2 network. Although all the encoders are transmitting a multicast stream, only the streams with a host in the (color coded) multicast group are forwarded to the distribution switch. From the distribution switch on, each multicast stream is only forwarded on any given segment if there is a downstream host joined to the multicast group.



Multicast on layer 3

PIM

Protocol-Independent Multicast (**PIM**) is a family of multicast routing protocols for IP networks that provide for distribution of multicast data between routers and across networks. PIM builds trees (multicast routes) which ensure shortest path and loop suppression. There are four varieties of PIM:

PIM Sparse Mode (PIM-SM)

PIM Sparse Mode builds trees rooted at a Rendezvous point. Trees are built before any multicast packets are sent. PIM-SM can create shortest path trees for each source. PIM-SM scales well and is the most commonly used PIM mode for video in enterprises.

PIM Dense Mode (PIM-DM)

PIM Dense Mode uses dense multicast routing. It creates trees by flooding the network with all the multicast traffic and pruning back routes that are not subscribed to the Multicast Group. This is used in applications that almost all hosts are subscribed to a Multicast, but the flooding can create issues in bandwidth heavy applications like streaming.

PIM Source-Specific Multicast (PIM-SSM)

PIM Source-Specific Multicast builds trees that are rooted in just one source. It can be more secure than other implementations because clients subscribe to specific sources. PIM-SSM requires IGMPV3 to be implemented.

Bidirectional PIM

Bidirectional PIM explicitly builds shared bi-directional trees and scales well for applications that communicate between device pools on multicast. Rarely used for streaming unless implemented for another application.

Note: PIM Sparse-Dense

PIM Sparse-Dense mode is not a type PIM, but rather a router setting on Cisco routers which allow the interface to act in both PIM Sparse mode and PIM Dense mode simultaneously (on different multicast groups) to support an network with both modes used for different multicast groups and applications.

PIM Sparse Mode (PIM-SM)


In PIM-SM Routers can take one of two roles:

Rendezvous Point (RP): A function on a router which is responsible for keeping track of multicast sources and building trees to distribute multicast to other routers. With a Rendezvous Point other routers do not need to know the addresses of the sources for every multicast group. All they need to know is the IP address of the RP router. The RP router discovers the sources for all multicast groups and forwards multicast packets to designated routers requesting them.

The Rendezvous Point function can be manually assigned to routers or can be automatically assigned. Depending on the streaming application it is typically best to manually assign the RP designation to the router immediately adjacent to the encoder VLAN. This minimizes encapsulated data and uses the least router processing power.

Designated Router (DR): Any multicast router in a network that is not the Rendezvous Point. The function of a DR is to send and forwards PIM Join messages to the Rendezvous Point to initiate a multicast to its networks and if attached to a multicast source, forward encapsulated multicast packets to the RP for distribution.

Multicast Routes and Messages



PIM Messages

PIM messages are very similar to IGMP messages in that there are only a few message types and they may be used for multiple purposes. PIM messages are sent with the routers source address and the destination address 224.0.0.13, which is the reserved address for PIM V2 messaging. The PIM messages are forwarded along the multicast tree.

PIM V2 Messages

Hello

Periodic message sent out all multicast router ports for neighbor discovery and Designated Router election.

Register

Multicast data packets forwarded Encapsulated in Unicast from a DR to the RP to indicate a multicast group is available.

Join/Prune

A dual purpose message sent upstream, towards the sender, which lists streams to be received (Join) and any streams being received to stop (Prune). Sent whenever a multicast group is joined or dropped, as well as periodically.

Register Stop

A specific Prune message to a DR to tell it to stop forwarding Register Packets.

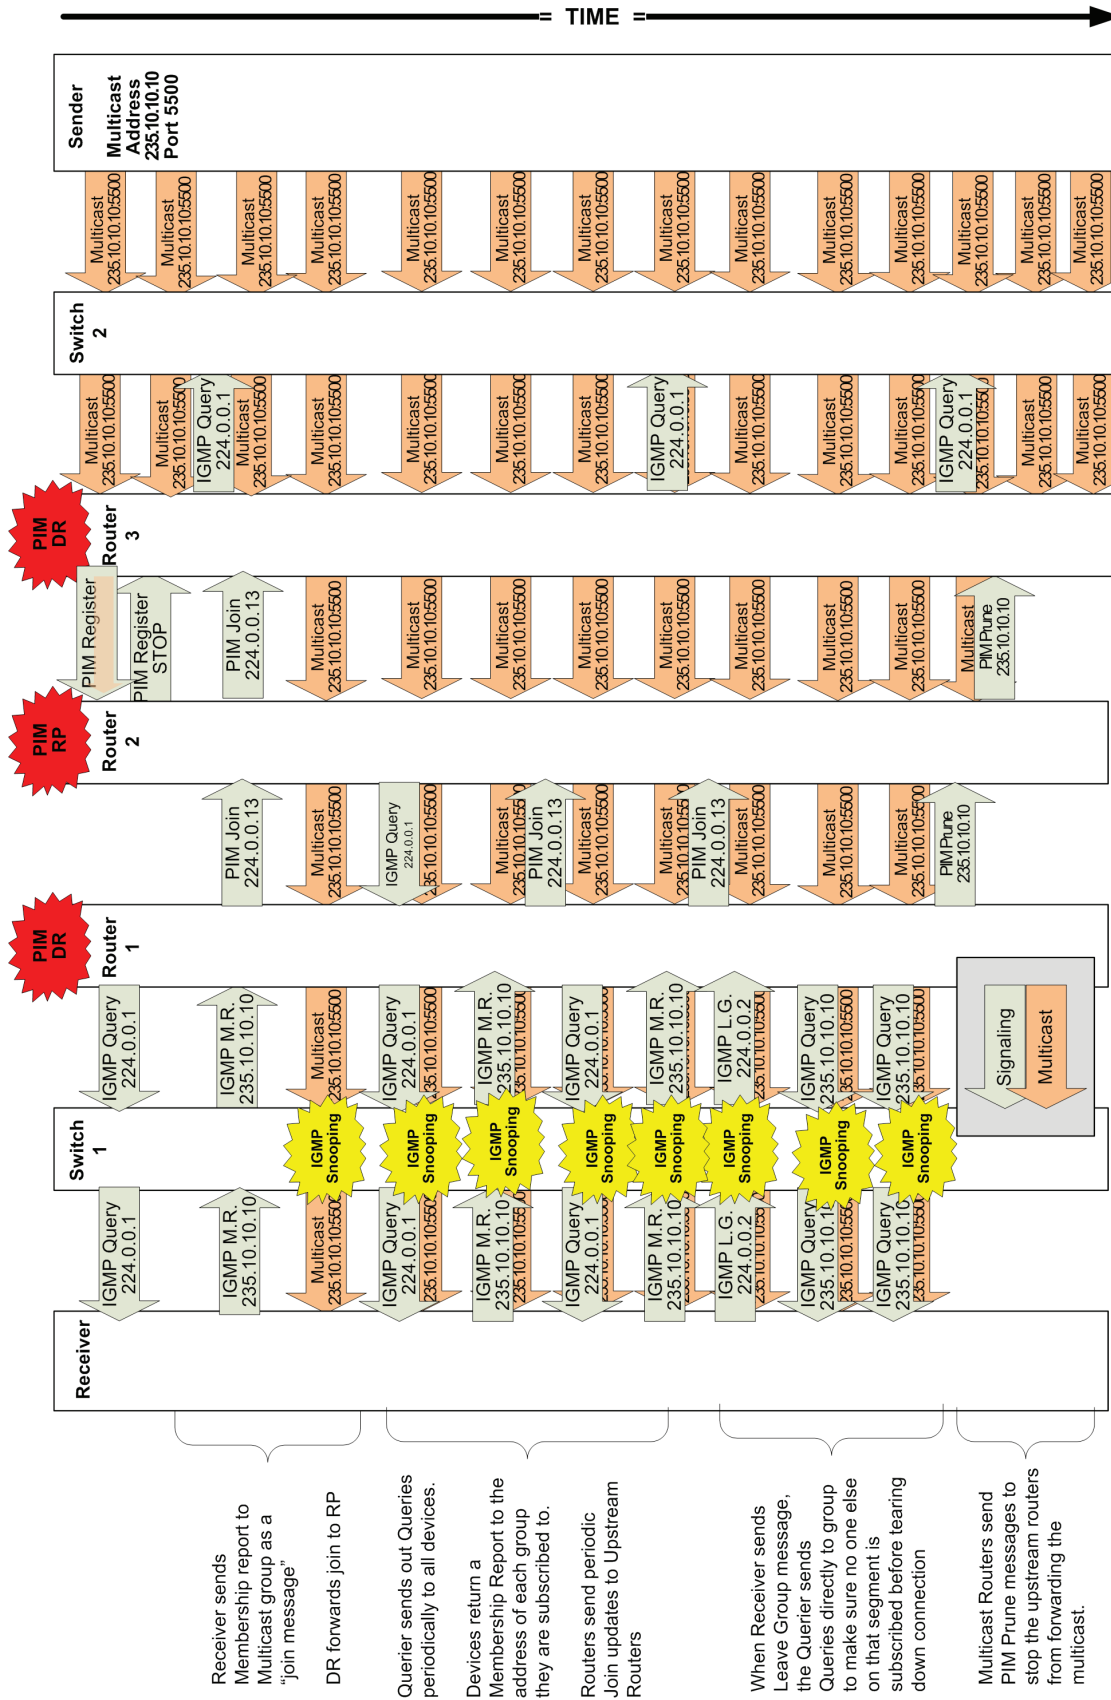
Multicast session example

The illustration on the next page depicts a complete multicast session between a multicast transmitter (Sender) and a multicast receiver (Receiver). Each host is attached to a switch and between the switches there are three routers. The session flow is top to bottom.

1. The Sender starts a multicast.
2. The multicast is forwarded to Router 3
 - Switch 2's IGMP snooping table forwards multicasts to multicast routers
3. Router 3 sends a PIM Register message periodically to the Rendezvous Point, Router 2.
 - The PIM Register message consist of a multicast packet encapsulated with a PIM header addressed to the Rendezvous Point
4. When Router 3 receives the PIM Register message it responds with a PIM Register Stop message.
5. Router 2 stops forwarding PIM Register messages.
 - It continues to receive the multicast from the sender.
6. The Receiver decides to receive the multicast and sends an IGMP Membership Report (MR) addressed to the multicast group (235.10.10.10) it wants to receive
7. Switch 1 receives the IGMP Membership Report and forwards it to the multicast router, Router 1
8. Router 1 sends a PIM Join message to the Rendezvous Point, Router 2 via the PIM V2 Multicast Address 224.0.0.2
9. The Rendezvous Point, Router 2 forwards the PIM Join to Router 1
10. Router 1 forwards the multicast traffic to the Rendezvous Point, Router 2
11. Router 2 forwards the multicast traffic to Router 1
12. Router 1 forwards the multicast traffic to Switch 1
13. Switch 1 looks up the multicast address in its IGMP Snooping Table and forward the multicast traffic to the Receiver
14. Router 1 continues to send periodic IGMP Membership Queries to the all host group (224.0.0.1)
 - There are no receivers so it doesn't get any replies.
15. Switch 1 forwards the IGMP Membership Queries to all ports
16. When the Receiver gets the IGMP Membership Queries it responds with a IGMP Membership Report address to the multicast group (235.10.10.10)
17. Router 1 receives the IGMP Membership Report and rests the IGMP timeout timer
18. Router 1 continues to send periodic PIM Join Messages to the Rendezvous Point, Router 2 to keep the session alive
19. When the Receiver wishes to stop receiving the multicast it sends a IGMP Leave Group message to the all routers group (224.0.0.2)
20. Switch 1 forwards the a IGMP Leave Group message to Router 1
21. Switch 1 removes the Receiver from the multicast group in the IGMP Snooping Table and stops forwarding the multicast traffic to the Receiver
22. Router 1 sends a series of IGMP Membership Queries to the multicast group (235.10.10.10) to see if there are any other hosts joined to the multicast group on the subnet
23. When Router 1 doesn't receive any membership reports for the multicast group (235.10.10.10) it sends a PIM Prune Message to the Rendezvous Point, Router 2
24. The Rendezvous Point, Router 2 sends a PIM Prune Message to Router 1
25. All Routers stop forwarding the multicast traffic

Not Shown:
Multicast
Routers send
out periodic
hello messages
out all their
interfaces.

MULTICAST STREAMING SETUP AND TEARDOWN



Other relevant network configurations

Spanning Tree

Spanning-tree is a layer 2 protocol designed to prevent broadcast storms caused by loops in the layer 2 topology. There are a couple of potential interactions between spanning-tree and multicast that should be kept in mind.

In many switches, when a switch receives spanning-tree Topology Change Notice (TCN), IGMP snooping will flood multicast traffic to all ports. This is based on the assumption that multicast traffic is critical and until the topology has converged multicast should be forwarded everywhere. If this is not desired or is causing a problem “TCN flooding” can usually be disabled.

When a switch receives spanning-tree Topology Change Notice (TCN) it puts all its ports in listening and then learning mode during which they do not forward packets. Depending on the spanning-tree protocol used, this can be from 10-50 seconds. This may not be a huge issue for traditional TCP traffic, but for a real time UDP service it is a major disruption. On at least the streaming source ports on the switch, an immediate forwarding setting like Cisco’s PortFast or Juniper’s Edge, should be set to avoid service interruption.

Storm Control

Storm control prevents traffic on a LAN from being disrupted by a broadcast, multicast, or unicast storm on a port. A LAN storm occurs when packets flood the LAN, creating excessive traffic and degrading network performance. Storm control uses thresholds to block and then restore the forwarding of broadcast, unicast, or multicast packets. You can also set the switch to shut down the port when the threshold is reached.

Storm control only works on inbound packets to a switch port so a careful application where there is a possibility of a multicast storm is important.

Because on layer 2, broadcasts are seen as a subset of multicast, when configuring storm control in a switch port, if you’re setting limits to both multicast and broadcast, you should set multicast limit higher than the broadcast limit, otherwise both broadcasts and multicasts will be limited by the multicast level. This is especially important in a spanning tree configuration to not block spanning tree BPDUs.

In AMX’s implementation, multicast clients do not request retransmission of multicast data and are therefore not likely to cause multicast storms. However, all multicast traffic is suppressed by multicast storm control, so if video is mission critical then consideration must be made.

IMPLEMENTING MULTITASK VIDEO STREAMING

Network preparation

Before any multicast streaming application can be installed a review of current multicast capabilities should be performed and any required configuration changes made.

Best Practices

1. Determine a multicast scope.
 - a. Is the multicast going to be available at all VLANs or limited to certain VLANs?
 - i. Will multicasts be selectively limited.
 1. You may not want employees watching TV at their desks but want them to be able to view meetings or HR videos.
 2. Certain channels or content may be sensitive in nature and not to be widely viewed.
 - ii. If this is the case, see the section Multicast Security
 - b. Look at the network topology to determine if there is a better location physically or logically for the multicast VLAN to minimize routed segments.
 - c. Determine the maximum number of layer 3 hops to intelligently set the TTL of the multicast stream.
 - d. Decide if you are going to use a static Rendezvous Point or Boot Strap Router (BSR) Rendezvous Point discovery.
 - e. Decide if you need to have separate Rendezvous Points different multicast groups.
2. Chose an unused multicast address range for the IPTV streaming application.
 - a. Chose an address range between 234.0.0.0-238.255.255.255, this is the largest unallocated block.
 - i. Since there is an overlap of IP Multicast addresses to Ethernet MAC multicast addresses, any multicast address in the [224-239].0.0.x and [224-239].128.0.x ranges should NOT be considered.
 - ii. Keep the first 3 octets of the ipv4 address the same and assign the last octet to individual Multicast streams. i.e. 235.10.1.xxx .
 - This will make configuring a dedicated RP easier.
3. Select a router to be used as the Rendezvous Point.
 - a. Look at the proposed traffic patterns and network topology.
 - i. If there are any network segments with bandwidth restrictions try to place the Rendezvous Point between the most possible sources and the restricted segment.
 - ii. Try to choose a point in the network to minimize the total number of hops between all the sources and the Rendezvous Point. (If there are many more sources at one location, put the rendezvous point there)
4. If you are using Boot Strap Router (BSR) RP discovery choose a backup Rendezvous Point.
5. Set a loopback port the chosen router as the rendezvous point.
 - a. **Option 1) (Preferred)**
 - i. Set a loopback port on the RP router as a Rendezvous Point candidate with the lowest priority number of all RP candidates. (lowest priority number is chosen)
 1. Allowed Range is 0-65535, default is 192
 2. In a complex network with multiple multicast applications, limit the RP to the IPTV Multicast group range
 3. Set the RP candidate announce to at least 15 seconds


- ii. Set a loopback port on the backup RP router as a Rendezvous Point candidate with the second lowest priority number of all RP candidates. (lowest priority number is chosen)
- iii. Make sure you have a Boot Strap Router (BSR) designated. If this is the only multicast application you can use the same router as a BSR, use a different loopback interface from the RP as BSR.

b. Option 2)

- i. Set a loopback port on the RP router as a static Rendezvous Point.
6. Enable IP multicasting on all layer 3 devices (routers and layer 3 switches) expected to pass multicast traffic.
 7. Enable Protocol Independent Multicast (PIM) on all layer 3 interfaces expected to pass multicast traffic.
 - a. Best practice for IPTV will employ a Rendezvous Point (RP) which will mean that the multicast will be managed in PIM-Sparse Mode.
 - b. Depending on the layer 3 device options you may configure the interface for PIM-Sparse-Dense mode if there is a PIM Dense mode application, otherwise configure the interfaces for PIM Sparse mode.
 8. Ensure the routers can find the Rendezvous Point.
 - a. If Boot Strap Router (BSR) RP discovery is used, enable the BSR listening and forwarding feature so that the BSR mechanism can dynamically inform all routers in the PIM domain of the RP set assigned to multicast group ranges.
 - b. If a static Rendezvous Point is going to be used configure all layer 3 devices with the static Rendezvous Point address.
 9. Ensure the multicast layer 3 interfaces are configured as IGMP Queriers.
 - Typically enabled by default when PIM is enabled.
 10. Check and note IGMP query interval and IGMP timeout interval on the Layer 3 devices.
 - a. The RFC default interval is 125 seconds; many manufacturers use a lower default interval. Typically the manufacturer default interval is fine.
 - b. The IGMP timeout interval should be at least 2 times the IGMP query interval. Some manufacturers use a timeout of ~2.1 times the IGMP query interval for a safety factor. Typically the manufacturer default interval is fine.
 11. Enable all switches expected to pass multicast traffic for IGMP snooping.
 - a. Typically this is a global configuration for all VLANS on a switch.
 - b. IGMP Snooping may, depending on manufacturer, be disabled in a per VLAN basis.
 12. Ensure the IGMP Snooping timeout on Layer 2 is at least 2 times the IGMP query interval.
 - a. On some switches it is set as a multiplier of the query interval as a default.
 - b. This is typically most important in a mixed manufacturer environment or when you depart from default settings.
 13. (Optional) Configure an IGMP Querier on a switch on each VLAN expected to pass multicast traffic. This will allow local multicast in the case of a router failure.
 - a. This is more applicable in a setup which does not include multicast routing.
 - b. If the switch has an IP address in a multicast VLAN (not a good idea unless the switch is a Layer 3 switch performing routing) the IP address should be higher than the Router IP address so the router is elected IGMP Querier.
 14. Test each VLAN for multicast access, IGMP joins and IGMP leaves.

Scenario 1 enterprise IPTV

In an enterprise IPTV installation, multiple encoder and/or producer channels stream multicast content continuously over the network. The multicast streams can be received by a variety of devices including desktop computers and set top boxes. The nature of IPTV is such that “channels” will be changed creating a changing multicast topology. In this situation it is important that multicast be properly configured to minimize network traffic where it is not required. These best practices apply when the multicast sources can be placed within a single VLAN.



Best Practices

1. Chose an unused multicast address range for the IPTV streaming application.
 - a. Chose an address range between 234.0.0.0-238.255.255.255, this is the largest unallocated block.
 - b. Since there is an overlap of IP Multicast addresses to Ethernet MAC multicast addresses, any multicast address in the [224-239].0.0.x and [224-239].128.0.x ranges should NOT be considered.
 - c. Keep the first 3 octets of the ipv4 address the same and assign the last octet to individual Multicast streams. i.e. 235.10.1.xxx.
 - This will make configuring a dedicated RP easier.
 - d. Remember that each Producer channel will require an individual multicast address.
2. Choose a Destination port and use the same port for every stream
 - e. For RTP the port number must be even, RTCP uses the next odd port. (not applicable for Transport stream)
3. Group the Multicast Sources on a single, dedicated VLAN
 - f. If the Multicast Source VLAN spans multiple switches, use a trunk dedicated to the Multicast Source VLAN between the switches.
 - i. Make the Multicast Source VLAN native on the trunk
 - ii. Make sure the Multicast Source VLAN is not included on other trunk ports
 - g. If you are using a streaming server with producer channels a separate multicast interface should be used on the Multicast Source VLAN. The control and Video on Demand functions should bind to a different VLAN.
4. Disable Spanning Tree or use a configuration like Cisco “PortFast” or Juniper “Edge”, on the access ports with streaming sources.

5. Use a dedicated router port, not a shared, trunked interface, for the Multicast Source VLAN
6. Set a loopback port on the same physical router dedicated to the Multicast Source VLAN as the rendezvous point.
7. Choose and configure a backup Rendezvous Point


Scenario 2 room overflow

Often in an enterprise environment, VLANs are terminated in a communications closet, known as an intermediate distribution frame (IDF) within 100 meters of all the hosts and there is no Layer 2 connectivity between that point and the rest of the network. The logical network topology is at least partially dictated by the physical layout. With this topology an application like a room overflow, where the multicast sources are distributed around the campus it is impractical to segregate the multicast sources on a single VLAN.

Best Practices

There are multiple strategies for approaching this topology. In this case the Best Practices are broken into large installations and small installations. The main differentiators for this application between a large installation and a small installation are the network engineer's analysis of traffic flow patterns, along with the number of devices and available bandwidth between locations. If the locations are joined by a VPN, consult the "Multicast Over VPN" section of this document.

Large Installation



These are the best practices with a large number of transmitters and receivers the concept of location can be abstracted to whatever unit (Floor, Building, Department, etc.) makes sense for the logical and physical topology.

1. Choose an unused multicast address range for each location for the streaming application.
 - a. Choose an address range between 234.0.0.0-238.255.255.255, this is the largest unallocated block.
 - b. Since there is an overlap of IP Multicast addresses to Ethernet MAC multicast addresses, any multicast address in the [224-239].0.0.x and [224-239].128.0.x ranges should NOT be considered.
 - c. Keep the first 3 octets of the ipv4 address the same and assign the last octet to individual Multicast streams. i.e. 235.10.1.xxx .
 - i. This will make configuring a dedicated RP easier.
 - ii. If supported by the Layer 3 manufacturer; smaller address ranges can be chosen, broken on binary boundaries. i.e. 235.10.1.16-235.10.1.31 (/28)
 - d. Remember that each Producer channel will require an individual multicast address.
2. Chose a Destination port and use the same port for every stream
 - For RTP the port number must be even, RTCP uses the next odd port (not applicable for Transport stream).
3. Determine if the physical and logical topology lends itself to a separate VLAN per Location for the multicast encoders (preferred).
 - a. If so; Group the multicast sources and dedicated receivers (usually set top boxes) on a single, dedicated VLAN per Location and add Multicast VLAN to trunks between switches.

Option 1

The multicast VLAN can be included on the 802.1Q trunks between the switches.

When using this method, make sure you understand any potential bandwidth issues for the traffic in any given trunk link.

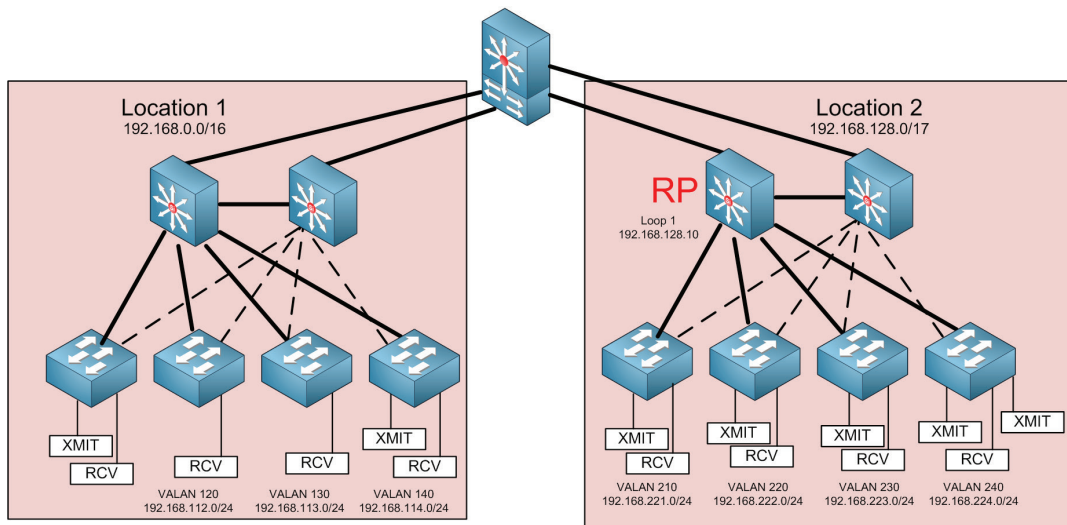
Option 2

Use a trunk dedicated to the Multicast VLAN between the switches.

Make the Multicast Source VLAN native on the trunk, and make sure the Multicast Source VLAN is not included on other trunk ports.

- b. If not; the multicast sources and receivers are located on the general data network.
4. Set a loopback port on each Location router as the rendezvous point.
 - a. Make the router the lowest rendezvous point candidate
 - b. Limit the RP to the IPTV Multicast group range chosen above
 - c. Ensure BSR is configured
5. Choose and configure a backup Rendezvous Point.

Small Installation



1. Choose an unused multicast address range for the streaming application..
 - a. Choose an address range between 234.0.0.0-238.255.255.255, this is the largest unallocated block.
 - b. Since there is an overlap of IP Multicast addresses to Ethernet MAC multicast addresses, any multicast address in the [224-239].0.0.x and [224-239].128.0.x ranges should NOT be considered.
 - c. Keep the first 3 octets of the ipv4 address the same and assign the last octet to individual Multicast streams. i.e. 235.10.1.xxx .
 - i. This will make configuring a dedicated RP easier.
 - ii. If supported by the Layer 3 manufacturer; smaller address ranges can be chosen, broken on binary boundaries. i.e. 235.10.1.16-235.10.1.31 (/28)
 - d. Remember that each Producer channel will require an individual multicast address.
2. Chose a Destination port and use the same port for every stream
 - a. For RTP the port number must be even, RTCP uses the next odd port. (not applicable for Transport stream)
3. Determine if the physical and logical topology lends itself to a separate VLAN per Location for the multicast encoders (preferred).
 - If so; Group the multicast sources and dedicated receivers (usually set top boxes) on a single, dedicated VLAN per Location and add Multicast VLAN to trunks between switches.

Option 1

The multicast VLAN can be included on the 802.1Q trunks between the switches.

When using this method, make sure you understand any potential bandwidth issues for the traffic in any given trunk link.

Option 2

Use a trunk dedicated to the Multicast VLAN between the switches.

Make the Multicast Source VLAN native on the trunk, and make sure the Multicast Source VLAN is not included on other trunk ports.


4. Select a router to be used as the Rendezvous Point.
 - a. Look at the proposed traffic patterns and network topology.
 - i. If there are any network segments with bandwidth restrictions try to place the Rendezvous Point between the most possible sources and the restricted segment.
 - ii. Try to choose a point in the network to minimize the total number of hops between all the sources and the Rendezvous Point. (If there are many more sources at one location, put the rendezvous point there)
5. Set a loopback port on router selected as the rendezvous point.
 - a. Make the router the lowest rendezvous point candidate
 - b. Limit the RP to the IPTV Multicast group range chosen above
 - c. Ensure BSR is configured
5. Choose and configure a backup Rendezvous Point.

Scenario 3 non-routed multicast

In many cases, such as Digital Signage, IPTV distribution to set top boxes, or room to room transmission no layer 3 transitions are required as all the multicast transmitters and receivers are on the same Subnet. The multicast can be on a totally isolated network as shown in Figure X1 or use shared infrastructure as shown in Figures X2 and X3.

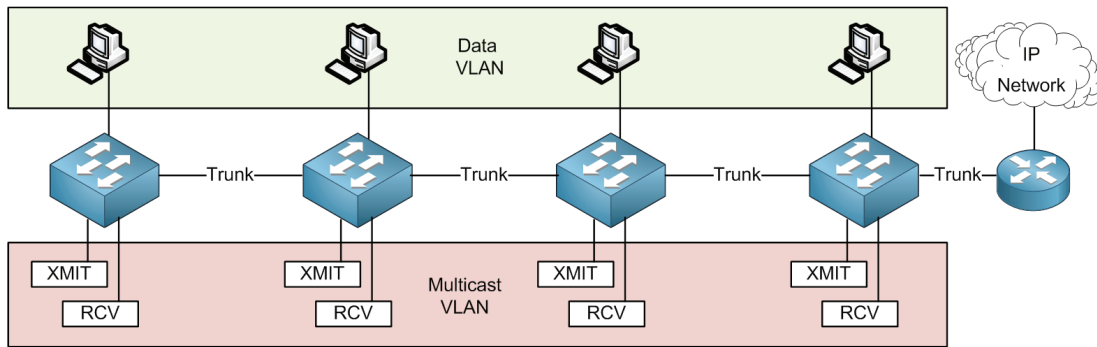
Best Practices

Stand Alone Infrastructure

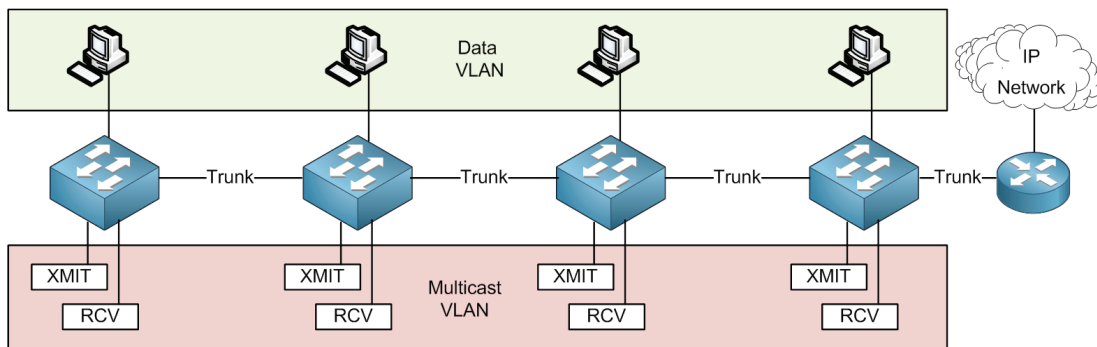


1. Choose a multicast address range between 234.0.0.0-238.255.255.255 for the streaming application.
 - Remember that each Producer channel will require an individual multicast address.
2. Choose a Destination port and use the same port for every stream
 - For RTP the port number must be even, RTCP uses the next odd port. (not applicable for Transport stream)
3. Enable all switches to pass multicast traffic for IGMP snooping.
 - Typically this is a global configuration for all VLANs on a switch.
4. Configure an IGMP Querier on at least one switch to pass multicast traffic.
 - If the switch has no IP address it may be required to set an IP address in the switch or specify a Querier Source address
5. (Optional) If the standalone Network spans multiple switches, use a trunk dedicated to the Multicast VLAN between the switches.

Shared Infrastructure — fig. 1



Shared Infrastructure — fig. 2



1. Choose a multicast address range between 234.0.0.0-238.255.255.255 for the streaming application.
 - Remember that each Producer channel will require an individual multicast address.
2. Choose a Destination port and use the same port for every stream
 - For RTP the port number must be even, RTCP uses the next odd port. (not applicable for Transport stream)
3. Create a separate VLAN on the shared infrastructure for Multicast.
4. Add Multicast VLAN to trunks between switches.

Option 1 (fig. 1)

The multicast VLAN can be included on the 802.1Q trunks between the switches.

When using this method, make sure you understand any potential bandwidth issues for the traffic in any given trunk link.

Option 2 (fig. 2)

Use a trunk dedicated to the Multicast VLAN between the switches.

Make the Multicast Source VLAN native on the trunk, and make sure the Multicast Source VLAN is not included on other trunk ports.

5. Enable all switches to pass multicast traffic for IGMP snooping.
 - Typically this is a global configuration for all VLANS on a switch.
6. Configure an IGMP Querier

Option 1

Configure an IGMP Querier on at least one switch to pass multicast traffic.

If the switch has no IP address it may be required to set an IP address in the switch or specify a Querier Source address.

Option 2

Configure PIM on the router interface for the Multicast VLAN

Add a sub-interface for the Multicast VLAN on a trunk connected to a router.
Enable Multicast on the router.
No Rendezvous needs to be configured.

APPENDIX 1, MULTICAST SECURITY

Blocking multicast to specific networks

By its nature, multicasts transmissions are insecure. In a fully enabled Any Source Multicast (ASM) network any multicast can be joined by any host. Even in networks with protected middleware obscuring the multicast addresses, if the multicast addresses are discovered the groups can be joined and the content viewed.

There are multiple strategies for limiting multicasts from being available on unauthorized networks. Various equipment manufacturers implement these strategies in different ways but most should be available on enterprise grade equipment.

Denying Multicast Traffic

Disabling multicast to selected VLANs entirely will eliminate the vulnerability, but may not be practical in cases where some multicasts should be available and others restricted.



Advantages

- Easy to configure, disable PIM on the interface
- Very secure, No multicast traffic is allowed to pass the boundary



Disadvantages

- Other applications may require multicast

Administratively Scoped Multicast

In a scenario where some multicasts need to be constrained within a defined contiguous network, an administratively scoped multicast can be used. In the illustrations below the network in Figure X1 can be administratively scoped, the network in figure X2 cannot.

Fig. X1





Fig. X2



In an administratively scoped multicast the Multicast groups are assigned within the range of 239.0.0.0-239.255.255 and administrative boundaries are set at the edge of the networks. No multicast traffic in the administratively scoped range is allowed to cross the boundary in or out.



Advantages

- Routers** only have to be configured at the edge of the network
- Very** secure, No administratively scoped traffic is allowed to pass the boundary
- Other** multicast traffic can pass the boundary



Disadvantages

- Multicast** source must be within the boundary
- Very** granular, no flexibility for exceptions
- Specific** network topology required
- Boundaries** cannot overlap

Access Control Lists (ACL)

Access control lists can be set on the router interfaces to deny specific multicast traffic to a VLAN. With a combination of filters a high degree of security can be maintained.

Access Control Lists (ACL)



IGMP Filters

When specific IGMP join messages are blocked to the layer 3 interface the router will never establish the connection to the multicast stream effectively blocking the multicast to the subnet. If the content is highly sensitive, IGMP filtering only may not provide enough security as network errors such as spanning tree recalculation may cause multicast to be flooded to the network.

Remember IGMP is a separate layer 3 protocol and will not be filtered with TCP or UDP. IGMP traffic should be filtered inbound to the layer 3 interface. Place a global permit statement at the end of the ACL to avoid implicit deny.

Some switches allow IGMP filtering at the switch port level. This can be useful if only a subset of the clients need to have specific multicasts blocked.



UDP Filters

UDP traffic destined to the multicast group can be filtered at the layer 3 outbound interface. This will block any multicast from the filtered groups from the subnet. UDP filters should not be used alone, because there could be a strain on the layer 3 processor from the large volume of packets if the router receives an IGMP join. Additionally unwanted traffic may be sent to the router only to be blocked.



IP Filters

Depending on the manufacturer filtering all IP traffic from any source to the multicast group both inbound and outbound at the interface can have the same outcome as the other two filters above.



Advantages

- Can** be very specifically configured
- Can** be very secure, blocking specific multicast groups
- Other** multicast traffic can pass the boundary



Disadvantages

- Complex** implementation, ACLs must be configured on each interface to be filtered
- ACLs** can be cumbersome to administrate

Rogue multicast

In some networks, such as college campuses, it may be desirable to limit individuals from setting up rogue multicasts. Rogue multicasts can consume bandwidth, or if misconfigured corrupt legitimate multicasts.

It is good practice to periodically check the multicast groups on your routers to see if there is any unauthorized multicast traffic. Rogue multicasts can be blocked by a variety of methods depending on the level of threat.

PIM Rendezvous Point Multipoint Group Restriction

If the PIM Rendezvous Points are configured for only the multicast groups authorized on the networks, rogue multicast transmitters attempting to transmit on an unauthorized IP address. The multicast can still transmit within the attached subnet.

IGMP and PIM operate using layer 3 addressing so a rogue multicast could still use an authorized multicast group on a different port.

Access Control List

An ACL can be configured on the inbound interface to the layer 3 interfaces filtering UDP multicast traffic to the network. The multicast can still transmit within the attached subnet.

IGMP Filters

IGMP filters applied as group policy on the access switch interfaces will keep hosts within the subnet from joining the rogue multicast group within the attached subnet.

Source Specific Multicast (SSM)

Any host can transmit on a valid multicast address. There are many applications where multiple hosts use the same multicast destination. This capability could present a vulnerability where a rogue host transmits to a multicast group assigned to another application with the intent to corrupt the media stream or exploit another vulnerability in the receiving host.

In IGMPv3 the ability for hosts to request to join a multicast group and only receive transmissions from a specific source was introduced. To implement Source Specific Multicast (Cisco sometimes terms it Single Source Multicast) the edge layer 3 device and the client device must support and be configured for IGMPv3. No other devices in the path need to support IGMPv3. Check the client application and middleware for IGMPv3 support.

APPENDIX 2, MULTICAST OVER 802.11 (WI-FI) WIRELESS

From an OSI model perspective 802.11 is executed on layer 1 and 2. All layer 3 protocols including UDP and IGMP used in multicast streaming are passed transparently as if they were on an Ethernet network. In a simple access point connected to the network scenario, all the potential issues with multicast are on the layer 1 and layer 2 (the Physical and MAC Layer). The primary differences between the various flavors of 802.11 (a, b, g, n, ac) are in radio frequency, bandwidth, and modulation type. Although there are other differences in signaling, from a multicast point of view they are very similar, so the discussion will be about generic 802.11 multicast.

Challenges to 802.11 Multicast

The 802.11 standard allows for multicast transmission as part of asynchronous services. Asynchronous services are provided by the 802.11 layer 2 protocols Carrier Sense Multiple Access / Collision Avoidance (CSMA/CA). CSMA/CA is a mechanism for multiple stations (hosts and access points) to time share the same radio frequency. This is accomplished by a sender listening to the frequency for a period of time and if it does not sense another transmission it is clear to transmit. In wireless networks it is not possible to detect a collision (two stations transmitting at the same time) so in unicast transmissions a layer 2 acknowledgement is sent by the receiving station to inform the sender that the frame has been received. In 802.11 multicasts once the multicast session established, no acknowledgement is returned, this increases the risk for lost frames due to collisions.

802.11 wireless is a shared, half duplex (one transmitter at a time) network. CSMA/CA is designed so that each station has roughly equal access to the radio channel so the more stations attached to the access point the higher the potential latency and lower effective throughput to an individual station or multicast group. Additionally 802.11 stations dynamically adjust the physical transmission rate and modulation used based on the signal quality (strength based primarily on distance) which in 802.11g can vary the data rate between 1 Mbps and 54Mbps. To solve this potential issue, most access points as a default setting, transmit broadcasts and multicasts at a 1Mbps Channel speed.

802.11 has a feature within it for power savings in the client stations which allow the wireless card to go to sleep and use very little power. If the wireless card is powered down then the client station would not be able to receive necessary broadcasts and multicasts. To solve this, when a client is in power saving mode it informs the access point which buffers the broadcast and multicast traffic. Periodic beacons sent out by the access point, typically about every 100ms. When an access point knows a client is in power savings mode, a Delivery Traffic Indication Message (DTIM) is sent with a periodic beacon. A DTIM interval is set within the access point to establish how often the DTIM is sent. If the DTIM interval is set to 5, the device in power savings mode only has to wake up every 5th beacon, to see if it has broadcast or multicast data it needs to listen to. This means that all multicast data would be buffered for ~500ms before it is forwarded. If no clients are in power savings mode, the multicast is forwarded through without being buffered.

802.11 Multicast implementation

Because of these challenges, the various equipment manufacturers have different strategies on how multicast is supported. In the simplest mode, with a single access point, connected directly to the IP subnet (not routed) The IEEE specification says the process should like below. Note that all communication is on layer 2 and the frames mention are 802.11 frames in wireless which are converted into Ethernet 802.3 frames at the access point.

- The client station sends a directed (Layer 2 unicast) frame, which requires an acknowledgement, containing multicast data (an IGMP Join) to the access point.
- The access point responds with an 802.11 acknowledgement frame.
- If the acknowledgement is not received the client station resends the frame.
- •The access point forwards the IGMP Join across the network.
- •When the access point receives the requested multicast, it forwards the packets on 802.11 group (broadcast or multicast) frames which require no acknowledgement and can be consumed by all client stations.
- •When the client station is done with the multicast it sends an IGMP Leave in a directed frame.
- The access point responds with an 802.11 acknowledgement frame.
- The access point forwards the IGMP leave across the network.
- The Ethernet network IGMP process determines if multicast continues to other clients.



Research your hardware

Each manufacturer has entered into multicast slightly differently, often from product to product. Features to research to determine the applicability of your wireless solution include:

- Is multicast supported, especially in managed systems?
- In access points with integrated routers
- Does the router support multicast routing?
- This is important if you are running separate SSIDs to separate networks.
- Can the wireless router be configured as an access point?
- What is the default and maximum wireless multicast rate? Can I set the rate?
- Are there multiple radios and can I set them to different modes? i.e. 1 auto, 1 802.11g, etc.
- Does the access point support IGMP snooping?
- Is there a built in multicast to unicast reflector?



Other components of an enterprise wireless system may change the way multicast works. Some of those issues may be:

- Often all IP traffic to managed access points is routed through wireless controllers
 - › The controller must support multicast routing and PIM
- Some wireless controllers do not support multicast
- Some controlled access points will only accept multicast traffic with a source address of the controller
 - › Controller has to act as a multicast gateway.


Optimizing 802.11 implementation

- Set the Delivery Traffic Indication Message (DTIM) Set DTIM interval to 1 to minimize delay for multicast traffic.
 - › This may make some devices have shorter battery life or potentially miss some broadcast traffic.
- If possible raise the Multicast Rate from 1Mbps
 - › Client stations connected at a lower bandwidth will not receive the multicast.
- Use multiple radio access point and segregate the 802.11b traffic from the other 802.11 networks to keep the 802.11b endpoint from lowering the bandwidth.
- Stream at a lower bandwidth.
 - › If you want the content streamed on the corporate network at 8 Mbps, consider a transcoded or second source stream at a much lower bandwidth.
- If you don't have a need for a large number of simultaneous viewers on wireless, consider reflecting the stream to unicast. This will allow for users connected at higher negotiated bandwidths.
- Experiment

APPENDIX 3, MULTICAST OVER VPN

Why VPN

In the Illustration below an organization has two locations, each with their own private IP address space. While each site has internet access, the networks are not routable so internal services are not available from one site to another. In this case the file server at Site 1 is not available to the clients at Site 2.



There are several possible solutions to the above situation, including:

Make the file server available through the internet by NAT forwarding.

- This is not very secure; the file server is potentially visible to everyone with an Internet connection.
- This is not very scalable; each service needs to be set up individually and the number of services available may be limited by the number of Public IP addresses available.
- This is not reliable
- This is not expensive; practically free. This is often used by small businesses or home users who have a low risk application like a web or FTP server they want available from another site.

Set up a private, point to point link between the sites with a carrier network like T1 or Frame Relay

- This is secure; data stays on the organizations private network and is not accessible from public networks.
- It is scalable; additional bandwidth can be purchased as needed and all the Private IP address space is routable.
- This is expensive; separate circuits have to be provisioned separately from the internet access the organization is already paying for.

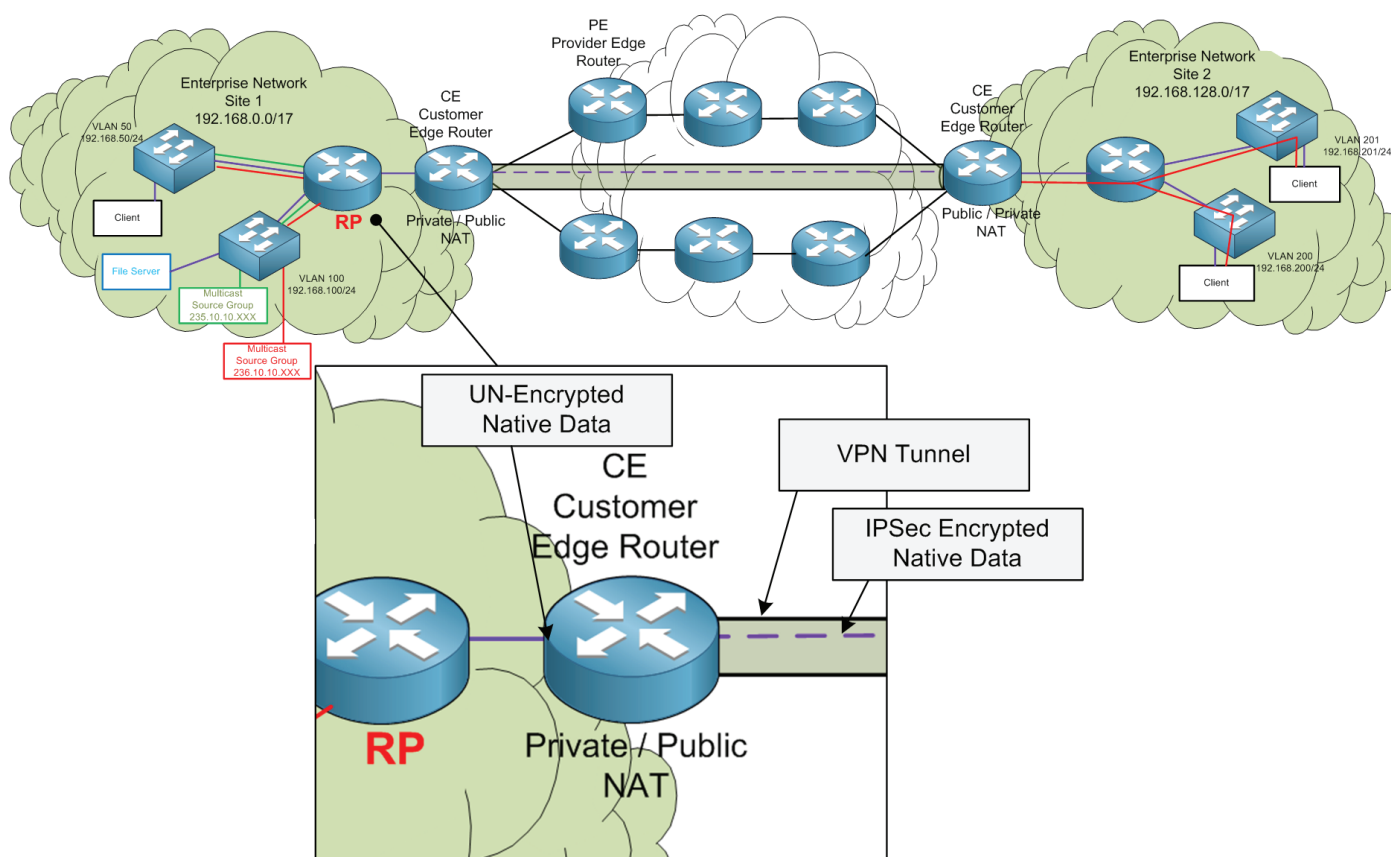
Set up a Virtual Private Network (VPN) between the sites over the existing internet connection.

- This is secure; although the information travels over the Internet, the traffic is encrypted and encapsulated so it is secure and the IP address space is not visible from outside the organization.
- It is scalable; bandwidth is limited only by the available public bandwidth and all the Private IP address space is routable.
- It is available; VPN can be scaled beyond point to point with multiple offices and home users.
- This is not expensive; often the same bandwidth and routing hardware used for internet access is used for the VPN.

Because of this Virtual Private Networks (VPNs) are the most common way that organizations link two or more sites. When adding a video streaming application to these organizations, chances are they already have some form of VPN between sites.

About VPN

A VPN is network connectivity across a shared infrastructure (such as the internet). It aims to provide the same policies and services as a private network, at a reduced cost of ownership. In the illustration below, the two sites from the previous Illustration have been joined by a VPN between the two Customer Edge Routers and shows a Virtual tunnel traversing the public network.



This allows a valid IP route between the Client at Site 2 and the file Server at Site one. Once a VPN is in place access to any permitted network service, such as Active Directory, email, printing, and instant messaging is available between the sites.

Pros and Cons

With VPNS, the biggest advantage is that they provide a more secure way for site-to-site traffic to communicate over the internet which is a public space. VPN tunnels encrypt traffic that is sent across the public using a protocol called IPsec (Internet Protocol Security), which encrypts each packet entirely, including the IP header. The entire Packet is the encapsulated into one of several VPN protocols and forwarded through the public network. If the packet is intercepted, all the information, including the original source and destination is encrypted and cannot be viewed.

VPN tunnels come with several disadvantages as well. VPN configurations must include statically maintained access lists to identify traffic allowed through the tunnel. This can become a tedious process for larger networks. Discontinuous subnets require separate tunnels. Due to limitations in

IPsec, VPNs do not allow multicast traffic to pass, therefore dynamic routing protocols, such as RIP and OSPF, are no longer options to use across VPN. Another consequence of this is that multicast video cannot be forwarded across a VPN natively.

Inter-site video streaming over VPN

The Internet and VPNs generally do not forward multicast traffic. There are starting to be some MPLS providers with a multicast option, but they are not widely deployed.

There are Three main methodologies for forwarding multicast traffic between two sites that that have one or more network segments that are not multicast compliant, Reflecting, Generic Routing Encapsulation (GRE) Tunnels and Automatic Multicast without explicit Tunnels (AMT).

Considerations in implementing multicast video streaming across VPN

Bandwidth

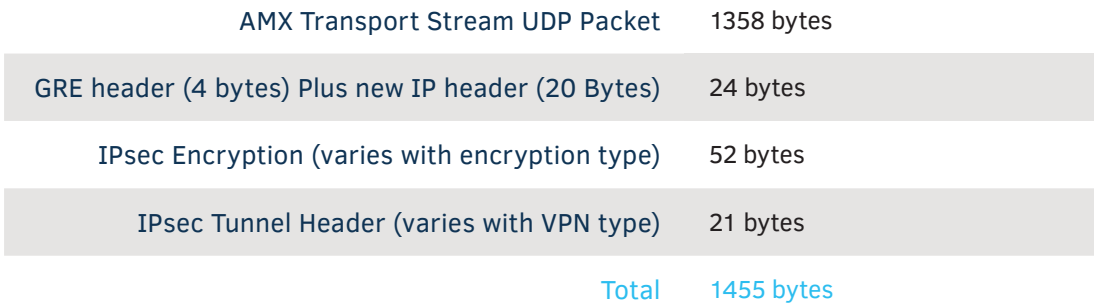
Streaming media can consume considerable bandwidth. With an application like IPTV with multiple channels at typical bandwidths of 6Mbps, bandwidth consumption can add up quickly.

Encapsulation, Overhead, and MTU Size

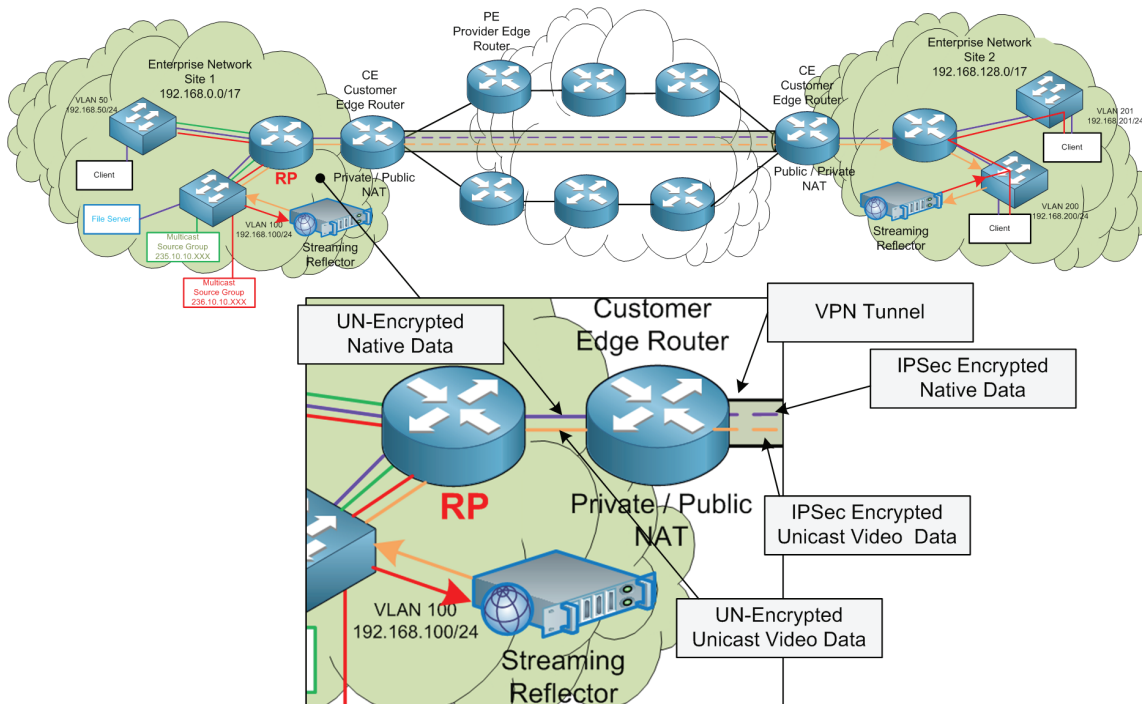
The process of forwarding multicast over a VPN Tunnel adds about a 7% head to the traffic. The additional overhead of TCP, in the case of reflection or the process of a 2 step encapsulation in the case of a GRE tunnel can potentially cause oversize MTUs which leads to packet fragmentation, which causes even more overhead, plus places increased processor load on the routers and receiving hosts.

Many manufacturers of video streaming products set the default MTU size on their equipment to Jumbo Frames to lower overhead. Jumbo Frames are any Ethernet packet with a packet size of more than 1500 bytes. While many LANs support Jumbo Frames, VPNs generally don't. AMX sets their packet size to 1358 bytes (7 transport stream packets) to allow for encapsulation overhead. The rule of thumb the maximum MTU size for a streaming application forwarded over a VPN is 1400 bytes, but the actual numbers should be checked.


© 2015 Pearson Education, Inc. or its affiliate(s). All rights reserved. Pearson Education, Inc., publishing as Pearson Benjamin Cummings, 101 Philip Drive, Assinippi Park, New York, NY 10984-2135. Pearson Education, Inc., publishing as Pearson Education, Inc., 501 Boylston Street, Boston, MA 02116-5093. Pearson Education, Inc., publishing as Pearson Education, Inc., 100 Brook Hill Drive, West Nyack, NY 10994-2133. Pearson Education, Inc., publishing as Pearson Education, Inc., 3501 Market Street, Philadelphia, PA 19104. Pearson Education, Inc., publishing as Pearson Education, Inc., 100 N. Zeeb Road, Ann Arbor, MI 48106-1546. Pearson Education, Inc., publishing as Pearson Education, Inc., 501 Boylston Street, Boston, MA 02116-5093. Pearson Education, Inc., publishing as Pearson Education, Inc., 100 Brook Hill Drive, West Nyack, NY 10994-2133. Pearson Education, Inc., publishing as Pearson Education, Inc., 3501 Market Street, Philadelphia, PA 19104. Pearson Education, Inc., publishing as Pearson Education, Inc., 100 N. Zeeb Road, Ann Arbor, MI 48106-1546.



A streaming video reflector (sometimes called a relay) subscribes to a video stream and re-transmits it to another address. This re-transmission can be any combination of multicast or unicast inputs and outputs. In the case of forwarding multicast across a VPN, a pair of reflectors can be used as shown below.



In the illustration, a streaming source, outputs a multicast stream. The reflector service subscribes to the stream and de-encapsulates layers 3 and 4 (IP and UDP headers). It then re-encapsulates the data with new TCP and IP headers and forwards the packet. The receiving end receives the unicast stream and performs the reverse process forwarding a multicast stream.



Considerations in Implementing Multicast Reflecting



Administration

The reflector is typically configured by the same administrator that is responsible for the streaming service. Beyond the initial configuration at installation, no network configuration is required to change the reflector service. For occasional use like company announcements a reflector at the source site can be configured and left active with the receive site reflectors for the duration of the event.



Bandwidth

Since the reflected streams are individually configured, there is no risk of inadvertently forwarding multicast streams.

The receiving side reflector requests the stream based on local administration. While the receive site reflector service is enabled, the bandwidth is used even if no one is subscribed to the multicast stream.

There is a small packet overhead increase with the conversion from UDP to TCP (> 1%).



Scalability

A single reflector at the source location can reflect to multiple receive site reflectors, but the bandwidth is unicast so each receive side reflector gets a separate stream.

Configuration Considerations

A separate PIM Rendezvous Point will need to be configured at the receiving site. Multicast addresses do not need to map between the sites. Each site can have its own addressing scheme. For IPTV applications a separate channel guide will have to be implemented if the multicast addresses do not map between the sites.

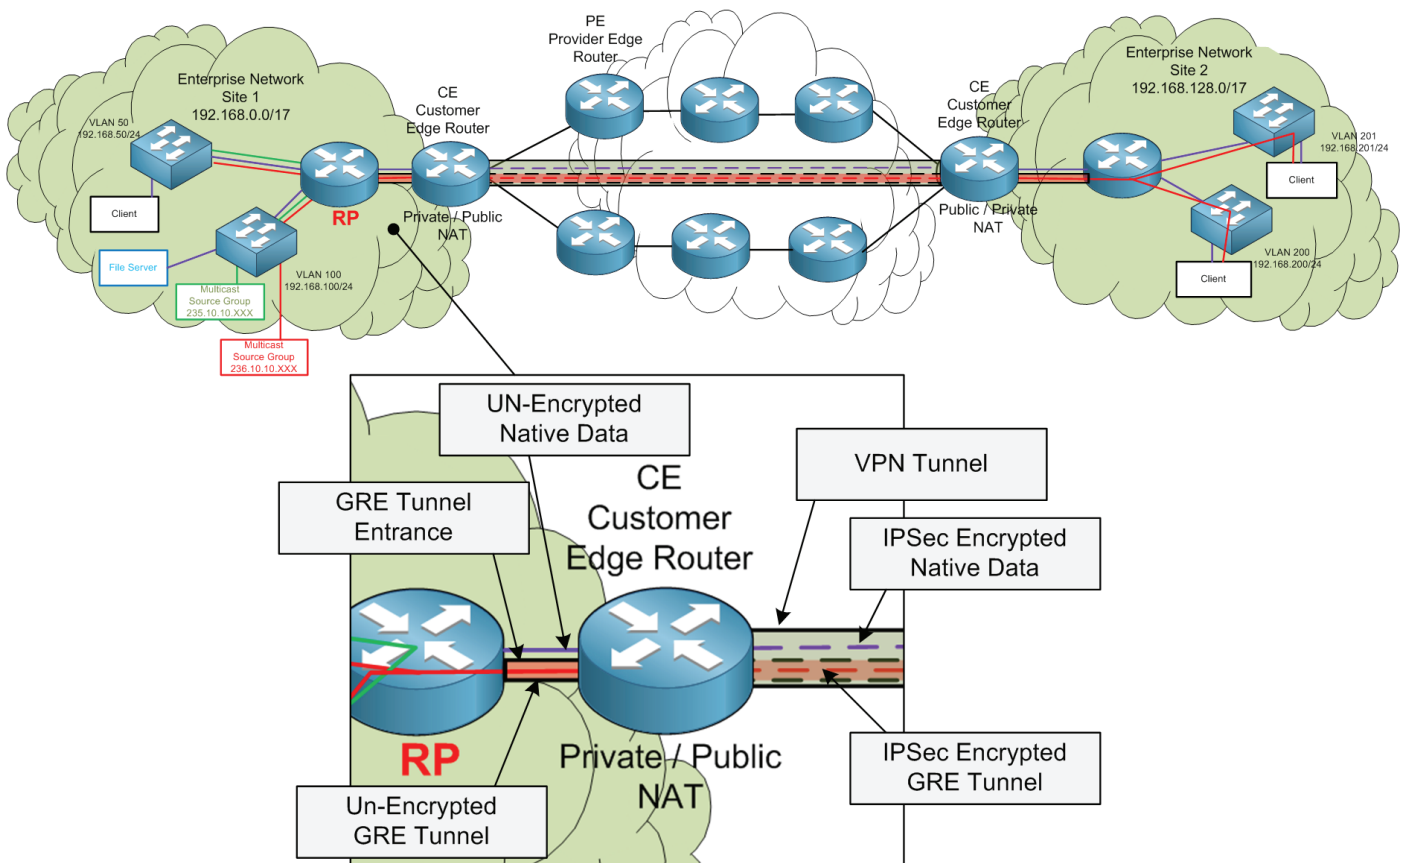
Generic Routing Encapsulation (GRE) tunnel

Generic Routing Encapsulation (GRE) is tunneling protocol that encapsulates a variety of Layer 3 protocols inside virtual point-to-point links over an IP internetwork. One upcoming example of an application for a GRE tunnel is to tunnel IPV6 traffic over a link which only supports IPV4. GRE tunnels are implemented Point-to-Point between two routers.

Unlike VPNs, GRE tunnels have no limitation on the types of traffic which can traverse it. It can route multiple subnets without multiple tunnels. GRE supports multicast so routing protocols and streaming video can be forwarded through a GRE tunnel. Users can merge VPNs and GRE tunnels together as a way to provide the security of VPN without running into the multicast limitations VPNs contain by configuring GRE over IPsec VPN tunnels. This allows GRE tunnel traffic to traverse across the VPN tunnel and creates a single IPsec association regardless of the number of multicast groups that need to get across.


When combined with Protocol Independent Multicast (PIM), a GRE tunnel will only forward multicast traffic when a host on the receiving side joins the multicast group.

A GRE tunnel simply encapsulate the layer 3 traffic (IP Packet) before sending it over the link, adding a 4 byte GRE header and a new 20 byte unicast IPV4 header, with the far side of the tunnel as the destination and forwards the packet.



The illustration above a GRE tunnel is shown implemented between the two interior routers and flow through the edge routers which create an IPsec Tunnel VLAN. The packets are encapsulated with the GRE header at the first router and are encrypted and the IPsec headers are added at the second router. At the far side the reverse is performed, the edge router decrypts the VLAN packet and forwards the GRE packet to the interior router which de-encapsulates it and forwards the original multicast packet.

The illustration below shows the two stage encapsulation.



Considerations in Implementing GRE Tunnels for Multicast



Administration

A GRE tunnel is configured by the network staff and runs on the same enterprise routers as the data network. Due to the mission critical nature of the data network this methodology is not suitable for configurations that will need to be changed frequently or on short notice.



Bandwidth

GRE tunnels can be set with bandwidth limitations to avoid overloading the VPN, but GRE has no ability to police the number of simultaneous streams. A number of simultaneous that exceeds the bandwidth limit can be established with the result being significant packet loss on all channels.



Scalability

An individual GRE tunnel must be established for each VPN that needs to be traversed. If multiple VPNs terminate through the same circuit, the potential bandwidth will be a multiple of the VPNs.

Configuration Considerations

To limit the multicast forwarded across the VPN, put the multicast groups allowed to be forwarded in a separate multicast group address space and limit the GRE tunnel to that range. Filter IGMP requests at the downstream inputs to the receiving side GRE router to only allow desired multicast groups.

Multicast address space must be coordinated between the two sites to avoid overlaps.

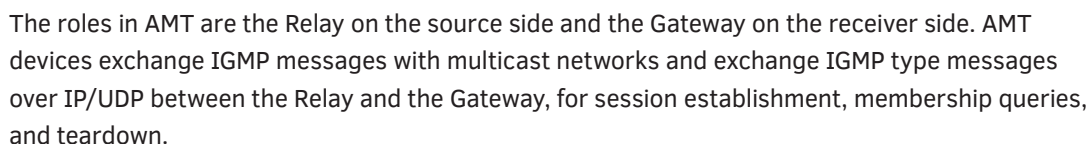
Ensure that there is a valid route for PIM traffic between all receiving side routers and the Rendezvous Point. Remember that PIM is its own protocol, not TCP or UDP.

If multicast is only streaming from one site, configure the GRE tunnel to block multicast from the receiving site to prevent inadvertent loops.

Make sure the multicast Time to Live (TTL) is large enough to cover all the router hops included in the new network.

Automatic Multicast without explicit Tunnels (AMT) is an emerging standard. It has been around as an Internet Engineering Task Force (IETF) draft since 2001, but it has been supported by router manufacturers only for a couple of years. It is still only available on high end routers, typically found on the edge of large enterprises. With the importance of multicast in IPV6, presumably AMT will become available in intermediate level routers.

AMT encapsulates the multicast packet with IP, UDP and AMT headers.



In the illustration on the previous page, the source side Edge Router acts as the Relay, encapsulating and forwarding the multicast traffic and the receiving side Edge Router acts as the Gateway de-encapsulating the traffic and forwarding the multicast onto the forwarding network.

Considerations in Implementing AMT



Administration

AMT is configured by the network staff and runs on the same enterprise routers as the data network. Due to the mission critical nature of the data network this methodology is not suitable for configurations that will need to be changed frequently or on short notice.



Bandwidth

AMT has no bandwidth limitations and no ability to police the number of simultaneous streams to avoid overloading the VPN. Other mechanisms or processes may be needed to limit and police multicast traffic across the VPN.



Scalability

AMT is highly scalable and can be potentially be used with software gateways for home or branch office VPNs that do not have AMT capable routers.

Configuration Considerations

AMT is still in the draft stage so multi-vendor interoperability is not ensured. This technology should be lab tested for suitability before it is implemented.

To limit the multicast forwarded across the VPN, put the multicast groups allowed to be forwarded in a separate multicast group address space and limit the AMT Relay to that range. Filter IGMP requests at the downstream inputs to the receiving side AMT Gateway to only allow desired multicast groups.

Multicast address space must be coordinated between the two sites to avoid overlaps.

Each side of the AMT session need a PIM Rendezvous Point if multicast is to be routed on both sides.

Make sure the multicast Time to Live (TTL) is large enough to cover all the router hops included in the new network.



About AMX by HARMAN

Founded in 1982 and acquired by HARMAN in 2014, AMX® is dedicated to providing AV solutions for an IT World. AMX solves the complexity of managing technology with reliable, consistent and scalable systems comprising control, video switching and distribution, digital signage and technology management. AMX systems are deployed worldwide in conference rooms, classrooms, network operation/command centers, homes, hotels, entertainment venues and broadcast facilities, among others. AMX is part of the HARMAN Professional Group, the only total audio, video, lighting, and control vendor in the professional AV market. HARMAN designs, manufactures and markets premier audio, video, infotainment and integrated control solutions for the automotive, consumer and professional markets.

©2015 Harman. All rights reserved. Specifications subject to change.

AMX.com | 800.222.0193 | 469.624.8000 | +1.469.624.7400 | fax 469.624.7153