

M&E Multicast Addressing Recommendations

Overview

Typical Live Production environments are heavily multicast orientated. Slightly different models are used in SMPTE ST2110, SMPTE ST2022 and AES67, but typically, each media flow is mapped to a unique multicast group, and there is typically only one source for this multicast group (S,G)

Differentiating flows by unique multicast group and source (S,G) is important to maximise the flexibility of the multicast based system. Multiple sources per multicast group is not used in this environment, as this results in restricting the flexibility of flow routing. Likewise, differentiating flows by layer 4 UDP port numbers is also not advised for the same reason.

In the SMPTE ST2110 environment each flow is an “essence” flow; video, audio or metadata, with the audio and metadata for example being able to contain multiple channels. Some media flows may be “bundles”, containing several audio channels for example.

Whereas in SMPTE ST2022-6 the flow represents all the elements present in the packetised SDI stream.

The multicast address range spans 224.0.0.0 through to 239.255.255.255, but for many reasons, it’s important to understand how to best make use of this address space.

Design Considerationsw

Multicast Address Ranges

The Multicast address range 224.0.0.0 -> 239.255.255.255 is split into sub ranges, as specified in RFC 5771 (<https://datatracker.ietf.org/doc/html/rfc5771>).

IP Multicast Address Range	Description	Routable	Usable for Live Production
224.0.0.0 -> 224.0.0.255	Local Network Control Block	No	Avoid
224.0.1.0 -> 224.0.1.255	Internetwork Control Block	Yes	Avoid
224.0.2.0 -> 224.0.255.255	Ad-Hoc Block 1	Yes	Avoid
224.1.0.0 -> 224.1.255.255	Reserved	Yes	Avoid
224.2.0.0 -> 224.2.255.255	SDP/SAP Block	Yes	Avoid
224.3.0.0 -> 224.4.255.255	Ad-Hoc Block 2	Yes	Avoid
224.5.0.0 -> 224.251.255.55	Reserved	Yes	Avoid
224.252.0.0 -> 224.255.255.255	DIS Transient Groups	Yes	Avoid
225.0.0.0 -> 231.255.255.255	Reserved	Yes	Possible
232.0.0.0 -> 232.255.255.255	Source Specific Multicast	Yes	Avoid
233.0.0.0 -> 233.251.255.255	GLOP addressing	Yes	Avoid
233.252.0.0 -> 233.255.255.255	Ad-Hoc Block 3	Yes	Avoid
234.0.0.0 -> 234.255.255.255	Unicast Prefix Based	Yes	Avoid
235.0.0.0 -> 238.255.255.255	Reserved	Yes	Possible
239.0.0.0 -> 239.255.255.255	Administratively Scoped	Yes	Recommend

224.0.0.0 -> 224.0.0.255

The local SubNetworking range is worth a closer look - these addresses are not routable, they remain bound to the subnet in which they are generated. They are used for a range of important local control traffic - IGMP, PIM, OSPF for example. mDNS (224.0.0.251) is also in this non-routable block, explaining why its use is not recommended for larger NMOS IS-04 implementations.

224.0.1.0 -> 224.0.1.255

The InterNetwork Control block are routable addresses, which are intended for specific network control type uses - for example PTP (224.0.1.129), NTP (224.0.1.1)

232.0.0.0 -> 232.255.255.255

Many Live Production capable end-points make use of Source Specific Multicast - initiating a join using IGMPv3 where the source of the desired multicast is specified (S,G). This method of joining is more precise and arguably more secure. However, limiting a Live Production facility to only end points that can initiate an IGMPv3 SSM join would be quite restricting.

239.0.0.0 -> 239.255.255.255

This range is assigned for private use within an organisation.

Within an organisation, most of the other multicast ranges are unused

L3 / L2 Aliasing

Most Live Production implementations transmit and receive flows as IP packets. This allows for routing, and this in turn allows for good resilience design practices, and scale.

Multicast IP packets (like other IP packets) have both a layer 3 destination IP address, and a layer 2 destination MAC address. The layer 3 destination IP address is known as its Multicast Group Address. The layer 2 destination MAC address is formed directly from the layer 3 IP address.

The diagram shows how the L2 destination address is formed from an L3 destination IP address - in this case 239.129.30.20

The corresponding layer2 (MAC) address is generated as follow:

- The top 24 bits of the MAC address are fixed as "01:00:5e" these are the same for all IPv4 multicast addresses (Green)
- The bottom 23 bits of the MAC address are mapped directly from the bottom 23 bits of the L3 IP address (Cyan)
- The remaining bit is fixed as "0" (Orange)




Figure 1 : L3/L2 Aliasing

From the diagram, we can see that none of the bits from the top octet of the L3 IP address make it into the L2 MAC address. The top bit of the next layer 3 octet is also not present in the L2 MAC address.

So, none of the top 9 bits of the layer3 destination address (yellow bits) make it into the layer 2 destination MAC address.

The top 4 bits of the top octet (binary 1110) are fixed for all multicast addresses, so are no great loss, but the remaining 5 bits that do not make it from the layer 3 address to the layer 2 address lead to there being 32 (2^5) possible layer 3 IP addresses that map to the same L2 MAC address - this is the L3 / L2 mapping alias.

For example, the following L3 IP addresses ALL map to the same L2 MAC address:

224.1.30.20, 224.129.30.20

225.1.30.20, 225.129.30.20

226.1.30.20, 226.129.30.20

227.1.30.20, 227.129.30.20

228.1.30.20, 228.129.30.20

229.1.30.20, 229.129.30.20

230.1.30.20, 230.129.30.20

231.1.30.20, 231.129.30.20

232.1.30.20, 232.129.30.20

233.1.30.20, 233.129.30.20

234.1.30.20, 234.129.30.20

235.1.30.20, 235.129.30.20

236.1.30.20, 236.129.30.20

237.1.30.20, 237.129.30.20

238.1.30.20, 238.129.30.20

239.1.30.20, 239.129.30.20

All 32 above L3 destination IP addresses will map to this layer 2 MAC address:

- 0x0100:5e40:30:20

Why does this L3/L2 alias matter?

This alias does not matter for the purposes of multicast routing, the routing of multicast flows between subnets, or routed ports, because routing functions use IP addresses to determine what forwarding should occur.

However, for the layer 2 forwarding operations needed when multicast flows are being delivered by snooping within the vlan, as a layer 2 operation, many implementations rely on the layer 2 destination MAC address, and do not take the layer 3 IP address into consideration.

If multiple aliasing multicast flows exist within a layer2 domain, either because they were sourced in this domain or have been delivered by layer 3 routing into the domain, a layer 2 only snooping process will deliver all of these flows to any end-point that requests any of these flows.

The delivery of unexpected flows in this example situation results in:

- Unexpected bandwidth being used in the network, which could result in oversubscription and packet loss
- Unexpected traffic being delivered to an end-point - which may result in errors for the receiver

For these reasons, it's best to design a multicast addressing scheme to avoid these unwanted aliases.

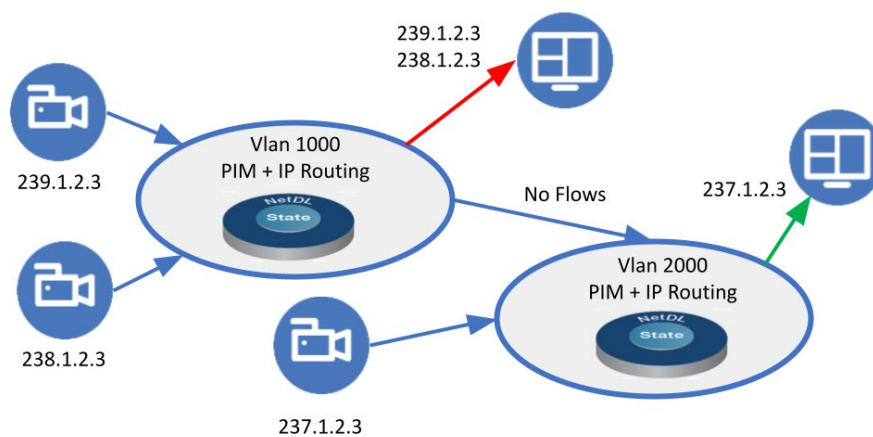


Figure 2: Example of L3/L2 aliasing problems

The diagram above provides an example of where L3/L2 aliasing can cause problems. All of the flows in the diagram are aliasing addresses.

In VLAN 1000, we see that when the Multiviewer requests the camera flow 239.1.2.3, it can receive all aliasing flows that exist in VLAN 1000 - in this case, both 239.1.2.3, and 238.1.2.3

However, in VLAN 2000, while 237.1.2.3 aliases with the other two cameras, these other flows are NOT pulled over the routed link to VLAN 1000 because Multicast IP Routing is L3 aware, and so the three addresses 239.1.2.3, 238.1.2.3 and 237.1.2.3 are distinguishable. So, no problems occur in this scenario.

How to avoid the L3/L2 Alias?

Avoiding these aliases is pretty straightforward - the addressing scheme used should avoid varying the 5 grey bits of the multicast group addresses in the diagram below.




Figure 3 : Avoiding L3/L2 Alias

In practice, this can be achieved by choosing a /9 multicast range - for example 239.0.0.0/9, and not straying from this range. The /9 mask means that you must keep both the top octet (239 in this case) and the top bit of the next octet the same.

Its also important to avoid any unintended aliasing with system multicast addresses - so avoid using addresses that would alias with (224.0.0.0 -> 224.0.1.255)

Summarization

Summarization is the process of being able to describe many smaller address ranges, using a single larger address and mask. It is an equally valid technique for IP source or destination IP addresses. For example, assume you have the following multicast address ranges in use in a facility:

- 239.0.0.0/16 Red Video flows
- 239.4.0.0/16 Red Audio flows
- 239.8.0.0/16 Red Metadata flows
- 239.64.0.0/16 Blue Video flows
- 239.68.0.0/16 Blue Audio flows
- 239.72.0.0/16 Blue Metadata flows

By choosing addresses that can be summarized, it becomes simpler to describe all the Red or Blue flows:

All Red flows = 239.0.0.0/10

All Blue flows = 239.64.0.0/10

Or for example, all Red Video flows = 239.0.0.0/14

These summarizations are useful when you want to configure the RP (Rendezvous Point) in an IGMP / PIM dynamic multicast system, or create security ACLs to prevent Red flows getting to Blue switches.

Choosing a Multicast Address Range for Live Production

A common choice for multicast address range in Live Production systems is 239.0.0.0/8. This range is available for private use, and while Source Specific Multicast (SSM) may be in use, avoids some implementation headaches of using the 232.0.0.0/8 range, if not all the anticipated hosts are capable of SSM.

It should be noted that some protocols direct the multicast ranges in use - for example, AES-67 describes that 239.0.0.0/8 shall be used.

The range 239.0.0.0 -> 239.127.255.255 provides a very large usable address range.

Other ranges can be used, but be aware of the intended use for any range you choose to use.

An Example Addressing Scheme

Its very tempting to design an addressing scheme for multicast media flows that allow for human interpretation at quite a fine granularity. For example, its nice to be able to interpret the following from a multicast group address

- Facility location
- Red or Blue
- Rack identifier
- Flow type - Video, audio, metadata, compressed, uncompressed
- Device identifier
- Flow on device identifier

In small lab situations, the ability to decode all this info from a multicast group is appealing, but as labs grow, and production environments are stood up, this approach becomes unwieldy, and ultimately will be impossible to maintain.

The chief reasons that these strategies fail are that as facilities grow, the impact of these encodings is to make address space use pretty inefficient, and as time goes on, changes in work patterns, extra SMPTE ST2110 flow types, and facility growth will render the encoding impossible to maintain.

However, it is still worth building in `_some_` structure to the addressing - for example, to ensure that the addresses can be effectively summarised (see above).

The level of structure, and human readability is always going to be a philosophical question, but it certainly makes sense to provide a structure that allows for simple summarisation and identification of Red vs Blue flows, and if there are several identifiable locations, these can benefit from being easily identifiable and summarisable.

Instead of building complex human decodable address schemes, it is recommended that human readable flow names are built into Live Production operational and monitoring toolkits, using standard IT technologies like DNS to provide centralised lookup and decoding.

To this end, there follows an example addressing scheme, that allows for the differentiation of flow origination location, Red or Blue network, and media type (Video, Audio, Metadata)

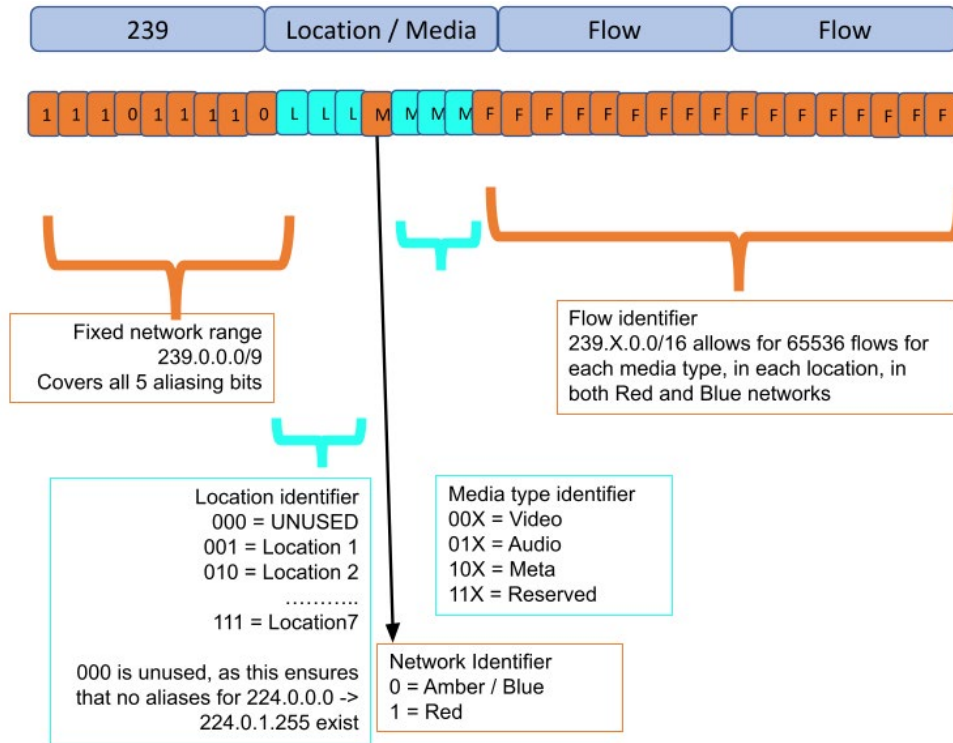


Figure 4: Example Addressing Scheme

Conclusion

Choosing multicast address ranges and addressing schemes is far from straightforward, but by understanding how multicast addresses work, and the ranges that have been made available via IANA, it's possible to arrive at a scalable solution that can provide top level human readability, without sacrificing scale and creating too much complexity.

Santa Clara—Corporate Headquarters

5453 Great America Parkway,
Santa Clara, CA 95054

Phone: +1-408-547-5500

Fax: +1-408-538-8920

Email: info@arista.com

Ireland—International Headquarters

3130 Atlantic Avenue
Westpark Business Campus
Shannon, Co. Clare
Ireland

Vancouver—R&D Office

9200 Glenlyon Pkwy, Unit 300
Burnaby, British Columbia
Canada V5J 5J8

San Francisco—R&D and Sales Office 1390

Market Street, Suite 800
San Francisco, CA 94102

India—R&D Office

Global Tech Park, Tower A, 11th Floor
Marathahalli Outer Ring Road
Devarabeesanahalli Village, Varthur Hobli
Bangalore, India 560103

Singapore—APAC Administrative Office

9 Temasek Boulevard
#29-01, Suntec Tower Two
Singapore 038989

Nashua—R&D Office

10 Tara Boulevard
Nashua, NH 03062

