

Reviewing New Multicast Designs with MLDv2



Nick Russo

NETWORK ENGINEER

@nickrusso42518 www.njrusmc.net



Agenda



Why do we need MLDv2?

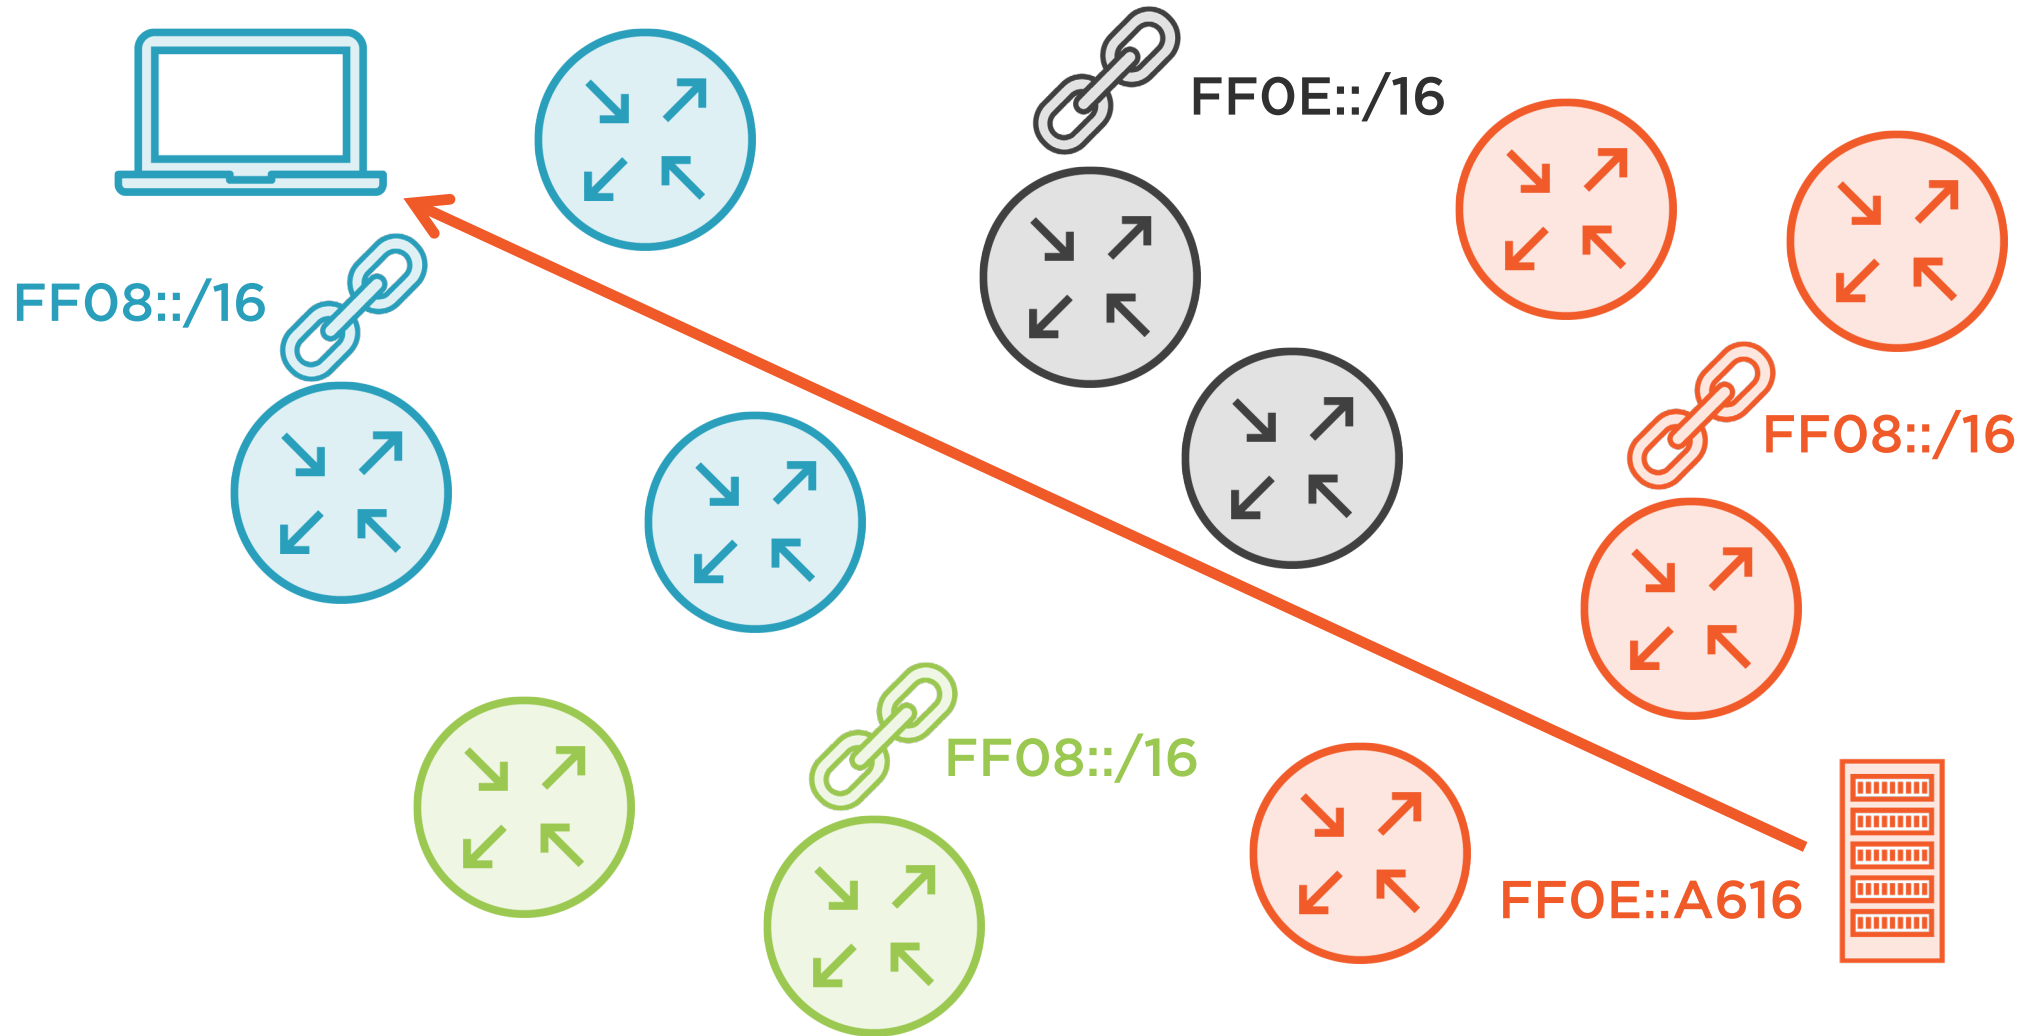
New MLDv2 capabilities

MLDv2 in action

Packet analysis



Big IPv6 Multicast Problem



New MLDv2 Capabilities

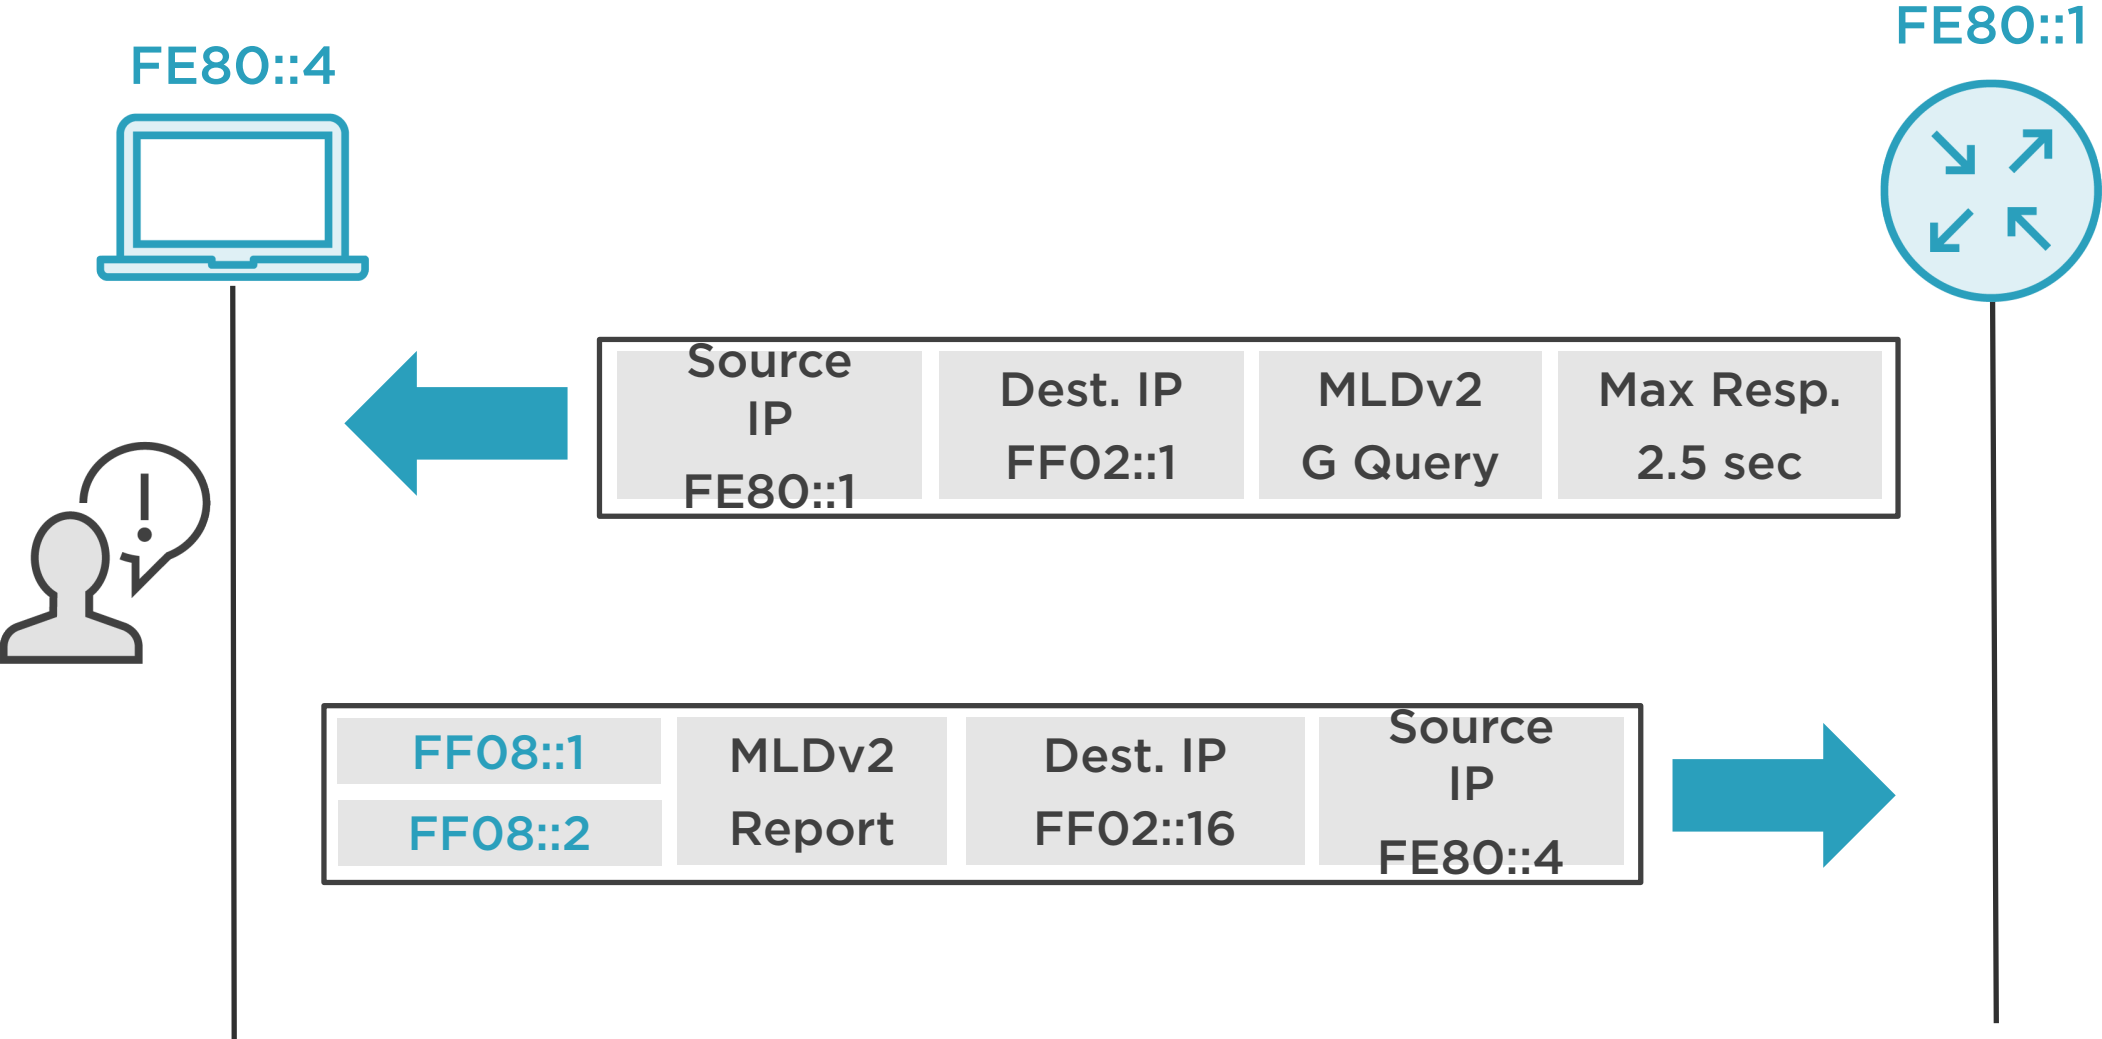
**Membership
Report Batching**

**Source-Specific
Multicast (SSM)
Support**

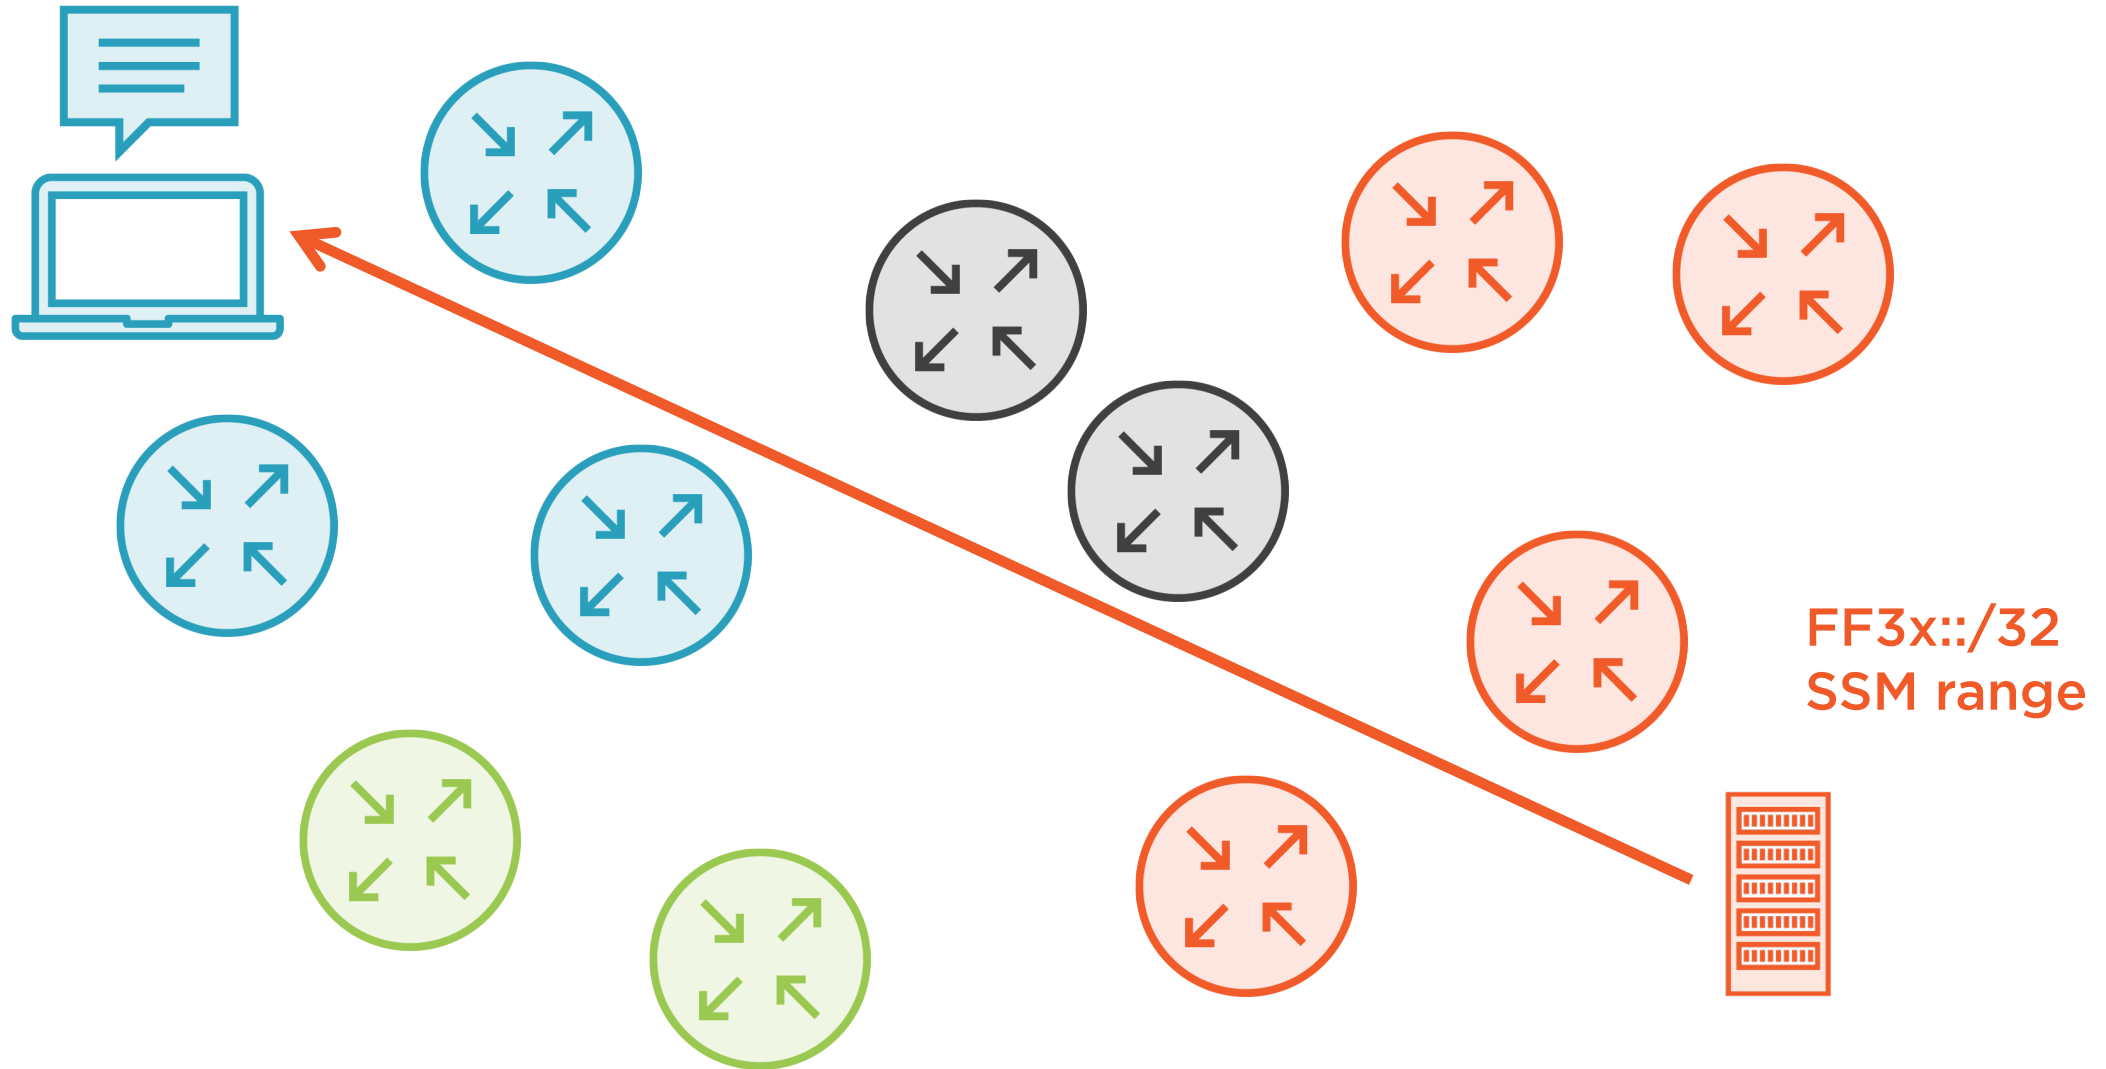
Group "Modes"




MLDv2 Membership Reports



SSM High-level View



Demo



MLDv2 in Action



MLDv2 General Query

No.	Time	Source	Destination	Protocol	Info
1	0.000000	fe80::1	ff02::1	ICMPv6	Multicast Listener Query
2	5.565540	fe80::4	ff02::16	ICMPv6	Multicast Listener Report Message v2
3	37.318891	fe80::4	ff02::16	ICMPv6	Multicast Listener Report Message v2
4	37.320010	fe80::1	ff08::1	ICMPv6	Multicast Listener Query
5	87.277115	fe80::4	ff02::16	ICMPv6	Multicast Listener Report Message v2
6	87.277881	fe80::1	ff38::3	ICMPv6	Multicast Listener Query

- ▶ Frame 1: 90 bytes on wire (720 bits), 90 bytes captured (720 bits) on interface 0
- ▶ Ethernet II, Src: 00:00:a6:16:00:01, Dst: 33:33:00:00:00:01
- ▶ Internet Protocol Version 6, Src: fe80::1, Dst: ff02::1
- ▼ Internet Control Message Protocol v6

Type: Multicast Listener Query (130)

Code: 0

Checksum: 0x5696 [correct]

[Checksum Status: Good]

Maximum Response Code: 10000

Reserved: 0000

Multicast Address: ::

▼ Flags: 0x02

.... 0... = Suppress Router-Side Processing: False

.... .010 = QRV (Querier's Robustness Variable): 2

0000 = Reserved: 0

QQIC (Querier's Query Interval Code): 125

Number of Sources: 0

← Type 130 means "Query"
Code is always 0

← No specific group means "General Query"

← Minor options already
discussed during IGMPv3

← Always 0 in general queries



MLDv2 Membership Report

No.	Time	Source	Destination	Protocol	Info
1	0.000000	fe80::1	ff02::1	ICMPv6	Multicast Listener Query
2	5.565540	fe80::4	ff02::16	ICMPv6	Multicast Listener Report Message v2
3	37.318891	fe80::4	ff02::16	ICMPv6	Multicast Listener Report Message v2
4	37.320010	fe80::1	ff08::1	ICMPv6	Multicast Listener Query
5	87.277115	fe80::4	ff02::16	ICMPv6	Multicast Listener Report Message v2
6	87.277881	fe80::1	ff38::3	ICMPv6	Multicast Listener Query

▶ Frame 2: 182 bytes on wire (1456 bits), 182 bytes captured (1456 bits) on interface 0
▶ Ethernet II, Src: 00:00:a6:16:00:04, Dst: 33:33:00:00:00:16
▶ Internet Protocol Version 6, Src: fe80::4, Dst: ff02::16
▼ Internet Control Message Protocol v6

Type: Multicast Listener Report Message v2 (143)
Code: 0
Checksum: 0x7814 [correct]
[Checksum Status: Good]
Reserved: 0000

Number of Multicast Address Records: 4

▶ Multicast Address Record Exclude: ff02::1:ff00:4

▼ Multicast Address Record Exclude: ff08::1

Record Type: Exclude (2)

Aux Data Len: 0

Number of Sources: 0

Multicast Address: ff08::1

▶ Multicast Address Record Exclude: ff08::2

▼ Multicast Address Record Include: ff38::3

Record Type: Include (1)

Aux Data Len: 0

Number of Sources: 2

Multicast Address: ff38::3

Source Address: fc00:10:2:5::5

Source Address: fc00:10:2:5::6

← Type 143 means "MLDv2 Report"
Code is always 0

← Solicited node address

← Record says "I want multicast
on FF08::1 from any source"

← Record says "I want multicast
on FF38::3 from sources
FC00:10:2:5::5 and
FC00:10:2:5::6 only"



MLDv2 "Done" Mechanism

No.	Time	Source	Destination	Protocol	Info
1	0.000000	fe80::1	ff02::1	ICMPv6	Multicast Listener Query
2	5.565540	fe80::4	ff02::16	ICMPv6	Multicast Listener Report Message v2
3	37.318891	fe80::4	ff02::16	ICMPv6	Multicast Listener Report Message v2
4	37.320010	fe80::1	ff08::1	ICMPv6	Multicast Listener Query
5	87.277115	fe80::4	ff02::16	ICMPv6	Multicast Listener Report Message v2
6	87.277881	fe80::1	ff38::3	ICMPv6	Multicast Listener Query

▶ Frame 3: 90 bytes on wire (720 bits), 90 bytes captured (720 bits) on interface 0
▶ Ethernet II, Src: 00:00:a6:16:00:04, Dst: 33:33:00:00:00:16
▶ Internet Protocol Version 6, Src: fe80::4, Dst: ff02::16
▼ Internet Control Message Protocol v6

Type: Multicast Listener Report Message v2 (143)
Code: 0
Checksum: 0x7100 [correct]
[Checksum Status: Good]
Reserved: 0000
Number of Multicast Address Records: 1

← Type 143 means "MLDv2 Report"
Code is always 0

▼ Multicast Address Record Changed to include: ff08::1
Record Type: Changed to include (3)
Aux Data Len: 0
Number of Sources: 0
Multicast Address: ff08::1

← Record means "I no longer want traffic on FF08::1"



MLDv2 Group-specific Query

No.	Time	Source	Destination	Protocol	Info
1	0.000000	fe80::1	ff02::1	ICMPv6	Multicast Listener Query
2	5.565540	fe80::4	ff02::16	ICMPv6	Multicast Listener Report Message v2
3	37.318891	fe80::4	ff02::16	ICMPv6	Multicast Listener Report Message v2
4	37.320010	fe80::1	ff08::1	ICMPv6	Multicast Listener Query
5	87.277115	fe80::4	ff02::16	ICMPv6	Multicast Listener Report Message v2
6	87.277881	fe80::1	ff38::3	ICMPv6	Multicast Listener Query

- ▶ Frame 4: 90 bytes on wire (720 bits), 90 bytes captured (720 bits) on interface 0
- ▶ Ethernet II, Src: 00:00:a6:16:00:01, Dst: 33:33:00:00:00:01
- ▶ Internet Protocol Version 6, Src: fe80::1, Dst: ff08::1
- ▼ Internet Control Message Protocol v6

Type: Multicast Listener Query (130)

Code: 0

Checksum: 0x7aae [correct]

[Checksum Status: Good]

Maximum Response Code: 1000

Reserved: 0000

Multicast Address: ff08::1

▼ Flags: 0x02

.... 0... = Suppress Router-Side Processing: False

.... .010 = QRV (Querier's Robustness Variable): 2

0000 = Reserved: 0

QQIC (Querier's Query Interval Code): 125

Number of Sources: 0

- ← Type 130 means "Query"
Code is always 0
- ← Faster max resp. time
- ← Specific group means
"Who wants traffic for FF08::1?"



MLDv2 Specific "Done" Mechanism

No.	Time	Source	Destination	Protocol	Info
1	0.000000	fe80::1	ff02::1	ICMPv6	Multicast Listener Query
2	5.565540	fe80::4	ff02::16	ICMPv6	Multicast Listener Report Message v2
3	37.318891	fe80::4	ff02::16	ICMPv6	Multicast Listener Report Message v2
4	37.320010	fe80::1	ff08::1	ICMPv6	Multicast Listener Query
5	87.277115	fe80::4	ff02::16	ICMPv6	Multicast Listener Report Message v2
6	87.277881	fe80::1	ff38::3	ICMPv6	Multicast Listener Query

- ▶ Frame 5: 106 bytes on wire (848 bits), 106 bytes captured (848 bits) on interface 0
- ▶ Ethernet II, Src: 00:00:a6:16:00:04, Dst: 33:33:00:00:00:16
- ▶ Internet Protocol Version 6, Src: fe80::4, Dst: ff02::16
- ▼ Internet Control Message Protocol v6

Type: Multicast Listener Report Message v2 (143)
Code: 0
Checksum: 0x71a0 [correct]
[Checksum Status: Good]
Reserved: 0000
Number of Multicast Address Records: 1

← Type 143 means "MLDv2 Report"
Code is always 0

- ▼ Multicast Address Record Block old sources: ff38::3
 - Record Type: Block old sources (6)
 - Aux Data Len: 0
 - Number of Sources: 1
 - Multicast Address: ff38::3
 - Source Address: fc00:10:2:5::5

← Specific source-group pair means "I no longer want traffic for FF38::3 specifically from FC00:10:2:5::5"



MLDv2 Source/Group-specific Query

No.	Time	Source	Destination	Protocol	Info
1	0.000000	fe80::1	ff02::1	ICMPv6	Multicast Listener Query
2	5.565540	fe80::4	ff02::16	ICMPv6	Multicast Listener Report Message v2
3	37.318891	fe80::4	ff02::16	ICMPv6	Multicast Listener Report Message v2
4	37.320010	fe80::1	ff08::1	ICMPv6	Multicast Listener Query
5	87.277115	fe80::4	ff02::16	ICMPv6	Multicast Listener Report Message v2
6	87.277881	fe80::1	ff38::3	ICMPv6	Multicast Listener Query

▶ Frame 6: 106 bytes on wire (848 bits), 106 bytes captured (848 bits) on interface 0

▶ Ethernet II, Src: 00:00:a6:16:00:01, Dst: 33:33:00:00:00:03

▶ Internet Protocol Version 6, Src: fe80::1, Dst: ff38::3

▼ Internet Control Message Protocol v6

- Type: Multicast Listener Query (130) ← **Type 130 means "Query"**
- Code: 0 ← **Code is always 0**
- Checksum: 0x7e1c [correct]
[Checksum Status: Good]
- Maximum Response Code: 1000
- Reserved: 0000
- Multicast Address: ff38::3 ← **Specific group means "Who wants traffic for FF38::3?"**
- ▼ Flags: 0x02
 - 0... = Suppress Router-Side Processing: False
 -010 = QRV (Querier's Robustness Variable): 2
 - 0000 = Reserved: 0
- QQIC (Querier's Query Interval Code): 125
- Number of Sources: 1
- Source Address: fc00:10:2:5::5 ← **Source/group-specific query**



Comparing MLDv1 and MLDv2

MLDv1

RFC 2710

Equivalent to IGMPv2

Inter-domain multicast is hard

Explicit Done message

MLDv2

RFC 3810

Equivalent to IGMPv3

SSM helps big time (No MSDP)

Implicit leaves via group mode logic

