Troubleshooting and Securing IGMP and MLD

Nick Russo
NETWORK ENGINEER

@nickrusso42518 www.njrusmc.net

Agenda

- Multicast traceroute overview
- Working traceroute with analysis
- Broken traceroute with analysis
- IGMP/MLD security design

Multicast Traceroute Explained

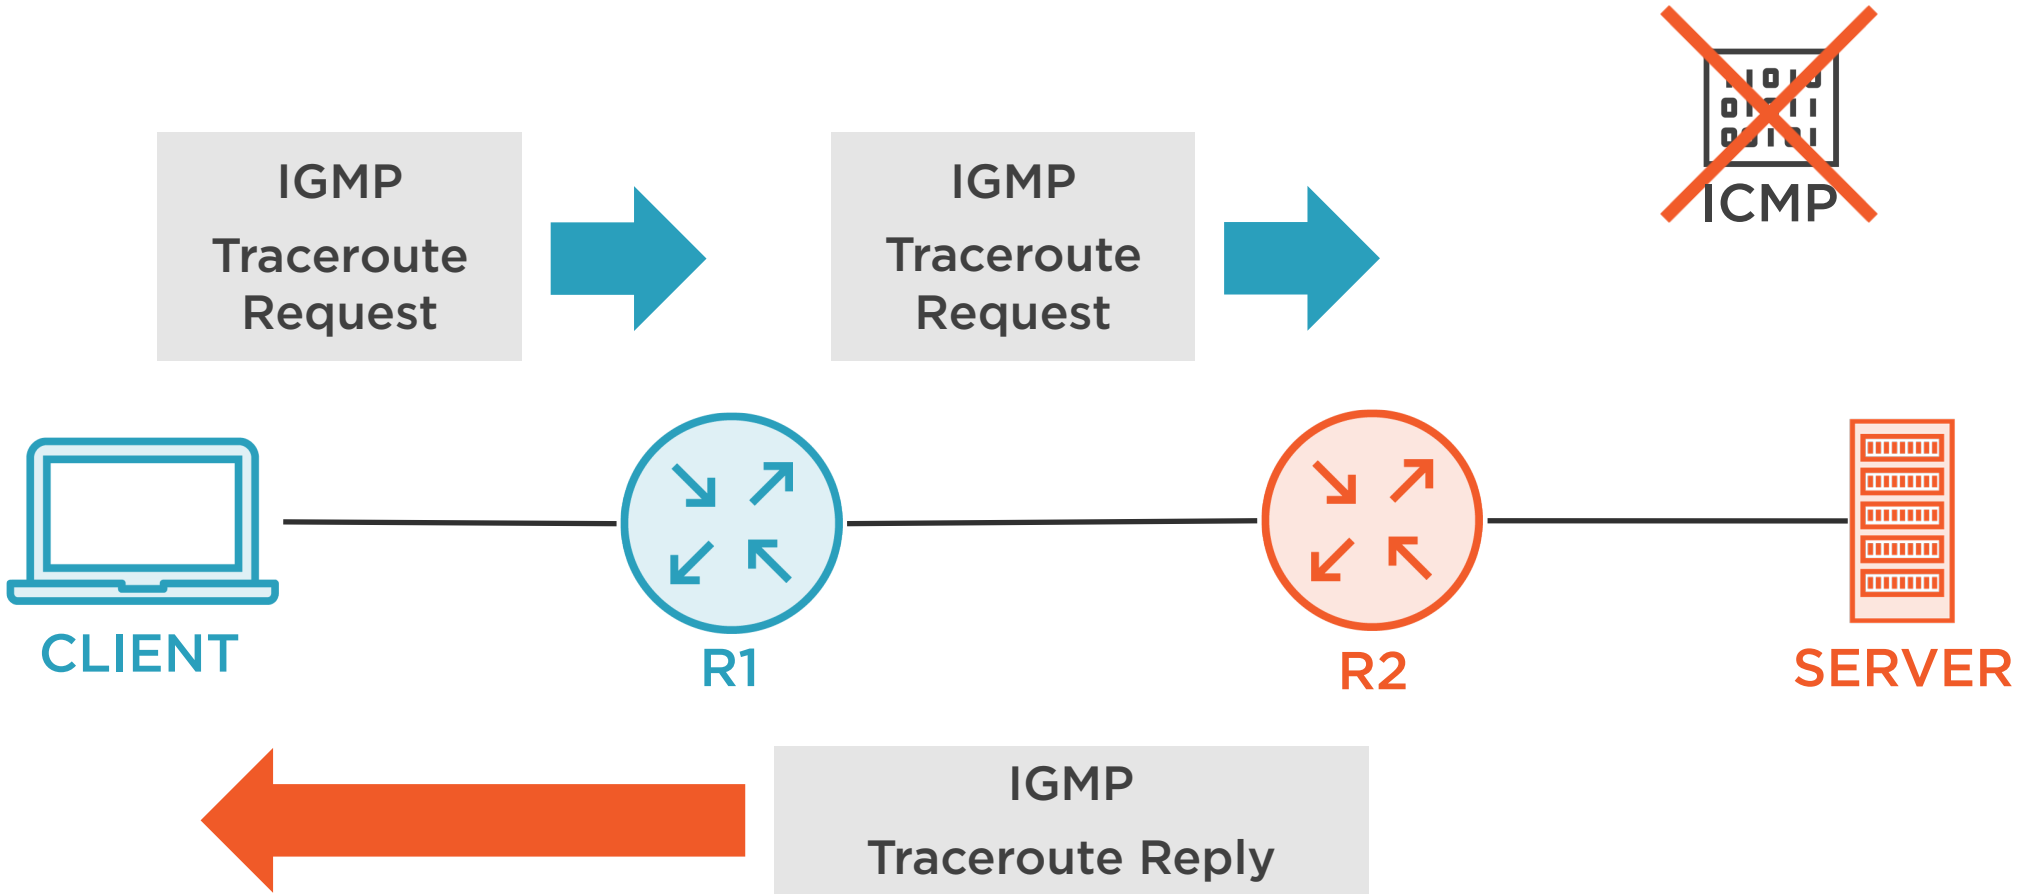

Demo

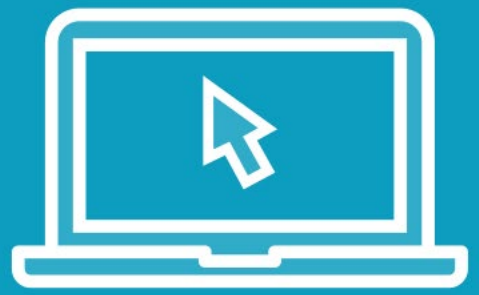

My favorite tool: Multicast traceroute

The Demo Network

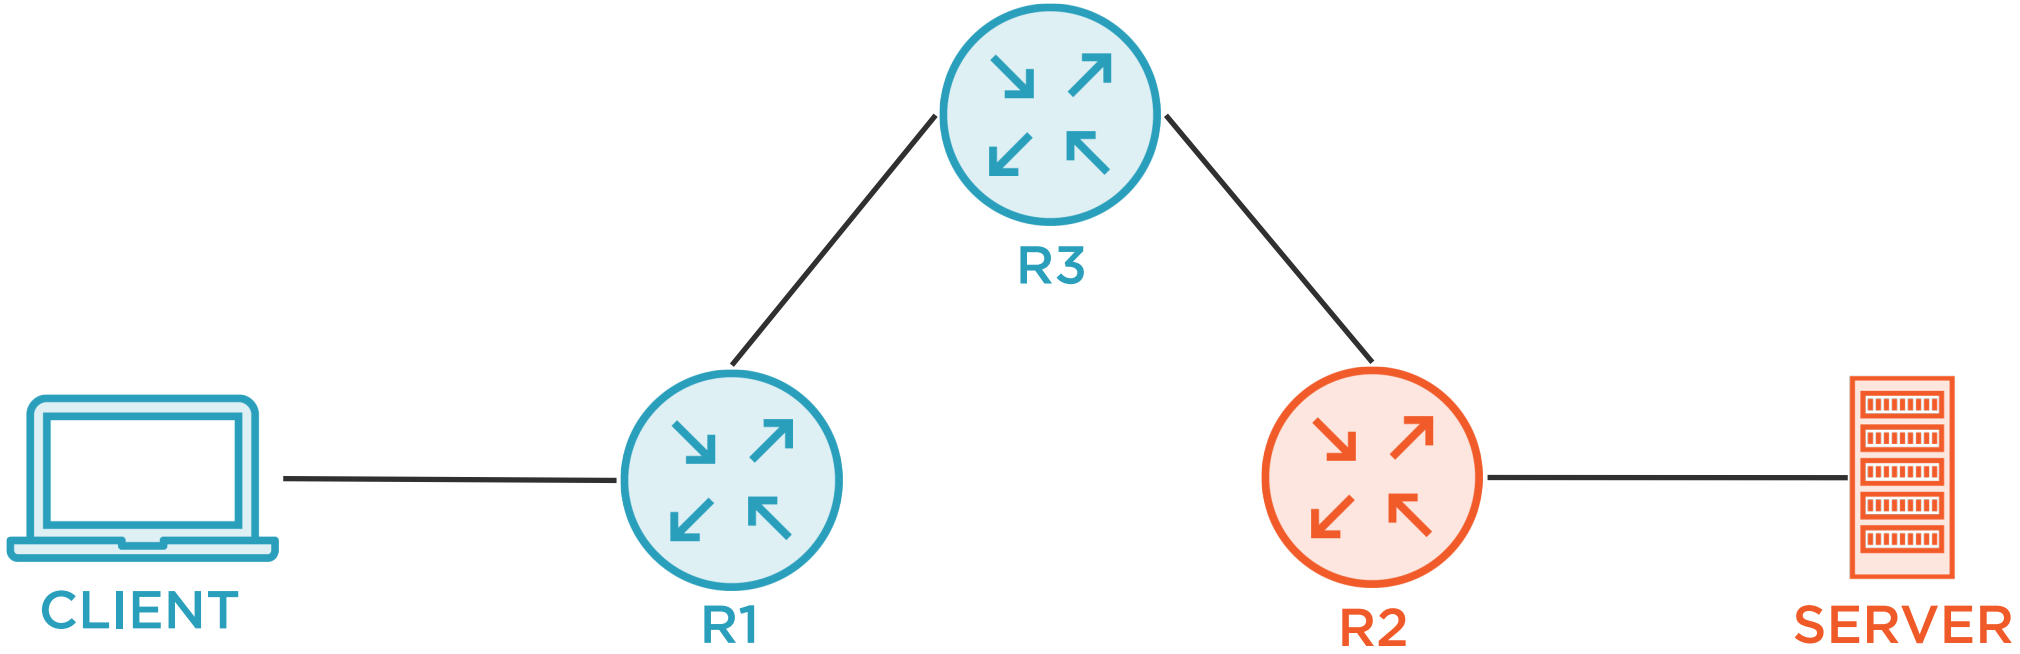

Multicast Traceroute Request

No.	Time	Source	Destination	Protocol	Info
1	0.000000	10.1.4.4	10.1.4.1	IGMP	Traceroute Request
2	0.001239	10.2.5.2	10.1.4.4	IGMP	Traceroute Response, 4 blocks

```
▶ Frame 1: 90 bytes on wire (720 bits), 90 bytes captured (720 bits) on interface 0
▶ Ethernet II, Src: 00:00:a6:16:00:04, Dst: 00:00:a6:16:00:01
▶ Internet Protocol Version 4, Src: 10.1.4.4, Dst: 10.1.4.1
▼ Internet Group Management Protocol
    Type: Traceroute Request (0x1f)
    # hops: 32
    Checksum: 0x7fec [correct]
    [Checksum Status: Good]
    Multicast Address: 0.0.0.0
    Source Address: 10.2.5.5
    Receiver Address: 10.1.4.4
    Response Address: 10.1.4.4
    Response TTL: 64
    Query ID: 8
  ▶ Response data block: 10.1.4.4 -> 10.1.4.4,  Proto: PIM using a static route,  Forwarding Code: NO_ERROR
```

Type 0x1F (31) means "Traceroute Request"
32 is upper-bound on mcast hops

Multicast: No specific group means "General Trace"
Source: Target of trace (mcast source)
Receiver: Unicast IP of client
Response: Used when client isn't unicast-reachable

What's this?

Multicast Traceroute Request Blocks

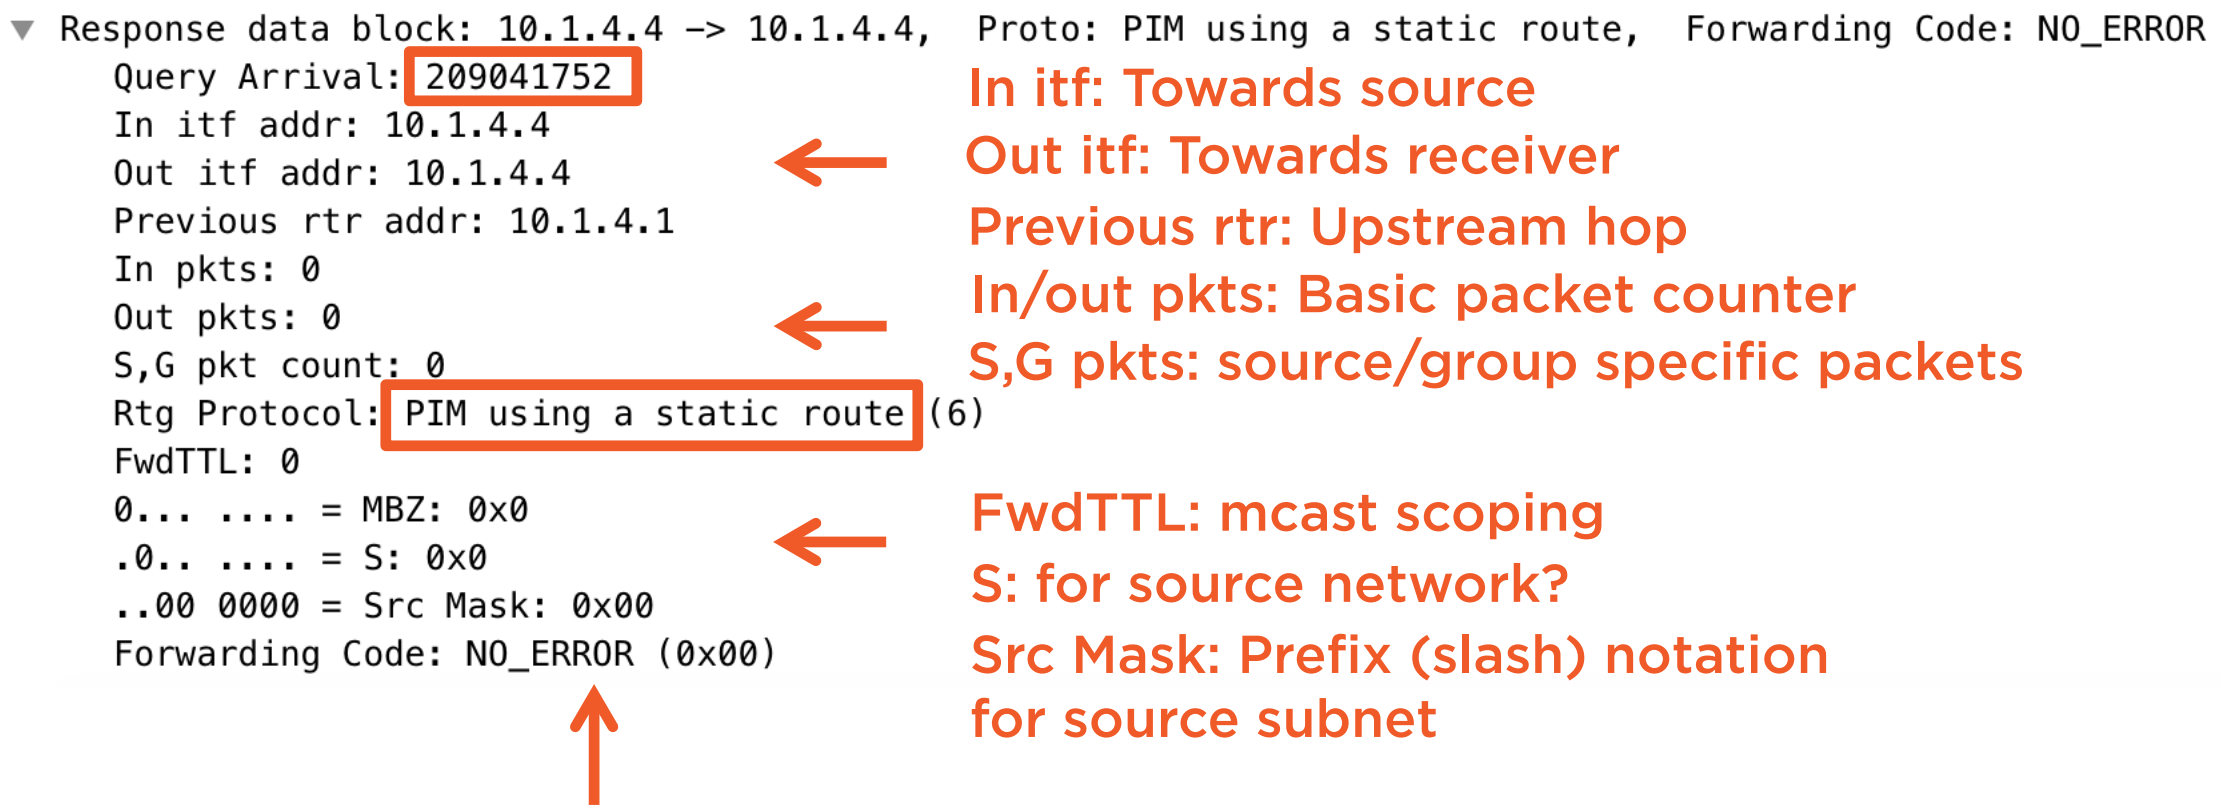

Multicast Traceroute Response

No.	Time	Source	Destination	Protocol	Info
1	0.000000	10.1.4.4	10.1.4.1	IGMP	Traceroute Request
2	0.001239	10.2.5.2	10.1.4.4	IGMP	Traceroute Response, 4 blocks

```
▶ Frame 2: 186 bytes on wire (1488 bits), 186 bytes captured (1488 bits) on interface 0
▶ Ethernet II, Src: 00:00:a6:16:00:01, Dst: 00:00:a6:16:00:04
▶ Internet Protocol Version 4, Src: 10.2.5.2, Dst: 10.1.4.4
▼ Internet Group Management Protocol
    Type: Traceroute Response (0x1e)
    # hops: 32
    Checksum: 0x6ca2 [correct]
    [Checksum Status: Good]
    Multicast Address: 0.0.0.0
    Source Address: 10.2.5.5
    Receiver Address: 10.1.4.4
    Response Address: 10.1.4.4
    Response TTL: 64
    Query ID: 8
  ▶ Response data block: 10.1.4.4 -> 10.1.4.4,  Proto: PIM using a static route,  Forwarding Code: NO_ERROR
  ▶ Response data block: 10.1.3.1 -> 10.1.4.1,  Proto: PIM,  Forwarding Code: NO_ERROR
  ▶ Response data block: 10.2.3.3 -> 10.1.3.3,  Proto: PIM,  Forwarding Code: NO_ERROR
  ▶ Response data block: 10.2.5.2 -> 10.2.3.2,  Proto: CBT using special routing table,  Forwarding Code: NO_ERROR
```

Type 0x1E (30) means "Traceroute Response"
32 is upper-bound on mcast hops

Multicast Traceroute Response First Block

```
▶ Response data block: 10.1.4.4 -> 10.1.4.4,  Proto: PIM using a static route,  Forwarding Code: NO_ERROR
▶ Response data block: 10.1.3.1 -> 10.1.4.1,  Proto: PIM,  Forwarding Code: NO_ERROR
▼ Response data block: 10.2.3.3 -> 10.1.3.3,  Proto: PIM,  Forwarding Code: NO_ERROR
    Query Arrival: 209041817
    In itf addr: 10.2.3.3
    Out itf addr: 10.1.3.3
    Previous rtr addr: 10.2.3.2
    In pkts: 0
    Out pkts: 0
    S,G pkt count: 0
    Rtg Protocol: PIM (3)
    FwdTTL: 0
    0... .... = MBZ: 0x0
    .0.. .... = S: 0x0
    ..01 1000 = Src Mask: 0x18
    Forwarding Code: NO_ERROR (0x00)
▶ Response data block: 10.2.5.2 -> 10.2.3.2,  Proto: CBT using special routing table,  Forwarding Code: NO_ERROR
```

Traffic inbound from R2
Traffic outbound to R1

Not actively sending mcast

Mask 0x18 is 24 in decimal

Most important thing!

Multicast Traceroute Response Last Block

```
▶ Response data block: 10.1.4.4 -> 10.1.4.4,  Proto: PIM using a static route,  Forwarding Code: NO_ERROR
▶ Response data block: 10.1.3.1 -> 10.1.4.1,  Proto: PIM,  Forwarding Code: NO_ERROR
▶ Response data block: 10.2.3.3 -> 10.1.3.3,  Proto: PIM,  Forwarding Code: NO_ERROR
▼ Response data block: 10.2.5.2 -> 10.2.3.2,  Proto: CBT using special routing table,  Forwarding Code: NO_ERROR
    Query Arrival: 209041817
    In itf addr: 10.2.5.2
    Out itf addr: 10.2.3.2
    Previous rtr addr: 0.0.0.0
    In pkts: 0
    Out pkts: 0
    S,G pkt count: 0
    Rtg Protocol: CBT using special routing table (9)
    FwdTTL: 0
    0... .... = MBZ: 0x0
    .0.. .... = S: 0x0
    ..01 1000 = Src Mask: 0x18
    Forwarding Code: NO_ERROR (0x00)
```

Traffic inbound from mcast source
Traffic outbound to R3

Not actively sending mcast

Mask 0x18 is 24 in decimal

Most important thing!

Demo

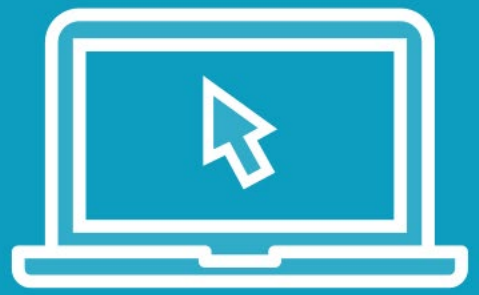

But does mtrace really help?

The Demo Network

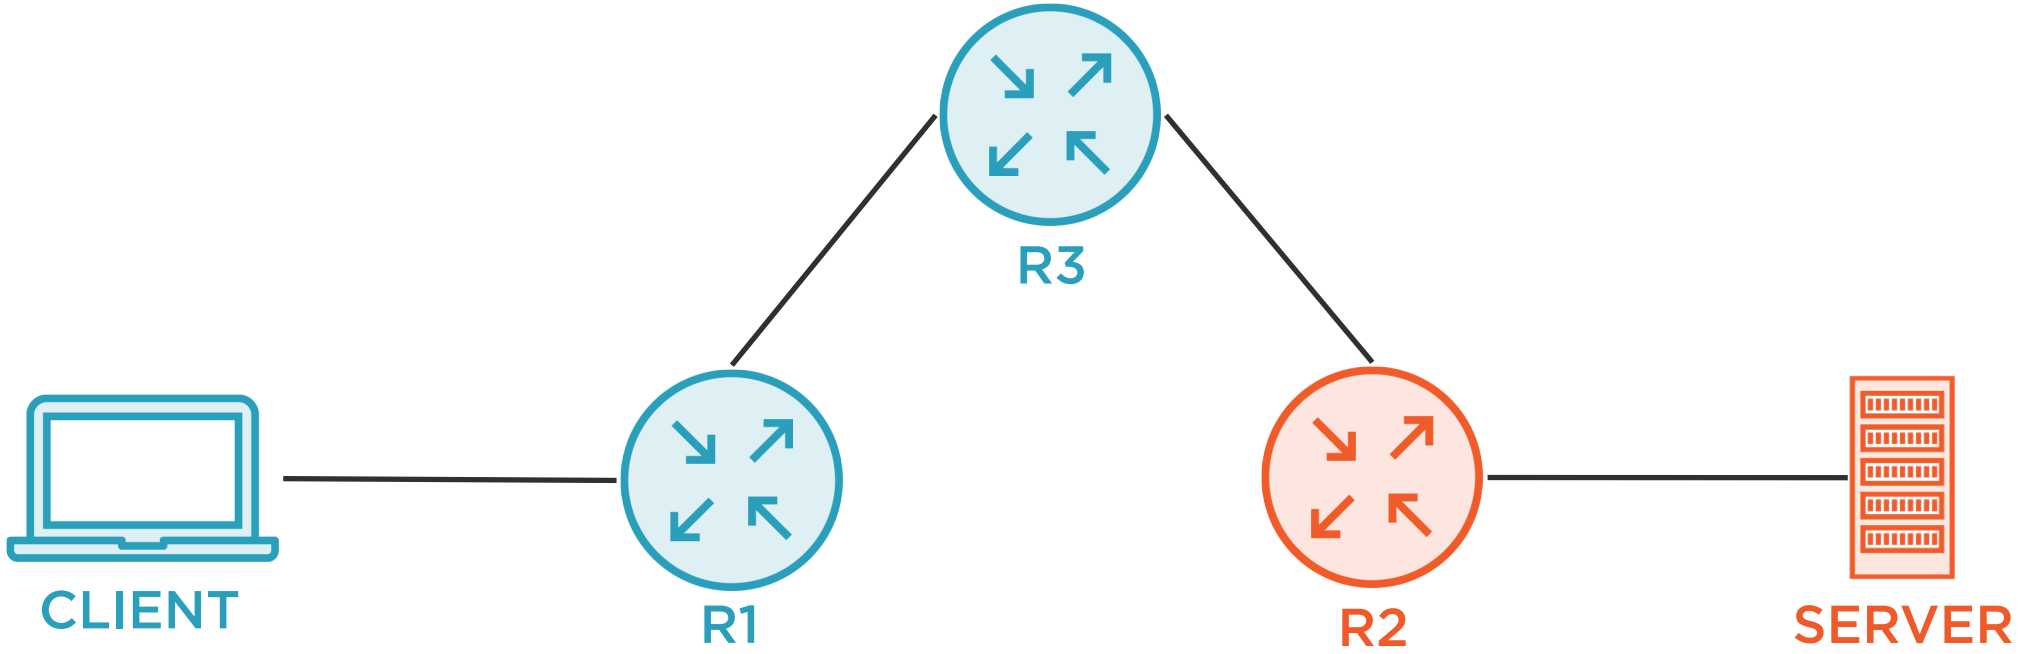

Troubleshooting Request

No.	Time	Source	Destination	Protocol	Info
1	0.000000	10.1.4.4	10.1.4.1	IGMP	Traceroute Request
2	0.000703	10.1.3.3	10.1.4.4	IGMP	Traceroute Response, 3 blocks

```
▶ Frame 1: 90 bytes on wire (720 bits), 90 bytes captured (720 bits) on interface 0
▶ Ethernet II, Src: 00:00:a6:16:00:04, Dst: 00:00:a6:16:00:01
▶ Internet Protocol Version 4, Src: 10.1.4.4, Dst: 10.1.4.1
▼ Internet Group Management Protocol
    Type: Traceroute Request (0x1f)
    # hops: 32
    Checksum: 0x40b2 [correct]
    [Checksum Status: Good]
    Multicast Address: 0.0.0.0
    Source Address: 10.2.5.5
    Receiver Address: 10.1.4.4
    Response Address: 10.1.4.4
    Response TTL: 64
    Query ID: 10
  ▶ Response data block: 10.1.4.4 -> 10.1.4.4,  Proto: PIM using a static route,  Forwarding Code: NO_ERROR
```

Troubleshooting Response

No.	Time	Source	Destination	Protocol	Info
1	0.000000	10.1.4.4	10.1.4.1	IGMP	Traceroute Request
2	0.000703	10.1.3.3	10.1.4.4	IGMP	Traceroute Response, 3 blocks

```
▶ Frame 2: 154 bytes on wire (1232 bits), 154 bytes captured (1232 bits) on interface 0
▶ Ethernet II, Src: 00:00:a6:16:00:01, Dst: 00:00:a6:16:00:04
▶ Internet Protocol Version 4, Src: 10.1.3.3, Dst: 10.1.4.4
▼ Internet Group Management Protocol
    Type: Traceroute Response (0x1e)
    # hops: 32
    Checksum: 0xe794 [correct]
    [Checksum Status: Good]
    Multicast Address: 0.0.0.0
    Source Address: 10.2.5.5
    Receiver Address: 10.1.4.4
    Response Address: 10.1.4.4
    Response TTL: 64
    Query ID: 10
  ▶ Response data block: 10.1.4.4 -> 10.1.4.4,  Proto: PIM using a static route   Forwarding Code: NO_ERROR
  ▶ Response data block: 10.1.3.1 -> 10.1.4.1,  Proto: PIM,  Forwarding Code: NO_ERROR
  ▶ Response data block: 0.0.0.0 -> 10.1.3.3,  Proto: Unknown,  Forwarding Code: NO_ROUTE
```

What happened?

Troubleshooting Response Last Block

```
▼ Response data block: 0.0.0.0 -> 10.1.3.3,  Proto: Unknown,  Forwarding Code: NO_ROUTE
    Query Arrival: 221771726
    In itf addr: 0.0.0.0
    Out itf addr: 10.1.3.3
    Previous rtr addr: 0.0.0.0
    In pkts: 0
    Out pkts: 0
    S,G pkt count: 0
    Rtg Protocol: Unknown (0)
    FwdTTL: 0
    0... .... = MBZ: 0x0
    .0.. .... = S: 0x0
    ..00 0000 = Src Mask: 0x00
    Forwarding Code: NO_ROUTE (0x05)
```

No idea where mcast comes in from

No idea who upstream router is

No idea how to RPF lookup

/0 is not the correct mask, /24 is

Problem: R3 doesn't have
RPF route for 10.2.5.5

Security Design for IGMP

igmp query (0x11 or 17)

mtrace reply (0x1E or 30)

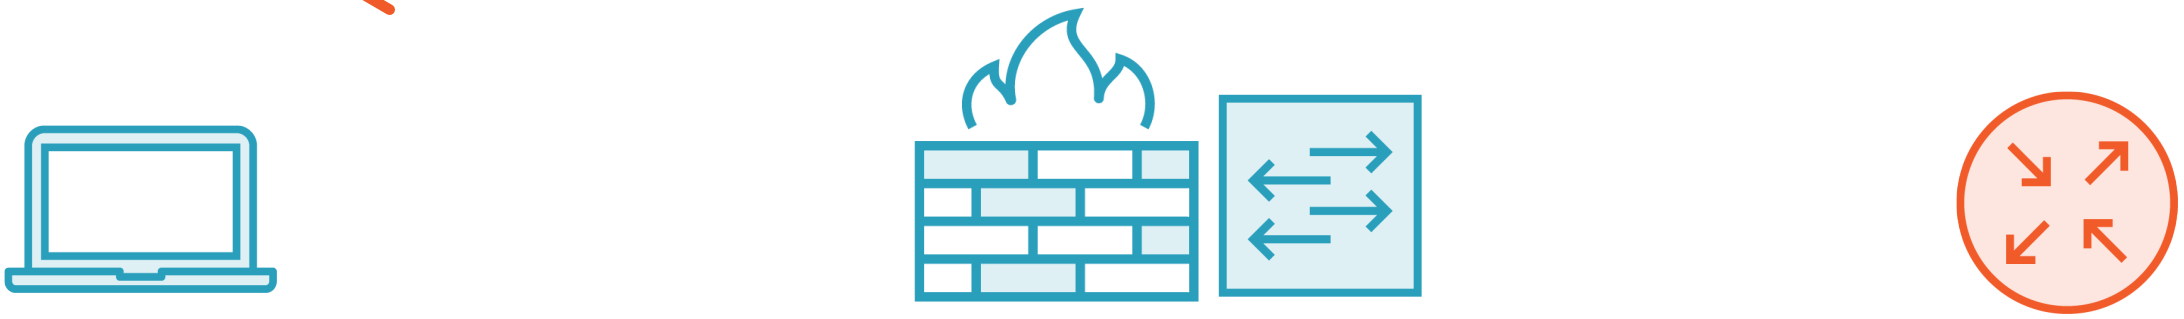

igmpv1 membership report (0x12 or 18)

igmpv2 membership report (0x16 or 22)

igmpv3 membership report (0x22 or 28)

igmpv2 leave group (0x17 or 23)

igmp mtrace request (0x1F or 31)

Security Design for MLD

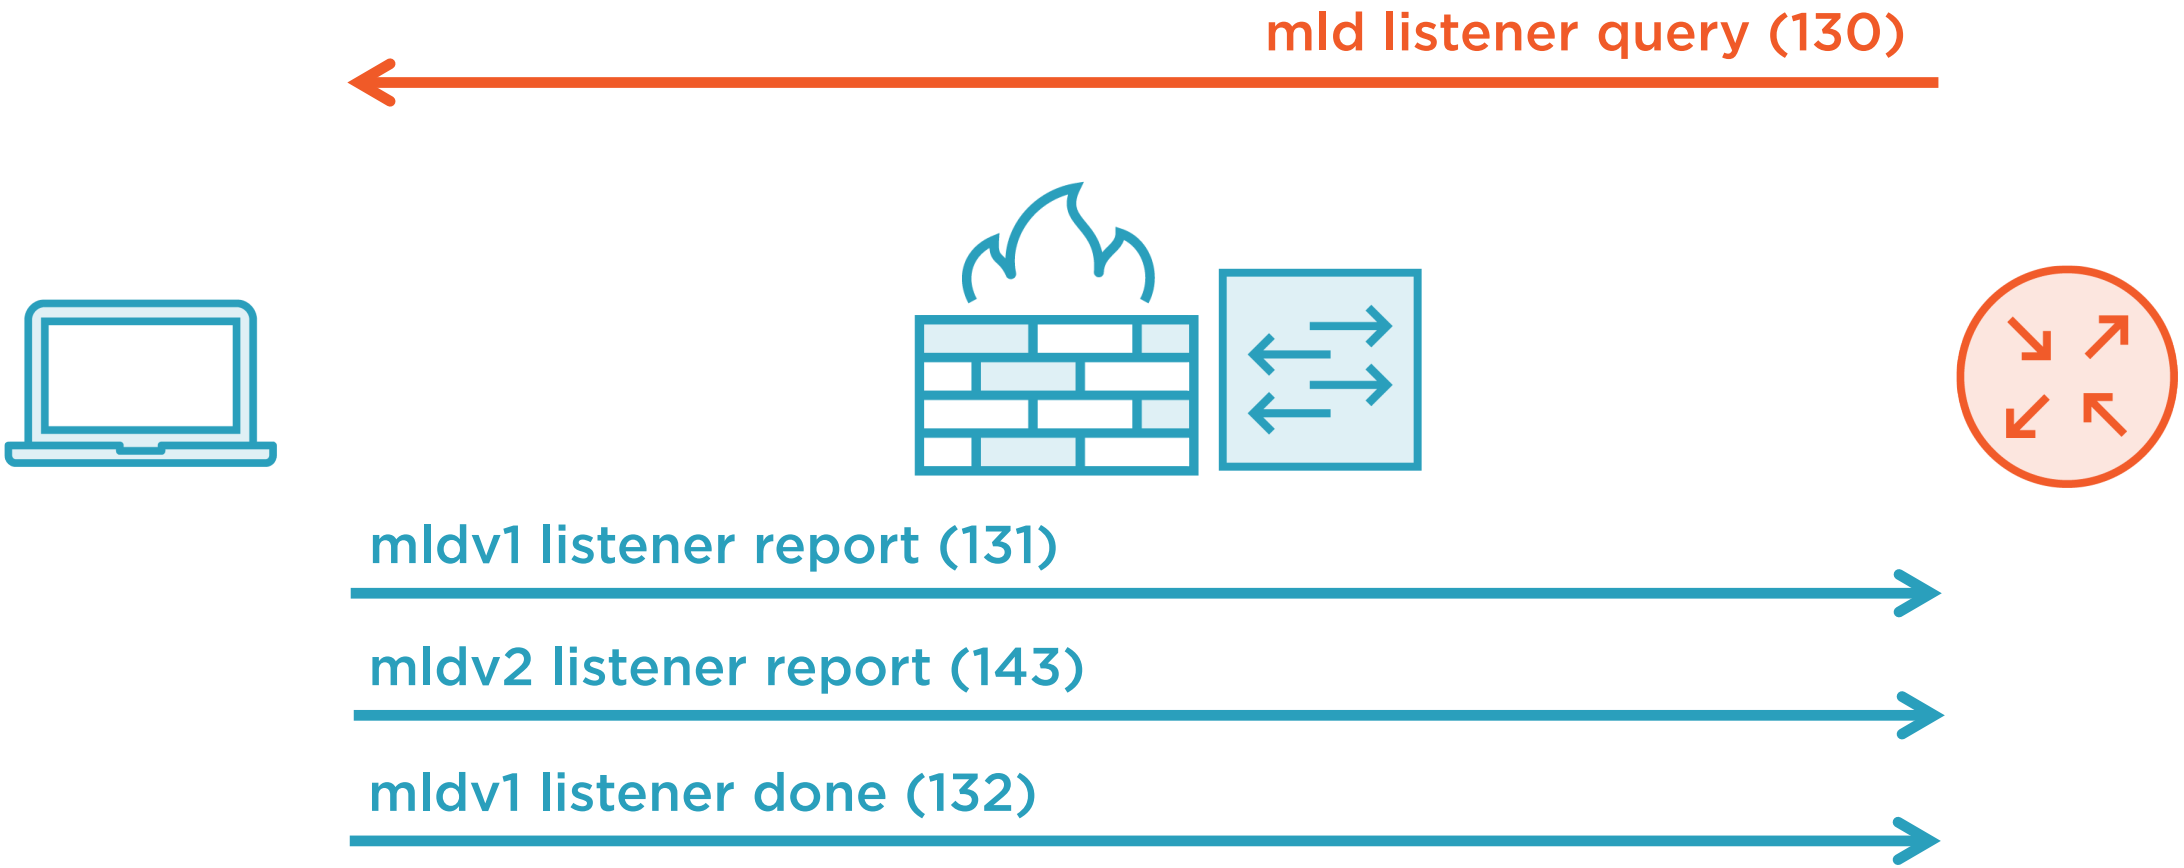

Don't forget about link-local traffic
for IPv6 neighbor discovery!

Final Thoughts

IGMP for IPv4

MLD for IPv6

Multicast traceroute is your friend

Thank you!