PIM Sparse Mode (IPv4)

Tim McConnaughy
Solutions Architect

@juangolbez carpe-dmvpn.com

Agenda

Topics:

- Basic Operation of PIM-SM

Demos:

- Rendezvous Point Election with Bootstrap Router
- Generate Multicast Traffic in a Lab
- PIM Join Process
- PIM Assert Process

Packet Analysis:

- PIM Candidate RP Advertisement/PIM Bootstrap
- PIM Register/Register Stop
- Shortest Path Switchover Join/Prune
- PIM Assert

Globomantics Multicast Deployment Continues

Branch Network

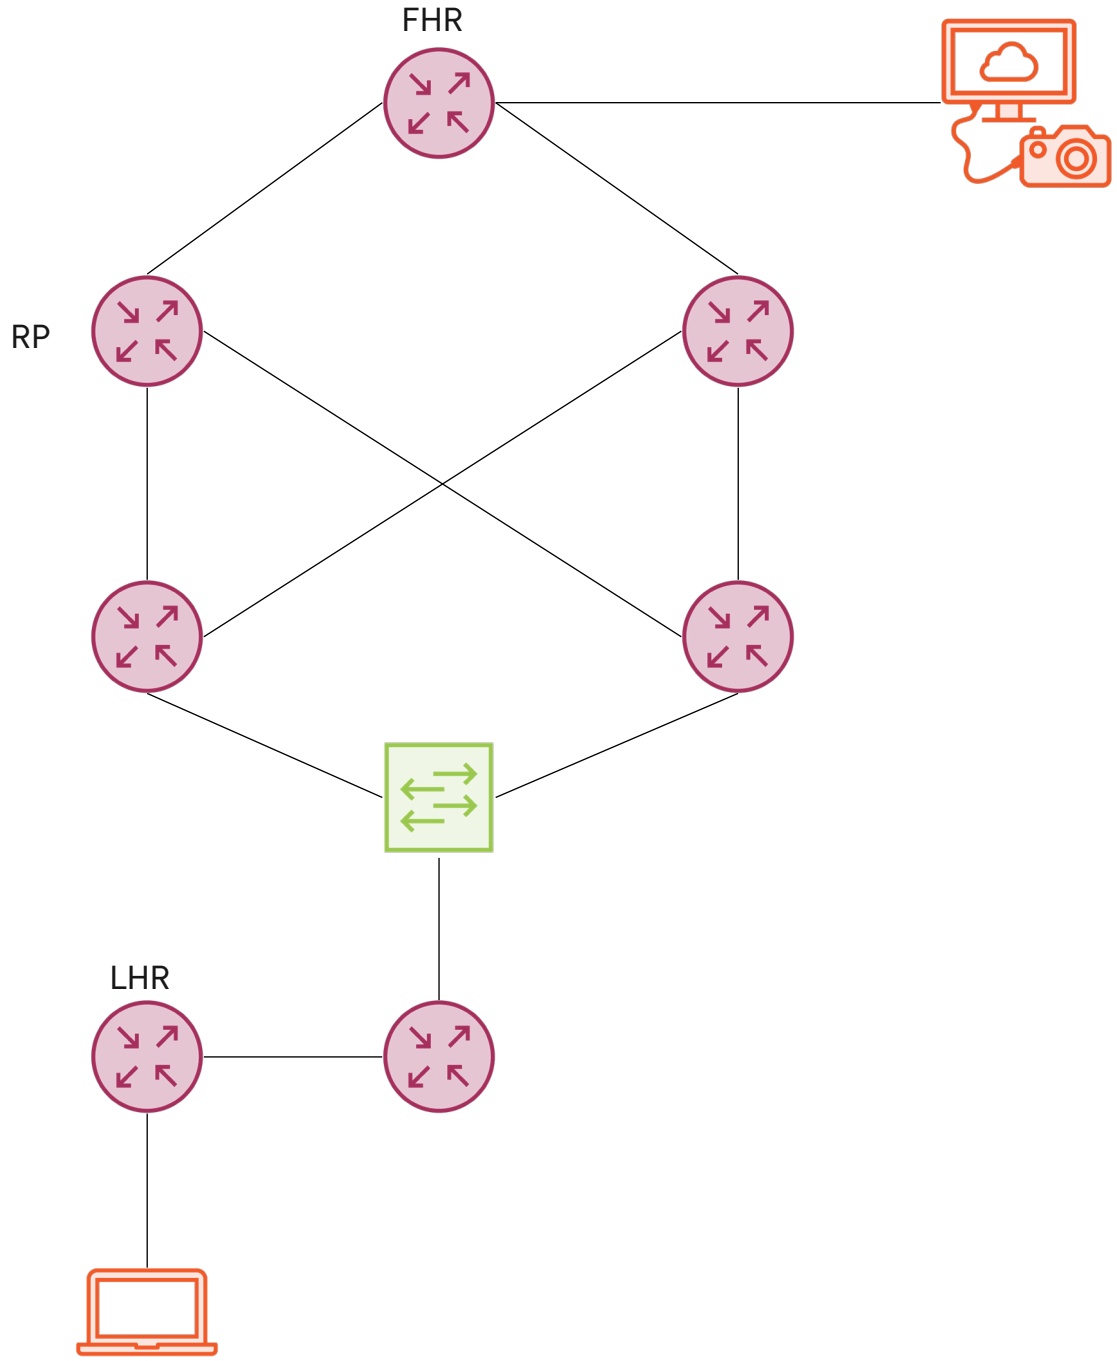

Rendezvous Point

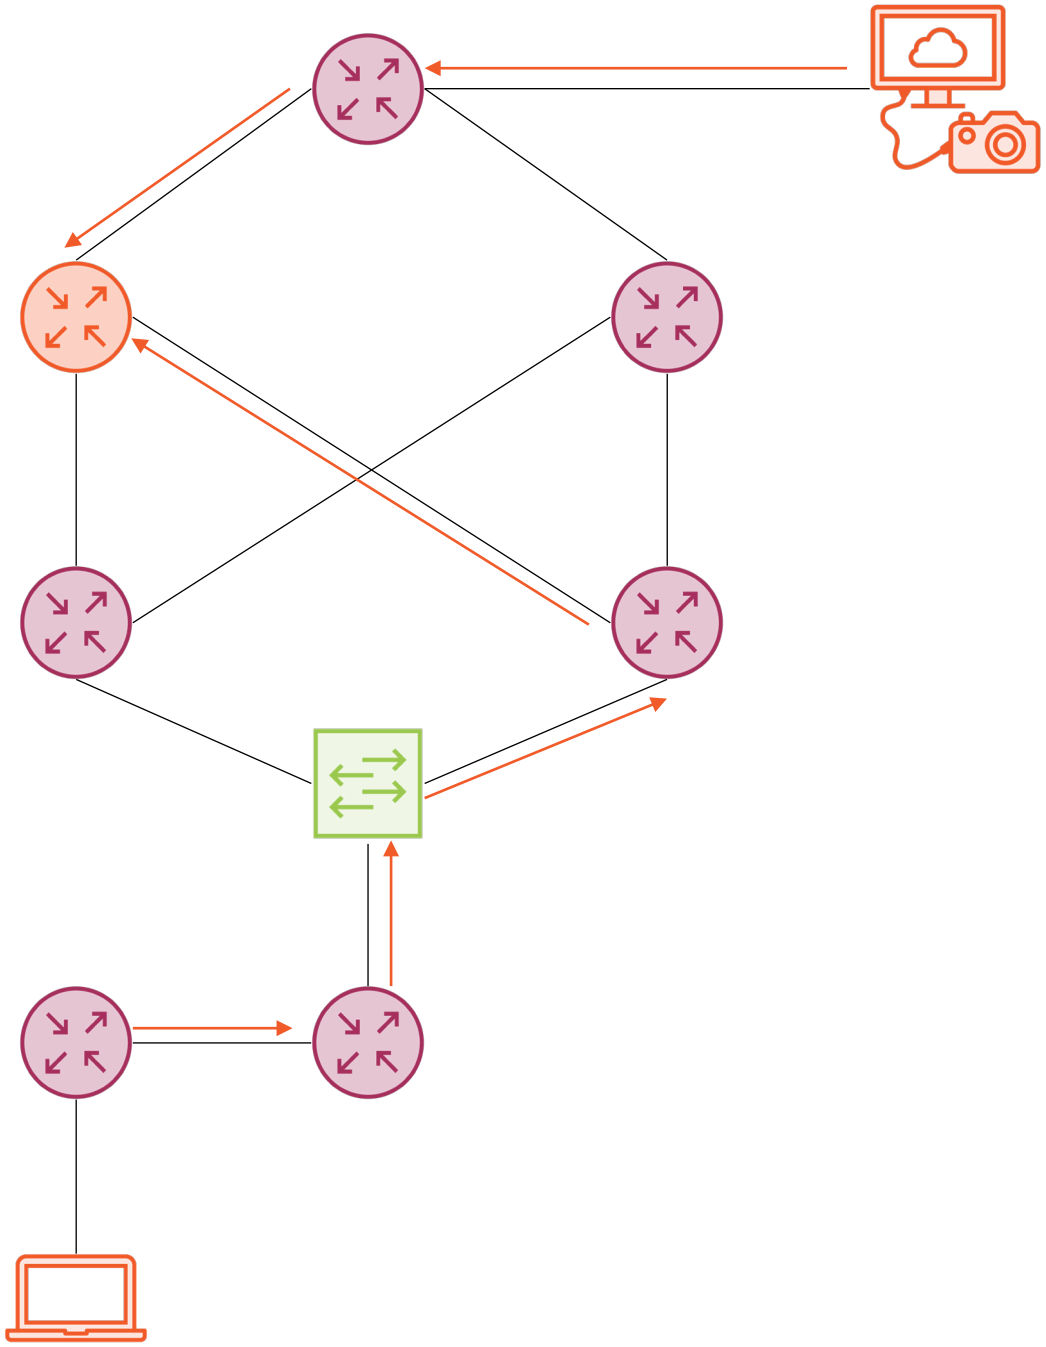

PIM Register

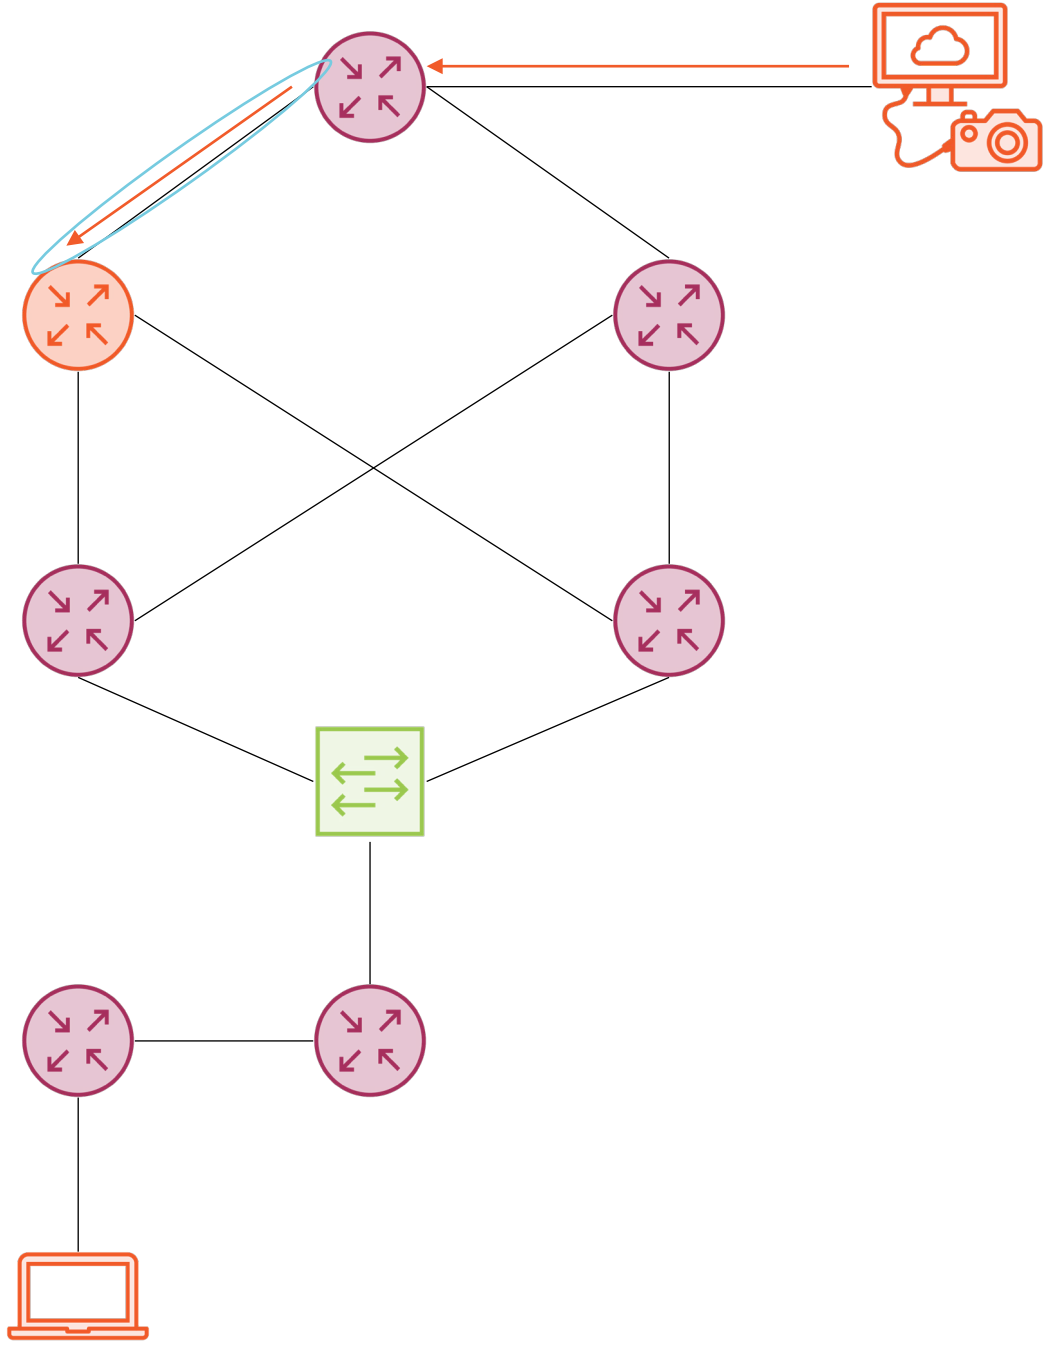

PIM Register

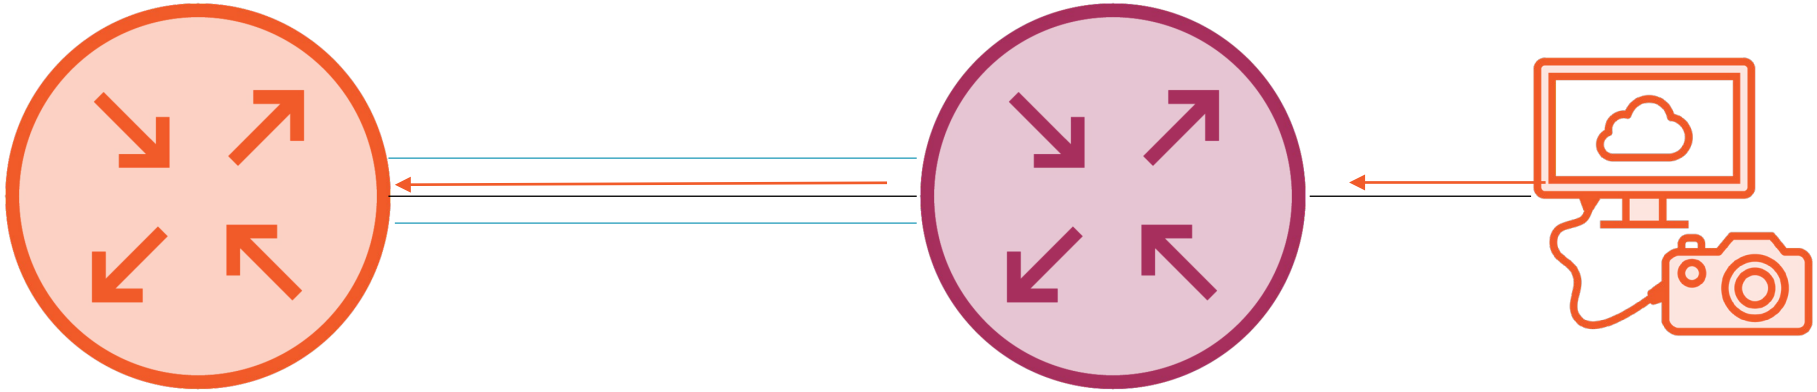

PIM Register Stop

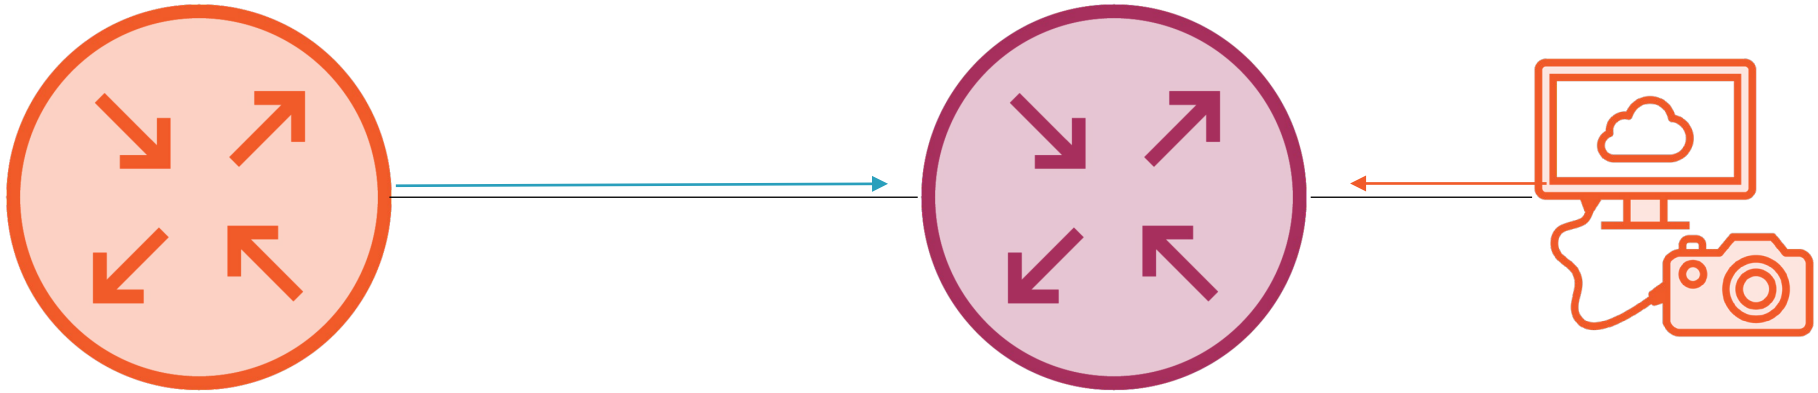

PIM Join From RP

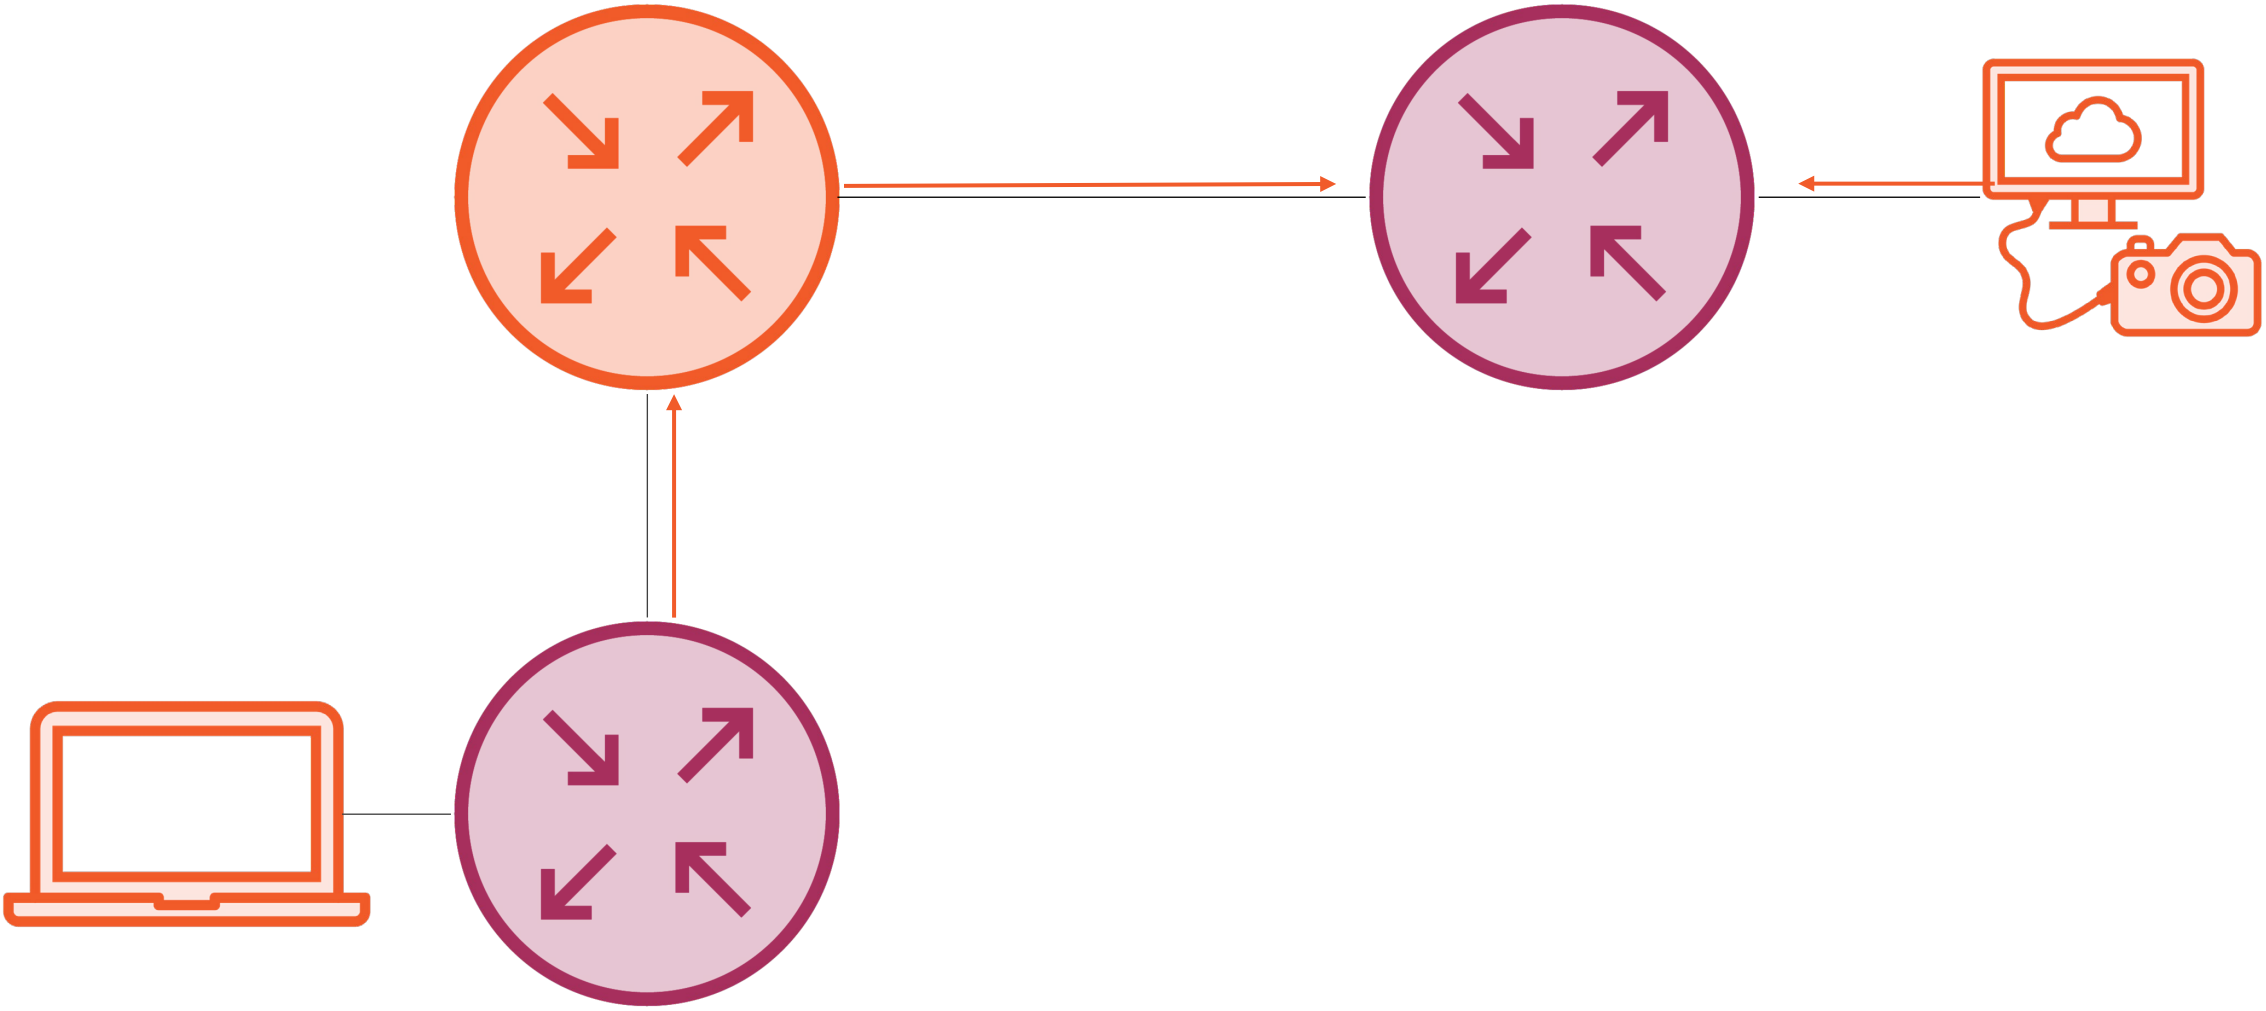

PIM Joins Build OIL

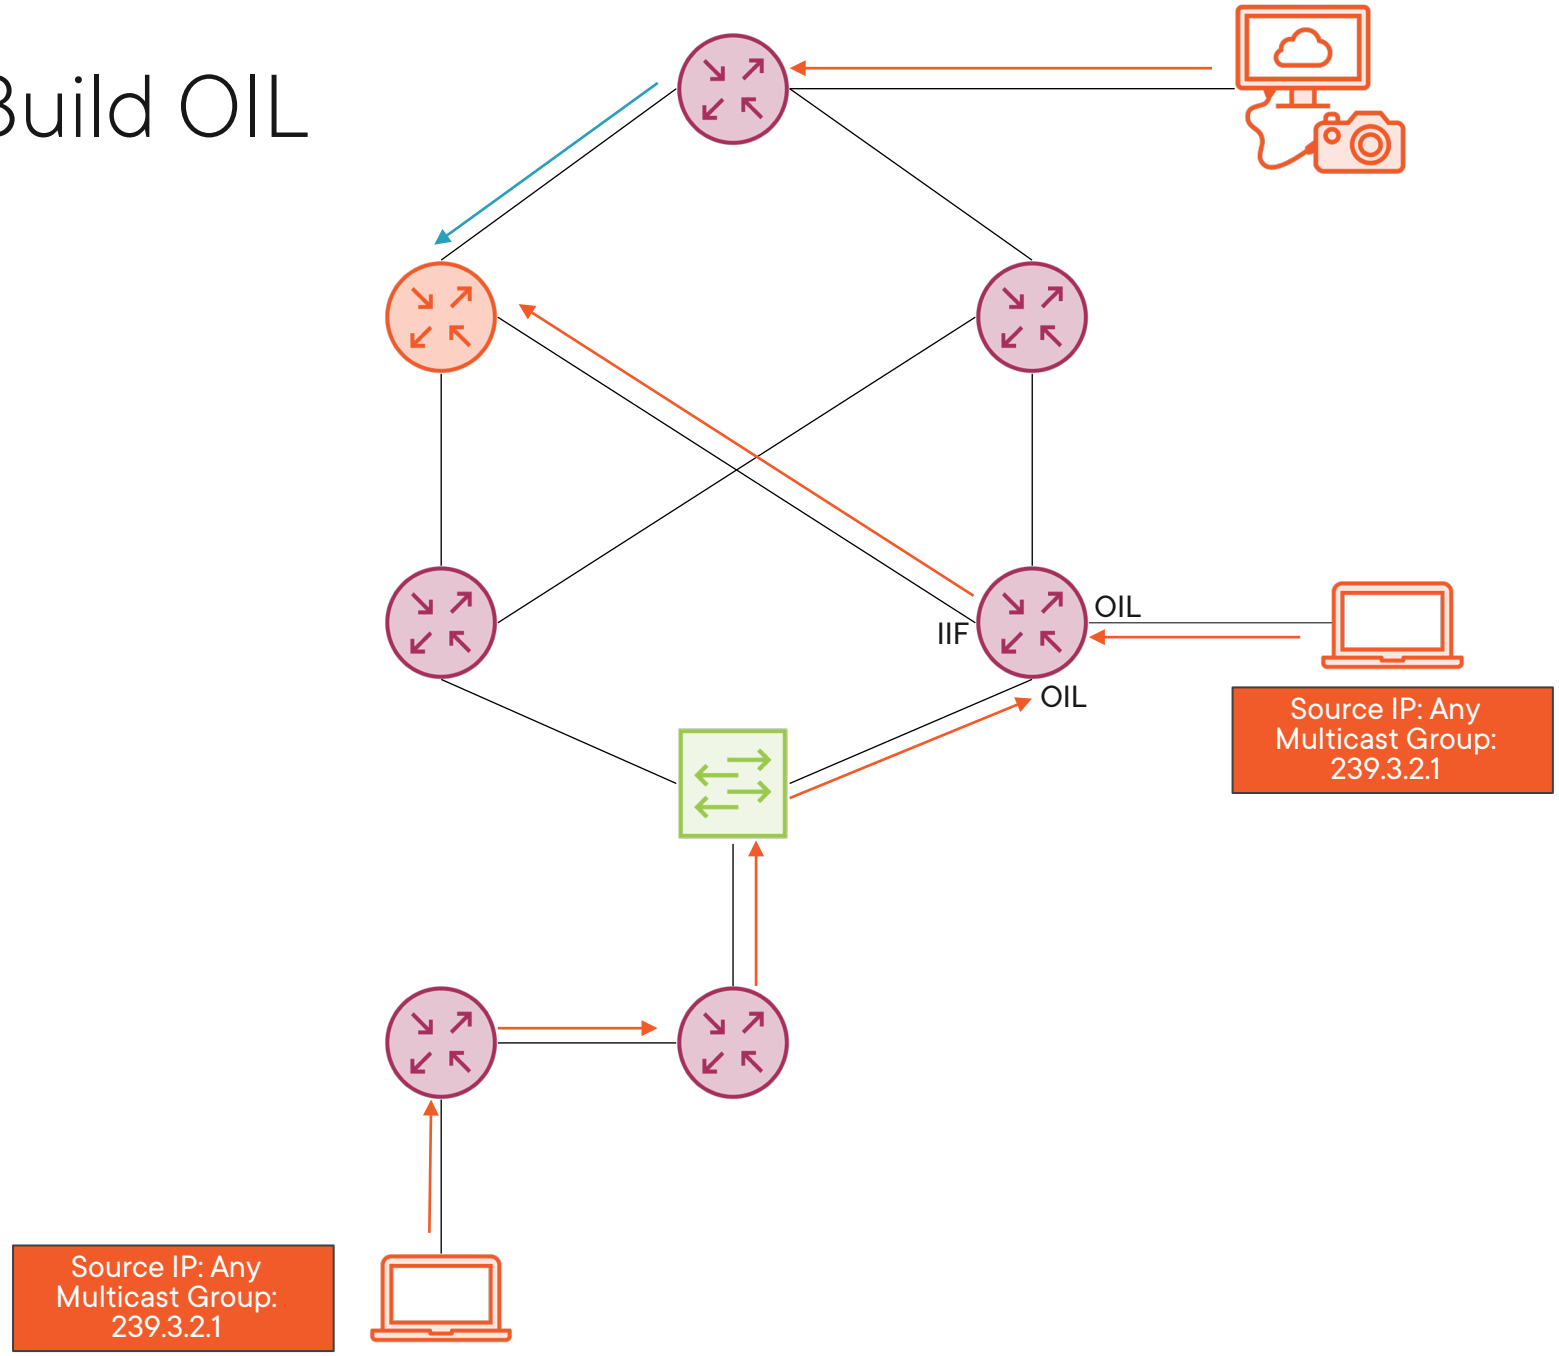

PIM Joins Build OIL

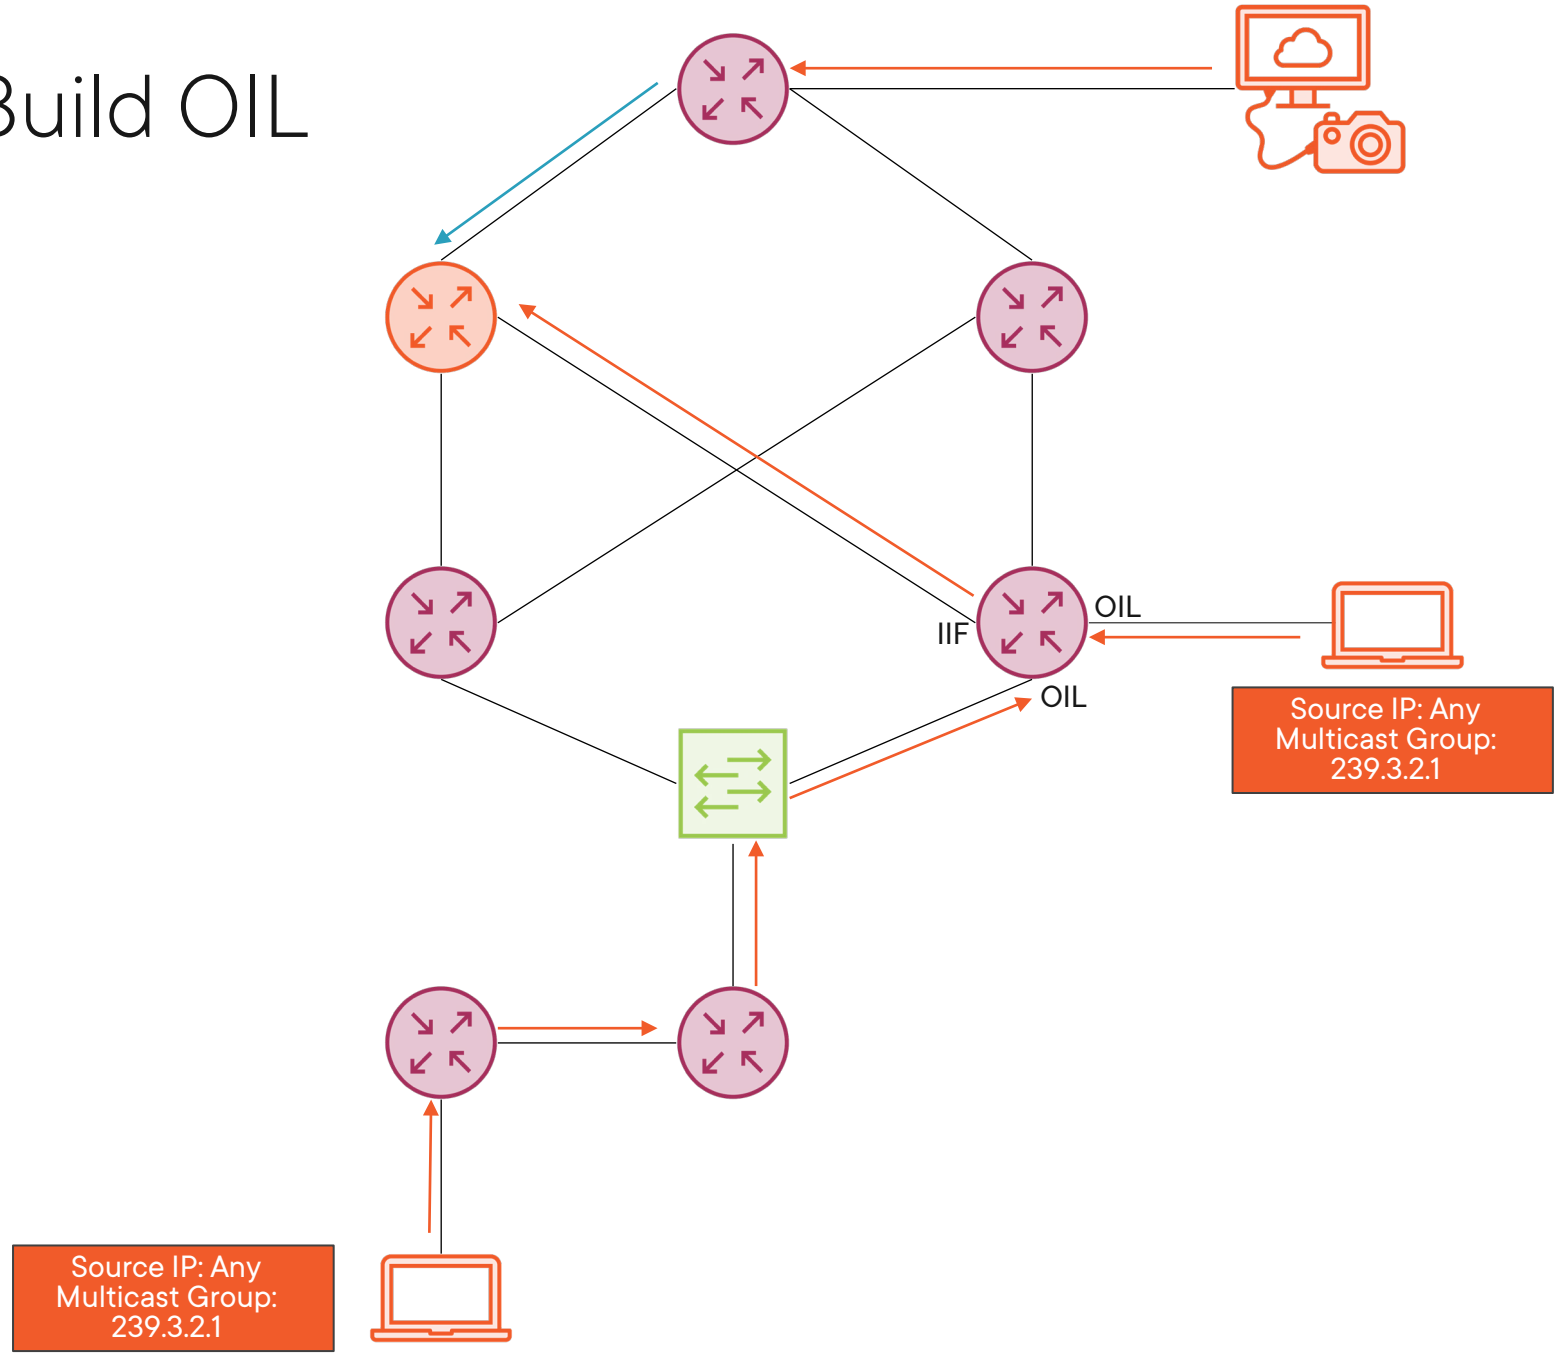

Shared Tree Flow

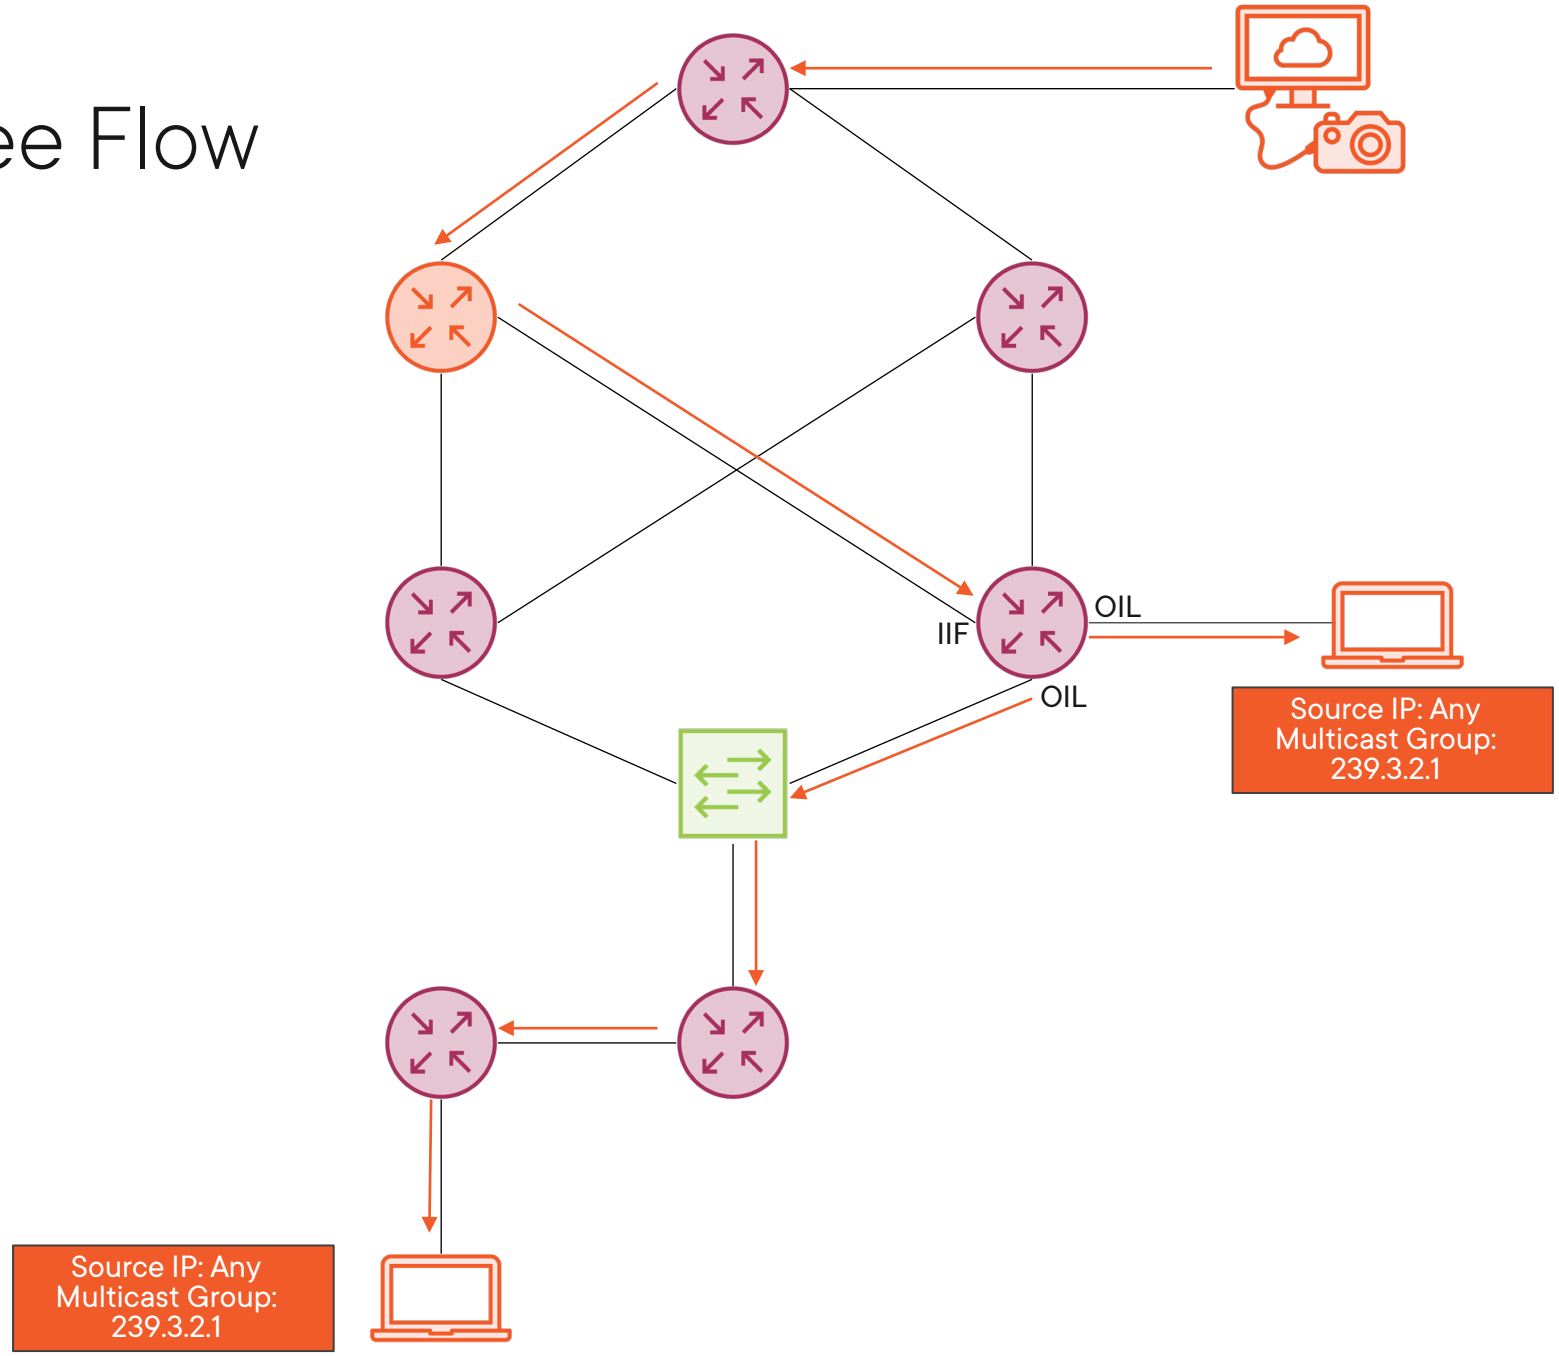

Source Tree Building

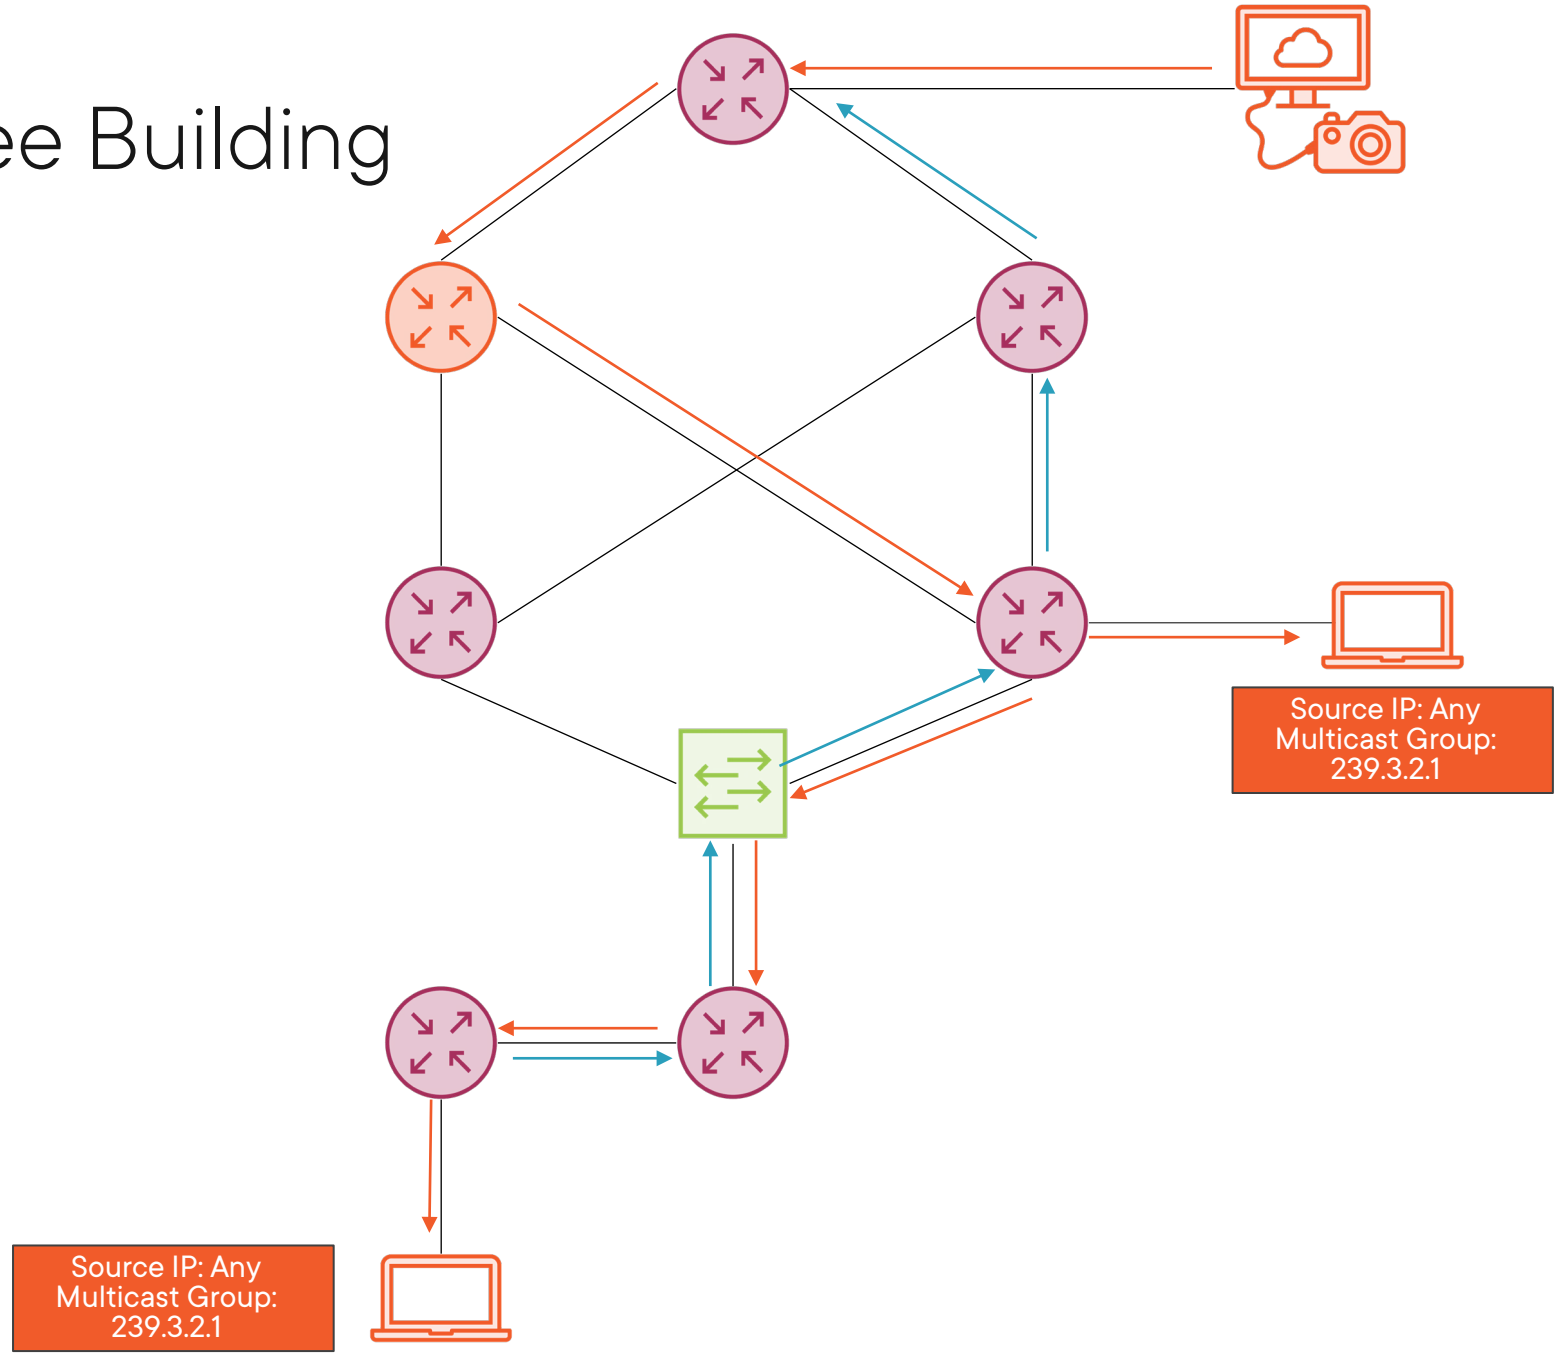

Source Tree Built

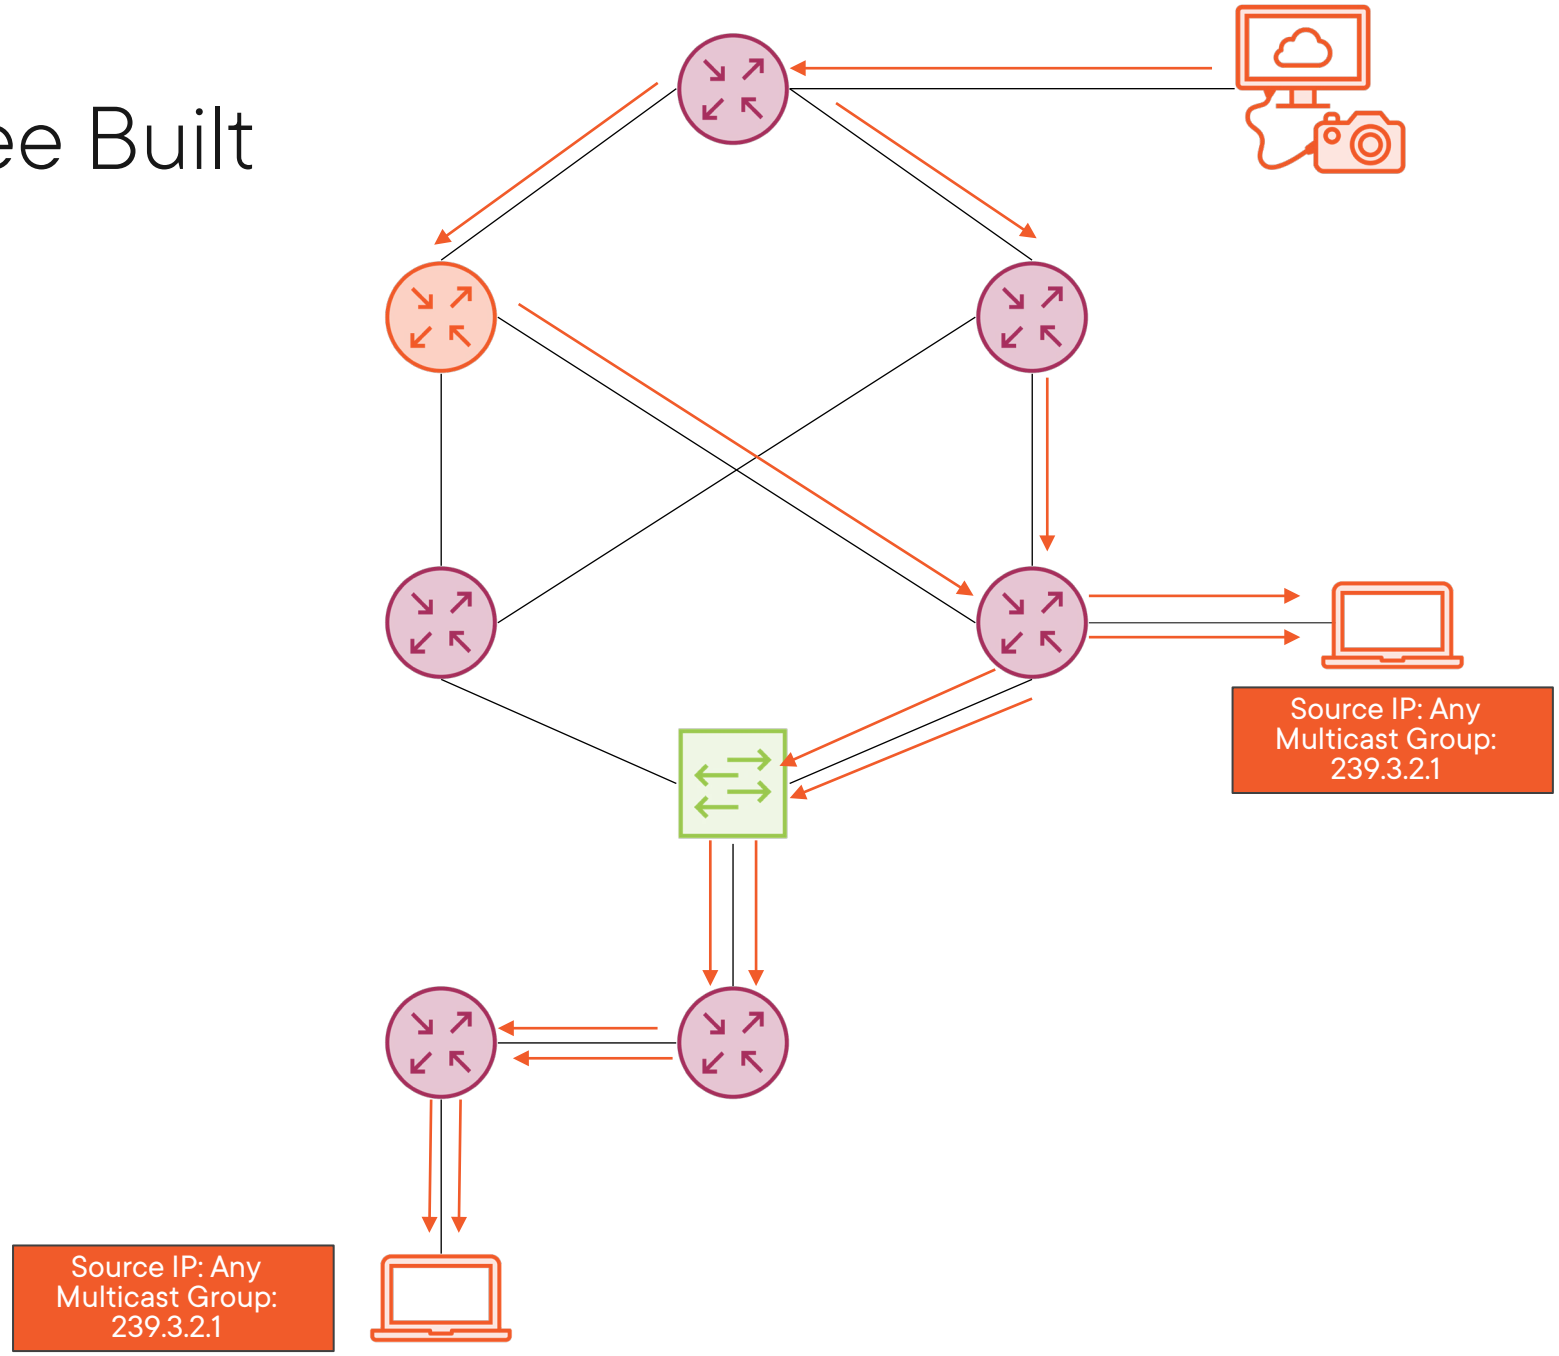

PIM Prune from RP

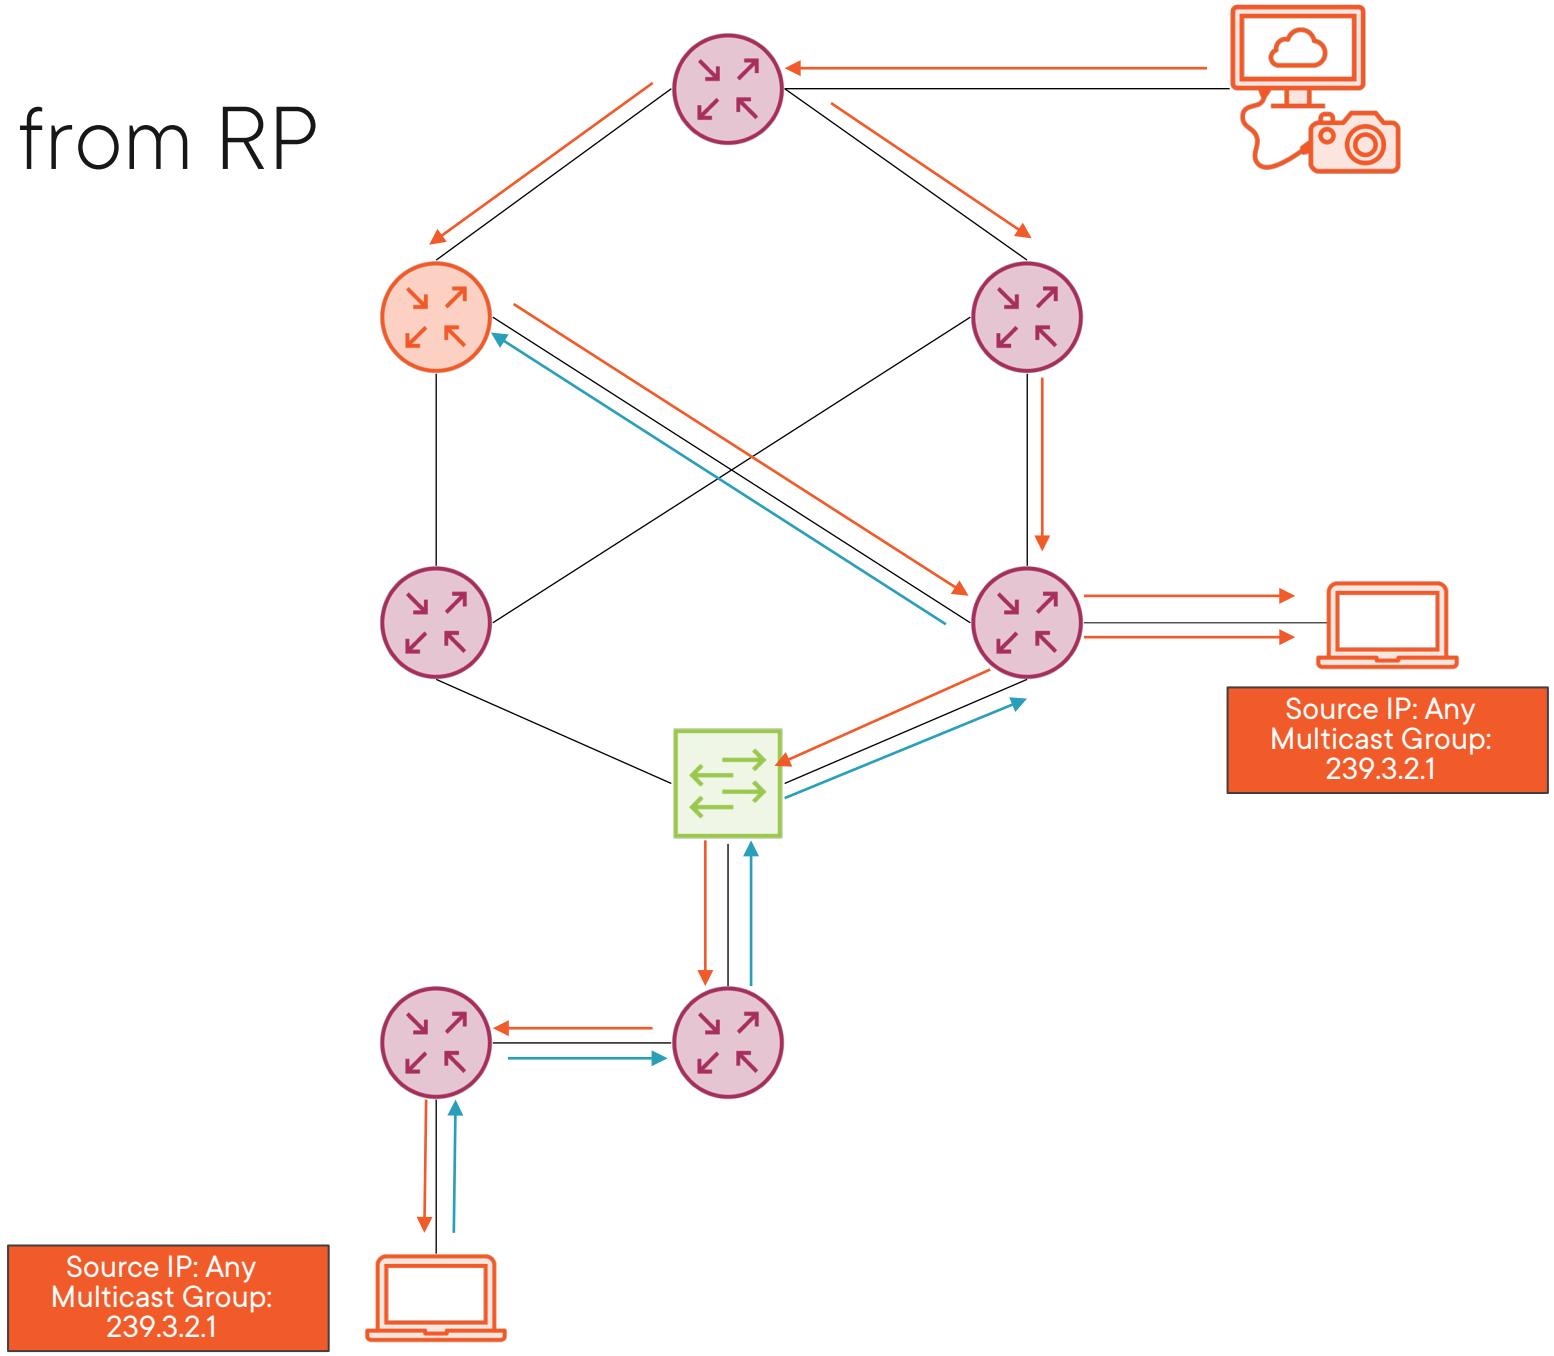

PIM Prune from FHR

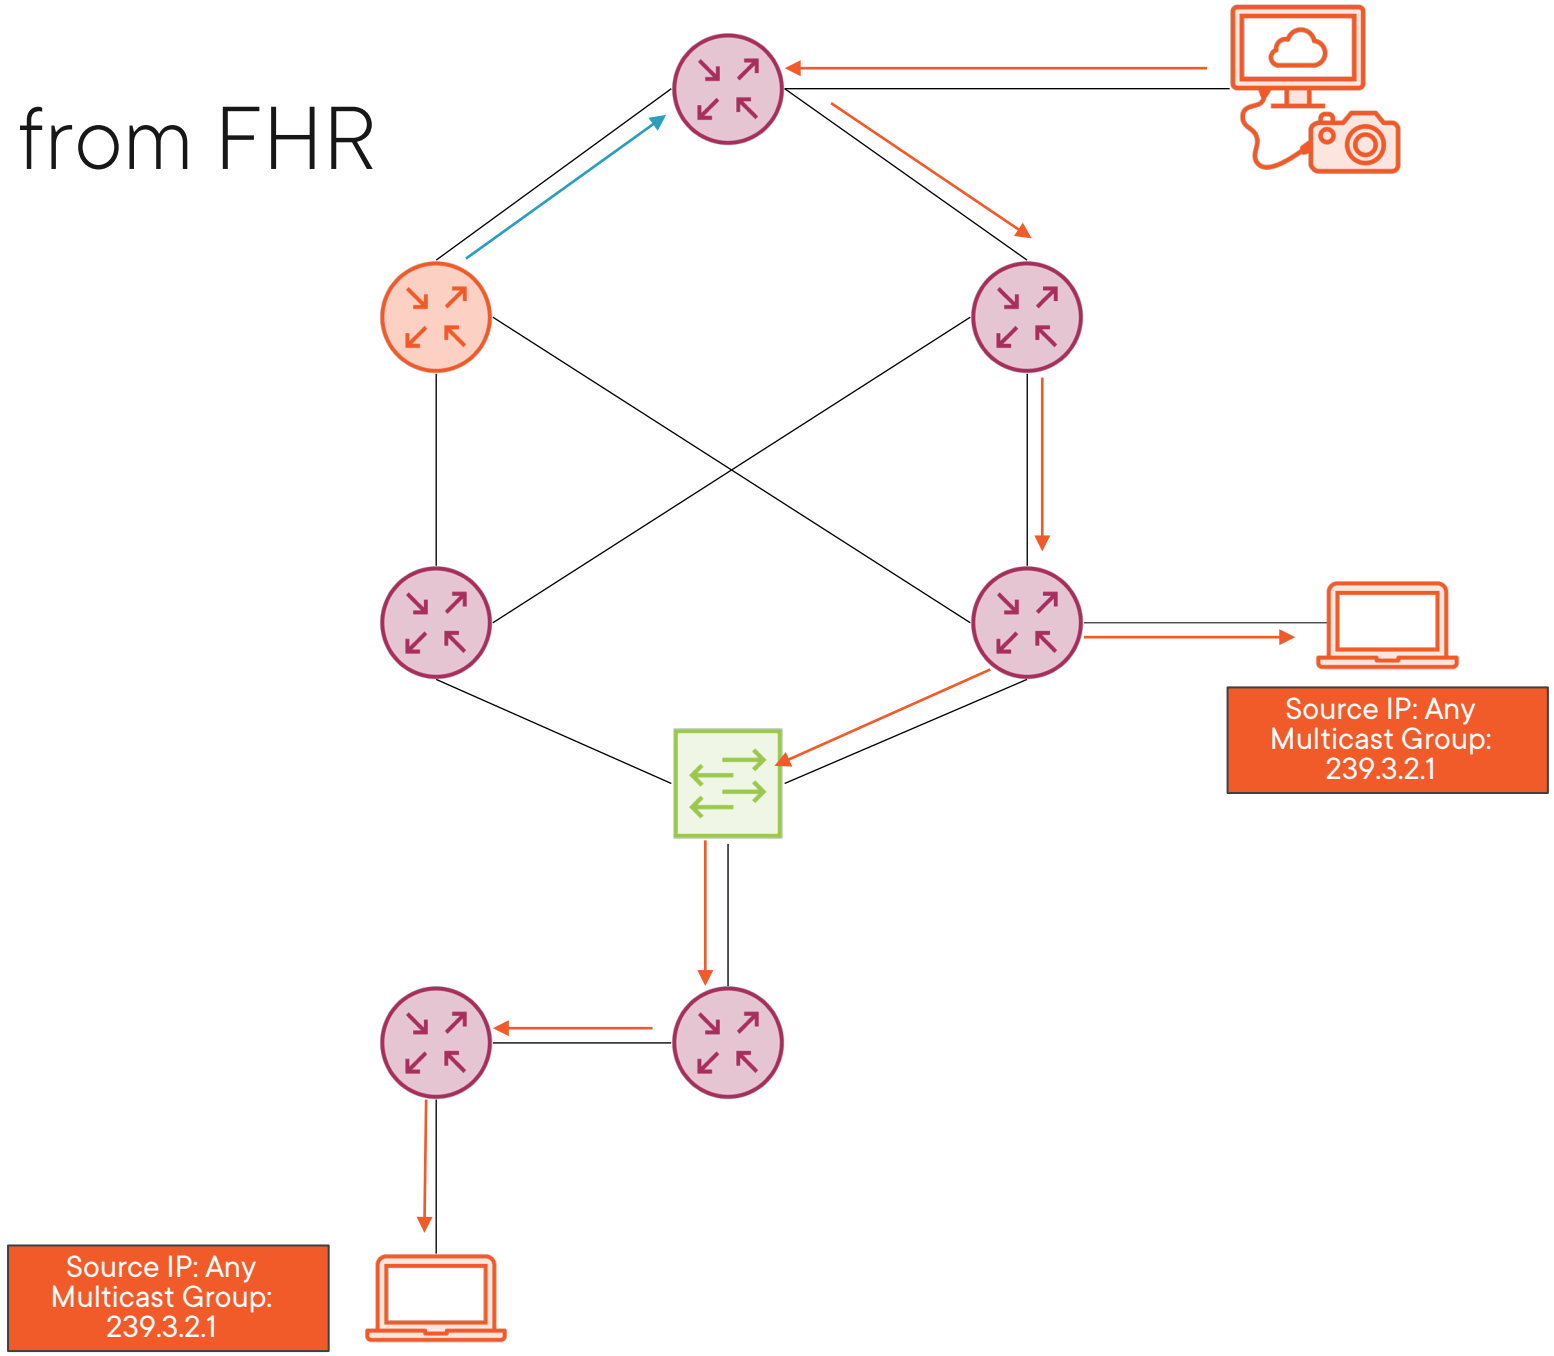

(S,G) Source Tree

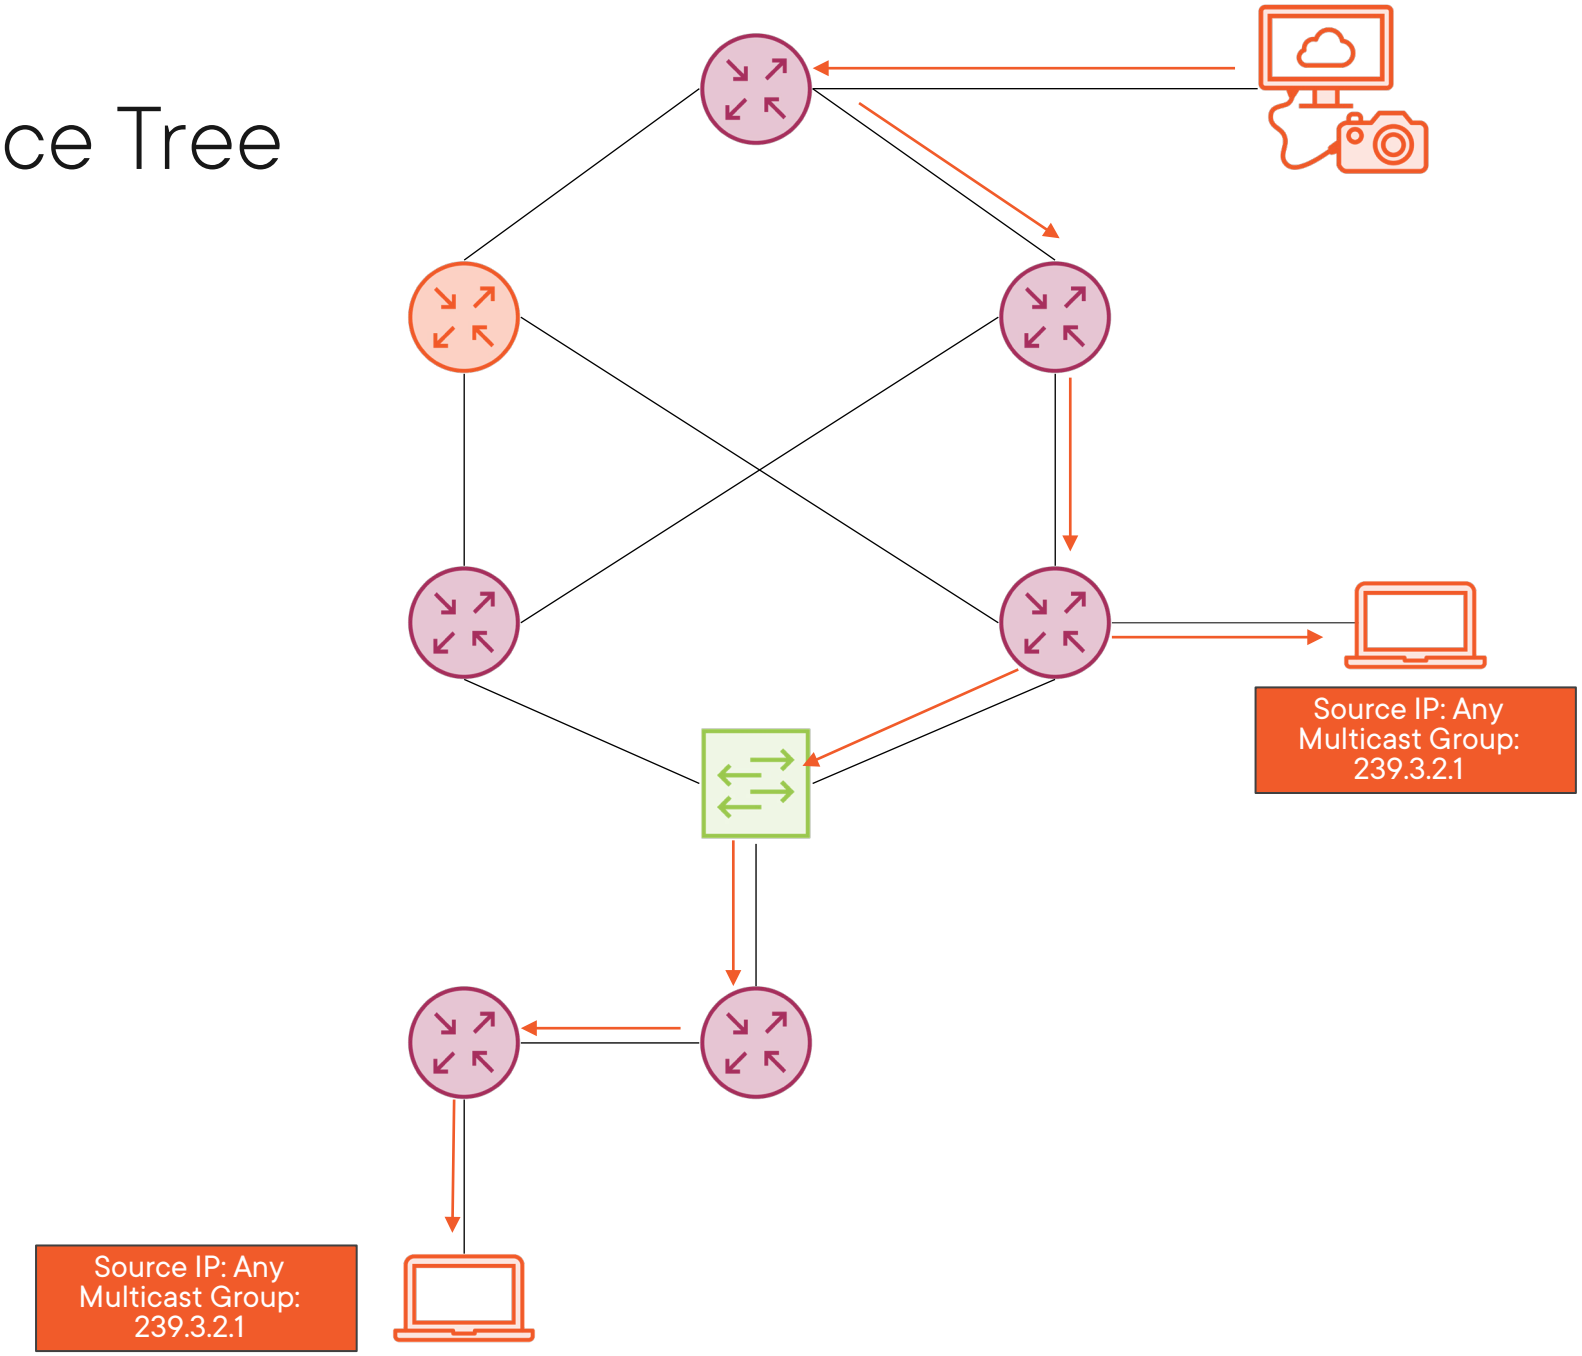

Demo

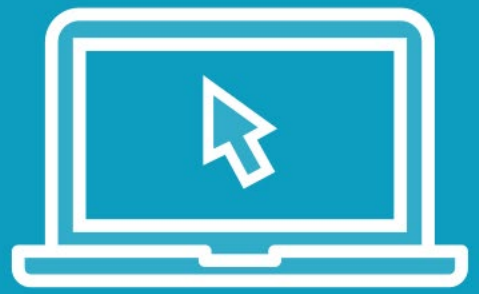

Rendezvous Point Election with BSR

- Configure a Bootstrap Router
- Use PIM Hellos to share BSR
- Dynamically elect an RP
- All routers learn RP address from PIM Bootstrap messages

RP and BSR

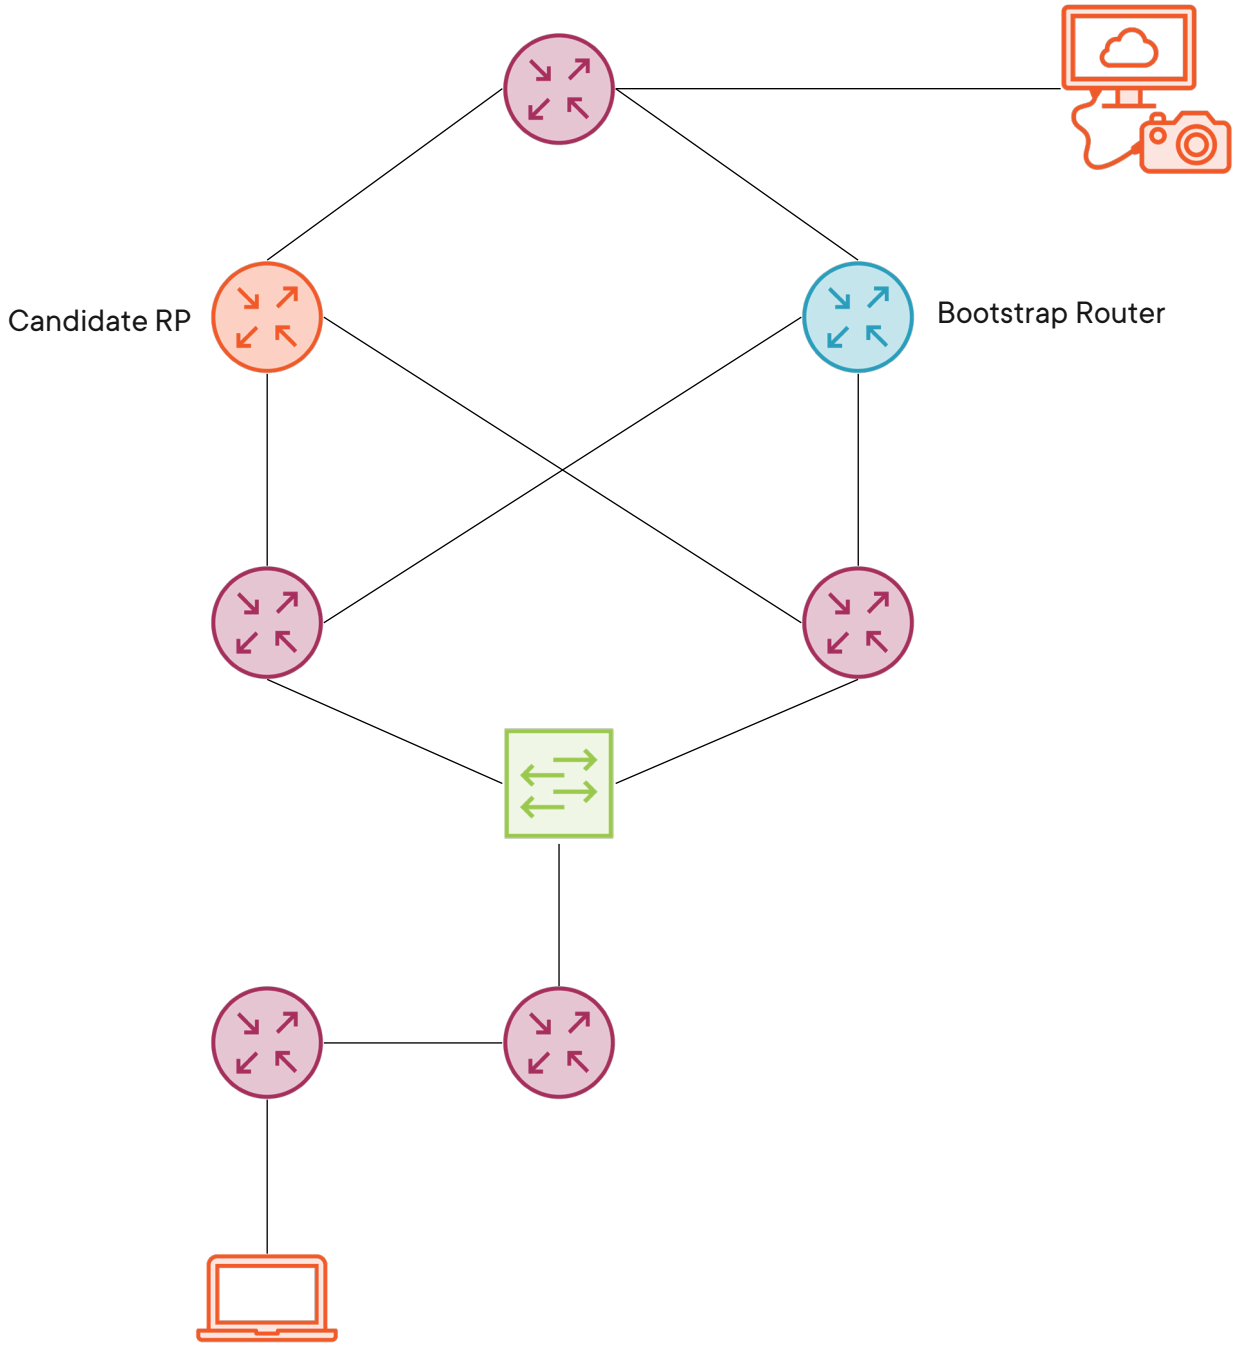

BSR Advertisement

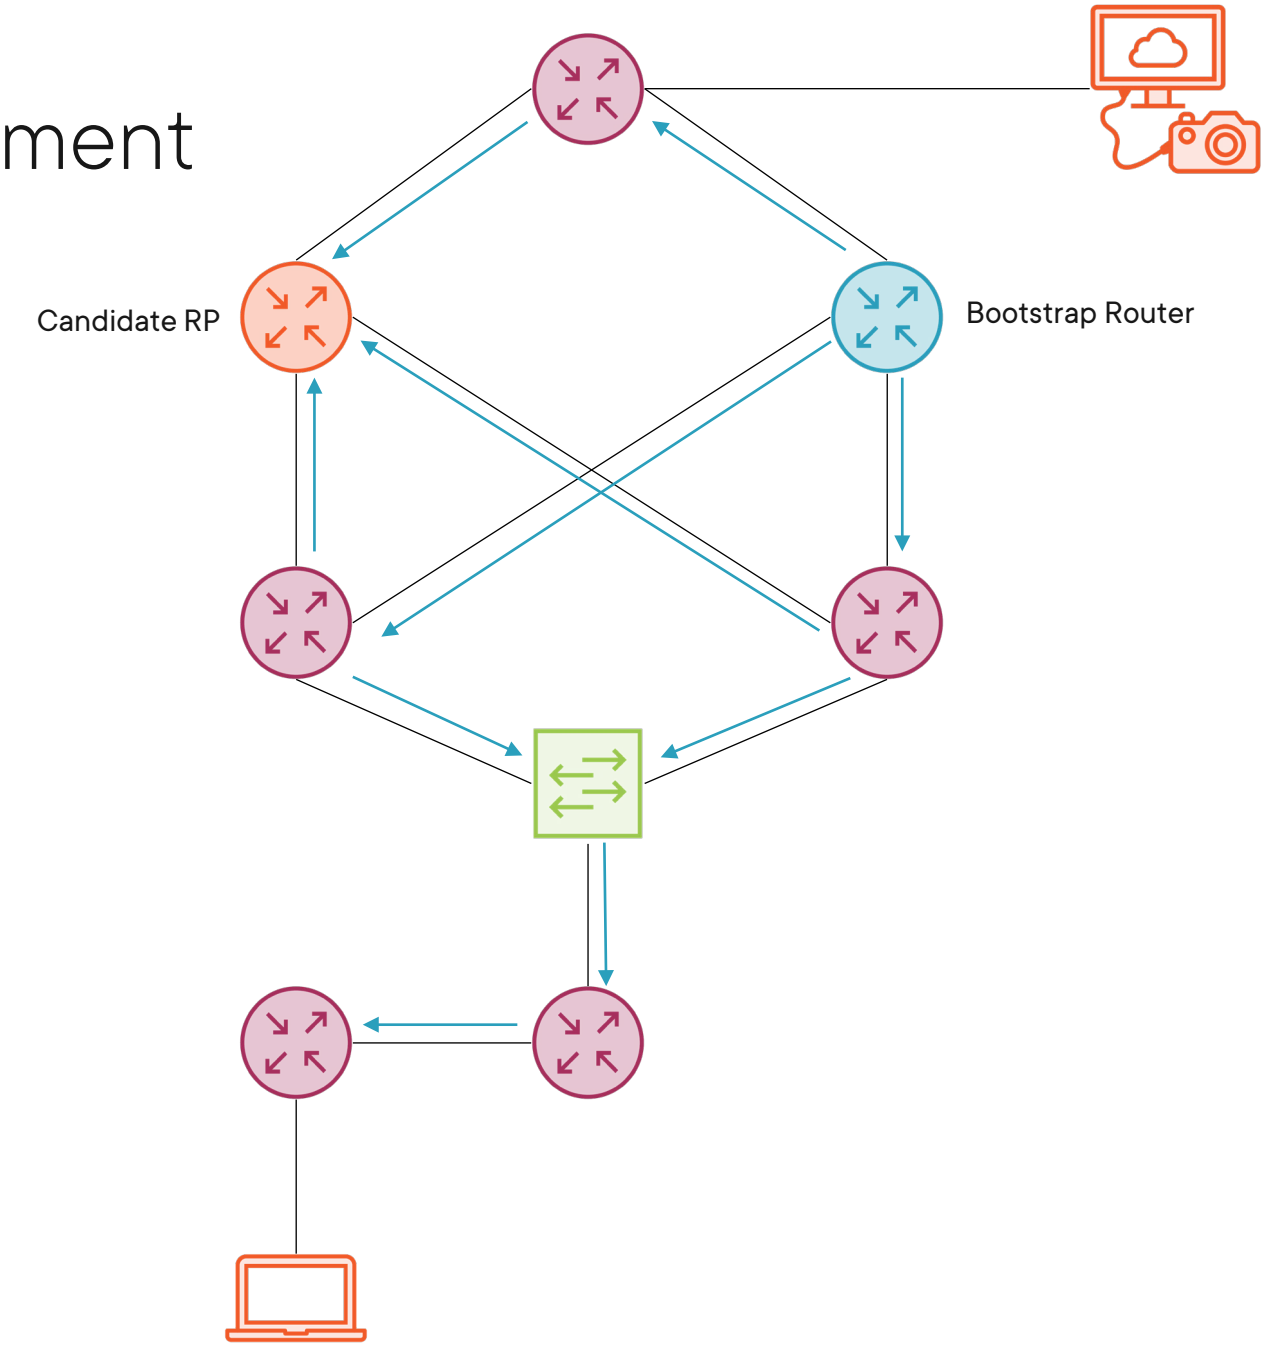

Candidate RP Advertisement

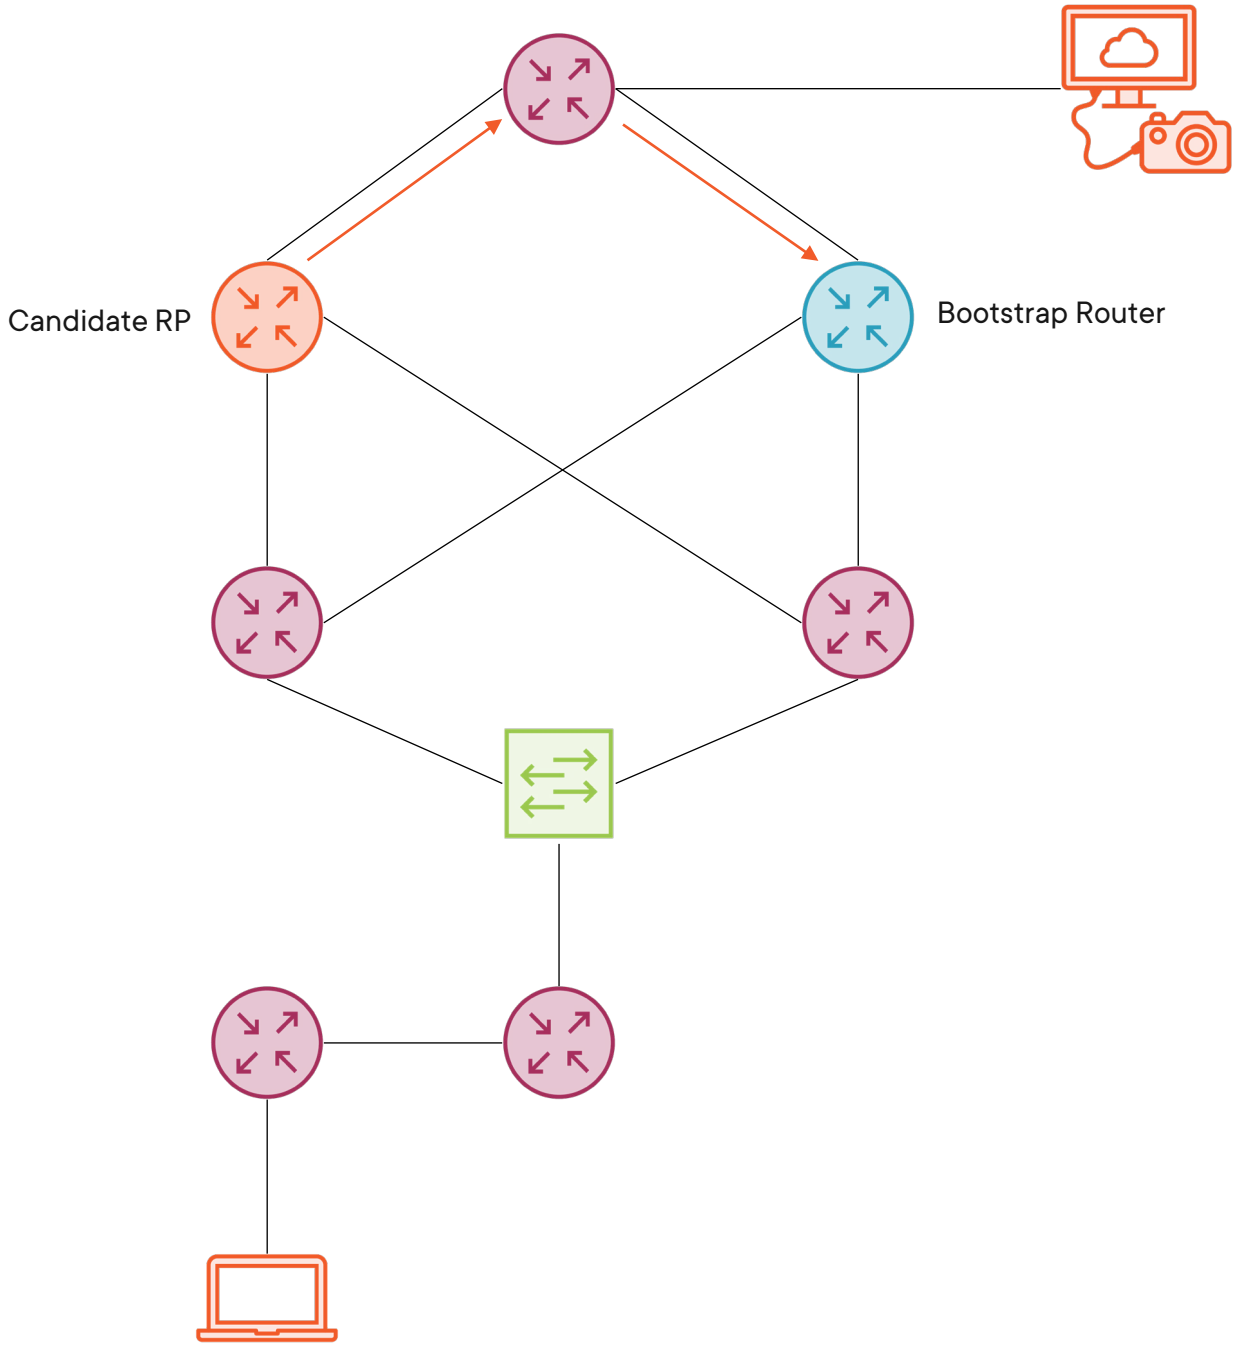

RP Selection and Advertisement

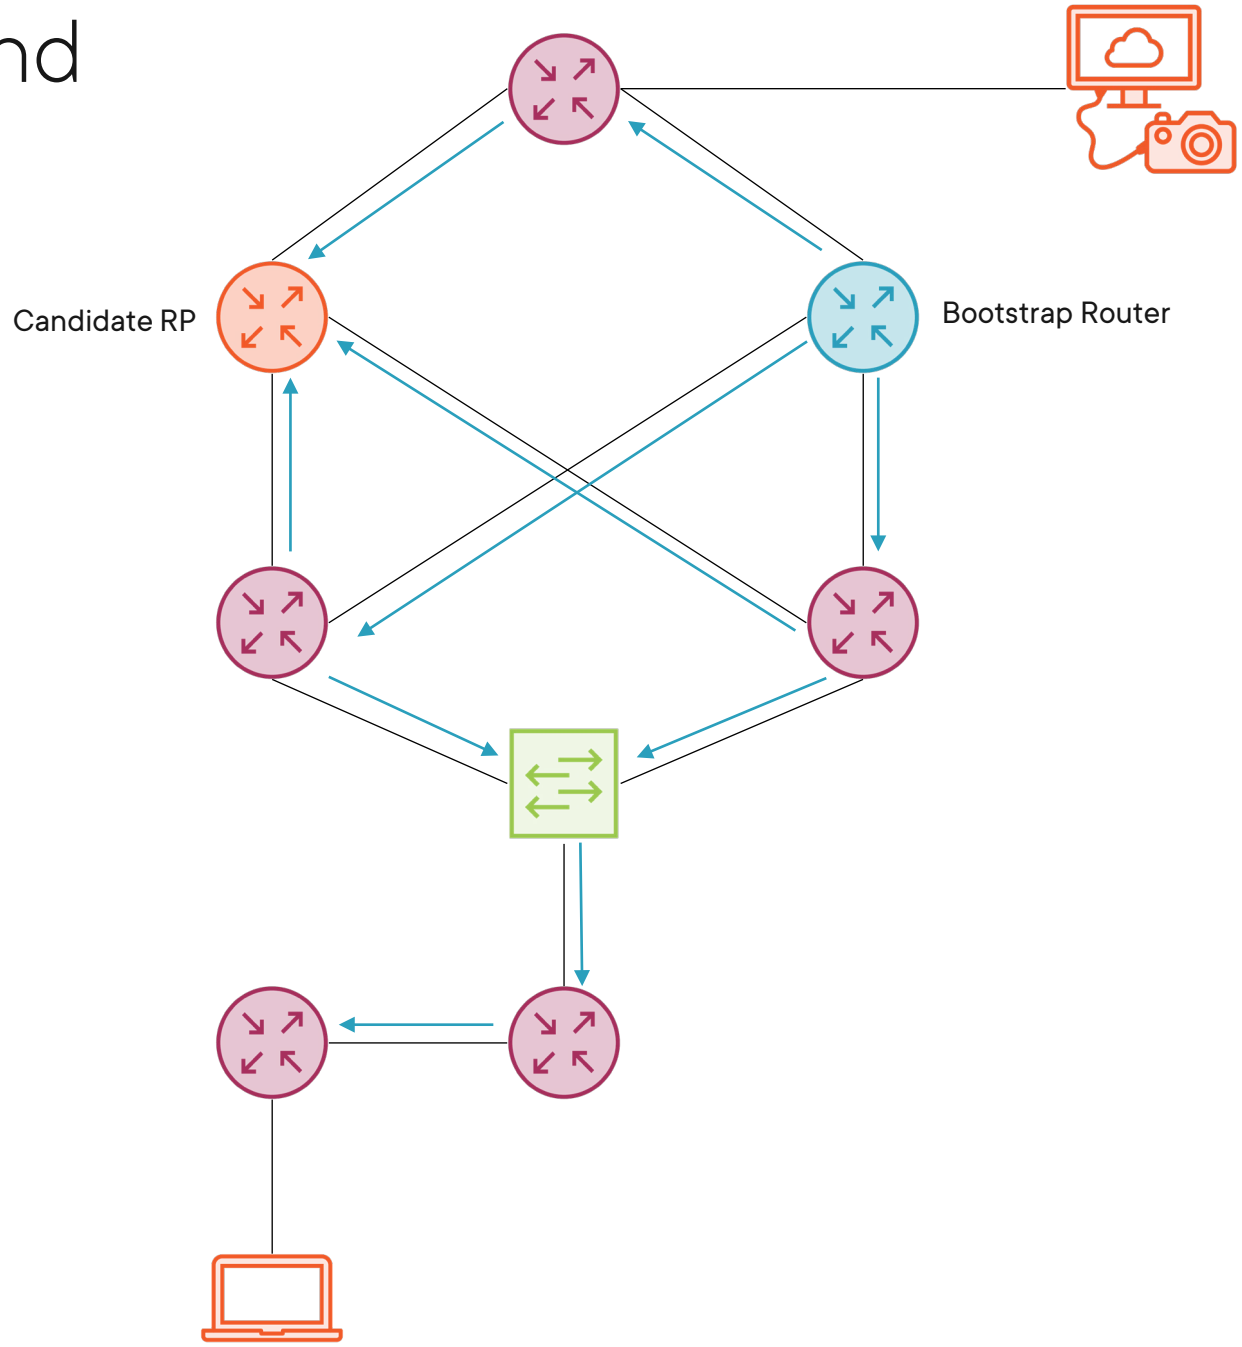

Demo

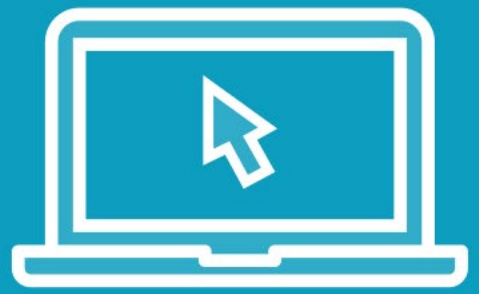

Send Multicast Source Traffic

- Originate multicast traffic in a lab setting to test multicast flows

Demo

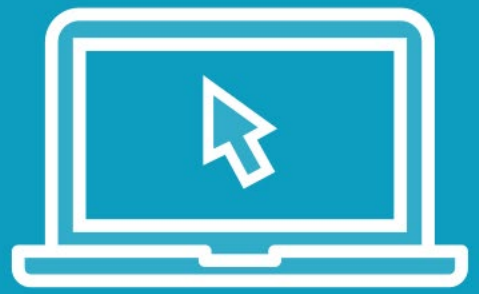

Create and Send a PIM Join

- Configure receiver to request multicast

Demo

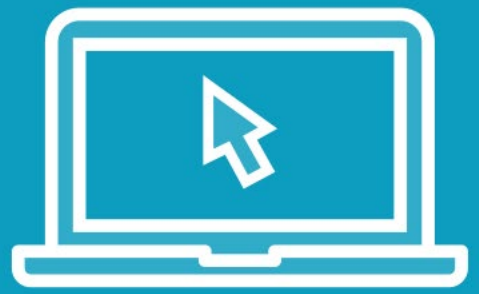

PIM Assert with Shortest Path Cutover

- Affects shared segments only
- When best path to the RP and source are through different routers

PIM Assert

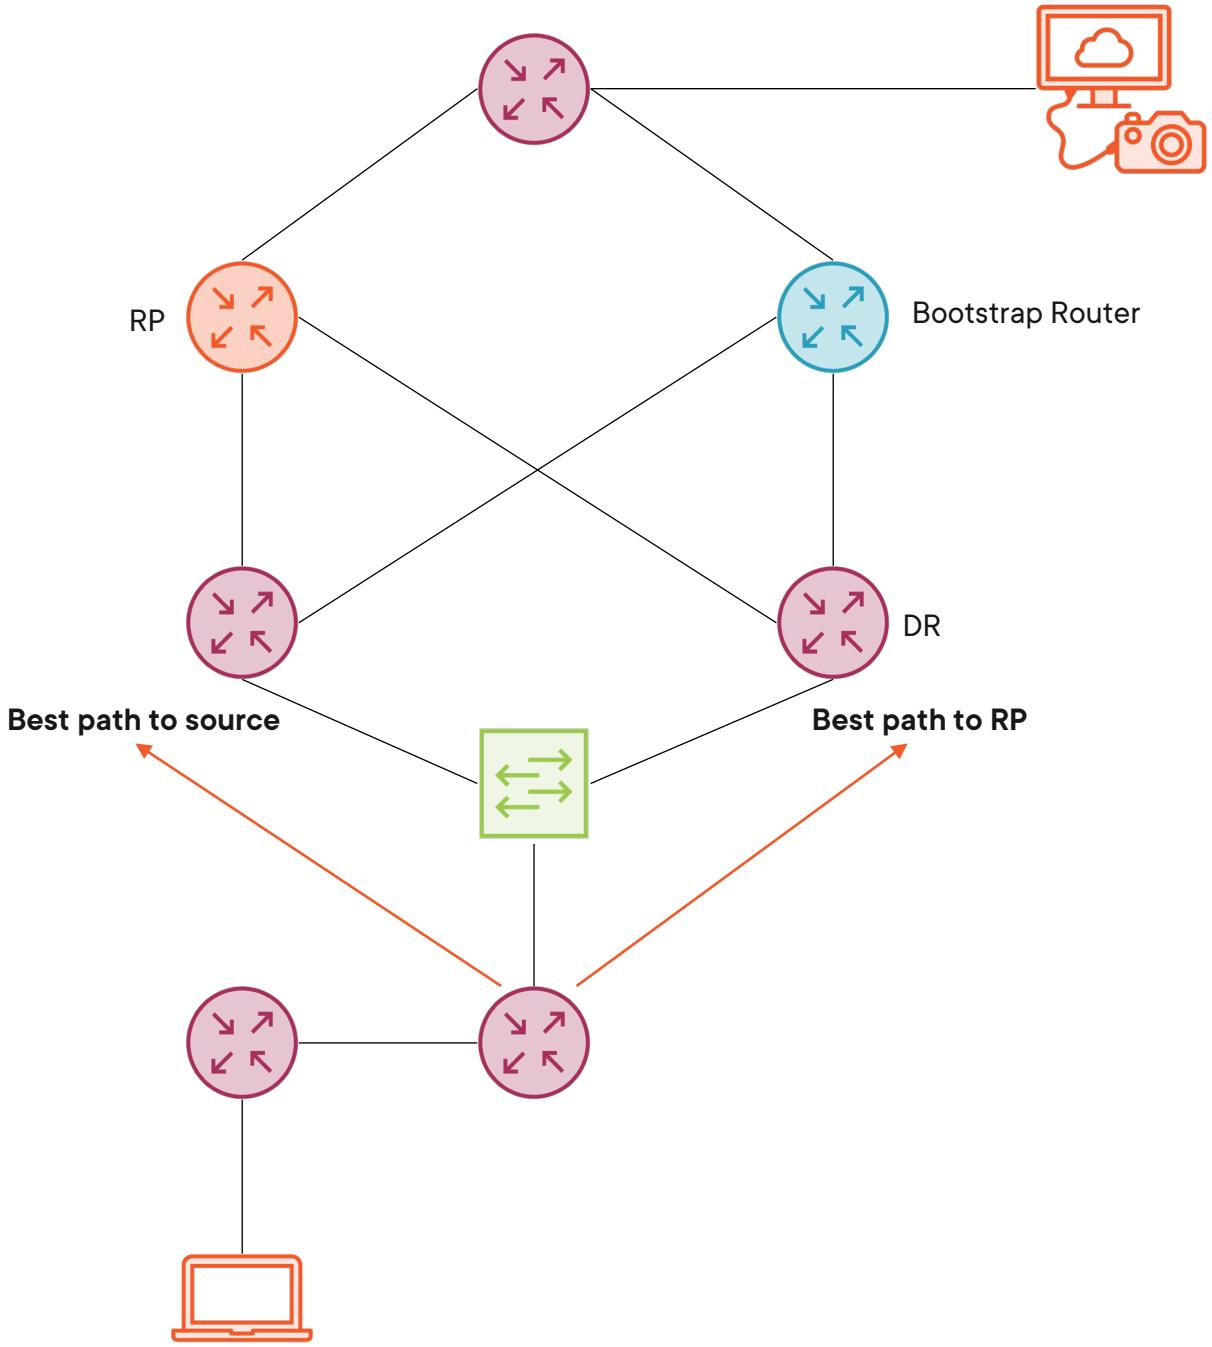

PIM Assert

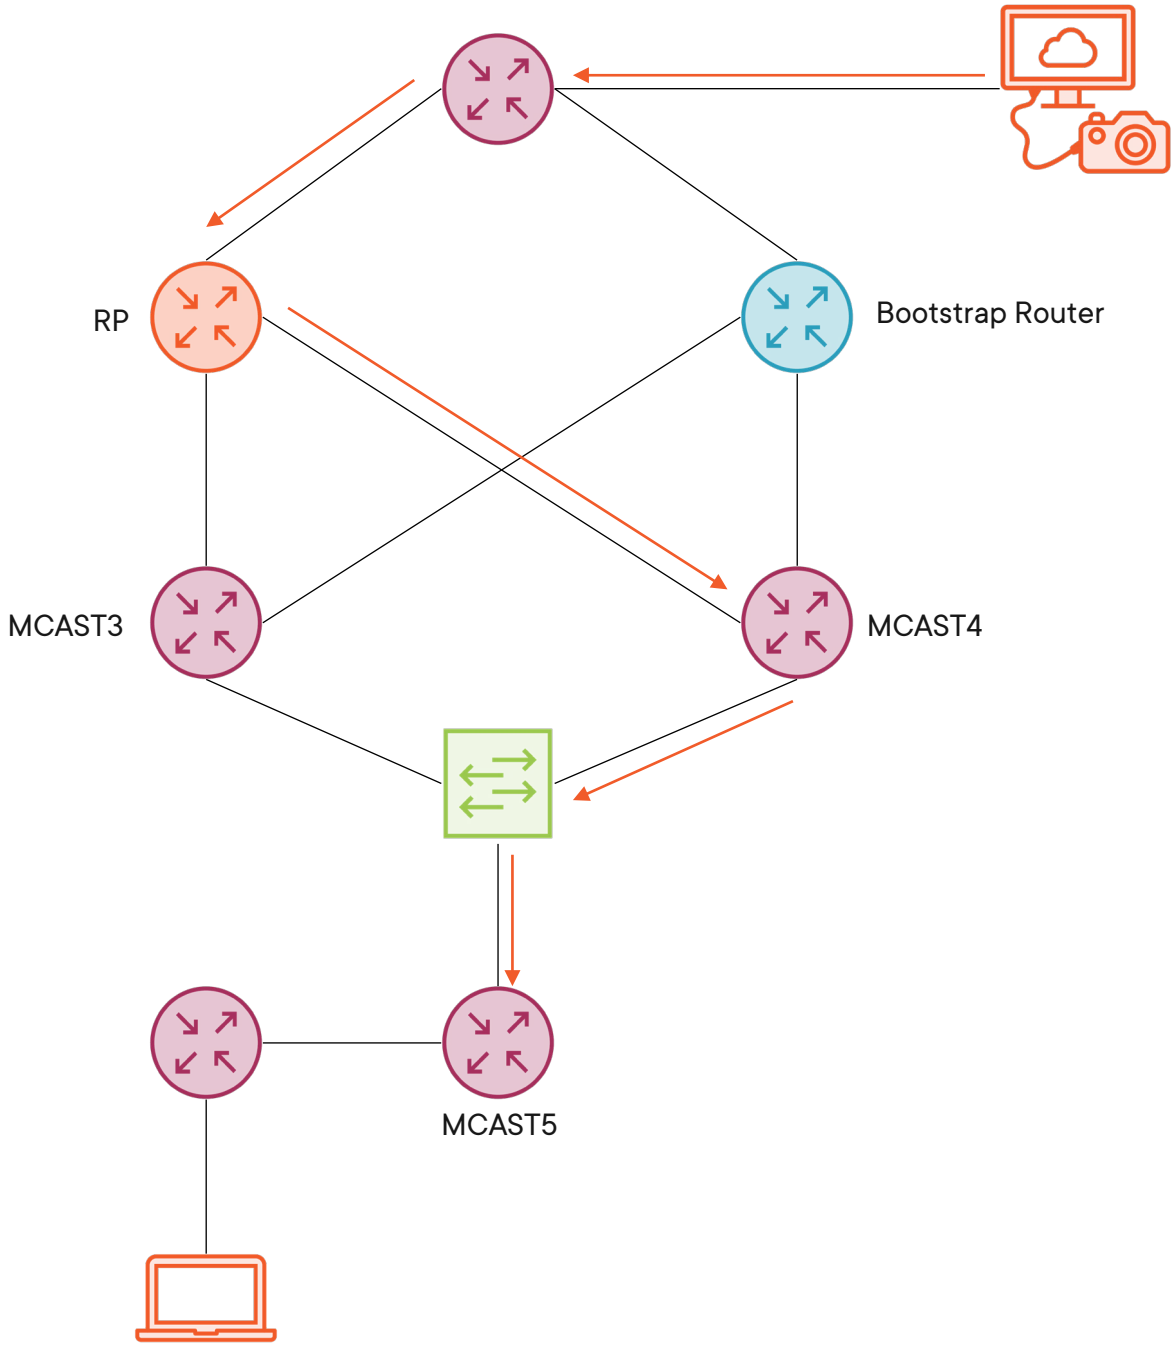

PIM Assert

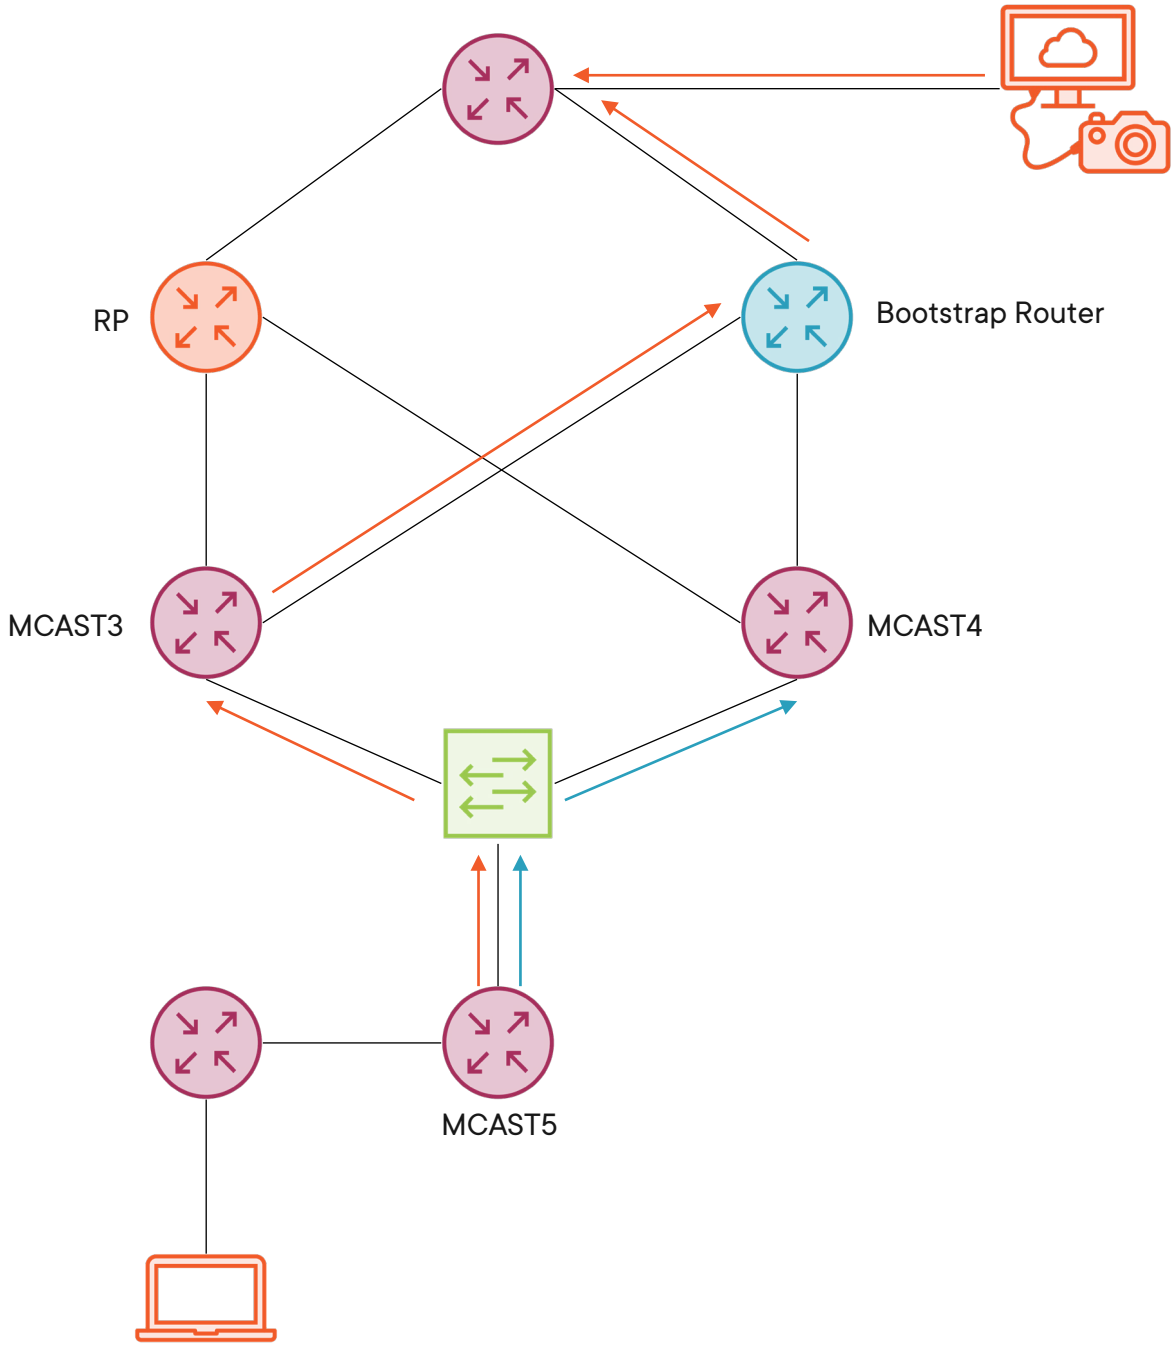

RP and BSR

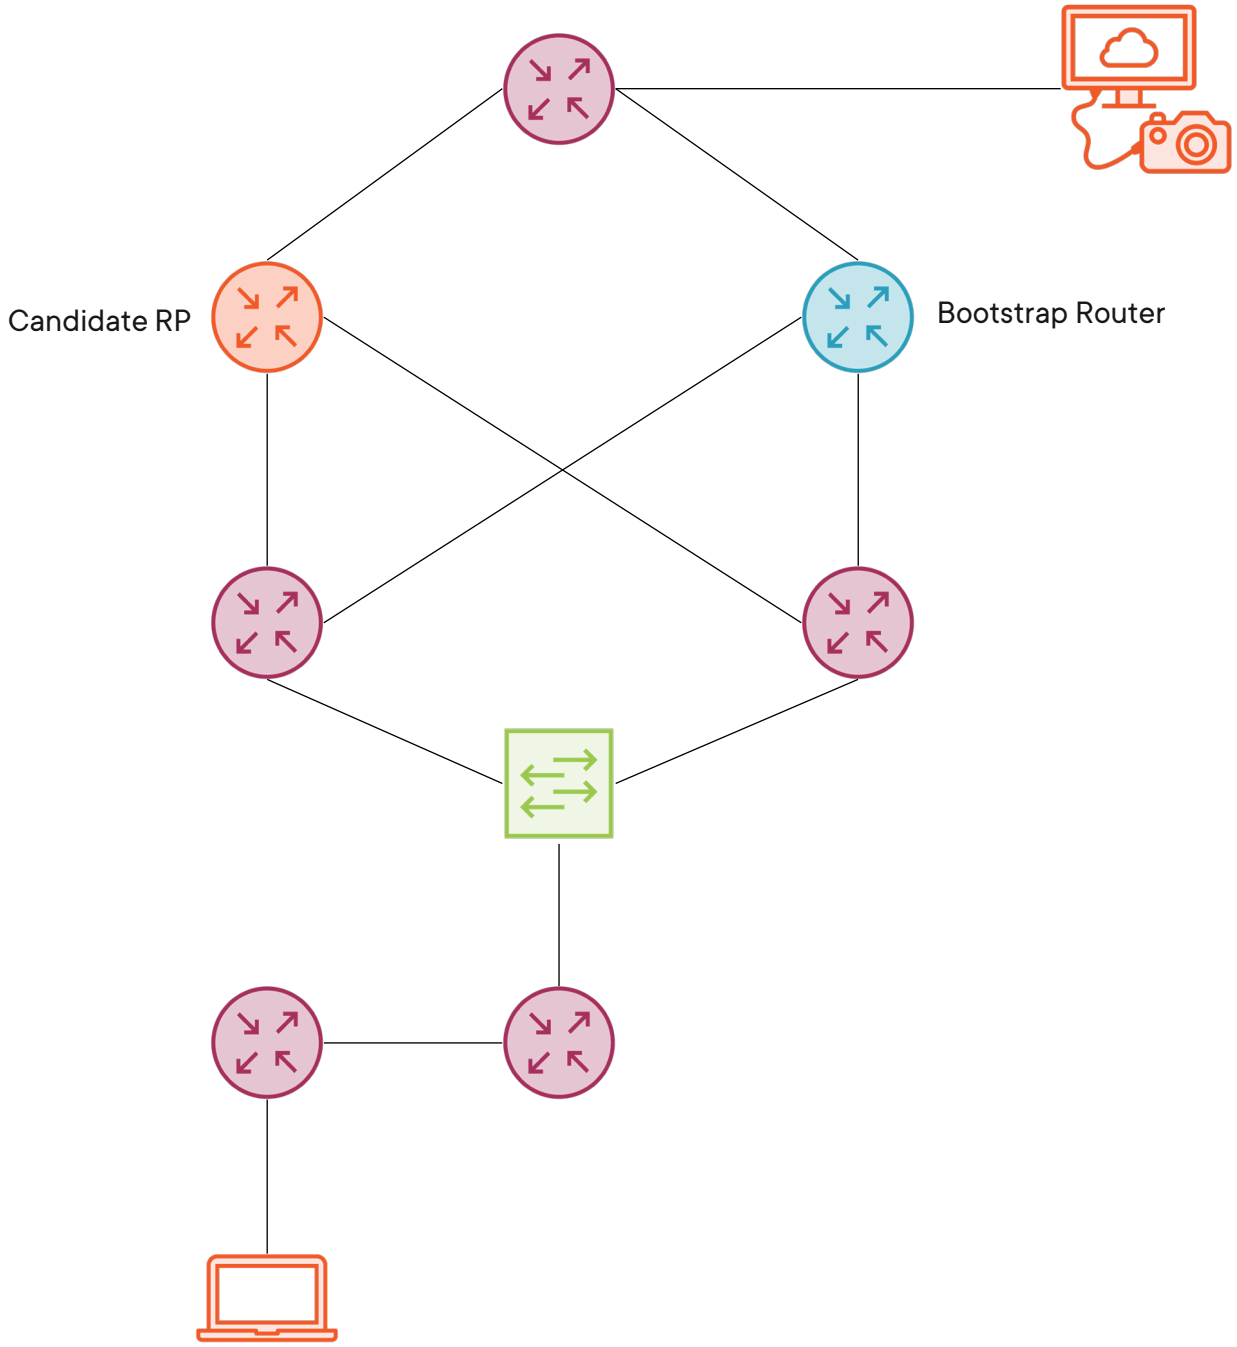

PIM Bootstrap Router

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000000	10.255.23.2	224.0.0.13	PIMv2	70	Bootstrap

```
> Frame 1: 70 bytes on wire (560 bits), 70 bytes captured (560 bits) on interface eth0, id 0
> Ethernet II, Src: 50:00:00:03:00:02 (50:00:00:03:00:02), Dst: IPv4mcast_0d (01:00:5e:00:00:0d)
> Internet Protocol Version 4, Src: 10.255.23.2, Dst: 224.0.0.13
v Protocol Independent Multicast
    0010 .... = Version: 2
    .... 0100 = Type: Bootstrap (4)
    Reserved byte(s): 00
    Checksum: 0xe3ba [correct]
    [Checksum Status: Good]
  v PIM Options
      Fragment tag: 0x0da3
      Hash mask len: 0
      BSR priority: 0
    v BSR: 2.2.2.2
        Address Family: IPv4 (1)
        Encoding Type: Native (0)
        Unicast: 2.2.2.2
    v Group 0: 224.0.0.0/4
        Address Family: IPv4 (1)
        Encoding Type: Native (0)
      > Flags: 0x00
        Masklen: 4
        Group: 224.0.0.0
        RP count: 1
        FRP count: 1
        Priority: 0
    v RP 0: 1.1.1.1
        Address Family: IPv4 (1)
        Encoding Type: Native (0)
        Unicast: 1.1.1.1
      Holdtime: 150
    Reserved byte(s): 00
```

PIM Candidate RP

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000000	1.1.1.1	2.2.2.2	PIMv2	56	Candidate-RP-Advertisement

```
> Frame 1: 56 bytes on wire (448 bits), 56 bytes captured (448 bits) on interface eth0, id 0
> Ethernet II, Src: 50:00:00:04:00:01 (50:00:00:04:00:01), Dst: 50:00:00:03:00:00 (50:00:00:03:00:00)
> Internet Protocol Version 4, Src: 1.1.1.1, Dst: 2.2.2.2
v Protocol Independent Multicast
    0010 .... = Version: 2
    .... 1000 = Type: Candidate-RP-Advertisement (8)
    Reserved byte(s): 00
    Checksum: 0xf262 [correct]
    [Checksum Status: Good]
  v PIM Options
      Prefix-count: 1
      Priority: 0
      Holdtime: 150
    v RP: 1.1.1.1
        Address Family: IPv4 (1)
        Encoding Type: Native (0)
        Unicast: 1.1.1.1
    v Group 0: 224.0.0.0/4
        Address Family: IPv4 (1)
        Encoding Type: Native (0)
      > Flags: 0x00
        Masklen: 4
        Group: 224.0.0.0
```

PIM Register

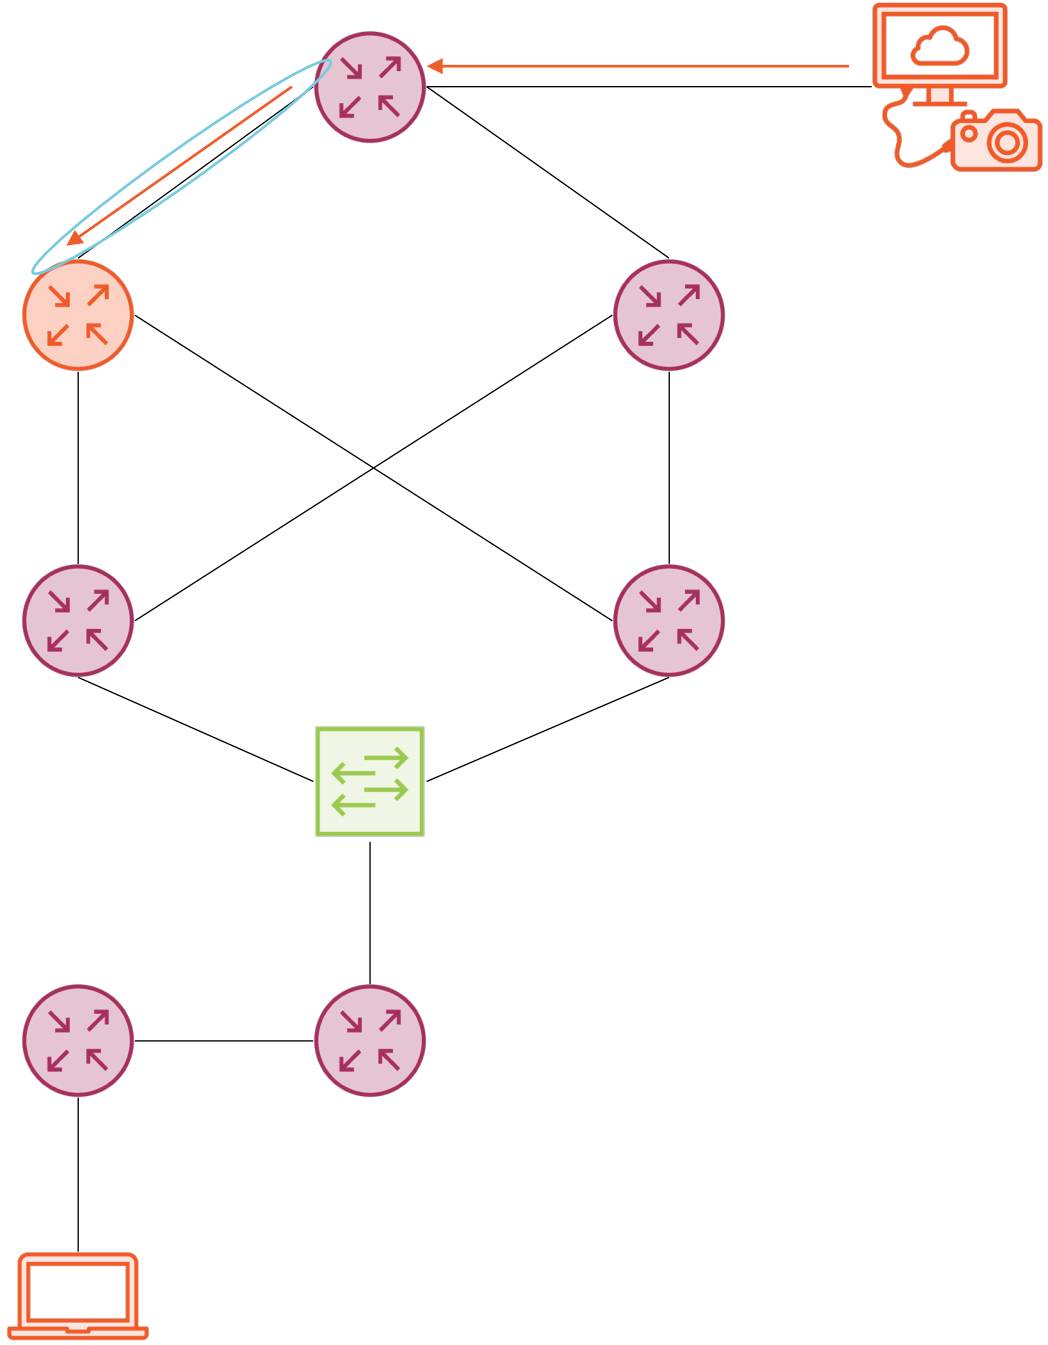

PIM Register Stop

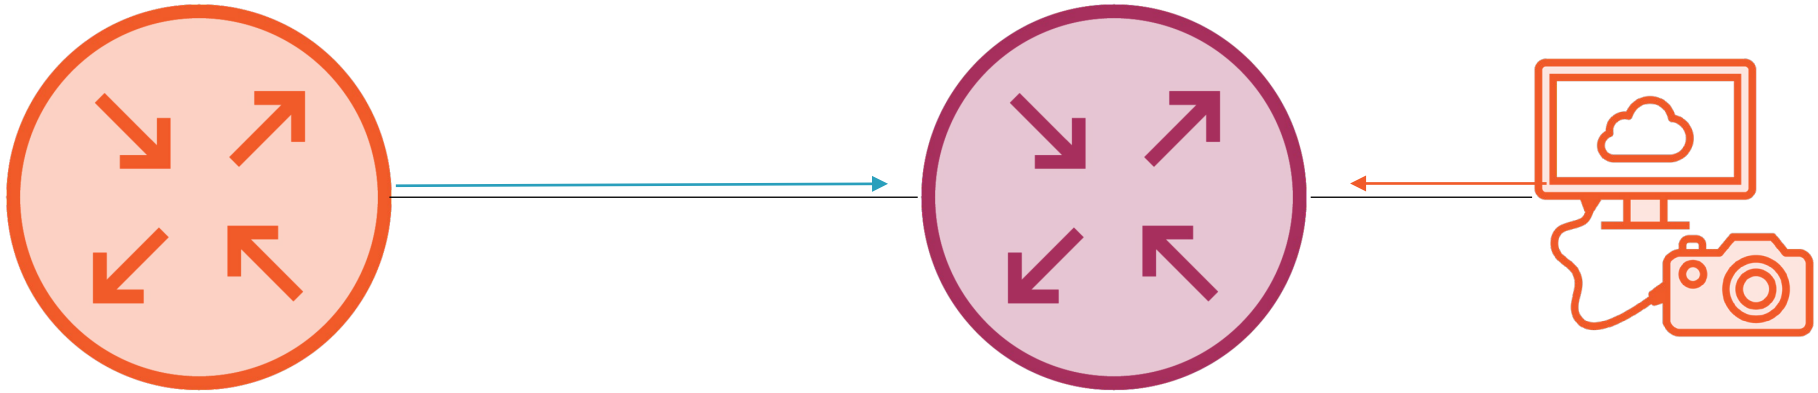

PIM Register

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000000	10.99.1.100	239.3.2.1	PIMv2	142	Register
2	0.009010484	1.1.1.1	10.255.10.254	PIMv2	52	Register-stop

```
> Frame 1: 142 bytes on wire (1136 bits), 142 bytes captured (1136 bits) on interface eth0, id 0
> Ethernet II, Src: 50:00:00:04:00:00 (50:00:00:04:00:00), Dst: 50:00:00:05:00:00 (50:00:00:05:00:00)
> Internet Protocol Version 4, Src: 10.255.10.254, Dst: 1.1.1.1
v Protocol Independent Multicast
    0010 .... = Version: 2
    .... 0001 = Type: Register (1)
    Reserved byte(s): 00
    Checksum: 0xdeff [correct]
    [Checksum Status: Good]
  v PIM Options
    v Flags: 0x00000000
        0... .... .... .... .... .... .... .... = Border: No
        .0.. .... .... .... .... .... .... .... = Null-Register: No
      0100 .... = IP Version: IPv4 (4)
> Internet Protocol Version 4, Src: 10.99.1.100, Dst: 239.3.2.1
v Internet Control Message Protocol
    Type: 8 (Echo (ping) request)
    Code: 0
    Checksum: 0x73d4 [correct]
    [Checksum Status: Good]
    Identifier (BE): 23 (0x0017)
    Identifier (LE): 5888 (0x1700)
    Sequence Number (BE): 0 (0x0000)
    Sequence Number (LE): 0 (0x0000)
  > Data (72 bytes)
```

PIM Register Stop

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000000	10.99.1.100	239.3.2.1	PIMv2	142	Register
2	0.009010484	1.1.1.1	10.255.10.254	PIMv2	52	Register-stop

```
> Frame 2: 52 bytes on wire (416 bits), 52 bytes captured (416 bits) on interface eth0, id 0
> Ethernet II, Src: 50:00:00:05:00:00 (50:00:00:05:00:00), Dst: 50:00:00:04:00:00 (50:00:00:04:00:00)
> Internet Protocol Version 4, Src: 1.1.1.1, Dst: 10.255.10.254
v Protocol Independent Multicast
    0010 .... = Version: 2
    .... 0010 = Type: Register-stop (2)
    Reserved byte(s): 00
    Checksum: 0xdf13 [correct]
    [Checksum Status: Good]
  v PIM Options
    v Group: 239.3.2.1/32
        Address Family: IPv4 (1)
        Encoding Type: Native (0)
      > Flags: 0x00
        Masklen: 32
        Group: 239.3.2.1
    v Source: 10.99.1.100
        Address Family: IPv4 (1)
        Encoding Type: Native (0)
        Unicast: 10.99.1.100
```

PIM Joins Build OIL

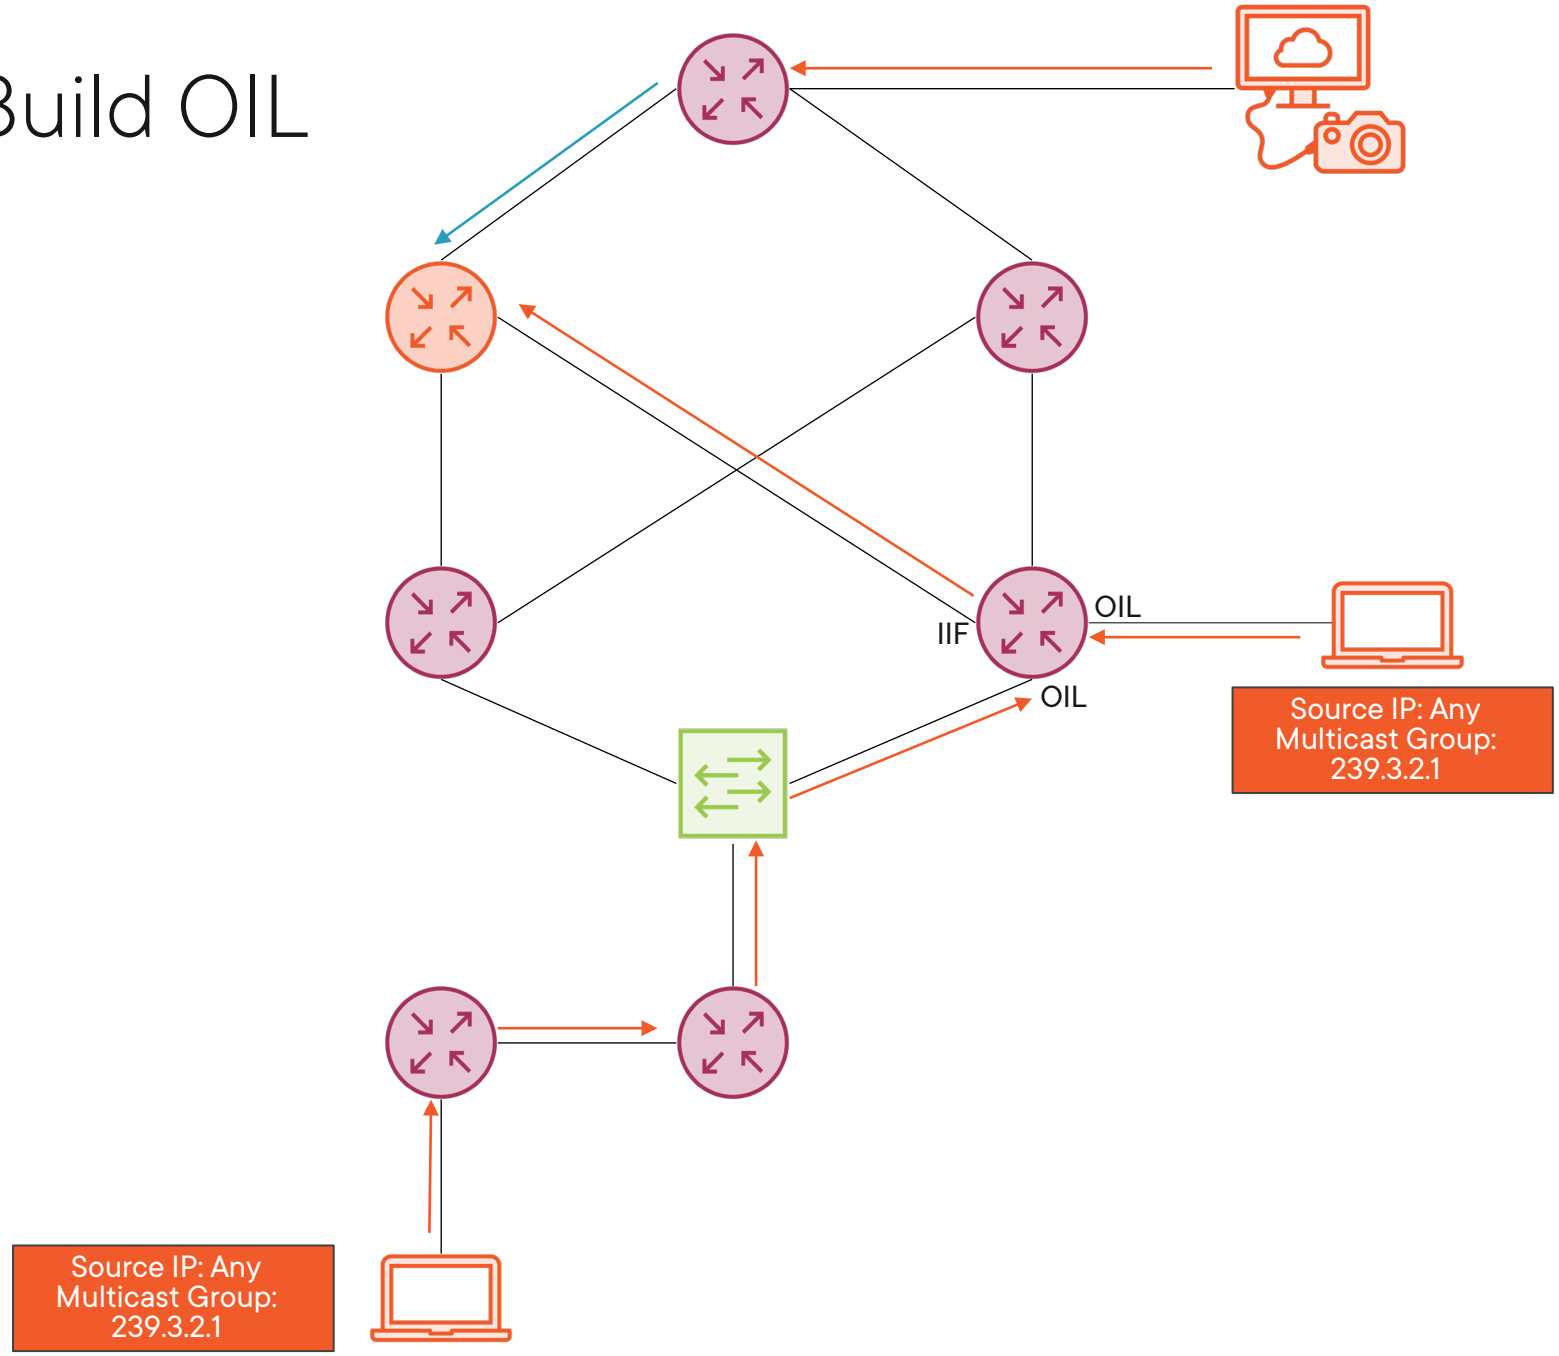

Source Tree Building

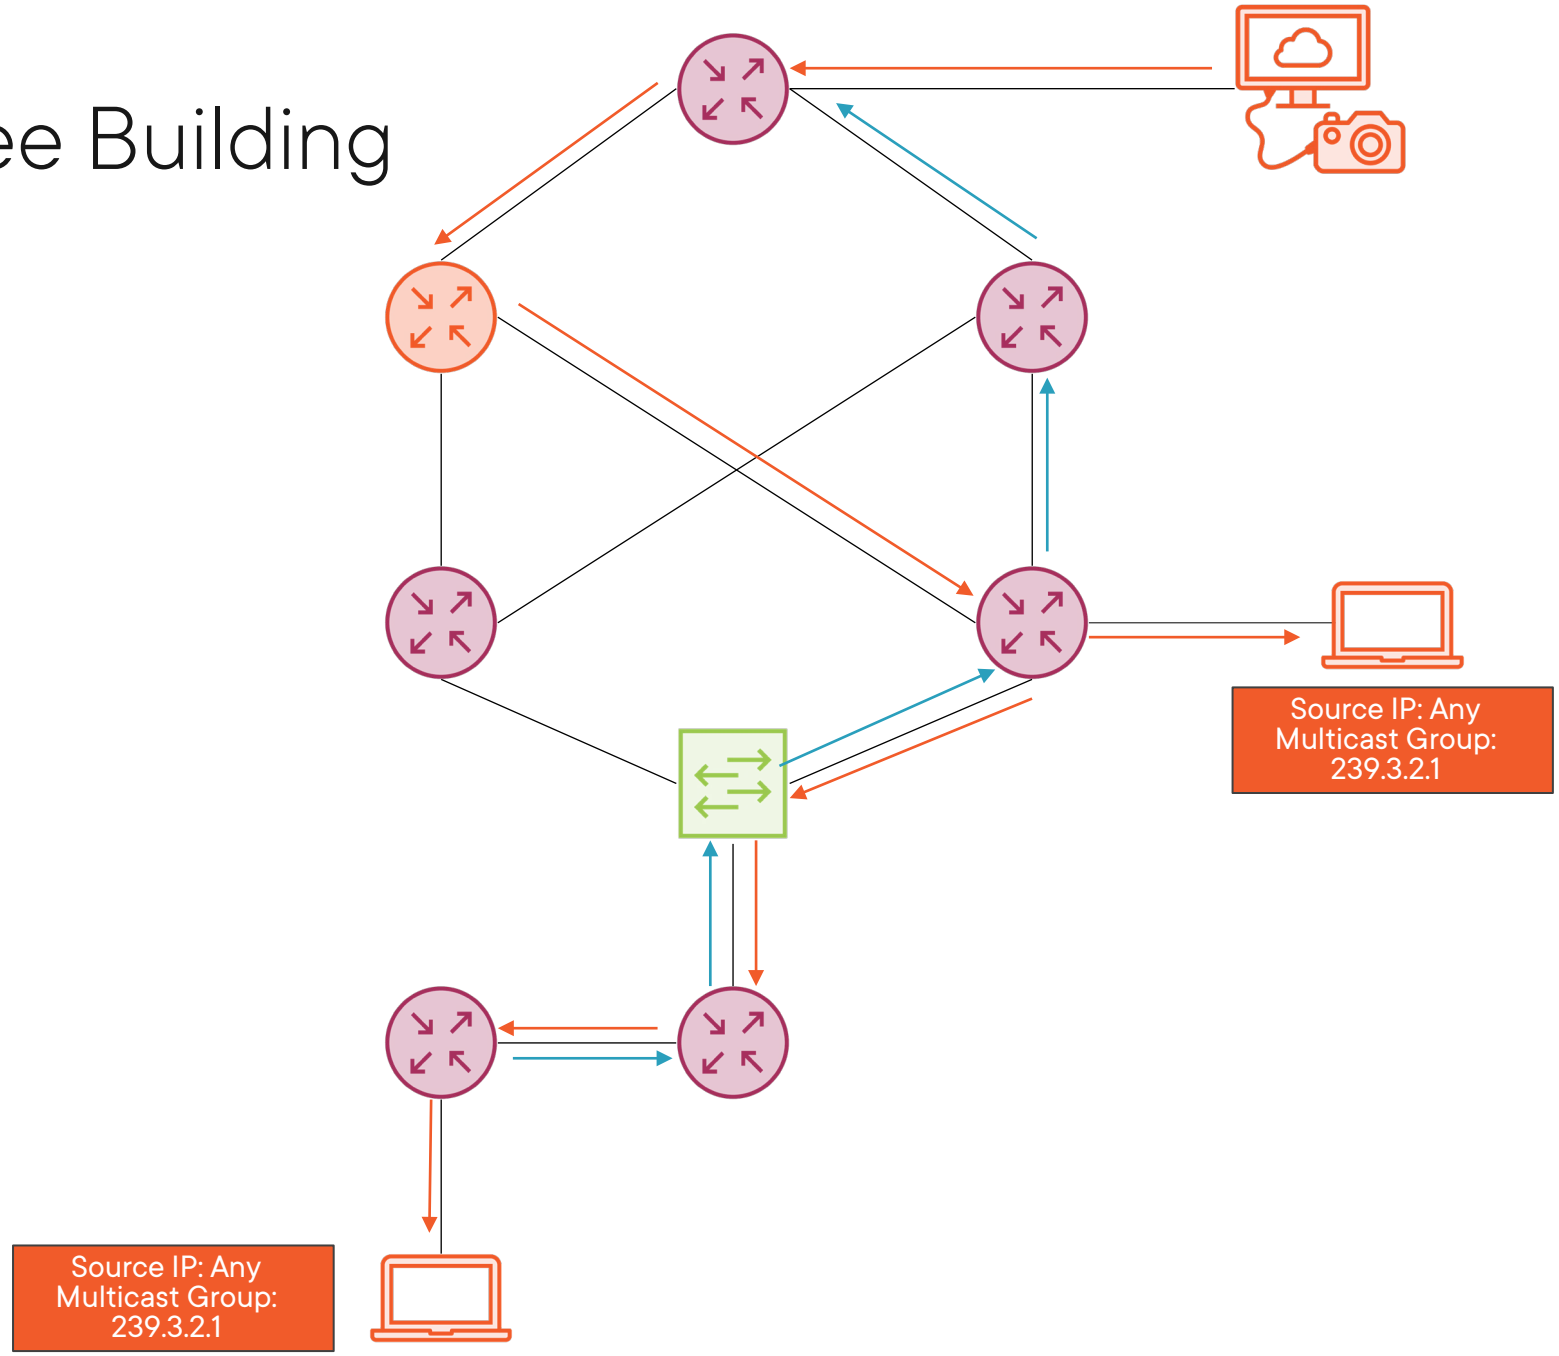

PIM Join (Shared)

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000000	10.255.254.5	224.0.0.13	PIMv2	68	Join/Prune

```
> Frame 1: 68 bytes on wire (544 bits), 68 bytes captured (544 bits) on interface eth0, id 0
> Ethernet II, Src: 50:00:00:06:00:00 (50:00:00:06:00:00), Dst: IPv4mcast_0d (01:00:5e:00:00:0d)
> Internet Protocol Version 4, Src: 10.255.254.5, Dst: 224.0.0.13
v Protocol Independent Multicast
    0010 .... = Version: 2
    .... 0011 = Type: Join/Prune (3)
    Reserved byte(s): 00
    Checksum: 0xd5e0 [correct]
    [Checksum Status: Good]
  v PIM Options
    v Upstream-neighbor: 10.255.254.4
        Address Family: IPv4 (1)
        Encoding Type: Native (0)
        Unicast: 10.255.254.4
      Reserved byte(s): 00
      Num Groups: 1
      Holdtime: 210
    v Group 0
      v Group 0: 239.3.2.1/32
          Address Family: IPv4 (1)
          Encoding Type: Native (0)
        > Flags: 0x00
          Masklen: 32
          Group: 239.3.2.1
      v Num Joins: 1
        v IP address: 1.1.1.1/32 (SWR)
            Address Family: IPv4 (1)
            Encoding Type: Native (0)
          > Flags: 0x07, Sparse, WildCard, Rendezvous Point Tree
            Masklen: 32
            Source: 1.1.1.1
      Num Prunes: 0
```

PIM Join (Source)

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000000	10.255.254.5	224.0.0.13	PIMv2	68	Join/Prune

```
> Frame 1: 68 bytes on wire (544 bits), 68 bytes captured (544 bits) on interface eth0, id 0
> Ethernet II, Src: 50:00:00:06:00:00 (50:00:00:06:00:00), Dst: IPv4mcast_0d (01:00:5e:00:00:0d)
> Internet Protocol Version 4, Src: 10.255.254.5, Dst: 224.0.0.13
v Protocol Independent Multicast
    0010 .... = Version: 2
    .... 0011 = Type: Join/Prune (3)
    Reserved byte(s): 00
    Checksum: 0xcf1b [correct]
    [Checksum Status: Good]
  v PIM Options
    v Upstream-neighbor: 10.255.254.4
        Address Family: IPv4 (1)
        Encoding Type: Native (0)
        Unicast: 10.255.254.4
      Reserved byte(s): 00
      Num Groups: 1
      Holdtime: 210
    v Group 0
      v Group 0: 239.3.2.1/32
          Address Family: IPv4 (1)
          Encoding Type: Native (0)
        > Flags: 0x00
          Masklen: 32
          Group: 239.3.2.1
      v Num Joins: 1
        v IP address: 10.99.1.100/32 (S)
            Address Family: IPv4 (1)
            Encoding Type: Native (0)
          > Flags: 0x04, Sparse
            Masklen: 32
            Source: 10.99.1.100
      Num Prunes: 0
```

PIM Prune (Shared)

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000000	10.255.254.5	224.0.0.13	PIMv2	68	Join/Prune

```
> Frame 1: 68 bytes on wire (544 bits), 68 bytes captured (544 bits) on interface eth0, id 0
> Ethernet II, Src: 50:00:00:06:00:00 (50:00:00:06:00:00), Dst: IPv4mcast_0d (01:00:5e:00:00:0d)
> Internet Protocol Version 4, Src: 10.255.254.5, Dst: 224.0.0.13
v Protocol Independent Multicast
    0010 .... = Version: 2
    .... 0011 = Type: Join/Prune (3)
    Reserved byte(s): 00
    Checksum: 0xd5e0 [correct]
    [Checksum Status: Good]
  v PIM Options
    v Upstream-neighbor: 10.255.254.4
        Address Family: IPv4 (1)
        Encoding Type: Native (0)
        Unicast: 10.255.254.4
      Reserved byte(s): 00
      Num Groups: 1
      Holdtime: 210
    v Group 0
      v Group 0: 239.3.2.1/32
          Address Family: IPv4 (1)
          Encoding Type: Native (0)
        > Flags: 0x00
          Masklen: 32
          Group: 239.3.2.1
        Num Joins: 0
      v Num Prunes: 1
        v IP address: 1.1.1.1/32 (SWR)
            Address Family: IPv4 (1)
            Encoding Type: Native (0)
          > Flags: 0x07, Sparse, WildCard, Rendezvous Point Tree
            Masklen: 32
            Source: 1.1.1.1
```

PIM Assert

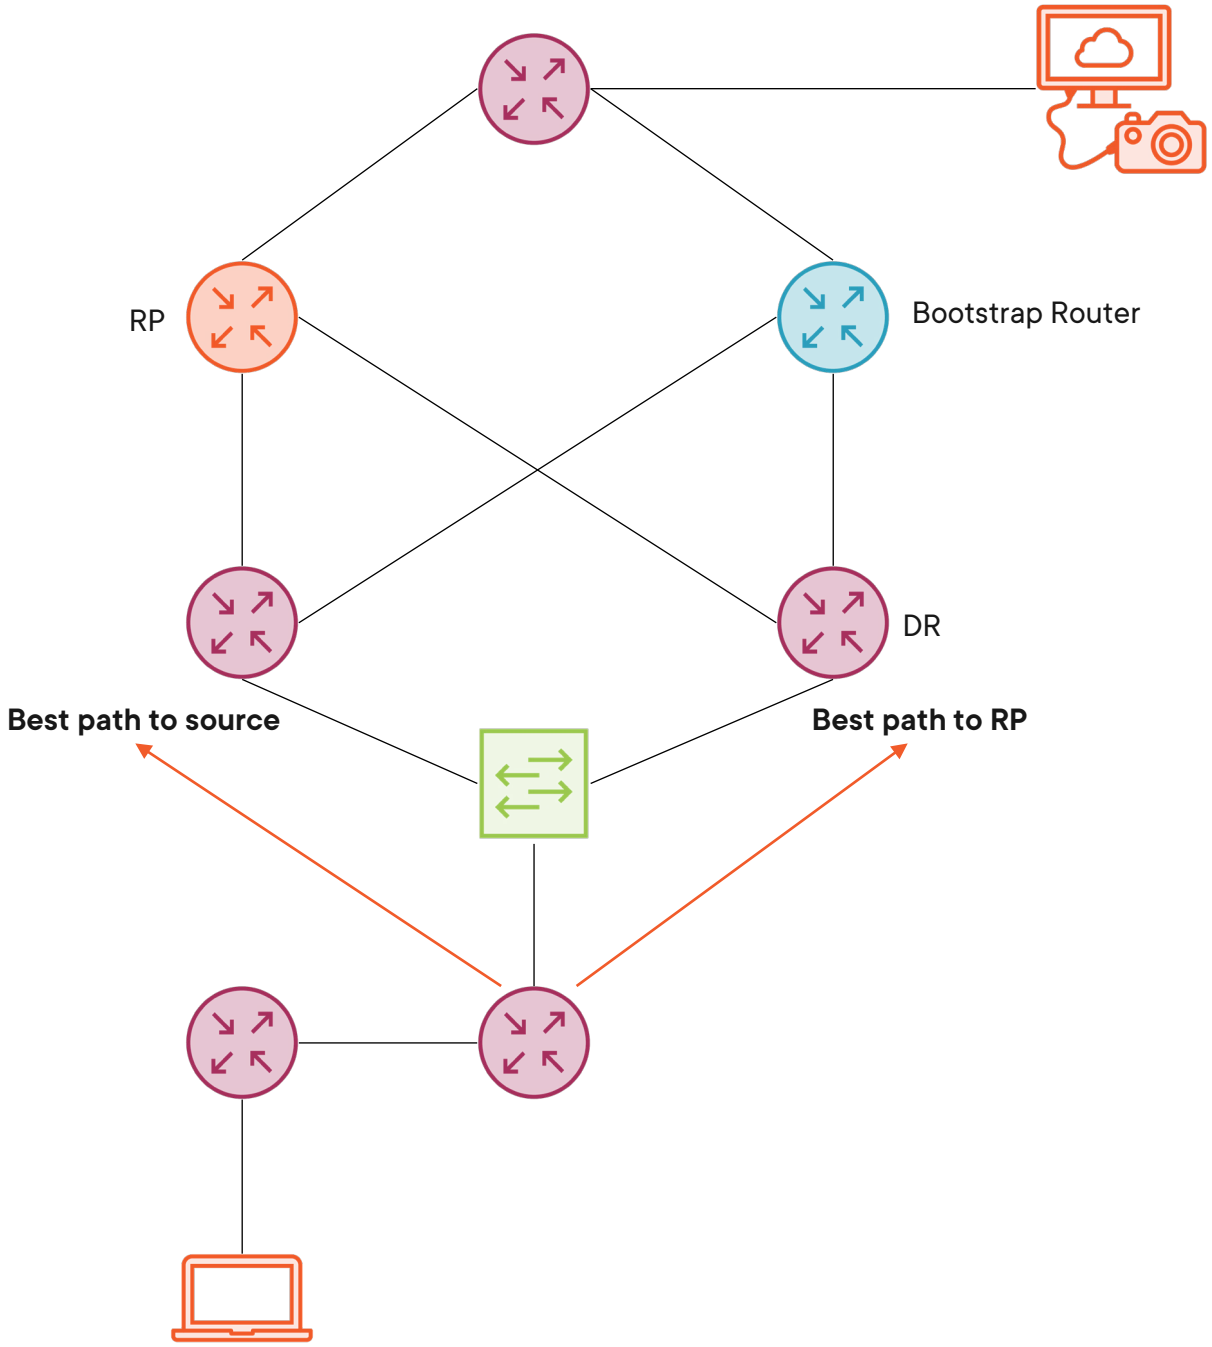

PIM Assert

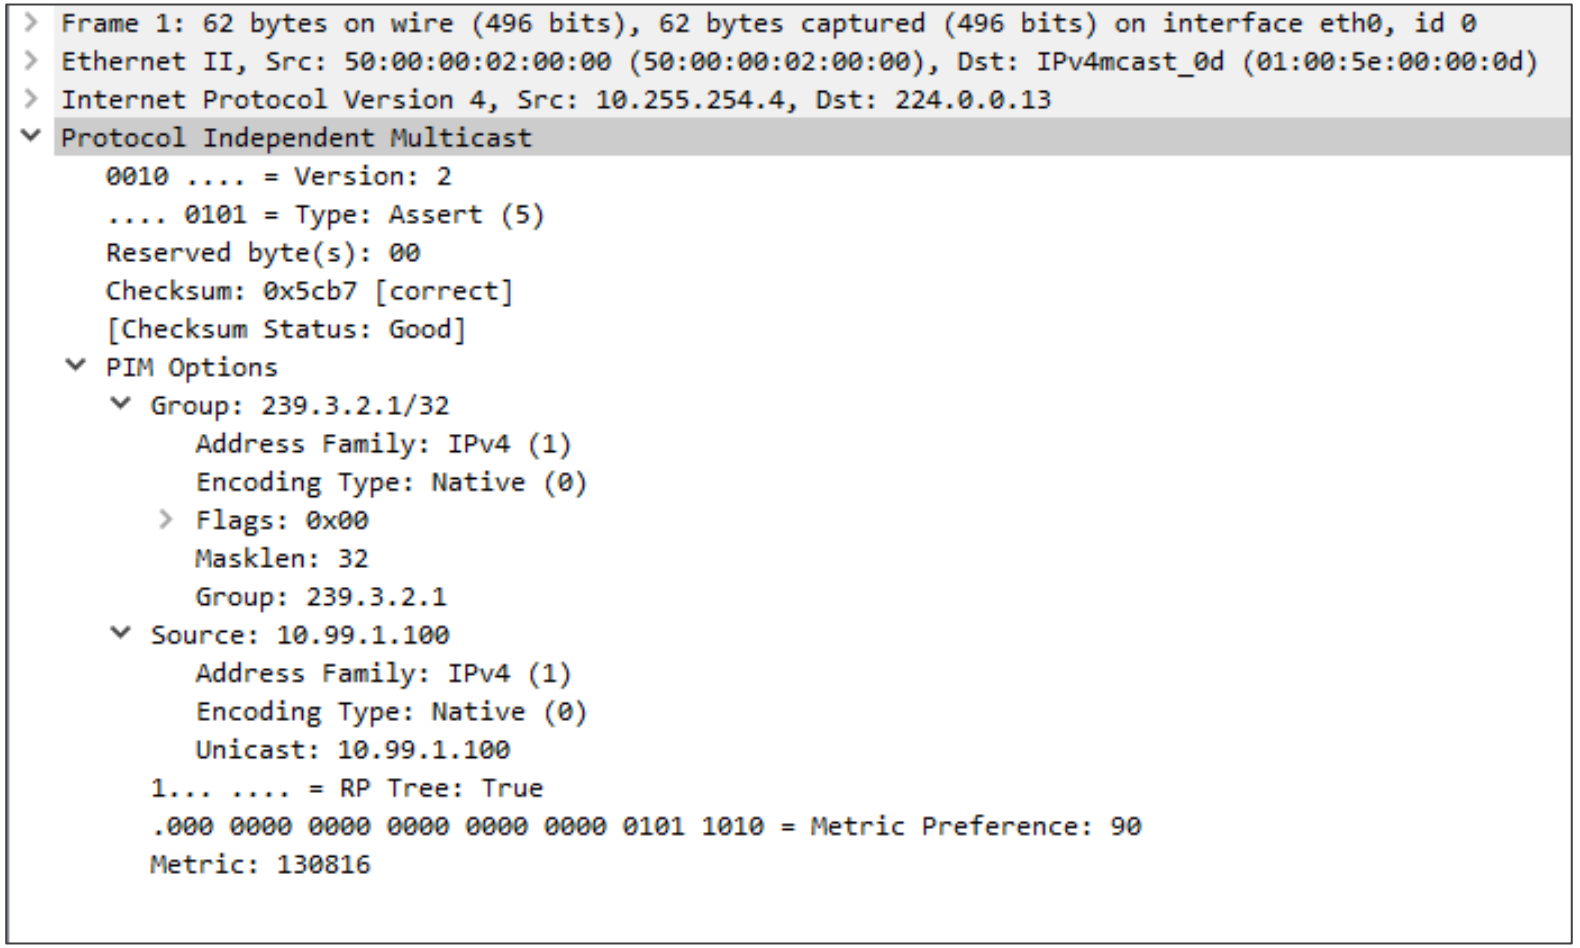
Frame 1: 62 bytes on wire (496 bits), 62 bytes captured (496 bits) on interface eth0, id 0
Ethernet II, Src: 50:00:00:02:00:00 (50:00:00:02:00:00), Dst: IPv4mcast_0d (01:00:5e:00:00:0d)
Internet Protocol Version 4, Src: 10.255.254.4, Dst: 224.0.0.13
Protocol Independent Multicast
0010 .... = Version: 2
.... 0101 = Type: Assert (5)
Reserved byte(s): 00
Checksum: 0x5cb7 [correct]
[Checksum Status: Good]
PIM Options
Group: 239.3.2.1/32
Address Family: IPv4 (1)
Encoding Type: Native (0)
Flags: 0x00
Masklen: 32
Group: 239.3.2.1
Source: 10.99.1.100
Address Family: IPv4 (1)
Encoding Type: Native (0)
Unicast: 10.99.1.100
1... .... = RP Tree: True
.000 0000 0000 0000 0000 0000 0101 1010 = Metric Preference: 90
Metric: 130816

PIM Assert

```
> Frame 2: 62 bytes on wire (496 bits), 62 bytes captured (496 bits) on interface eth0, id 0
> Ethernet II, Src: 50:00:00:01:00:00 (50:00:00:01:00:00), Dst: IPv4mcast_0d (01:00:5e:00:00:0d)
> Internet Protocol Version 4, Src: 10.255.254.3, Dst: 224.0.0.13
v Protocol Independent Multicast
    0010 .... = Version: 2
    .... 0101 = Type: Assert (5)
    Reserved byte(s): 00
    Checksum: 0xceb9 [correct]
    [Checksum Status: Good]
  v PIM Options
    v Group: 239.3.2.1/32
        Address Family: IPv4 (1)
        Encoding Type: Native (0)
      > Flags: 0x00
        Masklen: 32
        Group: 239.3.2.1
    v Source: 10.99.1.100
        Address Family: IPv4 (1)
        Encoding Type: Native (0)
        Unicast: 10.99.1.100
      0... .... = RP Tree: False
      .000 0000 0000 0000 0000 0000 0101 1010 = Metric Preference: 90
      Metric: 3328
```

Quick Detour: Sparse-Dense Mode

PIM Sparse-Dense Mode

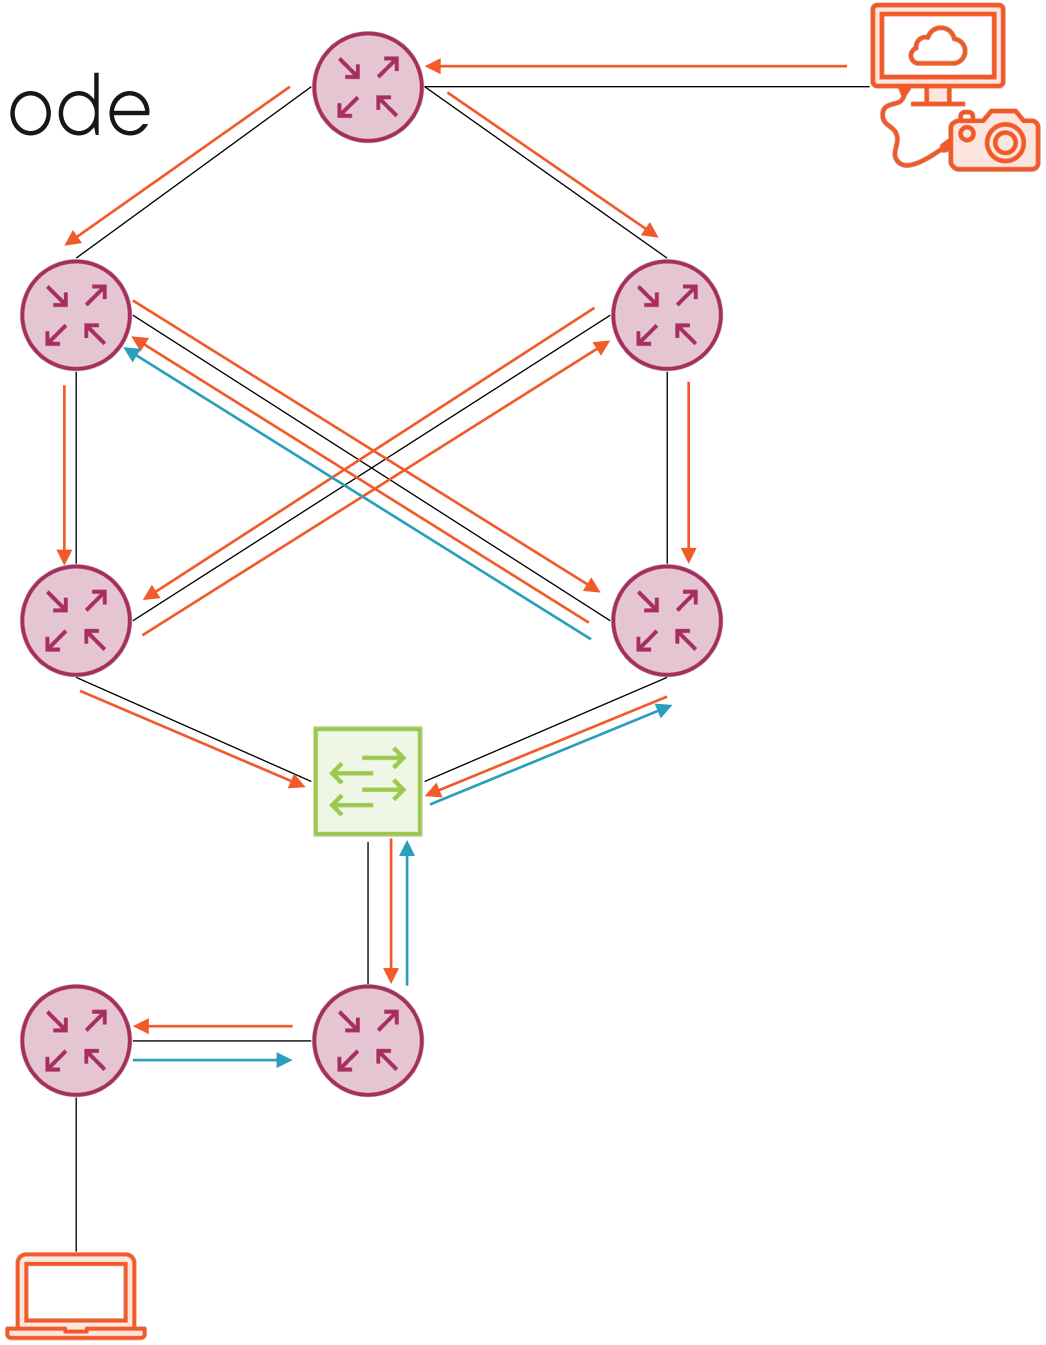

AutoRP Basics

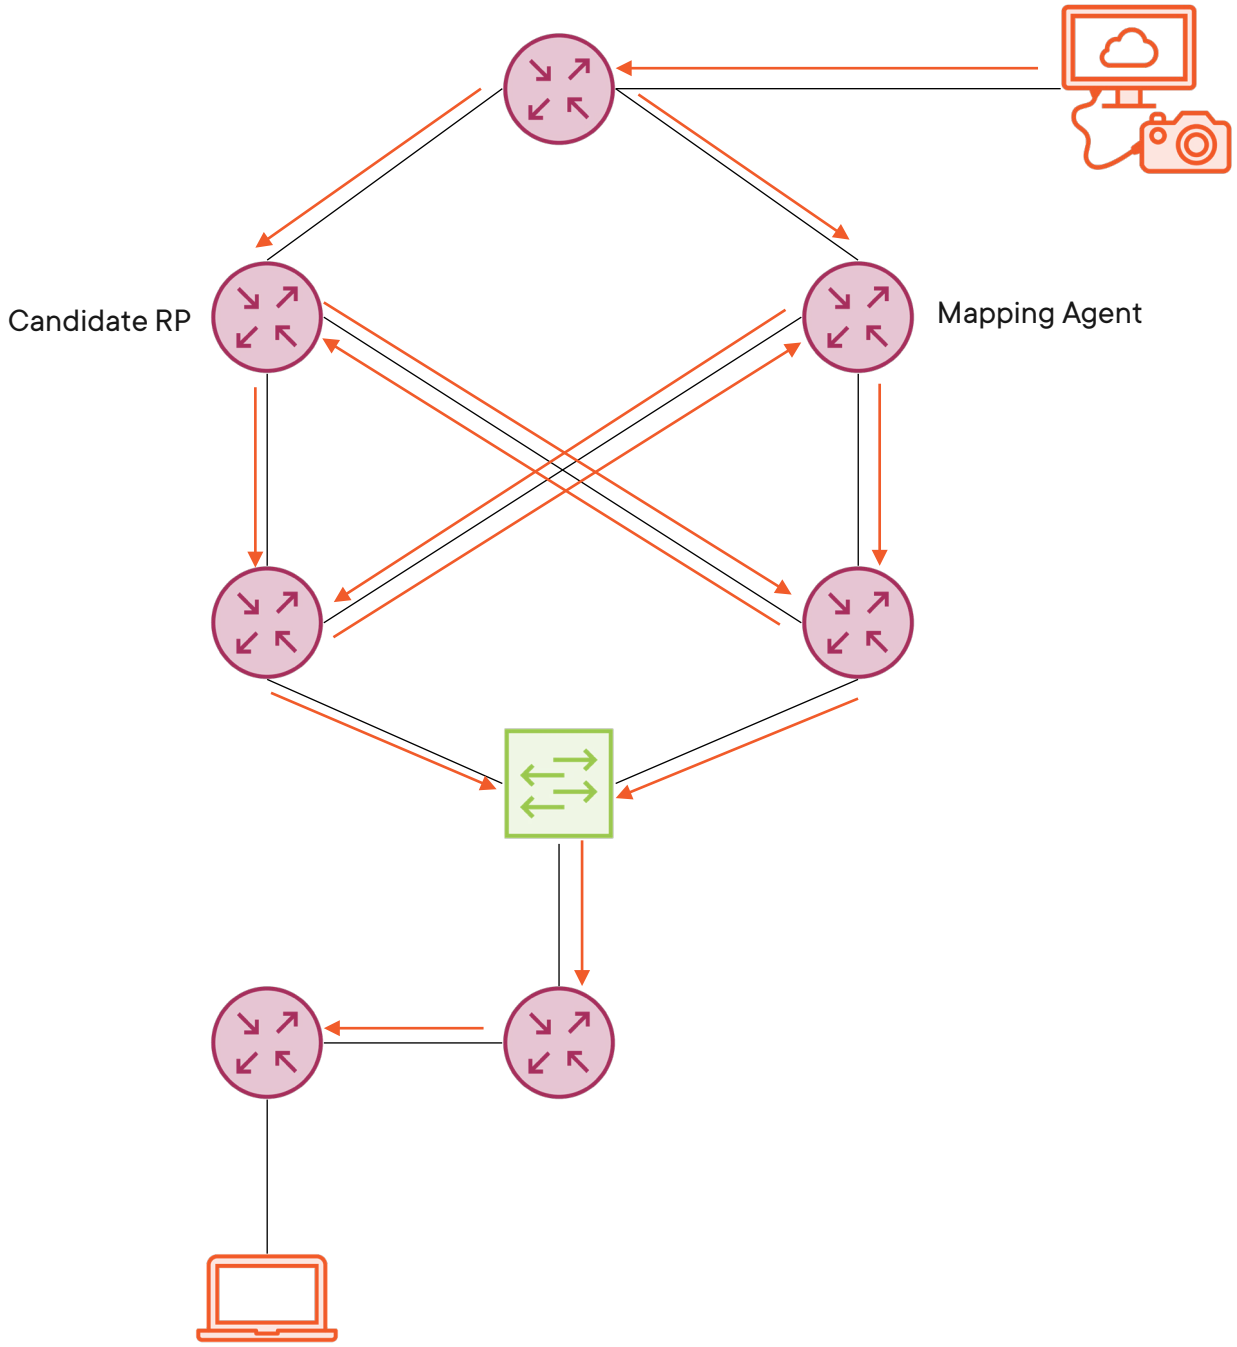

AutoRP Selection

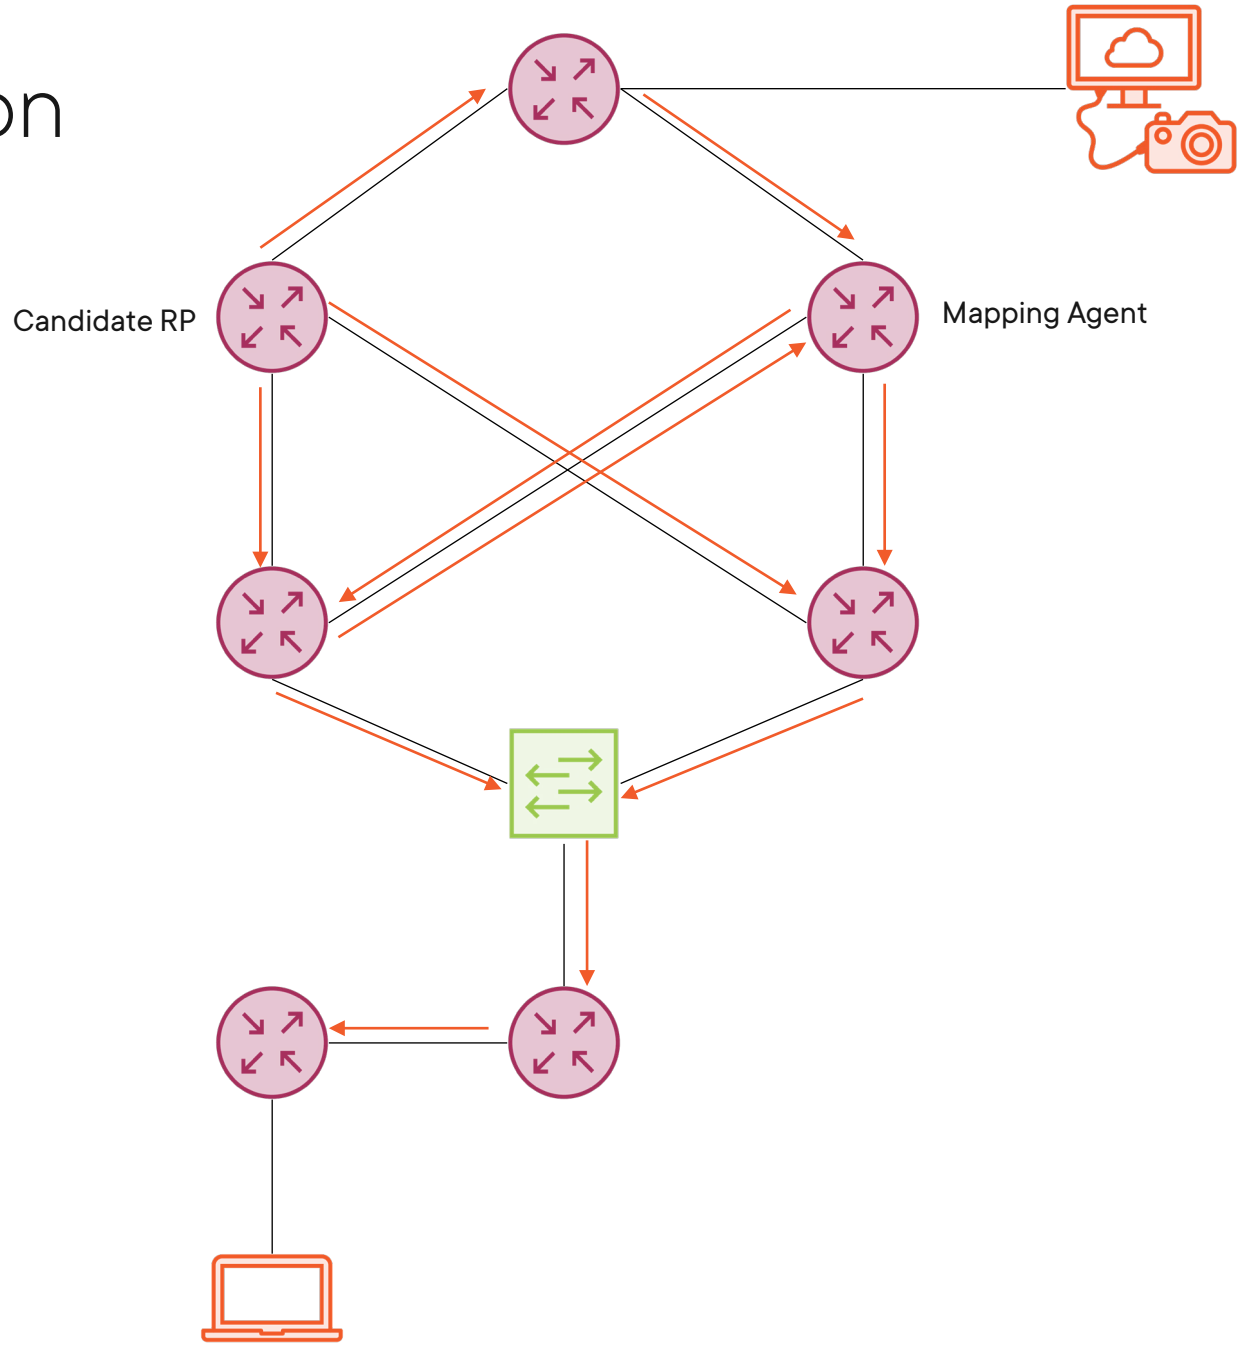

AutoRP Advertisement

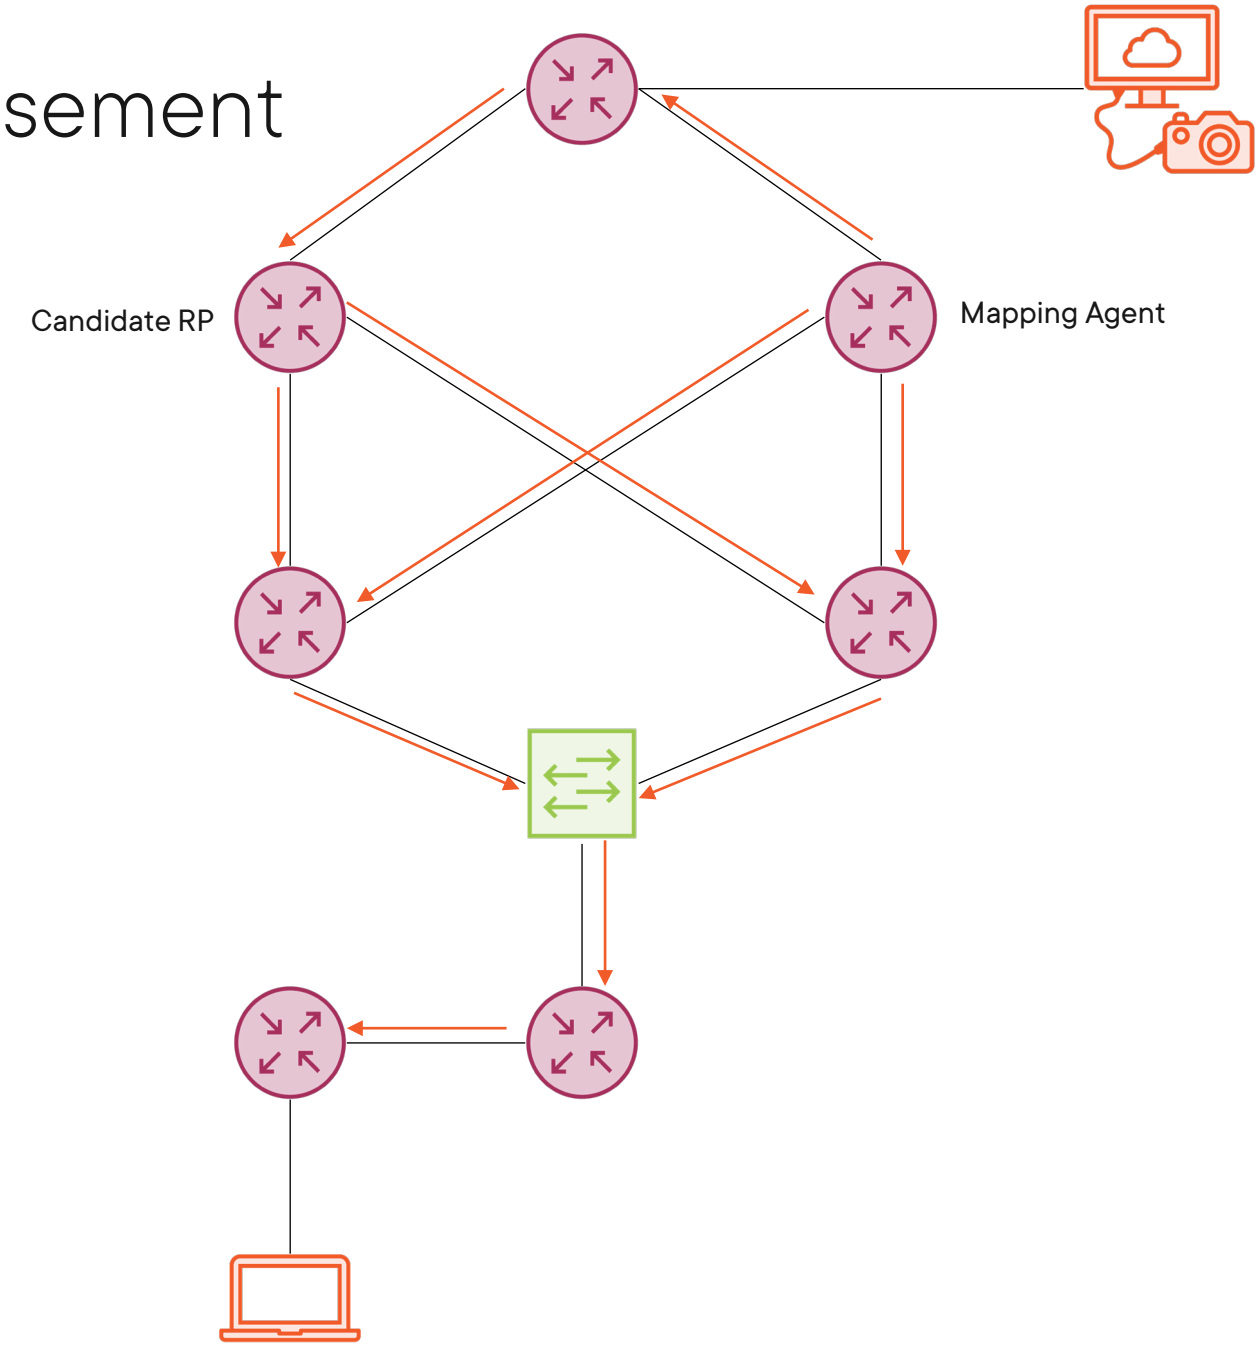

PIM-SDM Problems

Sparse-Dense Mode is inefficient

Dense Mode can cause routing loops and can act unpredictably in PIM-SDM

AutoRP Optimizations obsolete the need for SDM

AutoRP Listener functionality works fine with Sparse Mode

Summary

Topics:

- Basic Operation of PIM-SM

Demos:

- Rendezvous Point Election with Bootstrap Router
- Generate Multicast Traffic in a Lab
- PIM Join Process
- PIM Assert Process

Packet Analysis:

- PIM Candidate RP Advertisement/PIM Bootstrap
- PIM Register/Register Stop
- Shortest Path Switchover Join/Prune
- PIM Assert