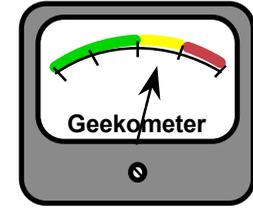


PIM Sparse Mode

Module Objectives

- **Identify the packet types used by PIM.**
- **Identify and explain the basic mechanisms of PIM Sparse Mode.**

Agenda



Cisco.com

- **PIM Packet Formats**
- **PIM Neighbor Discovery**
- **PIM State**
- **PIM SM Joining**
- **PIM SM Registering**
- **PIM SM SPT-Switchover**
- **PIM SM Pruning**
- **PIM Asserts**

PIM Packet Formats

- **PIM Packet Headers**
- **PIM Hello Messages**
- **PIM Join/Prunes**
- **PIM Grafts/Graft Acks**
- **PIM Asserts**
- **PIM Registers**
- **PIM Register-Stop**

PIM Packet Header

PIMv2 is assigned protocol number 103



Ver:

PIM Version = 2

Type:

0 = Hello

1 = Register (SM only)

2 = Register-Stop (SM only)

3 = Join/Prune

4 = Bootstrap (SM BSR only)

5 = Assert

6 = Graft (DM only)

7 = Graft-Ack (DM only)

8 = C-RP-Announcement (SM BSR only)

- All PIM messages (except Registers and Register-Stops) are multicast to the ALL-PIM-Routers (224.0.0.13) group with a TTL of 1.

PIM Hello Messages

3	7	15	31
Ver.	Type	Reserved	Checksum
Option Type		Option Length	
Option Value			
...			
Option Type		Option Length	
Option Value			

Option Types:

- 1 = Holdtime (Period of time in seconds before this PIM neighbor times out.)**
- 18 = DR Priority**
- 19 = Bidir PIM Support**
- 20 = Generation ID Support**

PIM Join/Prune Packets

3	7	15	31
Ver.	Type	Reserved	Checksum
Upstream Neighbor Address (Encoded-Unicast)			
Reserved	Num. Groups	Holdtime	
Group List			

Upstream Neighbor Address¹:
IP address of upstream neighbor

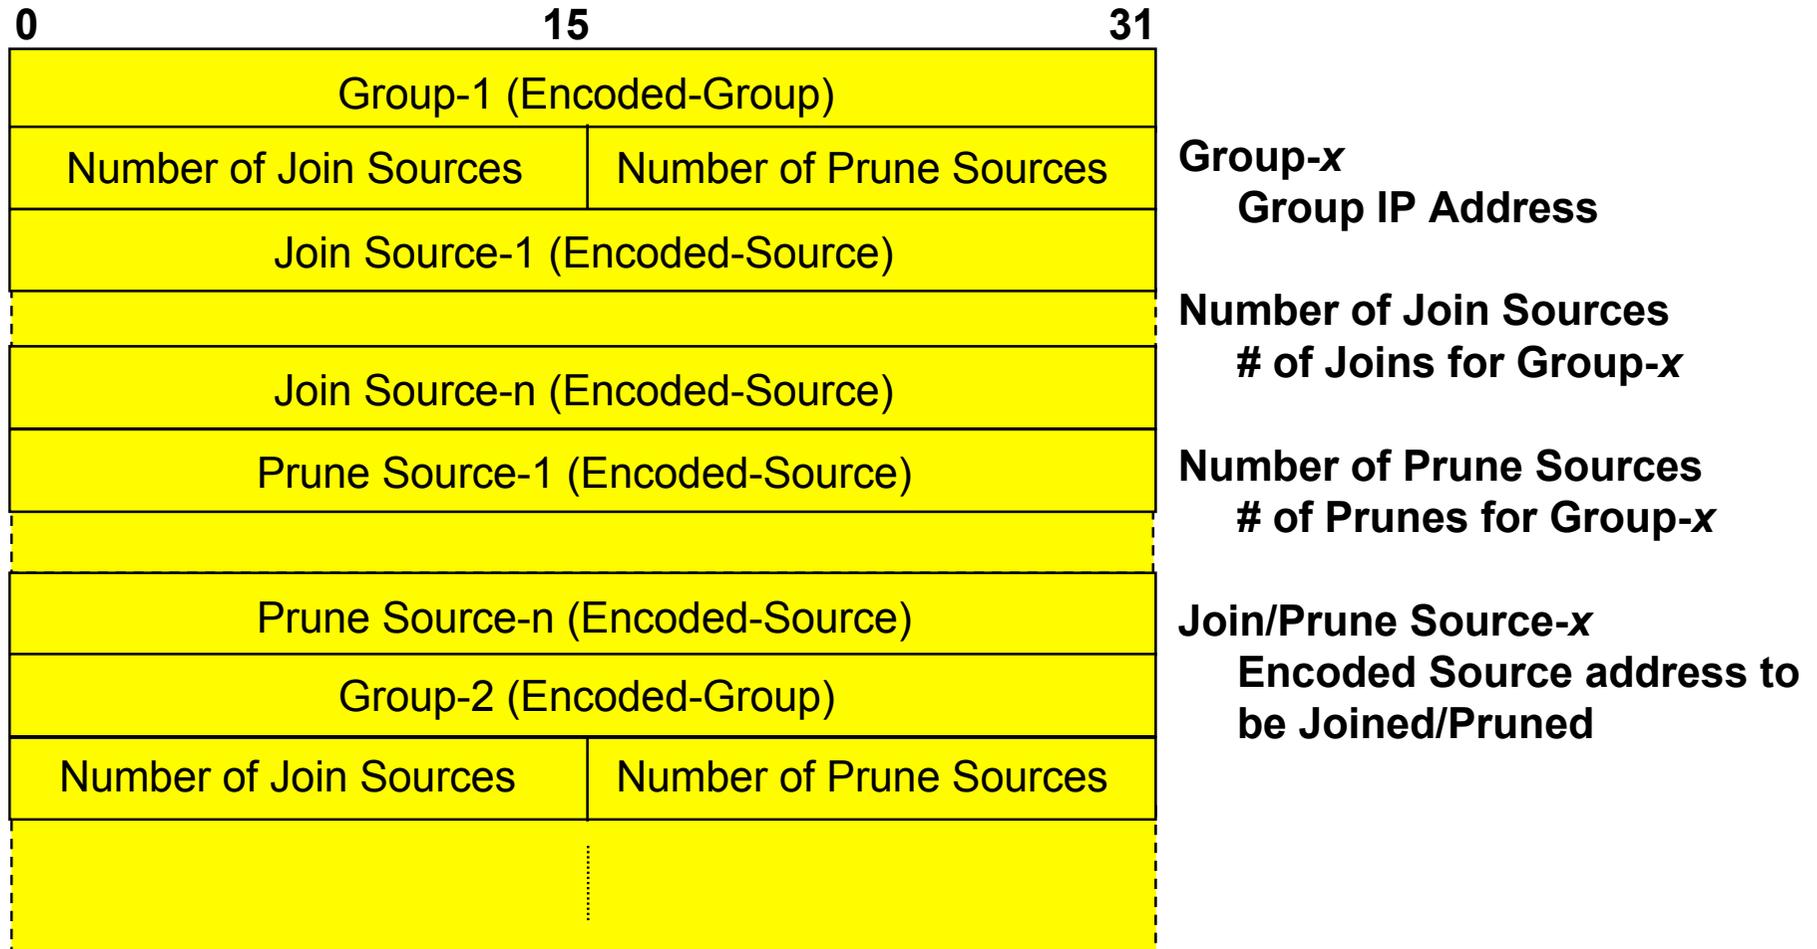
Num. Groups
of Groups in Group list

Holdtime:
Period of time in seconds before this join/prune times out.

Group List:
List (by group) of sources to Join and/or Prune

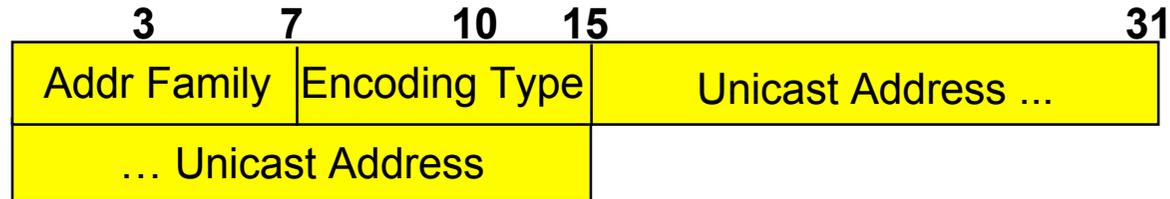
1. All PIM control packets are multicast on the wire to 224.0.0.13, thus all PIM routers on the wire receive and examine the Join/Prune packet. The Upstream Neighbor Address field in the Join/Prune packet specifies to which PIM router the Join/Prune packet is destined.

PIM Group Lists



- **Group Lists are used in Join/Prune and Graft/Graft-Ack messages.**

Encoded Unicast Addresses



Addr Family:

IANA Address Family Identifier (1=IPv4)

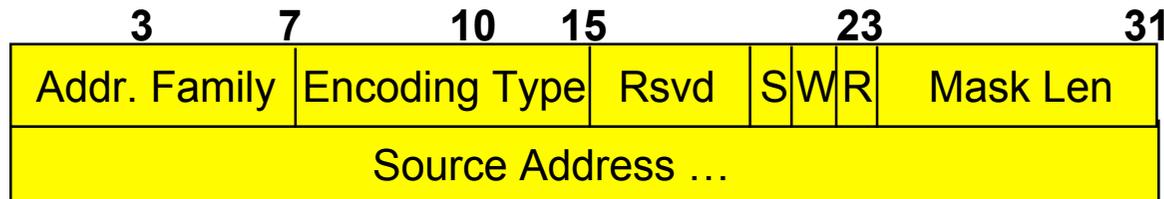
Encoding Type:

Type of encoding within Address Family

Unicast Address :

Unicast Address of the target device.

Encoded Source Addresses



Addr Family:

IANA Address Family Identifier (1=IPv4)

Encoding Type:

Type of encoding within Address Family

S = Sparse Mode bit :

Indicates sparse mode group.

W = Wildcard bit :

Indicates join/prune applies to (*, G) entry.

R = RP bit :

Indicates this join/prune should be sent up the Shared Tree towards the RP.

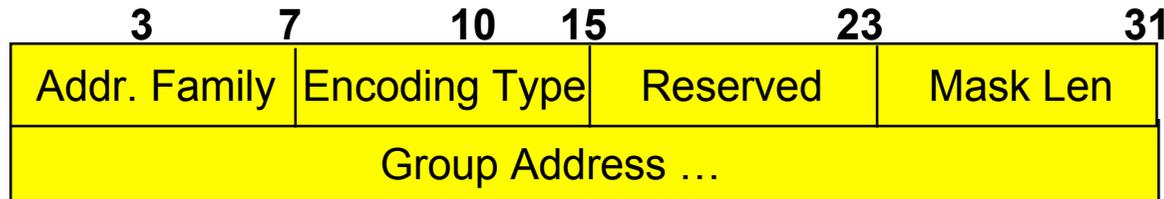
Mask Length:

Number of bits in the prefix of the Group Address.

Source Address :

Address of Multicast Source (or the RP Address if R bit is set).

Encoded Group Addresses



Addr Family:

IANA Address Family Identifier (1=IPv4)

Encoding Type:

Type of encoding within Address Family

Mask Length:

Number of bits in the prefix of the Group Address.

Group Address :

Multicast Group Address.

PIM Graft/Graft-Ack Packets (DM-only)

Cisco.com

3	7	15	31
Ver.	Type	Reserved	Checksum
Upstream Neighbor Address (Encoded-Unicast)			
Reserved	Num. Groups	Holdtime	
Group List			

Upstream Neighbor Address¹:
IP address of upstream neighbor

Num. Groups
of Groups in Group list

Holdtime:
Period of time in seconds before this join/prune times out.

Group List:
List (by group) of sources to Graft or Graft-Ack

1. All PIM control packets are multicast on the wire to 224.0.0.13, thus all PIM routers on the wire receive and examine Graft/Graft-Ack packets. The Upstream Neighbor Address field in the Graft/Graft-Ack packet specifies to which PIM router the Join/Prune packet is destined.

PIM Assert Packets

3	7	15	31
Ver.	Type	Reserved	Checksum
Group Address (Encoded-Group)			
Source Address (Encoded-Source)			
R	Metric Preference		
Metric			

Group Address:

Identifies Group of the Assert

Source Address:

Identifies Source of the Assert

R: (Sparse Mode)

1 = Assert down RP Tree; 0 = Assert Down SPT

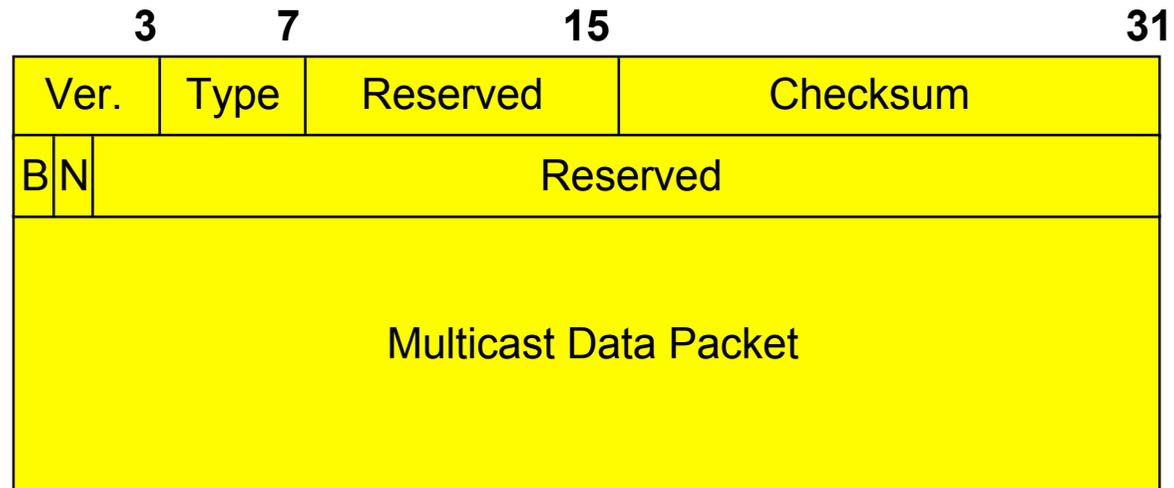
Metric Preference:

Admin. Distance of unicast routing protocol

Metric:

Unicast routing protocol metric

PIM Register Packets



B = Border Bit:

Indicates DR is a border router performing a proxy-register

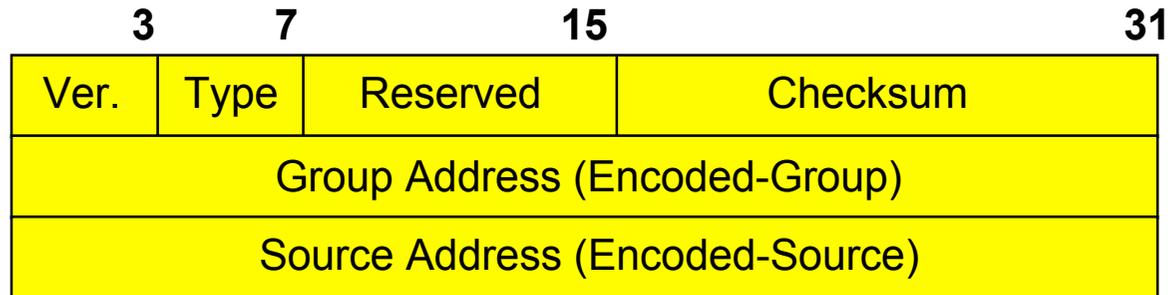
N = Null Register Bit:

Indicates DR is sending a Null-Register before expiring its register-suppression timer.

Multicast Data Packet:

The original packet sent by the source. For periodic sending of registers, this part is null.

PIM Register-Stop Packets



Group Address:

The group address from the register message.

Source Address:

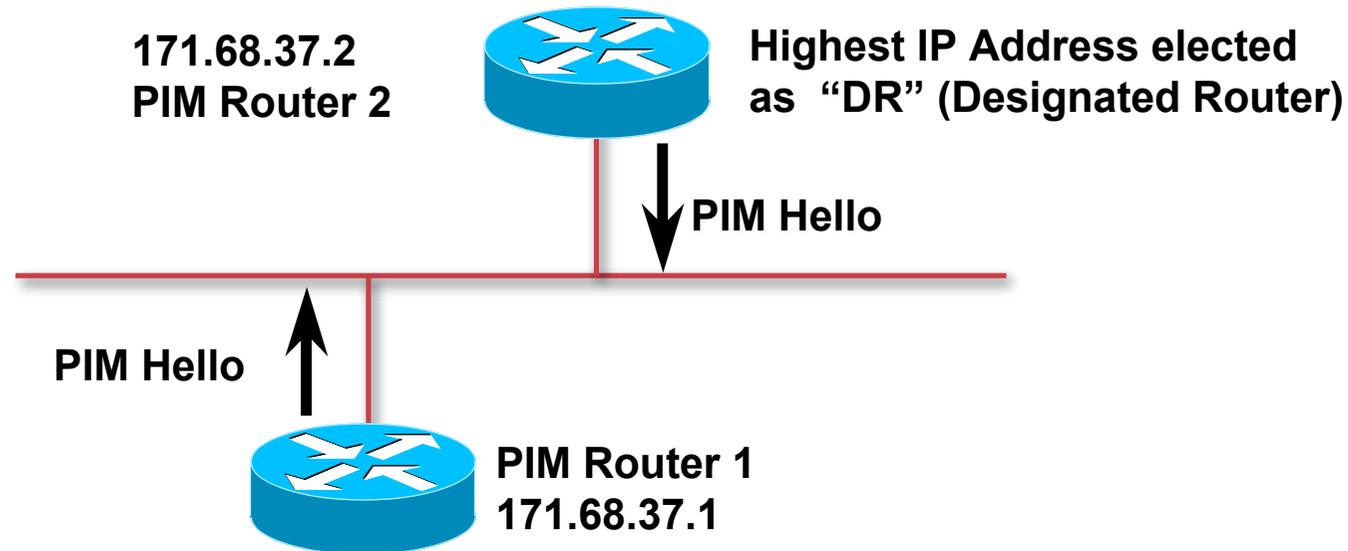
IP host address of source from multicast data packet in register.

Agenda

- PIM Packet Formats
- **PIM Neighbor Discovery**
- PIM State
- PIM SM Joining
- PIM SM Registering
- PIM SM SPT-Switchover
- PIM SM Pruning
- PIM Asserts

PIM Neighbor Discovery

Cisco.com



- PIMv2 Hellos are periodically multicast to the "All-PIM-Routers" (224.0.0.13) group address. (Default = 30 seconds)
- If the "DR" times-out, a new "DR" is elected.
- The "DR" is responsible for sending all Joins and Register messages for any receivers or senders on the network.

PIM Neighbor Discovery

```
wan-gw8>show ip pim neighbor
```

```
PIM Neighbor Table
```

Neighbor Address	Interface	Uptime/Expires	Ver	Mode Prio/Mode
171.68.0.70	FastEthernet0/0	2w1d/00:01:24	v2	1 / B S
171.68.0.91	FastEthernet0/0	2w6d/00:01:01	v2	1 / B S
171.68.0.82	FastEthernet0/0	7w0d/00:01:14	v2	5 / DR B S
171.68.0.86	FastEthernet0/0	7w0d/00:01:13	v2	1 / B S
171.68.0.80	FastEthernet0/0	7w0d/00:01:02	v2	1 / B S
171.68.28.70	Serial2.31	22:47:11/00:01:16	v2	1 / B S
171.68.28.50	Serial2.33	22:47:22/00:01:08	v2	1 / B S
171.68.27.74	Serial2.36	22:47:07/00:01:21	v2	N /
171.68.28.170	Serial0.70	1d4h/00:01:06	v2	N /
171.68.27.2	Serial1.51	1w4d/00:01:25	v2	1 / B S
171.68.28.110	Serial3.56	1d4h/00:01:20	v2	1 / B S
171.68.28.58	Serial3.102	12:53:25/00:01:03	v2	1 / B S

Agenda

- PIM Packet Formats
- PIM Neighbor Discovery
- **PIM State**
- PIM SM Joining
- PIM SM Registering
- PIM SM SPT-Switchover
- PIM SM Pruning
- PIM Asserts

PIM State

- **Describes the “state” of the multicast distribution trees as understood by the router at this point in the network.**
- **Represented by entries in the multicast routing (mroute) table**
 - **Used to make multicast traffic forwarding decisions**
 - **Composed of (*, G) and (S, G) entries**
 - **Each entry contains RPF information**
 - **Incoming (i.e. RPF) interface**
 - **RPF Neighbor (upstream)**
 - **Each entry contains an Outgoing Interface List (OIL)**
 - **OIL may be NULL**

PIM-SM State Example

```
sj-mbone> show ip mroute
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
      L - Local, P - Pruned, R - RP-bit set, F - Register flag,
      T - SPT-bit set, J - Join SPT, M - MSDP created entry,
      X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
      U - URD, I - Received Source Specific Host Report
Outgoing interface flags: H - Hardware switched
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.1.1.1), 2w1d/00:00:00, RP 172.16.25.1, flags: SJC
  Incoming interface: Serial0/1, RPF nbr 172.16.4.1
  Outgoing interface list:
    Ethernet0/1, Forward/Sparse-Dense, 2w1d/00:01:40
    Serial0/0, Forward/Sparse-Dense, 00:4:52/00:02:08

(172.16.8.2, 224.1.1.1), 00:00:10/00:02:59, flags: CJT
  Incoming interface: Serial0/1, RPF nbr 172.16.4.1
  Outgoing interface list:
    Ethernet0/1, Forward/Sparse-Dense, 00:00:10/00:02:49
    Serial0/0, Forward/Sparse-Dense, 00:4:52/00:02:08
```

PIM-SM (*,G) State Rules

- **(*,G) creation**
 - Receipt of a (*,G) Join or IGMP Report
 - Automatically if (S,G) must be created
- **(*,G) reflects default group forwarding**
 - IIF = RPF interface toward RP
 - OIL = interfaces
 - that received a (*,G) Join or
 - with directly connected members or
 - manually configured
- **(*,G) deletion**
 - When OIL = NULL and
 - no child (S,G) state exists

PIM-SM (S,G) State Rules

- **(S,G) creation**
 - By receipt of (S,G) Join or Prune or
 - By “Register” process
 - Parent (*,G) created (if doesn't exist)
- **(S,G) reflects forwarding of “S” to “G”**
 - IIF = RPF Interface normally toward source
 - RPF toward RP if “RP-bit” set
 - OIL = Initially, copy of (*,G) OIL minus IIF
- **(S,G) deletion**
 - By normal (S,G) entry timeout

PIM-SM OIL Rules

- **Interfaces in OIL added**
 - **By receipt of Join message**
 - **Interfaces added to (*,G) are added to all (S,G)'s**
- **Interfaces in OIL removed**
 - **By receipt of Prune message**
 - **Interfaces removed from (*,G) are removed from all (S,G)'s**
 - **Interface Expire timer counts down to zero**
 - **Timer reset (to 3 min.) by receipt of periodic Join**
or
 - **By IGMP membership report**

PIM-SM Triggered Join/Prune Rules

- **Triggering Join/Prune Messages**
 - **(*,G) Joins are triggered when:**
 - The (*,G) OIL transitions from Null to non-Null
 - **(*,G) Prunes are triggered when:**
 - The (*,G) OIL transitions from non-Null to Null
 - **(S,G) Joins are triggered when:**
 - The (S,G) OIL transitions from Null to non-Null
 - **(S,G) Prunes are triggered when:**
 - The (S,G) OIL transitions from non-Null to Null
 - **(S,G)RP-bit Prunes are triggered when:**
 - The (S,G) RPF info \neq the (*,G) RPF info

PIM-SM State Flags

- **S = Sparse**
- **C = Directly Connected Host**
- **L = Local (Router is member)**
- **P = Pruned (All intfcs in OIL = Prune)**
- **T = Forwarding via SPT**
 - **Indicates at least one packet was forwarded**

PIM-SM State Flags (cont.)

- **J = Join SPT**
 - In (*, G) entry
 - Indicates SPT-Threshold is being exceeded
 - Next (S,G) received will trigger join of SPT
 - In (S, G) entry
 - Indicates SPT joined due to SPT-Threshold
 - If rate < SPT-Threshold, switch back to Shared Tree
- **F = Register/First-Hop**
 - In (S,G) entry
 - “S” is a directly connected source
 - Triggers the Register Process
 - In (*, G) entry
 - Set when “F” set in at least one child (S,G)

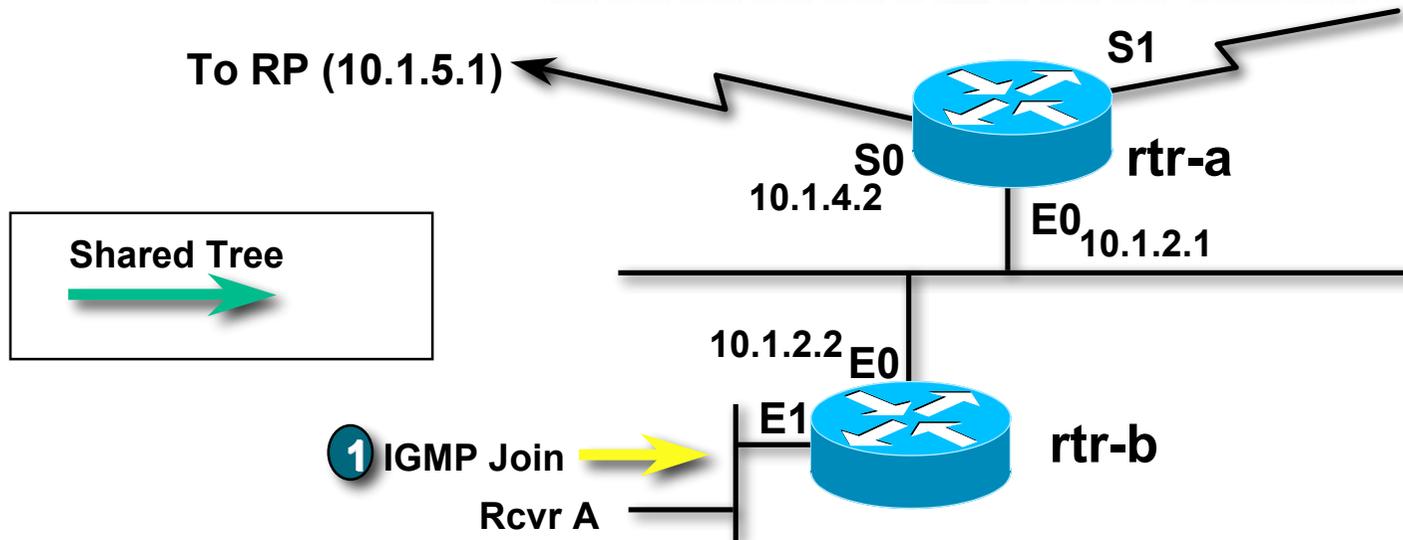
PIM-SM State Flags (cont.)

- **R = RP bit**
 - **(S, G) entries only**
 - **Set by (S,G)RP-bit Prune**
 - **Indicates info is applicable to Shared Tree**
 - **Used to prune (S,G) traffic from Shared Tree**
 - **Initiated by Last-hop router after switch to SPT**
 - **Modifies (S,G) forwarding behavior**
 - **IIF = RPF toward RP (I.e. up the Shared Tree)**
 - **OIL = Pruned accordingly**

Agenda

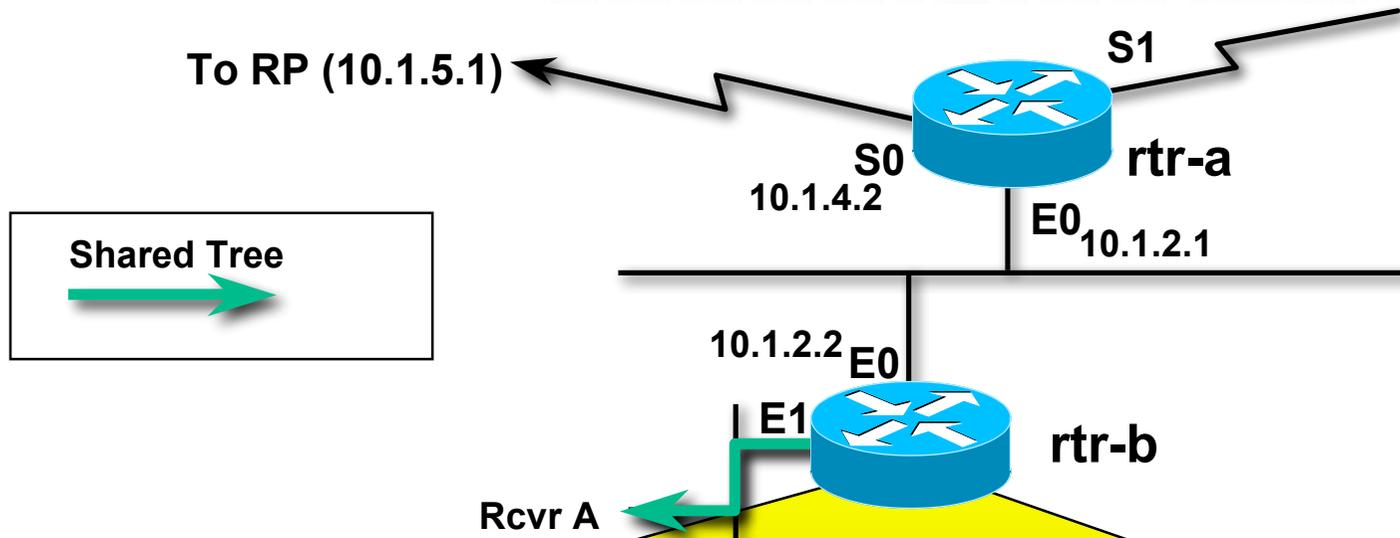
- PIM Packet Formats
- PIM Neighbor Discovery
- PIM State
- **PIM SM Joining**
- PIM SM Registering
- PIM SM SPT-Switchover
- PIM SM Pruning
- PIM Asserts

PIM SM Joining



- 1 “Rcvr A” wishes to receive group G traffic. Sends IGMP Join for G.

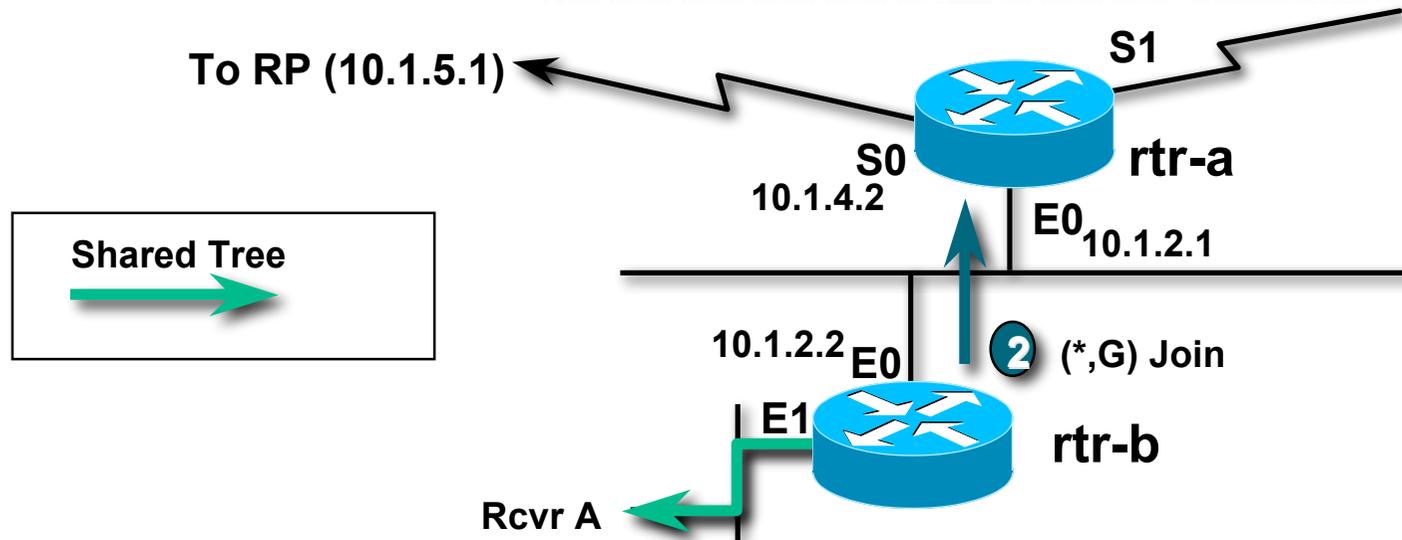
PIM SM Joining



```
(*, 224.1.1.1), 00:00:05/00:00:00, RP 10.1.5.1, flags: SC  
Incoming interface: Ethernet0, RPF nbr 10.1.2.1  
Outgoing interface list:  
Ethernet1, Forward/Sparse-Dense, 00:00:05/00:02:54
```

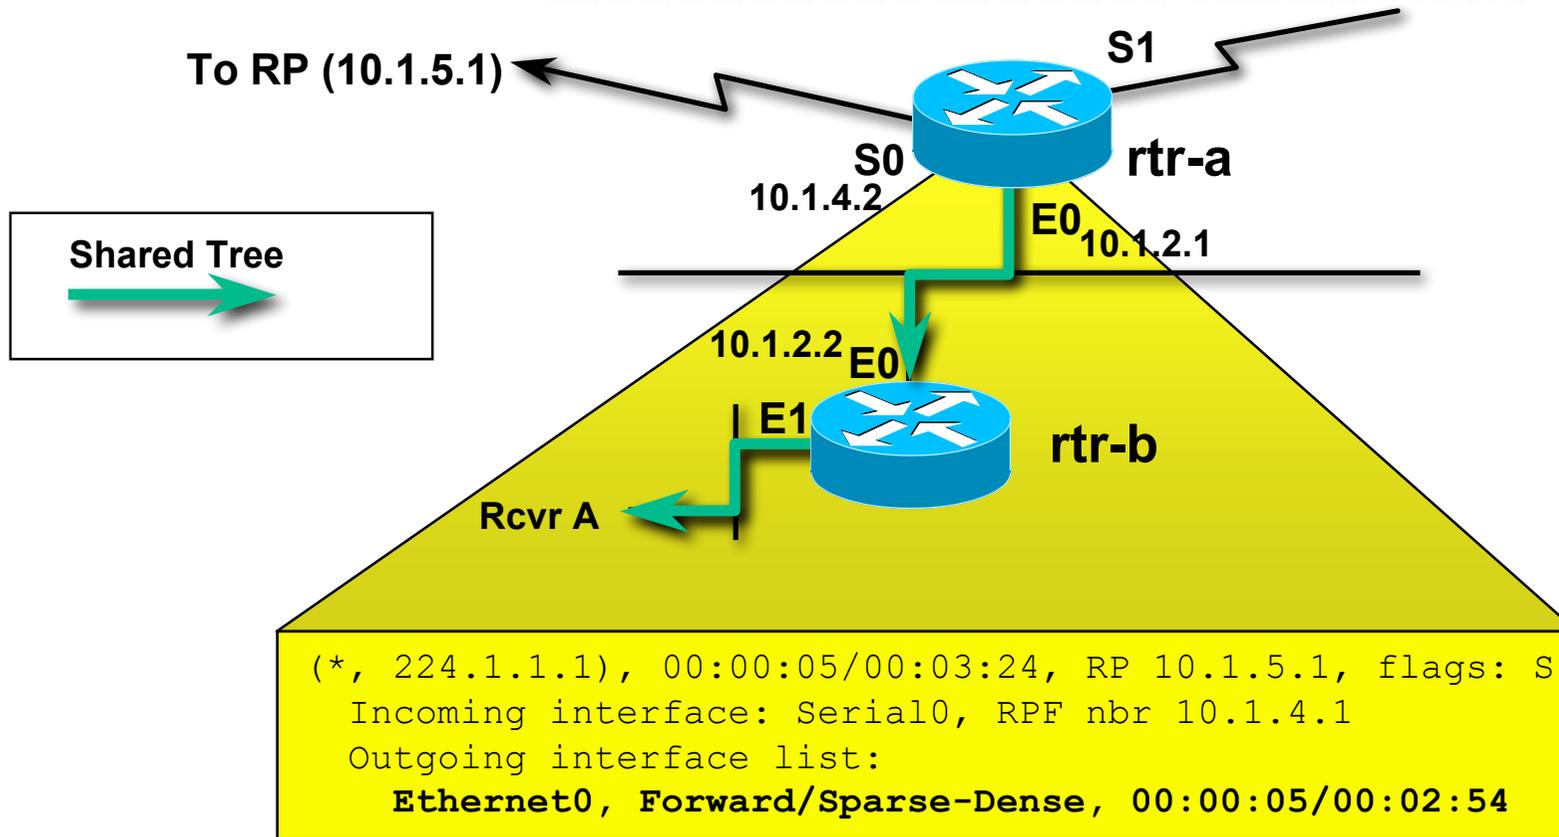
“rtr-b” creates (*, 224.1.1.1) state

PIM SM Joining



- 1 "Rcvr A" wishes to receive group G traffic. Sends IGMP Join for G.
- 2 "rtr-b" sends (*,G) Join towards RP.

PIM SM Joining



“rtr-a” creates (*, 224.1.1.1) state.

Agenda

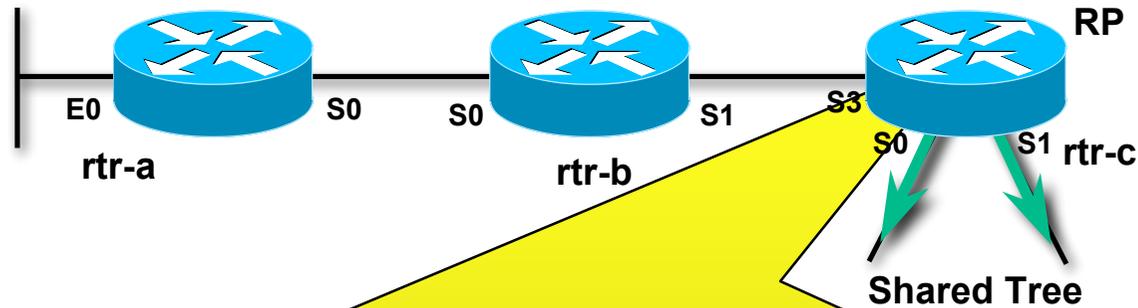
- PIM Packet Formats
- PIM Neighbor Discovery
- PIM State
- PIM SM Joining
- **PIM SM Registering**
- PIM SM SPT-Switchover
- PIM SM Pruning
- PIM Asserts

PIM SM Register Examples

- **Receivers Join Group First**
- **Source Registers First**
- **Receivers along the SPT**

PIM SM Registering

Receiver Joins Group First

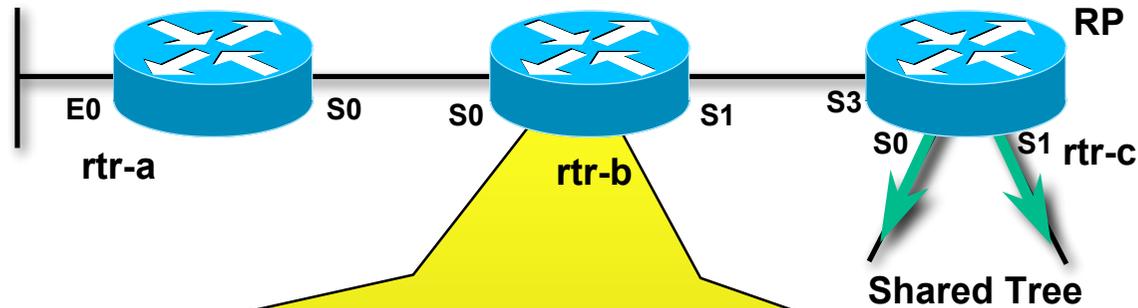


```
(*, 224.1.1.1), 00:03:14/00:02:59, RP 171.68.28.140, flags:S  
Incoming interface: Null, RPF nbr 0.0.0.0,  
Outgoing interface list:  
  Serial0, Forward/Sparse-Dense, 00:03:14/00:03:15  
  Serial1, Forward/Sparse-Dense, 00:03:14/00:03:15
```

**State in “RP” before any source registers
(with receivers on Shared Tree)**

PIM SM Registering

Receiver Joins Group First



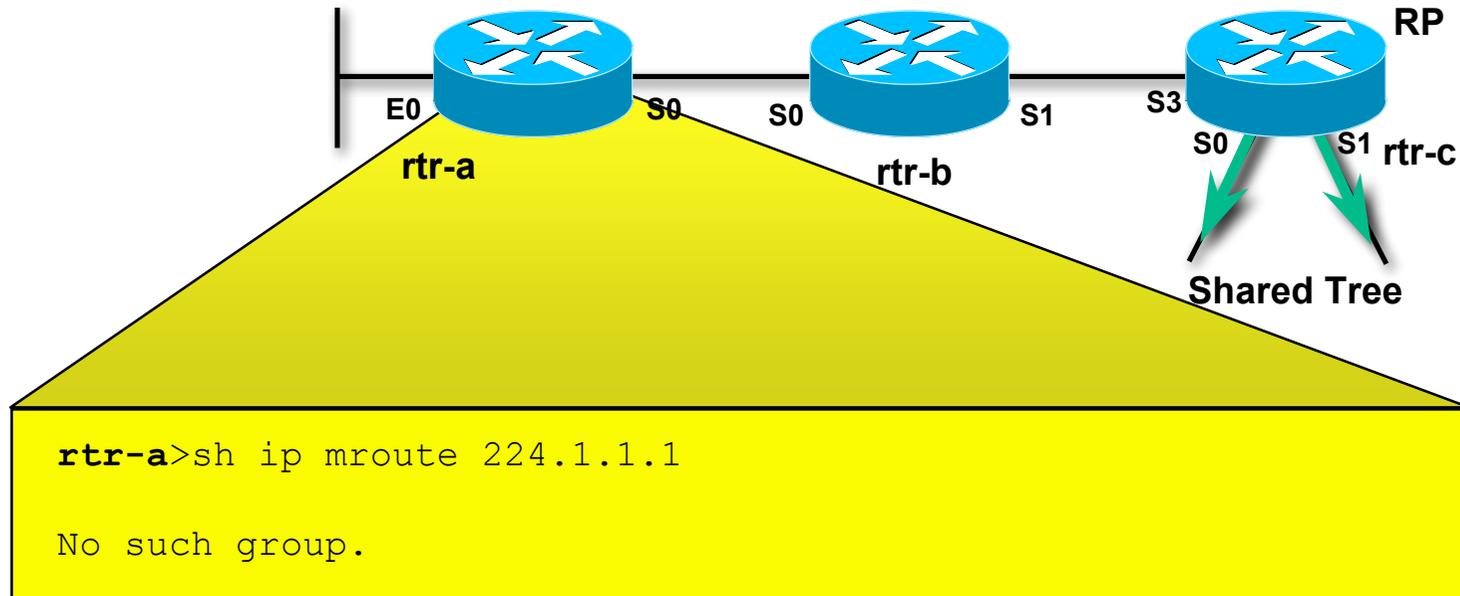
```
rtr-b>sh ip mroute 224.1.1.1
```

```
No such group
```

State in “rtr-b” before any source registers
(with receivers on Shared Tree)

PIM SM Registering

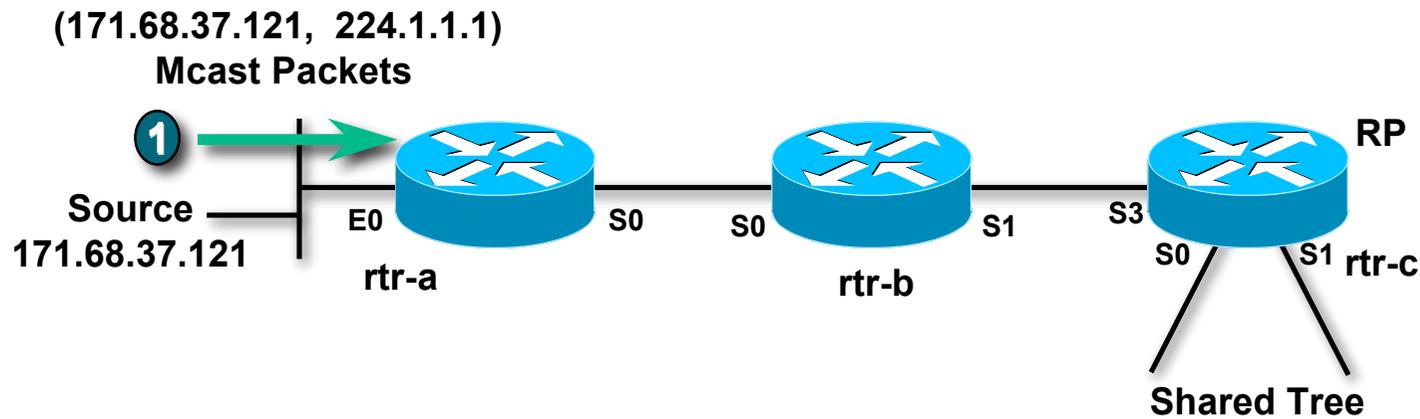
Receiver Joins Group First



**State in “rtr-a” before any source registers
(with receivers on Shared Tree)**

PIM SM Registering

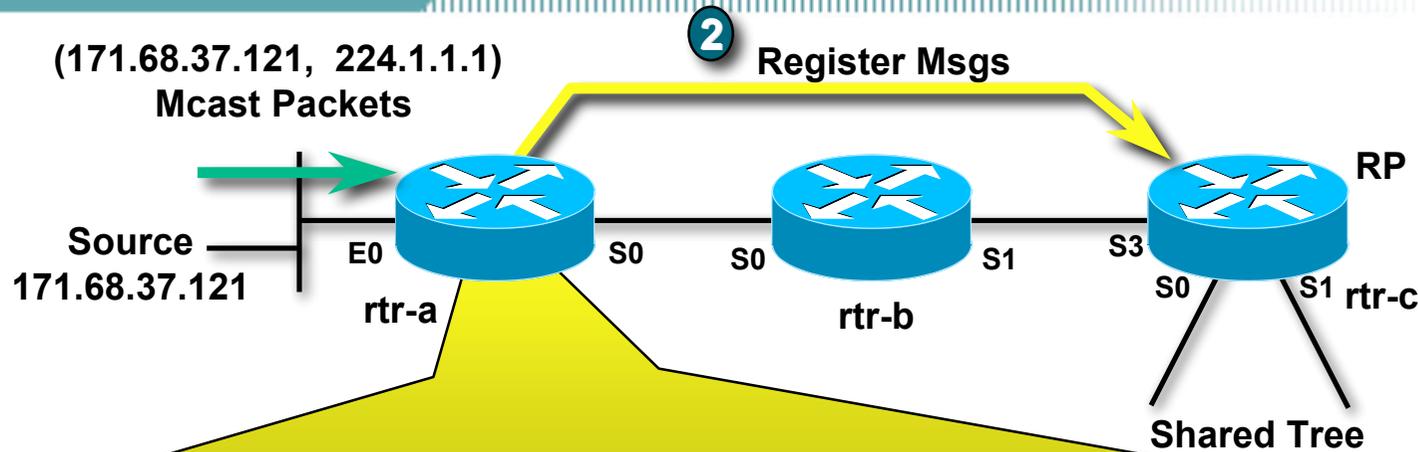
Receiver Joins Group First



- 1 “Source” begins sending group G traffic.

PIM SM Registering

Receiver Joins Group First



```
(*, 224.1.1.1), 00:00:03/00:02:59, RP 171.68.28.140, flags: SP  
Incoming interface: Serial0, RPF nbr 171.68.28.191,  
Outgoing interface list: Null
```

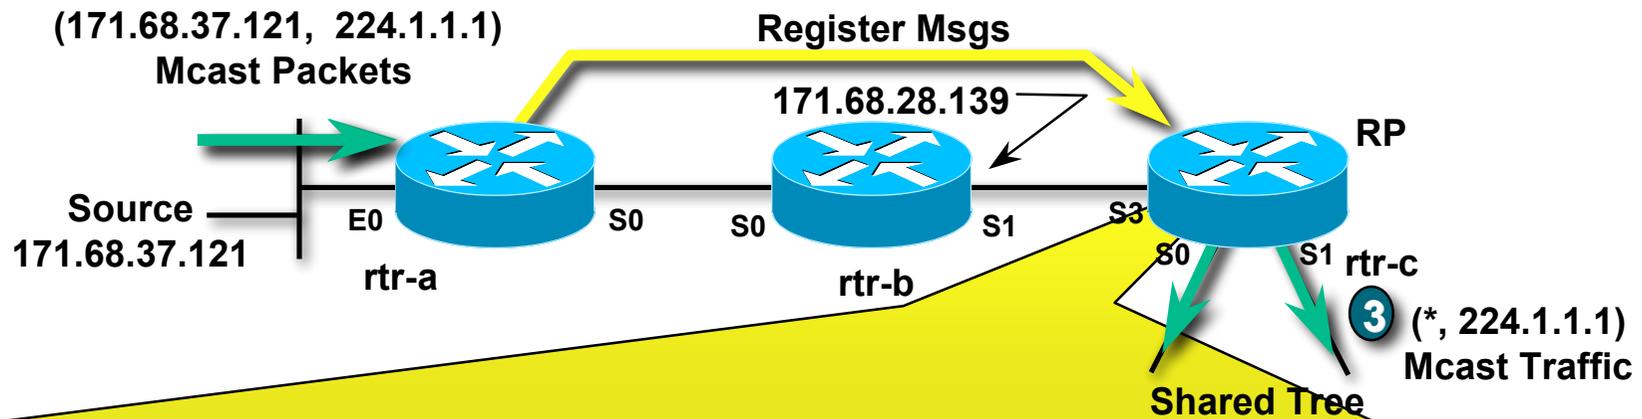
```
(171.68.37.121/32, 224.1.1.1), 00:00:03/00:02:56, flags: FPT  
Incoming interface: Ethernet0, RPF nbr 0.0.0.0, Registering  
Outgoing interface list: Null
```

“rtr-a” creates (S, G) state for source
(After automatically creating a (*, G) entry)

- 1 “Source” begins sending group G traffic.
- 2 “rtr-a” encapsulates packets in Registers; unicasts to RP.

PIM SM Registering

Receiver Joins Group First



```
(*, 224.1.1.1), 00:09:21/00:02:59, RP 171.68.28.140, flags: S
Incoming interface: Null, RPF nbr 0.0.0.0,
Outgoing interface list:
  Serial0, Forward/Sparse-Dense, 00:09:21/00:02:38
  Serial1, Forward/Sparse-Dense, 00:03:14/00:02:46

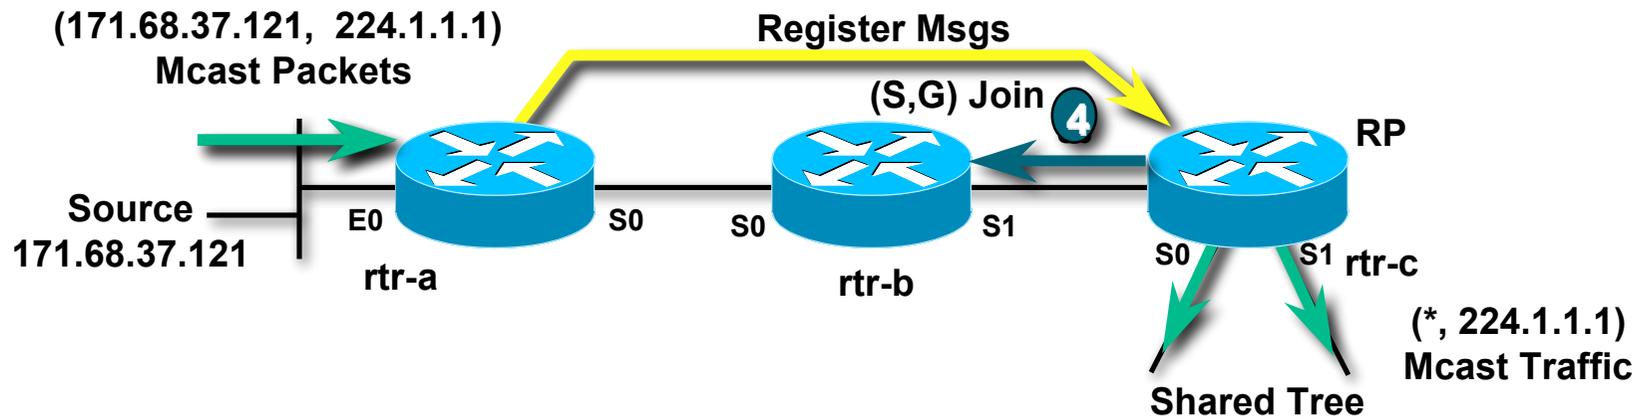
(171.68.37.121, 224.1.1.1, 00:01:15/00:02:46, flags:
Incoming interface: Serial3, RPF nbr 171.68.28.139,
Outgoing interface list:
  Serial0, Forward/Sparse-Dense, 00:00:49/00:02:11
  Serial1, Forward/Sparse-Dense, 00:00:49/00:02:11
```

“RP” processes Register; creates (S, G) state

③ “rtr-c” (RP) de-encapsulates packets; forwards down Shared tree.

PIM SM Registering

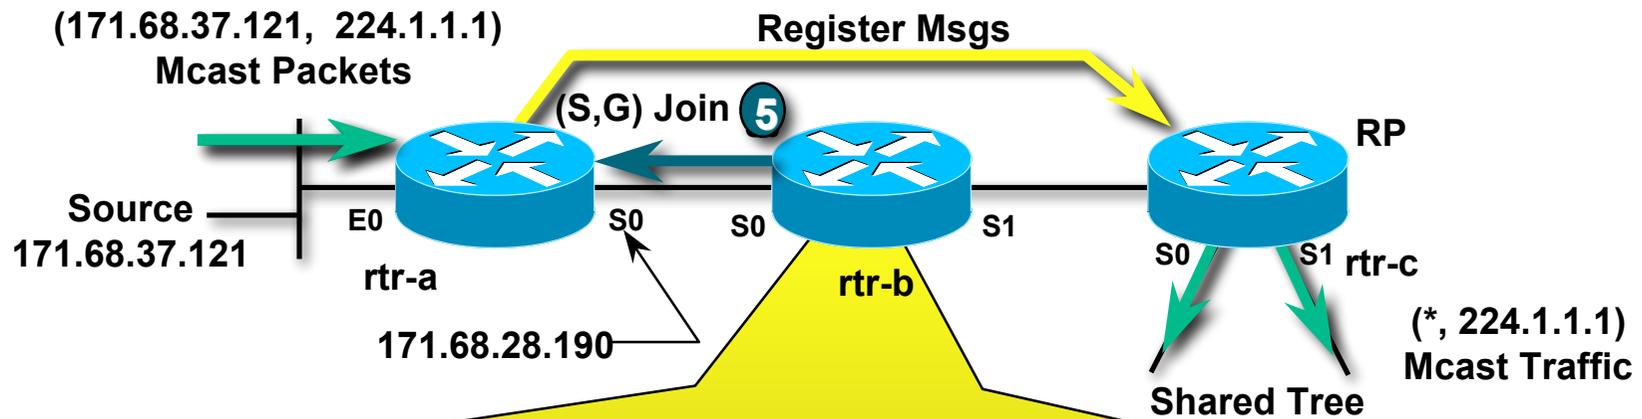
Receiver Joins Group First



- 4 RP sends (S,G) Join toward Source to build SPT.

PIM SM Registering

Receiver Joins Group First



```
(*, 224.1.1.1), 00:04:28/00:02:59, RP 171.68.28.140, flags: SP
Incoming interface: Serial1, RPF nbr 171.68.28.140,
Outgoing interface list: Null

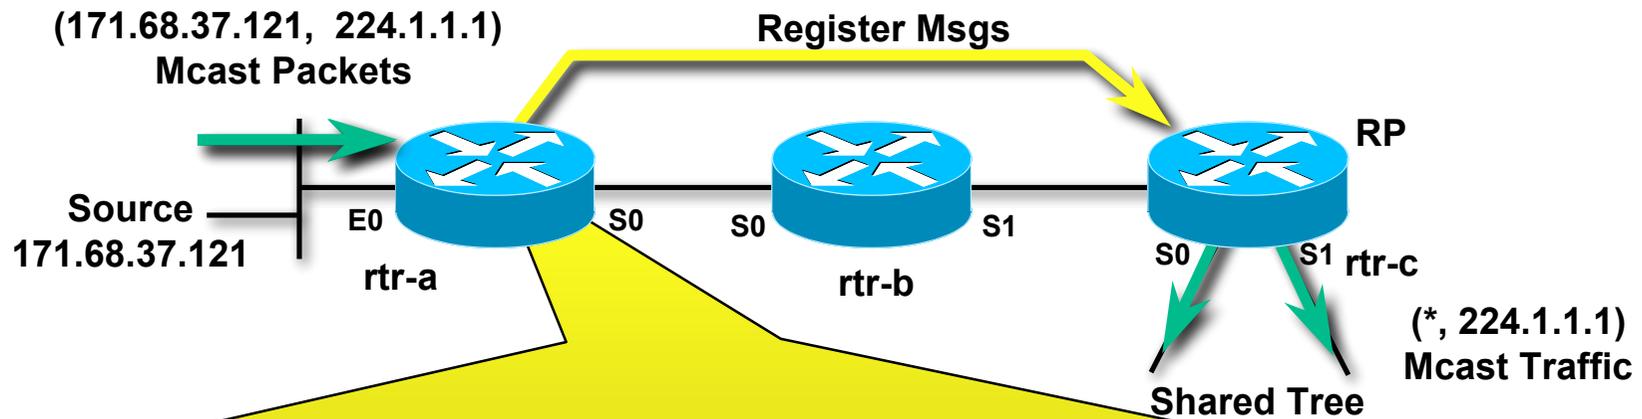
(171.68.37.121/32, 224.1.1.1), 00:04:28/00:01:32, flags:
Incoming interface: Serial0, RPF nbr 171.68.28.190
Outgoing interface list:
Serial1, Forward/Sparse-Dense, 00:04:28/00:01:32
```

“rtr-b” processes Join, creates (S, G) state
 (After automatically creating the (*, G) entry)

5 “rtr-b” sends (S,G) Join toward Source to continue building SPT.

PIM SM Registering

Receiver Joins Group First



```
(*, 224.1.1.1), 00:04:28/00:02:59, RP 171.68.28.140, flags: SP
Incoming interface: Serial0, RPF nbr 171.68.28.191,
Outgoing interface list: Null

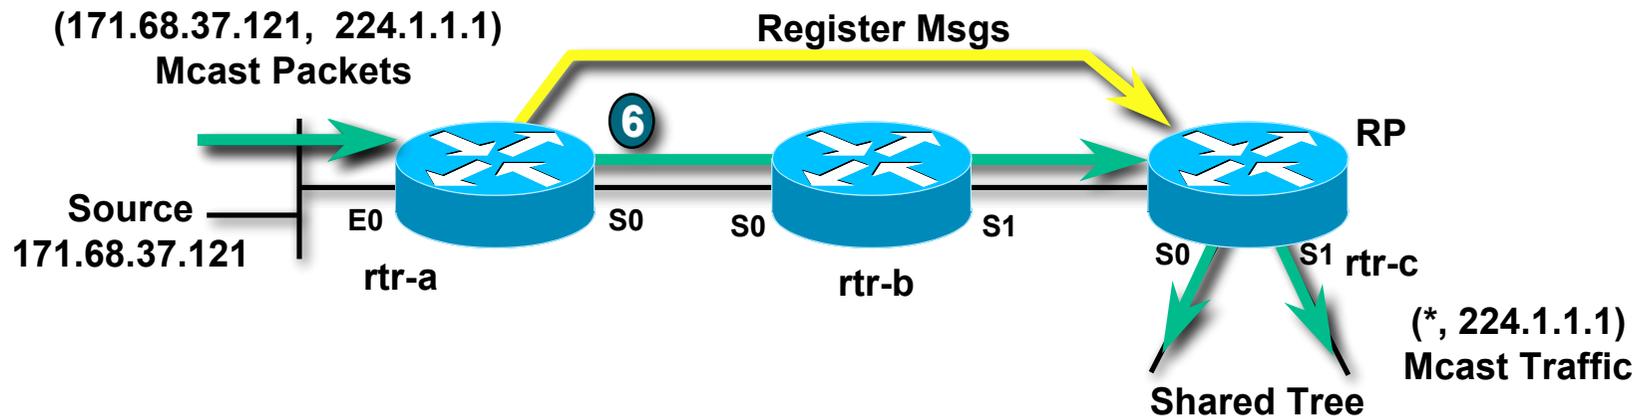
(171.68.37.121/32, 224.1.1.1), 00:04:28/00:01:32, flags: FT
Incoming interface: Ethernet0, RPF nbr 0.0.0.0, Registering
Outgoing interface list:
  Serial0, Forward/Sparse-Dense, 00:04:28/00:01:32
```

“rtr-a” processes the (S, G) Join; adds Serial 0 to OIL

PIM SM Registering

Receiver Joins Group First

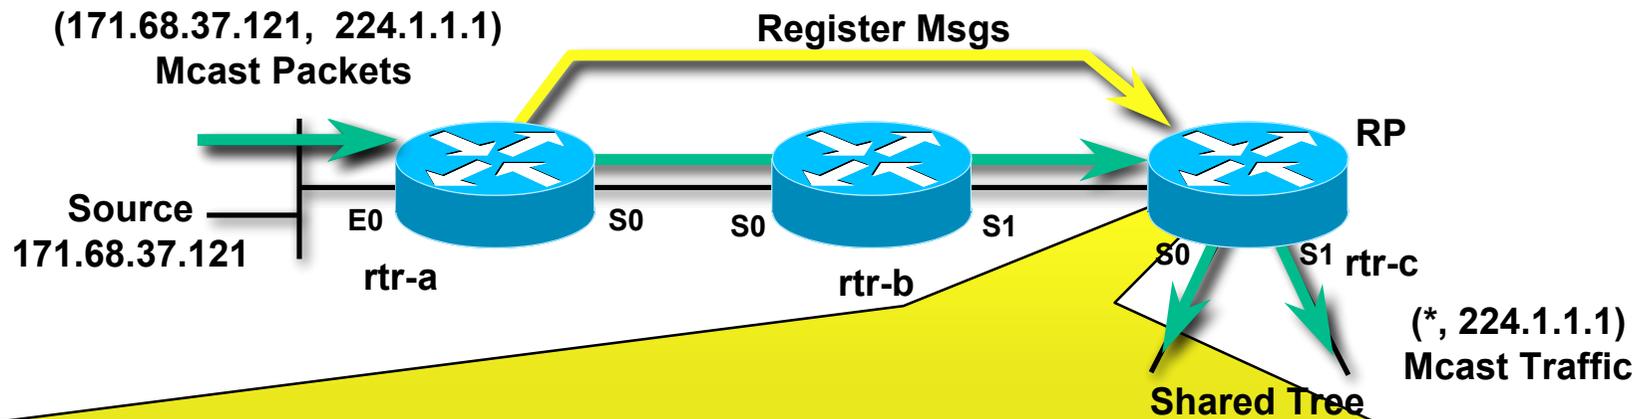
Cisco.com



- 6 RP begins receiving (S,G) traffic down SPT.

PIM SM Registering

Receiver Joins Group First



```
(*, 224.1.1.1), 00:09:21/00:02:38, RP 171.68.28.140, flags: S
Incoming interface: Null, RPF nbr 0.0.0.0,
Outgoing interface list:
  Serial0, Forward/Sparse-Dense, 00:09:21/00:02:38
  Serial1, Forward/Sparse-Dense, 00:03:14/00:02:46
```

```
(171.68.37.121, 224.1.1.1, 00:01:15/00:02:46, flags:T
Incoming interface: Serial3, RPF nbr 171.68.28.139,
Outgoing interface list:
  Serial0, Forward/Sparse-Dense, 00:00:49/00:02:11
  Serial1, Forward/Sparse-Dense, 00:00:49/00:02:11
```

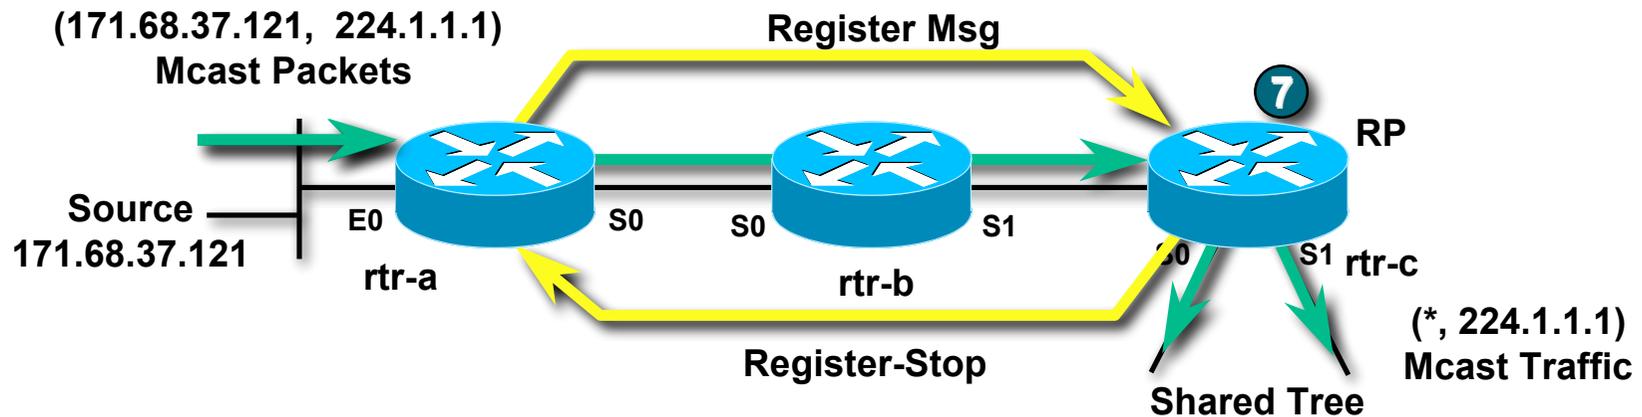
Note "T" Flag is now set

**Traffic arriving via SPT is forwarded down Shared Tree
(This causes the "T" flag to be set.)**

PIM SM Registering

Receiver Joins Group First

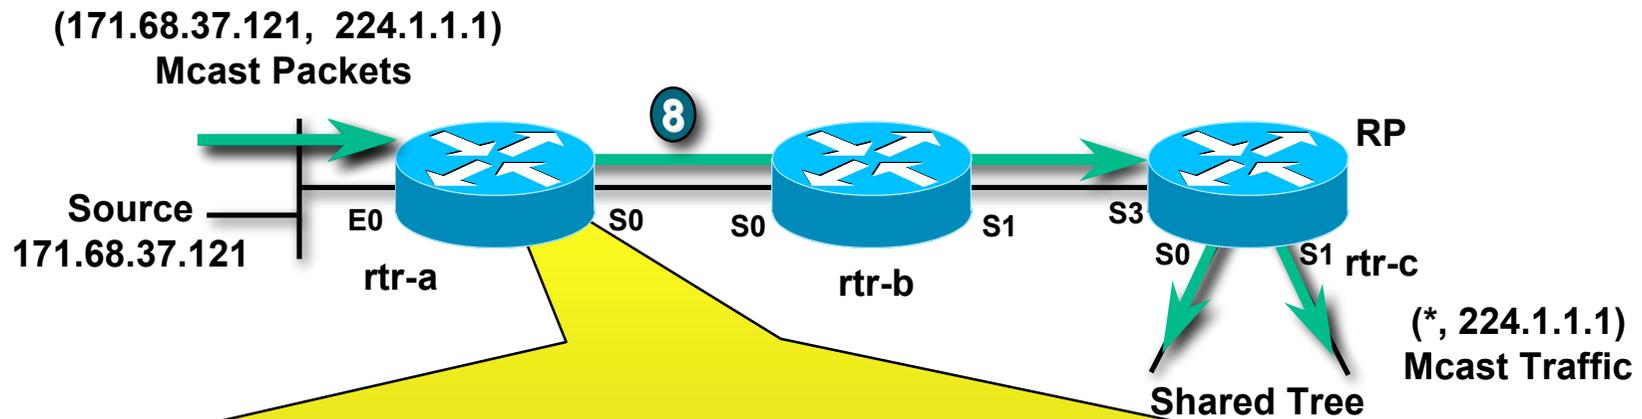
Cisco.com



- 7 Once “T” Flag is set, next “Register” causes RP to send back a “Register-Stop” to “rtr-a”.

PIM SM Registering

Receiver Joins Group First



```
(* , 224.1.1.1), 00:04:28/00:01:32, RP 171.68.28.140, flags: SP  
Incoming interface: Serial0, RPF nbr 171.68.28.191,  
Outgoing interface list: Null
```

```
(171.68.37.121/32, 224.1.1.1), 00:04:28/00:01:32, flags: FT  
Incoming interface: Ethernet0, RPF nbr 0.0.0.0,  
Outgoing interface list:  
Serial0, Forward/Sparse-Dense, 00:04:28/00:01:32
```

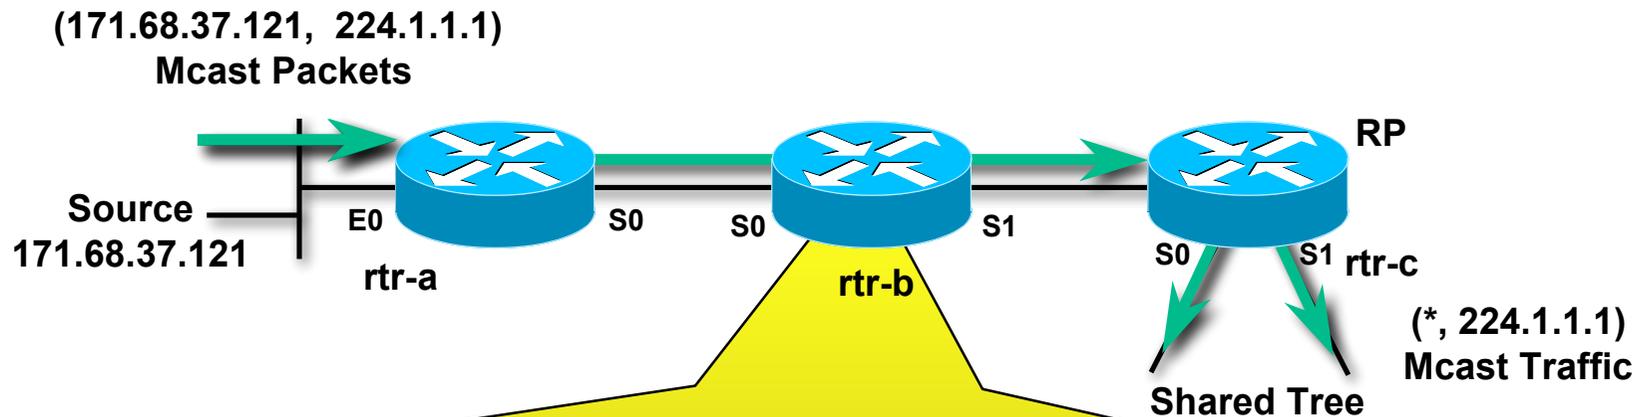
**“rtr-a” stops sending Register messages
(Final State in “rtr-a”)**

8 (S,G) Traffic now flowing down a single path (SPT) to RP.

PIM SM Registering

Receiver Joins Group First

Cisco.com

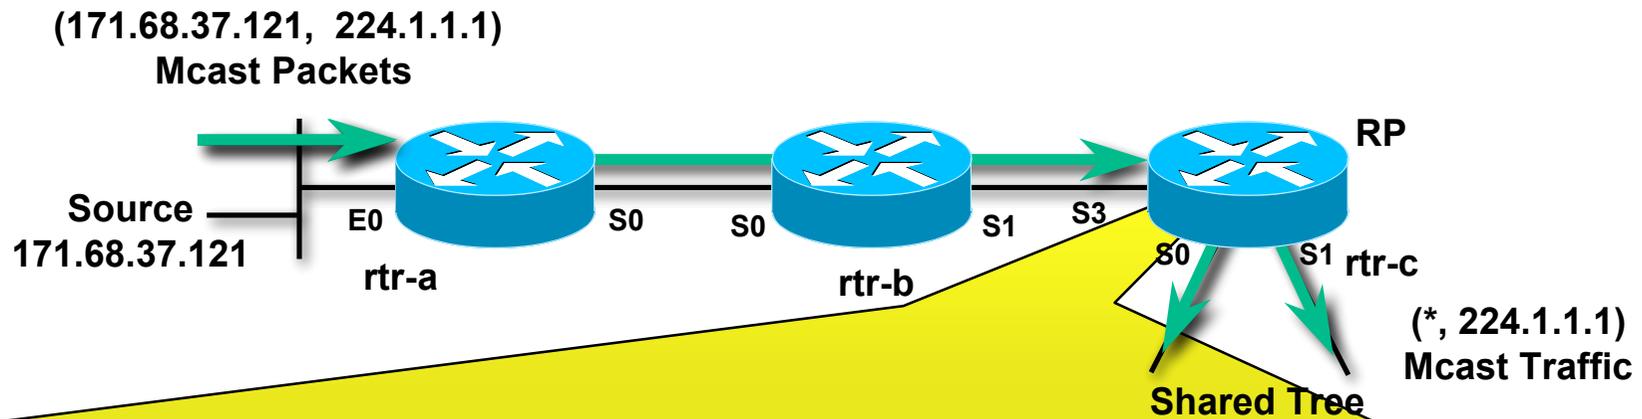


```
(*, 224.1.1.1), 00:04:28/00:01:32, RP 171.68.28.140, flags: SP  
Incoming interface: Serial1, RPF nbr 171.68.28.140,  
Outgoing interface list: Null  
  
(171.68.37.121/32, 224.1.1.1), 00:04:28/00:01:32, flags: T  
Incoming interface: Serial0, RPF nbr 171.68.28.190  
Outgoing interface list:  
Serial1, Forward/Sparse-Dense, 00:04:28/00:01:32
```

Final state in "rtr-b"

PIM SM Registering

Receiver Joins Group First



```
(*, 224.1.1.1), 00:09:21/00:02:38, RP 171.68.28.140, flags: S  
Incoming interface: Null, RPF nbr 0.0.0.0,  
Outgoing interface list:  
  Serial0, Forward/Sparse-Dense, 00:09:21/00:02:38  
  Serial11, Forward/Sparse-Dense, 00:03:14/00:02:46
```

```
(171.68.37.121, 224.1.1.1, 00:01:15/00:02:46, flags: T  
Incoming interface: Serial3, RPF nbr 171.68.28.139,  
Outgoing interface list:  
  Serial0, Forward/Sparse-Dense, 00:00:49/00:02:11  
  Serial11, Forward/Sparse-Dense, 00:00:49/00:02:11
```

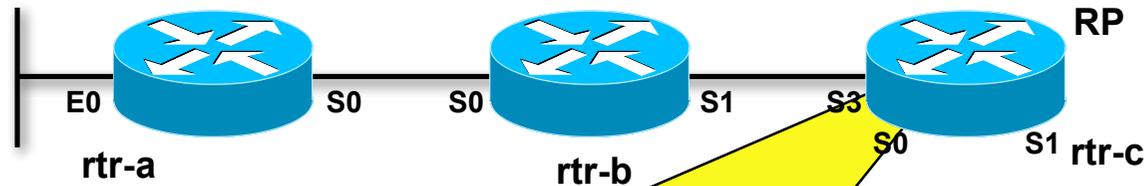
**Final state in the “RP”
(with receivers on Shared Tree)**

PIM SM Register Examples

- Receivers Join Group First
- **Source Registers First**
- Receivers along the SPT

PIM SM Registering

Source Registers First

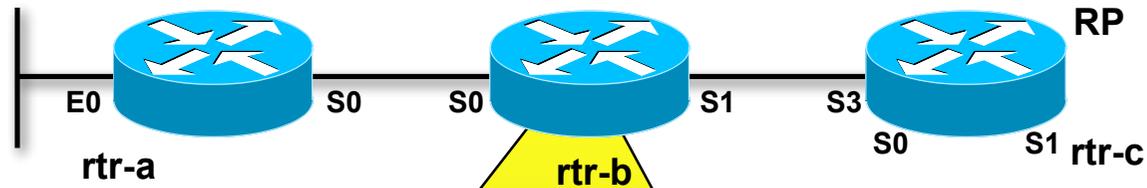


```
rtr-c>show ip mroute 224.1.1.1  
  
Group 224.1.1.1 not found.
```

**State in “RP” before Registering
(without receivers on Shared Tree)**

PIM SM Registering

Source Registers First



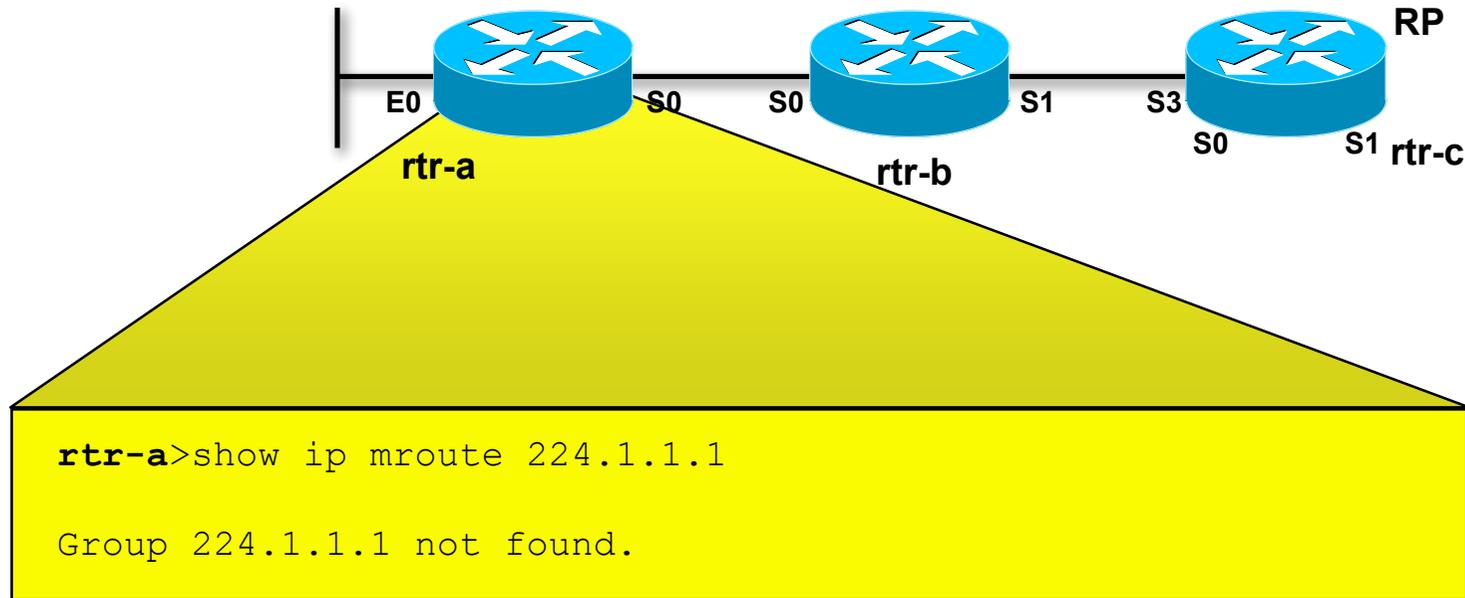
```
rtr-b>show ip mroute 224.1.1.1
```

```
Group 224.1.1.1 not found.
```

**State in “rtr-b” before any source registers
(with receivers on Shared Tree)**

PIM SM Registering

Source Registers First

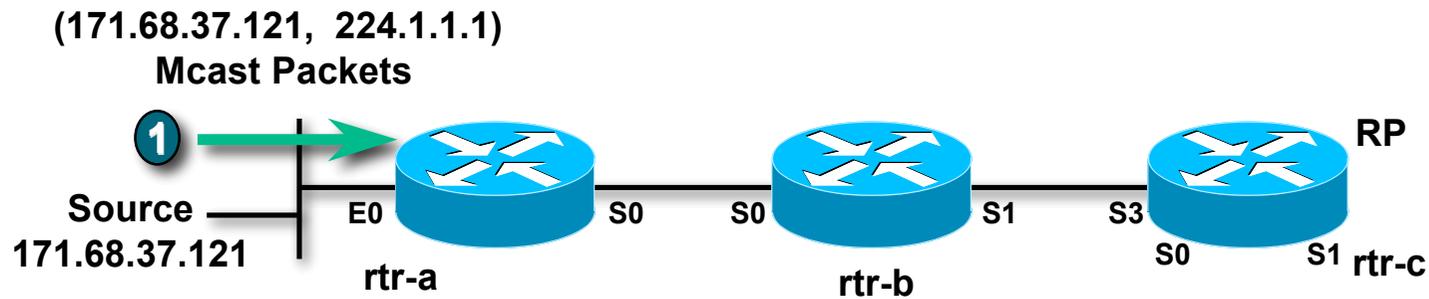


**State in “rtr-a” before any source registers
(with receivers on Shared Tree)**

PIM SM Registering

Source Registers First

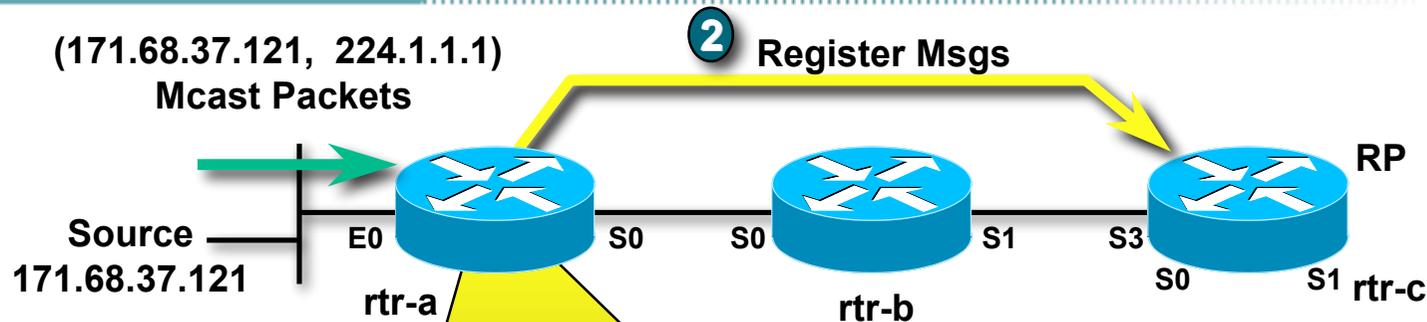
Cisco.com



- 1 “Source” begins sending group G traffic.

PIM SM Registering

Source Registers First



```
(* , 224.1.1.1), 00:00:03/00:02:56, RP 171.68.28.140, flags: SP  
Incoming interface: Serial0, RPF nbr 171.68.28.191,  
Outgoing interface list: Null
```

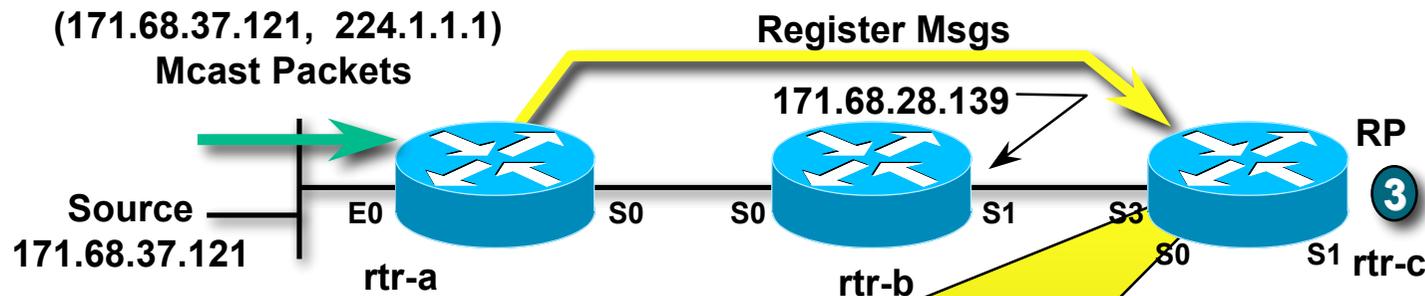
```
(171.68.37.121/32, 224.1.1.1), 00:00:03/00:02:56, flags: FPT  
Incoming interface: Ethernet0, RPF nbr 0.0.0.0, Registering  
Outgoing interface list: Null
```

“rtr-a” creates (S, G) state for source
(After automatically creating a (*, G) entry)

- 1 “Source” begins sending group G traffic.
- 2 “rtr-a” encapsulates packets in Registers; unicasts to RP.

PIM SM Registering

Source Registers First



```
(* , 224.1.1.1), 00:01:15/00:01:45, RP 171.68.28.140, flags: SP  
Incoming interface: Null, RPF nbr 0.0.0.0,  
Outgoing interface list: Null
```

```
(171.68.37.121, 224.1.1.1), 00:01:15/00:01:45, flags: P  
Incoming interface: Serial3, RPF nbr 171.68.28.139,  
Outgoing interface list: Null
```

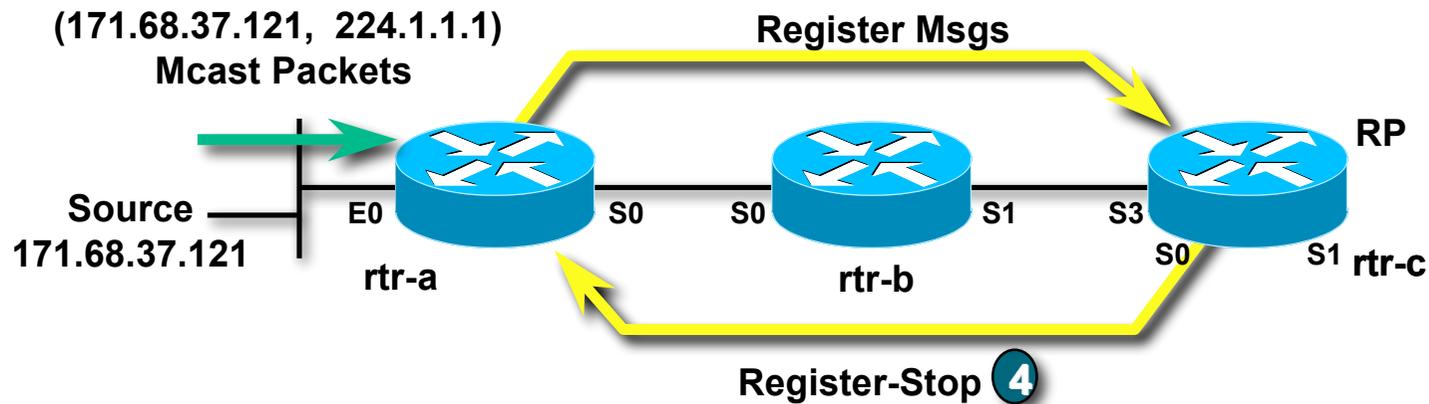
“RP” processes Register; creates (S, G) state
(After automatically creating the (*, G) entry)

3 “rtr-c” (RP) has no receivers on Shared Tree; discards packet.

PIM SM Registering

Source Registers First

Cisco.com

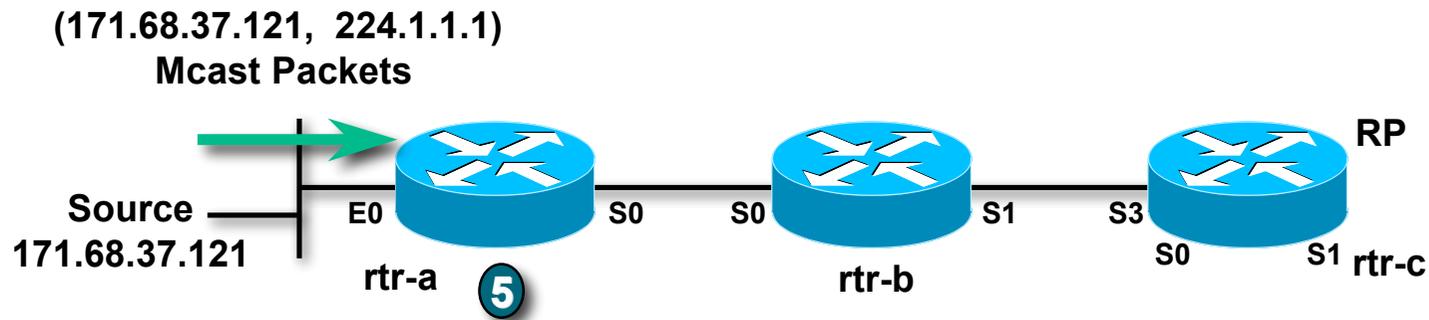


4 RP sends "Register-Stop" to "rtr-a".

PIM SM Registering

Source Registers First

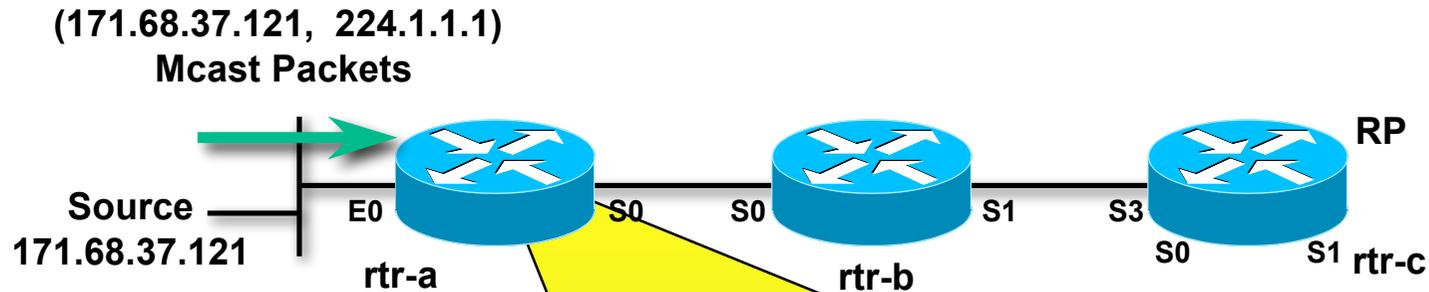
Cisco.com



- 5 “rtr-a” stops encapsulating traffic in Register Messages; drops packets from Source.

PIM SM Registering

Source Registers First



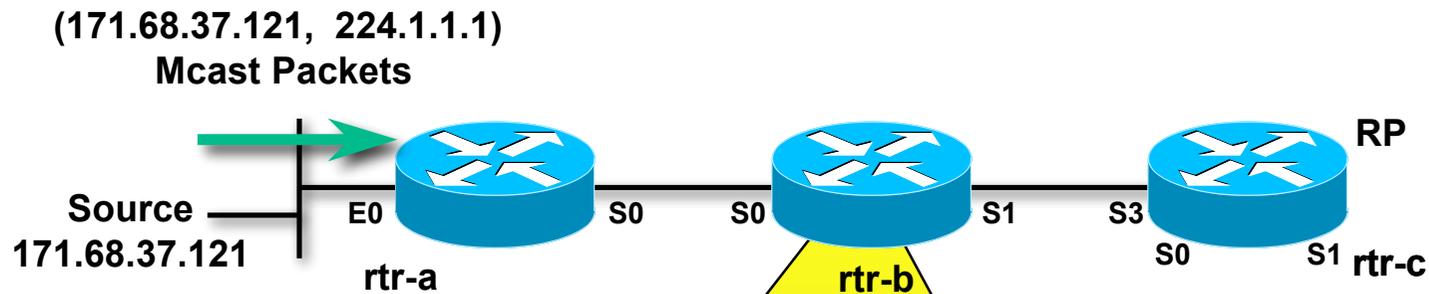
```
(*, 224.1.1.1), 00:01:28/00:01:32, RP 171.68.28.140, flags: SP  
Incoming interface: Serial0, RPF nbr 171.68.28.191,  
Outgoing interface list: Null
```

```
(171.68.37.121/32, 224.1.1.1), 00:01:28/00:01:32, flags: FPT  
Incoming interface: Ethernet0, RPF nbr 0.0.0.0  
Outgoing interface list: Null
```

**State in “rtr-a” after Registering
(without receivers on Shared Tree)**

PIM SM Registering

Source Registers First



```
rtr-b>show ip mroute 224.1.1.1
```

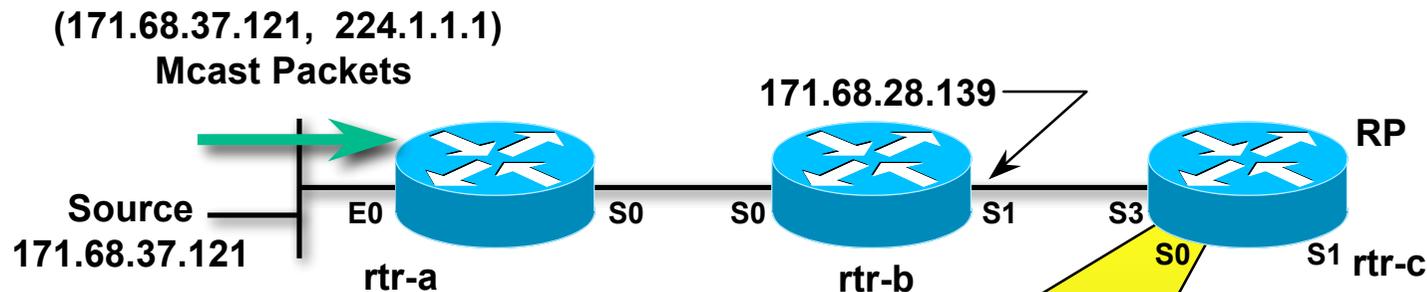
```
Group 224.1.1.1 not found.
```

State in “rtr-b” after “rtr-a” Registers
(without receivers on Shared Tree)

PIM SM Registering

Source Registers First

Cisco.com



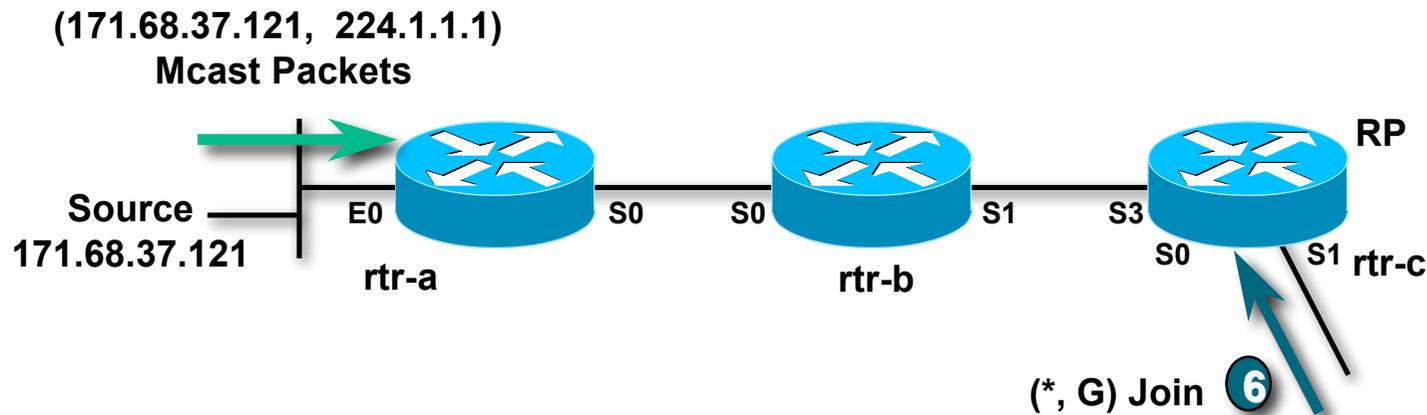
```
(* , 224.1.1.1), 00:01:15/00:01:45, RP 171.68.28.140, flags: SP  
Incoming interface: Null, RPF nbr 0.0.0.0,  
Outgoing interface list: Null  
  
(171.68.37.121, 224.1.1.1), 00:01:15/00:01:45, flags: P  
Incoming interface: Serial3, RPF nbr 171.68.28.139,  
Outgoing interface list: Null
```

State in “RP” after “rtr-a” Registers
(without receivers on Shared Tree)

PIM SM Registering

Source Registers First

Cisco.com

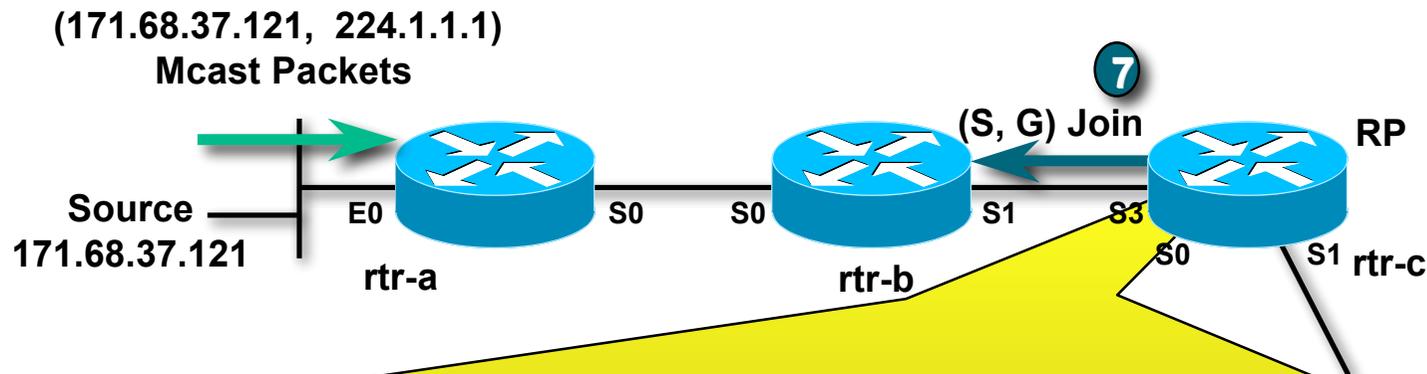


Receivers begin joining the Shared Tree

- 6 RP (“rtr-c”) receives (*, G) Join from a receiver on Shared Tree.

PIM SM Registering

Source Registers First



```
(*, 224.1.1.1), 00:09:21/00:02:38, RP 171.68.28.140, flags: S
Incoming interface: Null, RPF nbr 0.0.0.0,
Outgoing interface list:
  Serial1, Forward/Sparse-Dense, 00:00:14/00:02:46

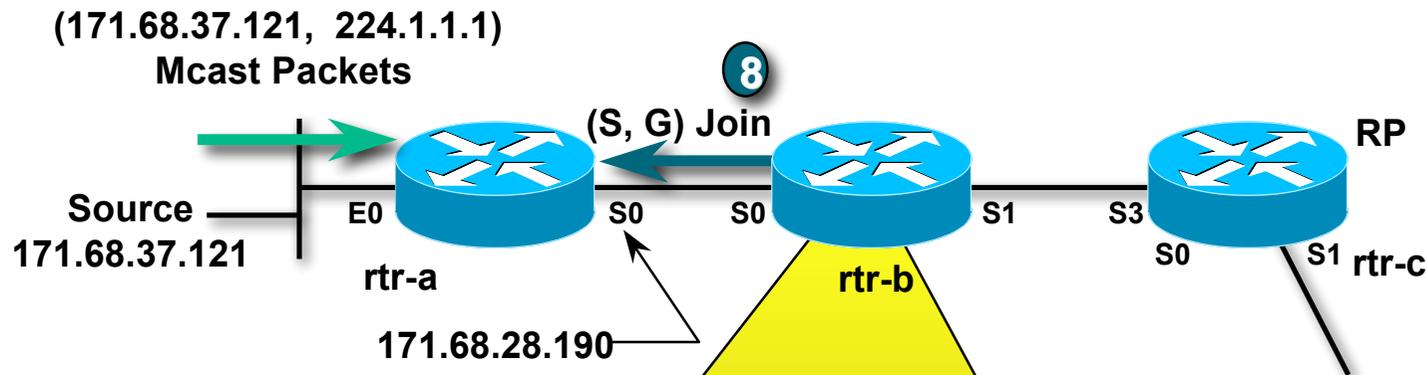
(171.68.37.121/32, 224.1.1.1, 00:01:15/00:02:46, flags: T
Incoming interface: Serial3, RPF nbr 171.68.28.139,
Outgoing interface list:
  Serial1, Forward/Sparse-Dense, 00:00:14/00:02:46
```

**“RP” processes (*,G) Join
(Adds Serial1 to Outgoing Interface Lists)**

7 RP sends (S,G) Joins for all known Sources in Group.

PIM SM Registering

Source Registers First



```
(* , 224.1.1.1), 00:04:28/00:01:32, RP 171.68.28.140, flags: SP  
Incoming interface: Serial1, RPF nbr 171.68.28.140,  
Outgoing interface list: Null
```

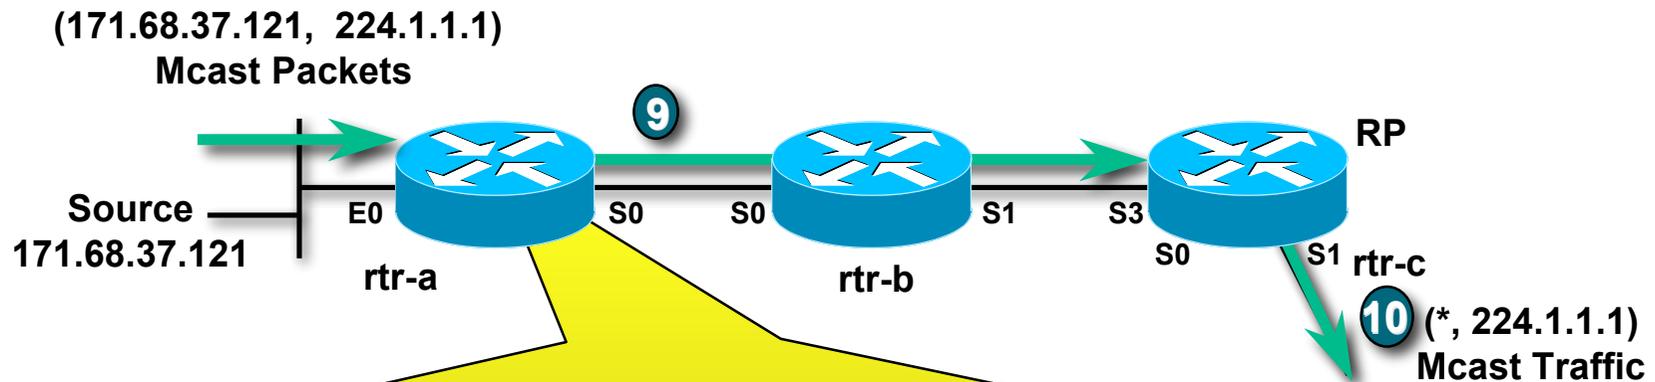
```
(171.68.37.121/32, 224.1.1.1), 00:04:28/00:01:32, flags:  
Incoming interface: Serial0, RPF nbr 171.68.28.190  
Outgoing interface list:  
Serial1, Forward/Sparse-Dense, 00:04:28/00:01:32
```

“rtr-b” processes Join, creates (S, G) state
(After automatically creating the (*, G) entry)

8 “rtr-b” sends (S,G) Join toward Source to continue building SPT.

PIM SM Registering

Source Registers First



```
(*, 224.1.1.1), 00:04:28/00:01:32, RP 171.68.28.140, flags: SP  
Incoming interface: Serial0, RPF nbr 171.68.28.191,  
Outgoing interface list: Null
```

```
(171.68.37.121/32, 224.1.1.1), 00:04:28/00:01:32, flags: FT  
Incoming interface: Ethernet0, RPF nbr 0.0.0.0,  
Outgoing interface list:  
Serial0, Forward/Sparse-Dense, 00:04:28/00:01:32
```

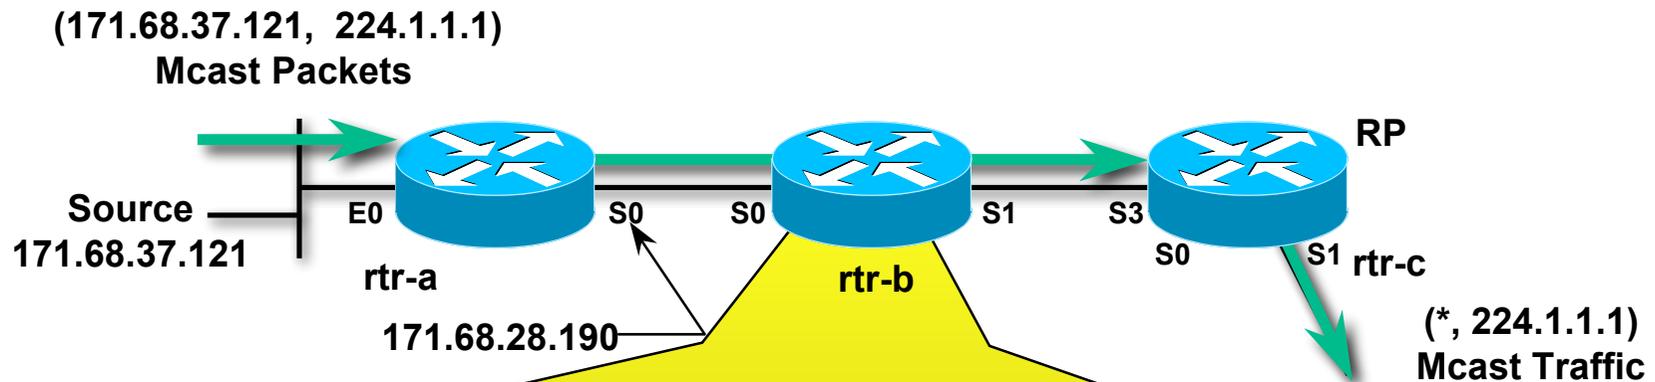
“rtr-a” processes the (S, G) Join; adds Serial0 to OIL

- 9 RP begins receiving (S,G) traffic down SPT.
- 10 RP forwards (S,G) traffic down Shared Tree to receivers.

PIM SM Registering

Source Registers First

Cisco.com



```
(*, 224.1.1.1), 00:04:28/00:01:32, RP 171.68.28.140, flags: SP  
Incoming interface: Serial1, RPF nbr 171.68.28.140,  
Outgoing interface list: Null
```

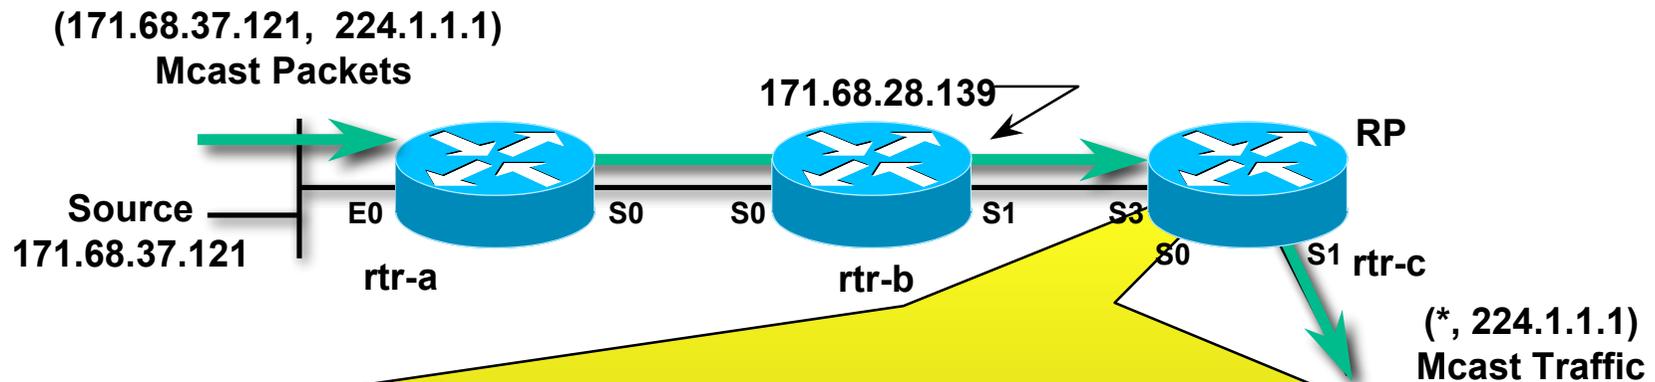
```
(171.68.37.121/32, 224.1.1.1), 00:04:28/00:01:32, flags: T  
Incoming interface: Serial0, RPF nbr 171.68.28.190  
Outgoing interface list:  
Serial1, Forward/Sparse-Dense, 00:04:28/00:01:32
```

Final state in "rtr-b" after Receivers Join

PIM SM Registering

Source Registers First

Cisco.com



```
(*, 224.1.1.1), 00:09:21/00:02:38, RP 171.68.28.140, flags: S
Incoming interface: Null, RPF nbr 0.0.0.0,
Outgoing interface list:
  Serial1, Forward/Sparse-Dense, 00:03:14/00:02:46

(171.68.37.121/32, 224.1.1.1, 00:01:15/00:02:46, flags: T
Incoming interface: Serial3, RPF nbr 171.68.28.139,
Outgoing interface list:
  Serial1, Forward/Sparse-Dense, 00:00:49/00:02:11
```

Final state in "RP" after Receivers Join

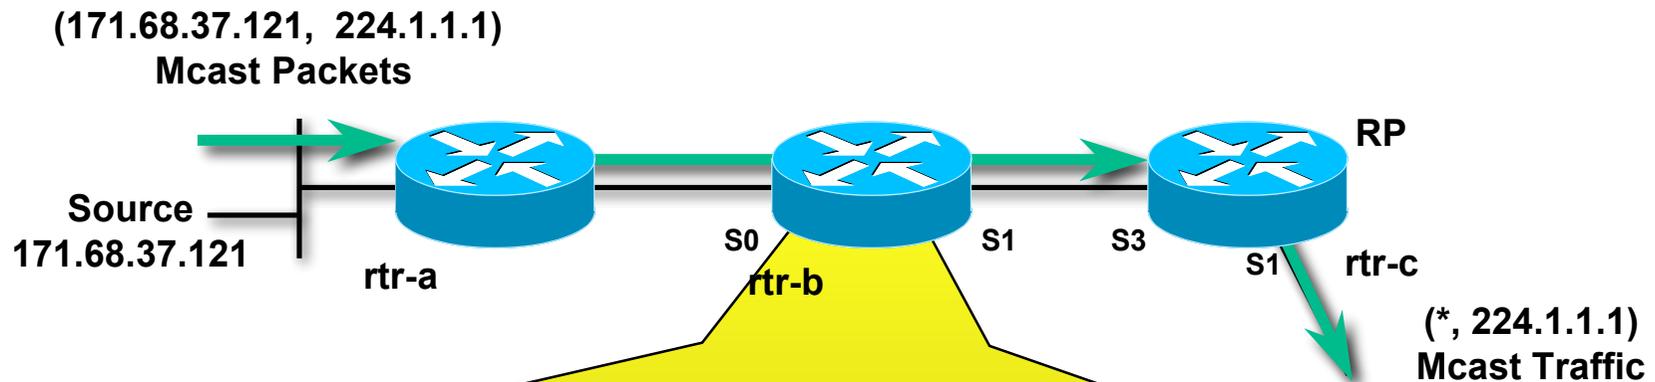
PIM SM Register Examples

- **Receivers Join Group First**
- **Source Registers First**
- **Receivers along the SPT**

PIM SM Registering

Receivers along the SPT

Cisco.com



```
(*, 224.1.1.1), 00:04:28/00:01:32, RP 171.68.28.140, flags: SP  
Incoming interface: Serial1, RPF nbr 171.68.28.140,  
Outgoing interface list: Null
```

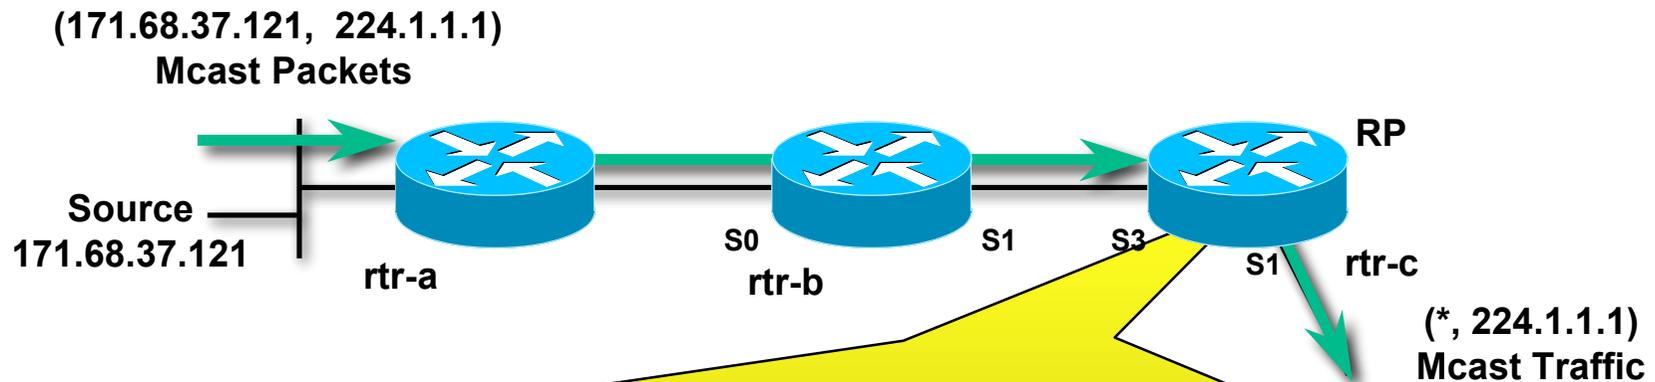
```
(171.68.37.121/32, 224.1.1.1), 00:04:28/00:01:32, flags: T  
Incoming interface: Serial0, RPF nbr 171.68.28.190  
Outgoing interface list:  
Serial1, Forward/Sparse-Dense, 00:04:28/00:01:32
```

Current state in "rtr-b"

PIM SM Registering

Receivers along the SPT

Cisco.com



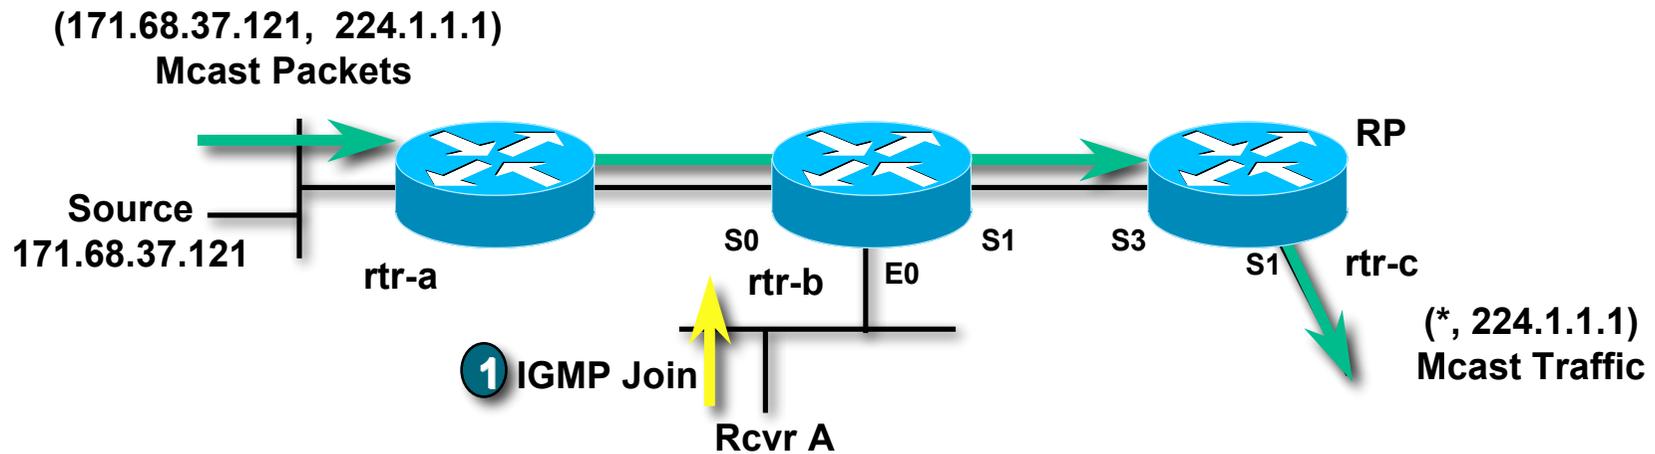
```
(*, 224.1.1.1), 00:09:21/00:02:38, RP 171.68.28.140, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0,
  Outgoing interface list:
    Serial1, Forward/Sparse-Dense, 00:03:14/00:02:46

(171.68.37.121/32, 224.1.1.1, 00:01:15/00:02:46, flags: T
  Incoming interface: Serial3, RPF nbr 171.68.28.139,
  Outgoing interface list:
    Serial1, Forward/Sparse-Dense, 00:00:49/00:02:11
```

Current state in the RP

PIM SM Registering

Receivers along the SPT

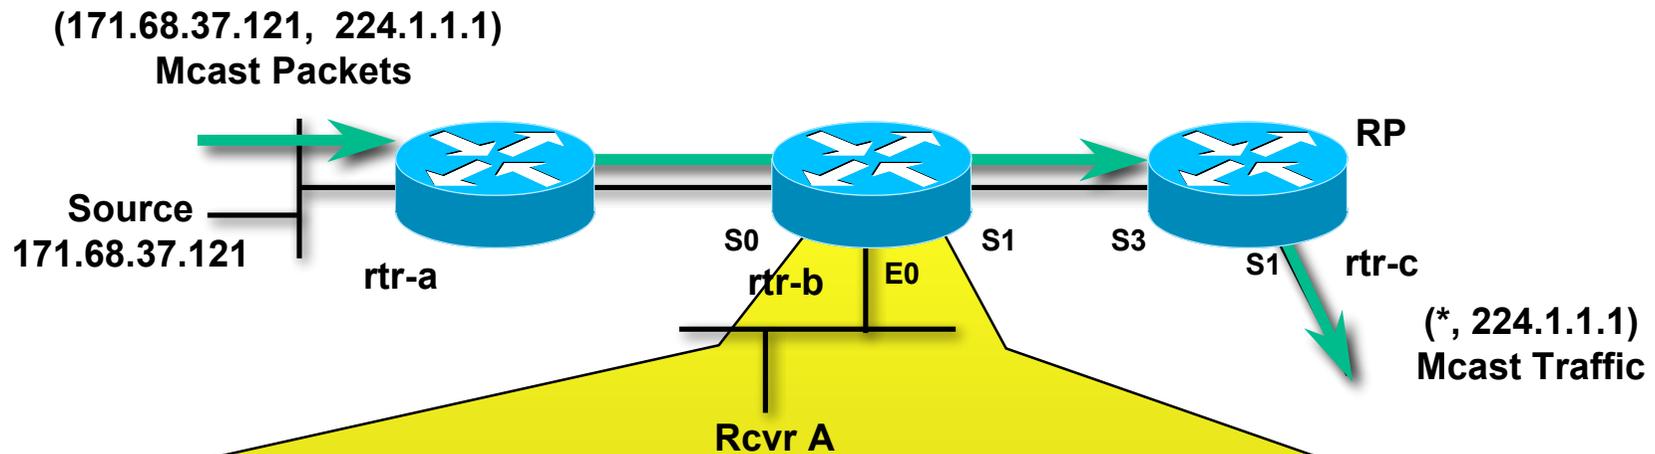


- 1 “Rcvr A” wishes to receive group G traffic. Sends IGMP Join for G.

PIM SM Registering

Receivers along the SPT

Cisco.com



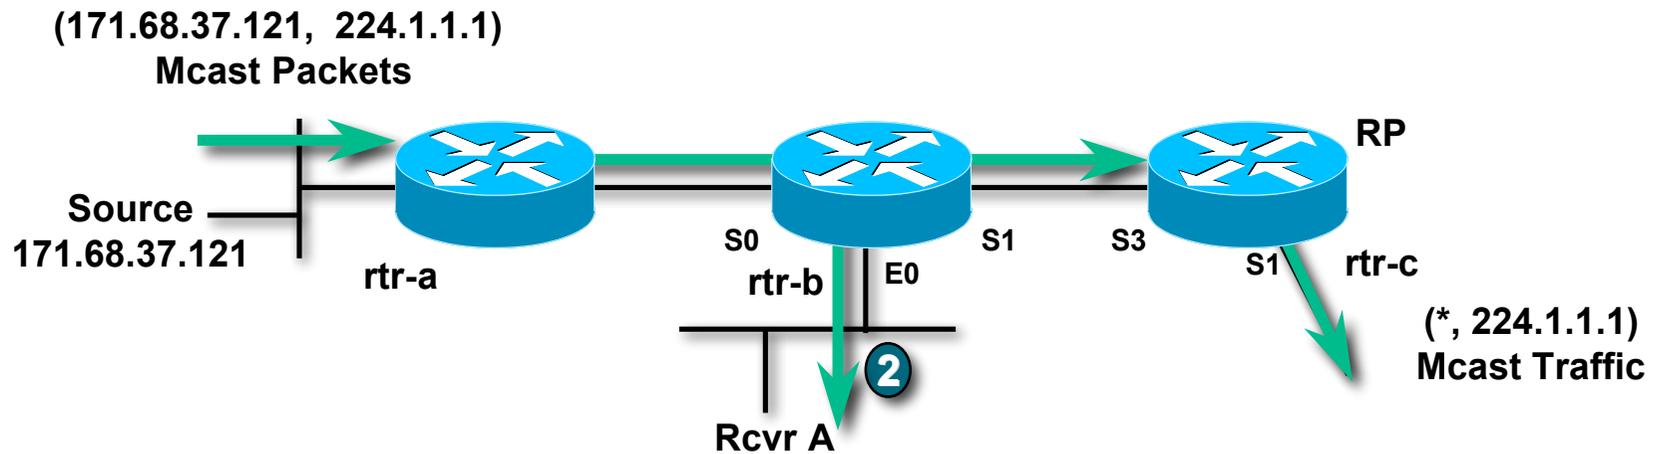
```
(* , 224.1.1.1), 00:04:28/00:01:32, RP 171.68.28.140, flags: SC  
Incoming interface: Serial1, RPF nbr 171.68.28.140,  
Outgoing interface list:  
  Ethernet0, Forward/Sparse-Dense, 00:00:30/00:02:30  
  
(171.68.37.121/32, 224.1.1.1), 00:04:28/00:01:32, flags: CT  
Incoming interface: Serial0, RPF nbr 171.68.28.190  
Outgoing interface list:  
  Serial1, Forward/Sparse-Dense, 00:04:28/00:01:32  
  Ethernet0, Forward/Sparse-Dense, 00:00:30/00:02:30
```

State in "rtr-b" after "Rcvr A" joins group

PIM SM Registering

Receivers along the SPT

Cisco.com



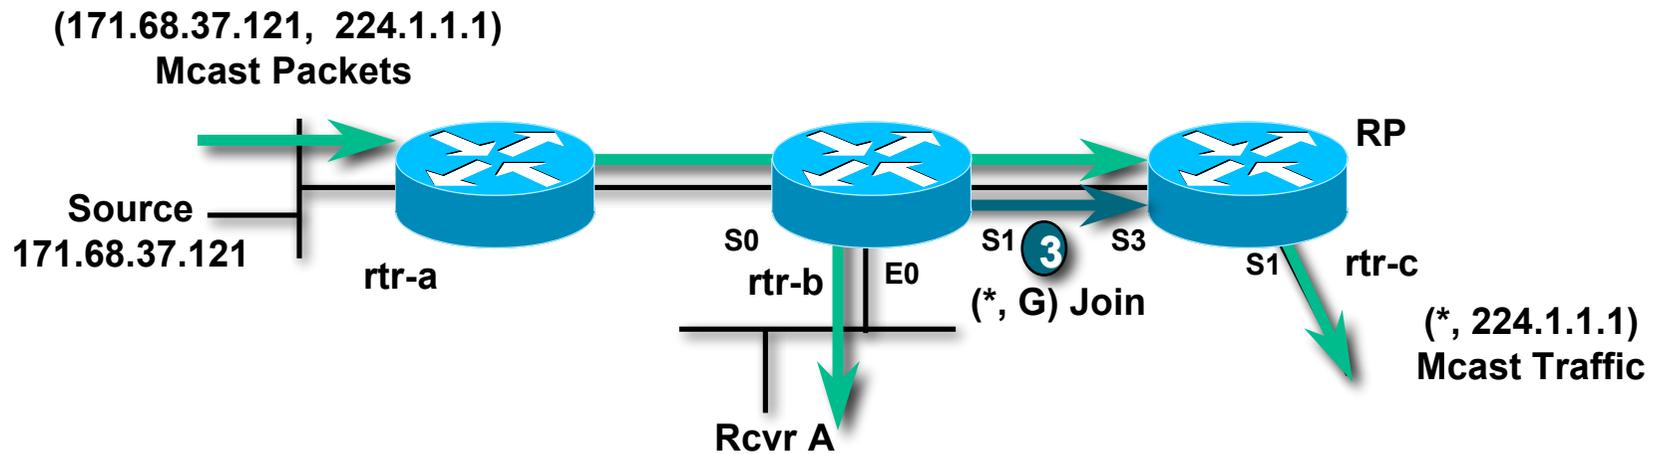
2 Group G traffic begins to flow to “Rcvr A”.

(Note: 171.68.37.121 traffic doesn't flow to RP then back down to rtr-b)

PIM SM Registering

Receivers along the SPT

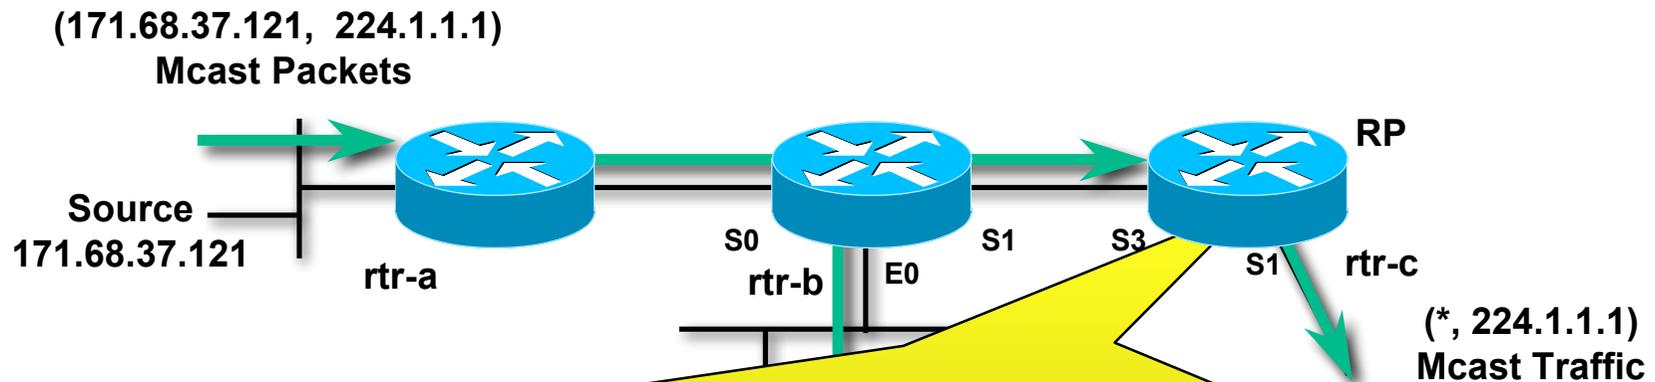
Cisco.com



- ③ “rtr-b” triggers a (*,G) Join to join the Shared Tree

PIM SM Registering

Receivers along the SPT



```
(*, 224.1.1.1), 00:09:21/00:02:38, RP 171.68.28.140, flags: S
Incoming interface: Null, RPF nbr 0.0.0.0,
Outgoing interface list:
  Serial1, Forward/Sparse-Dense, 00:03:14/00:02:46
  Serial3, Forward/Sparse-Dense, 00:00:10/00:02:50

(171.68.37.121/32, 224.1.1.1, 00:01:15/00:02:46, flags: T
Incoming interface: Serial3, RPF nbr 171.68.28.139,
Outgoing interface list:
  Serial1, Forward/Sparse-Dense, 00:00:49/00:02:11
```

State in "RP" after "rtr-b" joins Shared Tree

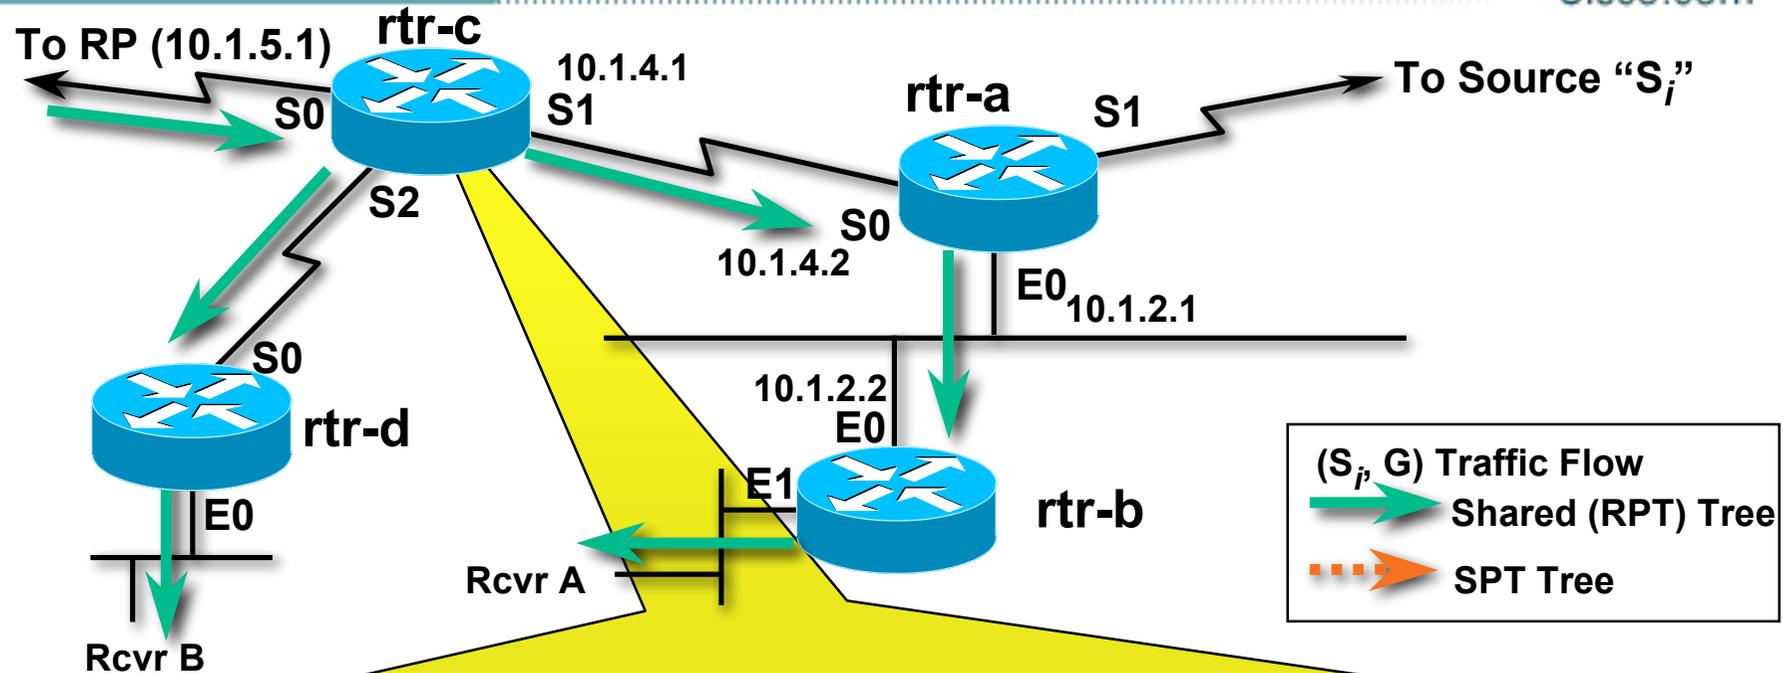
Agenda

- PIM Packet Formats
- PIM Neighbor Discovery
- PIM State
- PIM SM Joining
- PIM SM Registering
- **PIM SM SPT-Switchover**
- PIM SM Pruning
- PIM Asserts

PIM SM SPT-Switchover

- **SPT Thresholds may be set for any Group**
 - Access Lists may be used to specify which Groups
 - Default Threshold = 0kbps (I.e. immediately join SPT)
 - Threshold = “infinity” means “never join SPT”.
 - *Don't use values in between “0” and “infinity”.*
- **Threshold triggers Join of Source Tree**
 - Sends an (S,G) Join up SPT for next “S” in “G” packet received.

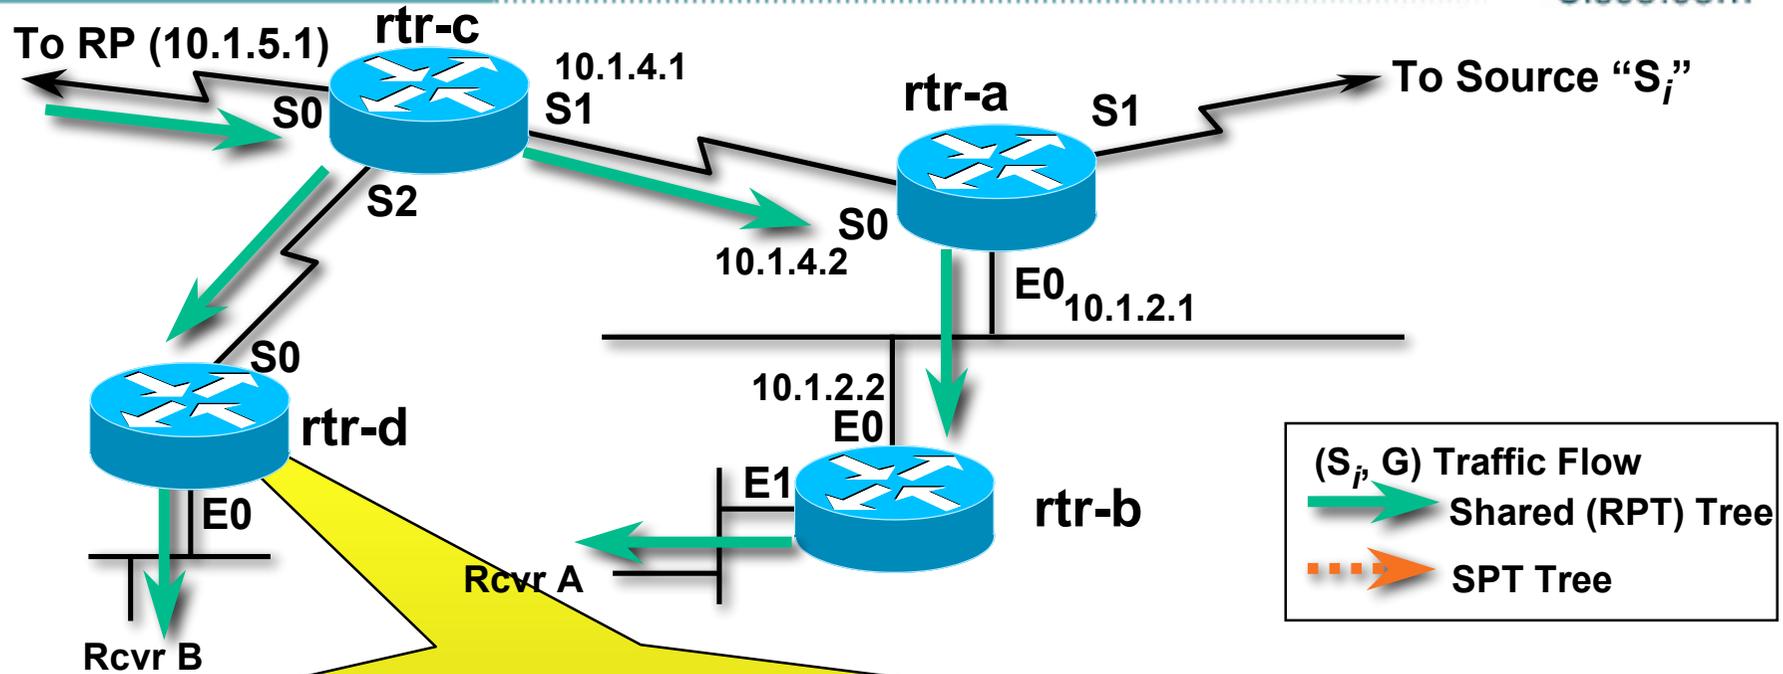
PIM SM SPT-Switchover



```
(*, 224.1.1.1), 00:01:43/00:02:13, RP 10.1.5.1, flags: S
Incoming interface: Serial0, RPF nbr 10.1.5.1,
Outgoing interface list:
  Serial1, Forward/Sparse-Dense, 00:01:43/00:02:11
  Serial2, Forward/Sparse-Dense, 00:00:32/00:02:28
```

State in "rtr-c" before switch

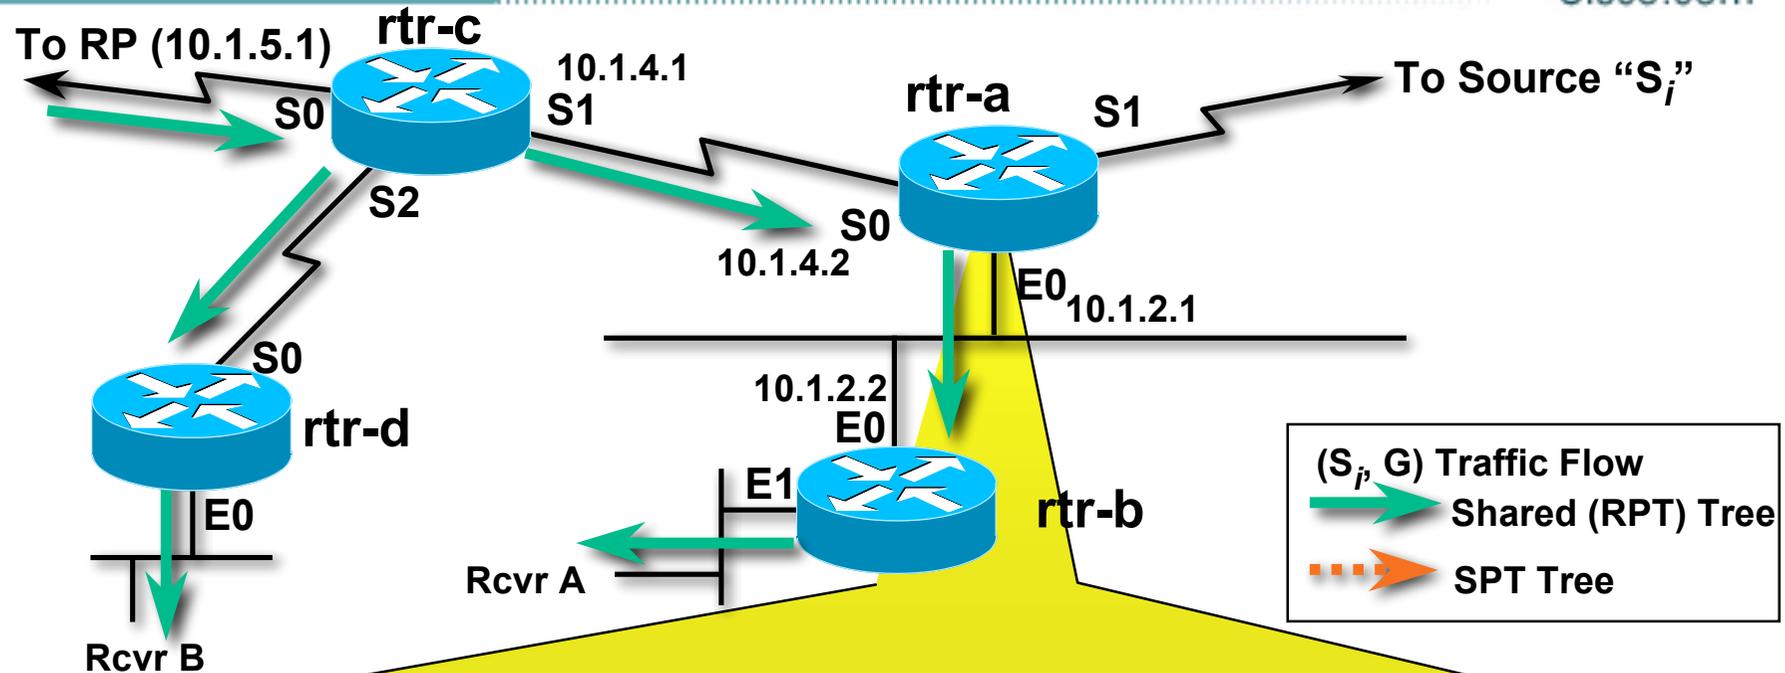
PIM SM SPT-Switchover



```
(*, 224.1.1.1), 00:01:43/00:02:13, RP 10.1.5.1, flags: SC
Incoming interface: Serial0, RPF nbr 10.1.4.8,
Outgoing interface list:
  Ethernet0, Forward/Sparse-Dense, 00:01:43/00:02:11
```

State in "rtr-d" before switch

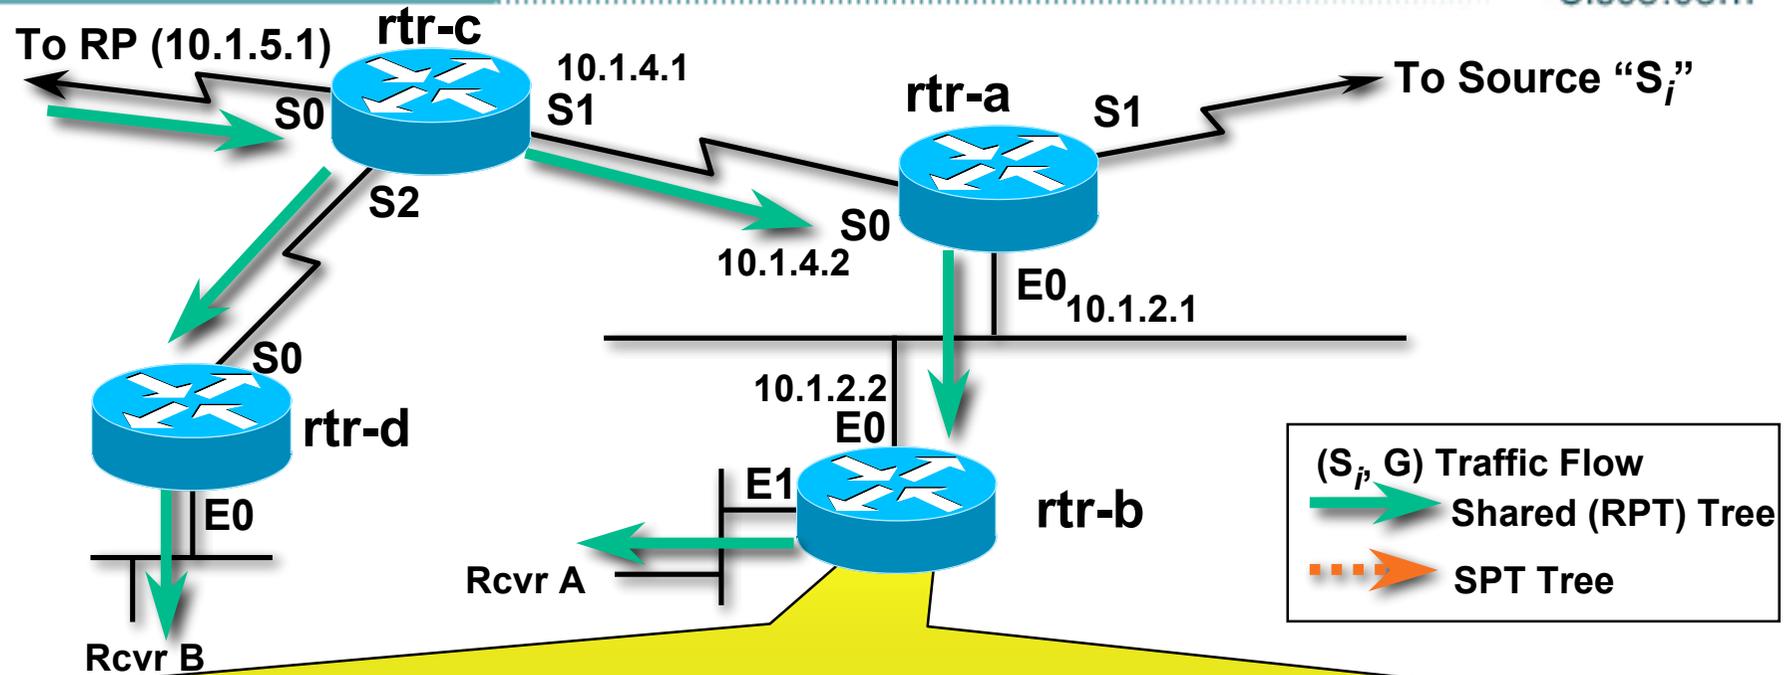
PIM SM SPT-Switchover



```
(*, 224.1.1.1), 00:01:43/00:02:13, RP 10.1.5.1, flags: S
Incoming interface: Serial0, RPF nbr 10.1.4.1,
Outgoing interface list:
Ethernet0, Forward/Sparse-Dense, 00:01:43/00:02:11
```

State in "rtr-a" before switch

PIM SM SPT-Switchover

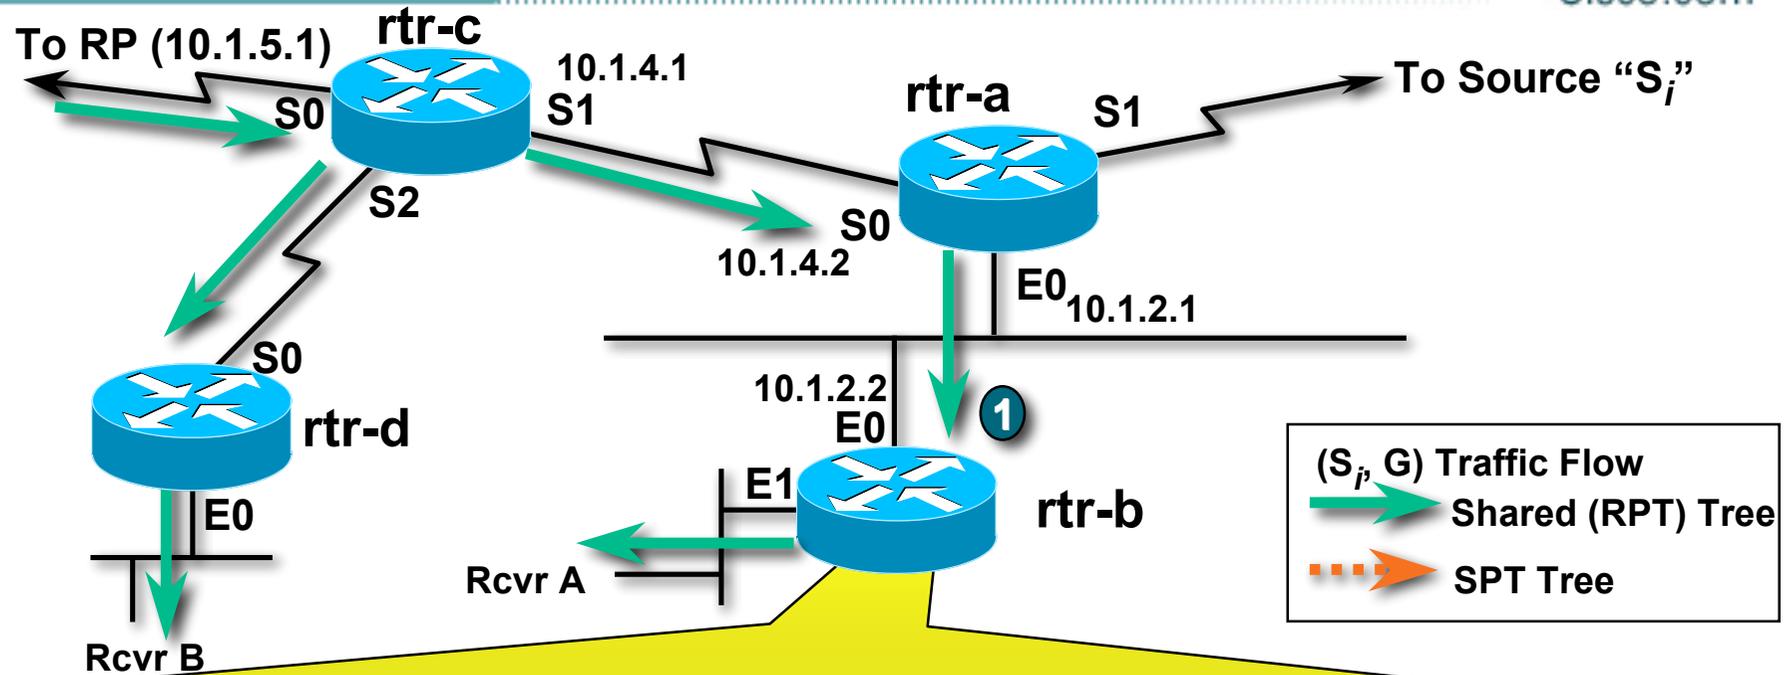


```
(*, 224.1.1.1), 00:01:43/00:02:13, RP 10.1.5.1, flags: SCJ
Incoming interface: Ethernet0, RPF nbr 10.1.2.1,
Outgoing interface list:
Ethernet1, Forward/Sparse-Dense, 00:01:43/00:02:11
```

Note "J"
Flag is set

State in "rtr-b" before switch

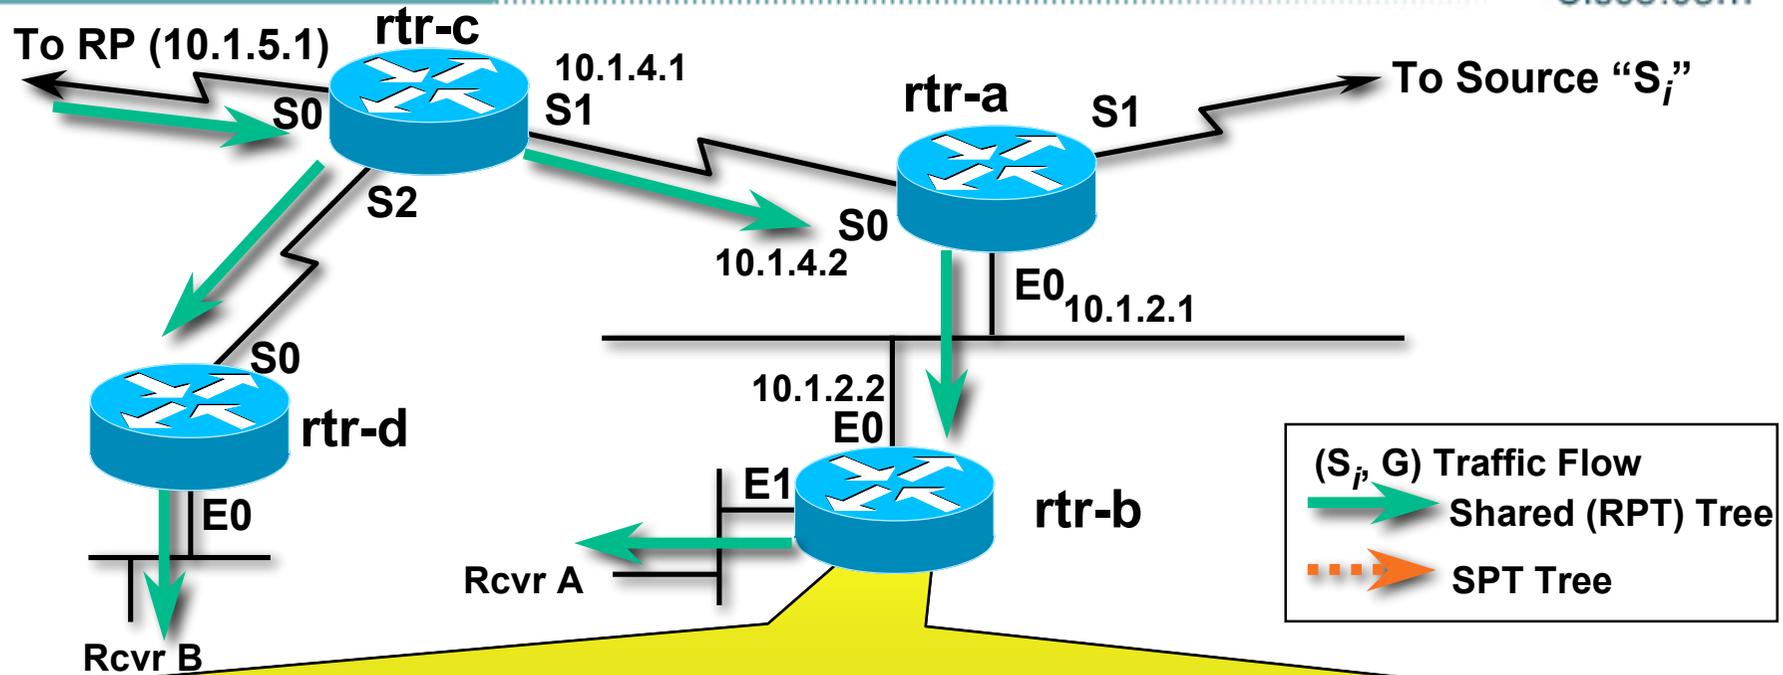
PIM SM SPT-Switchover



```
(*, 224.1.1.1), 00:01:43/00:02:13, RP 10.1.5.1, flags: SCJ
Incoming interface: Ethernet0, RPF nbr 10.1.2.1,
Outgoing interface list:
Ethernet1, Forward/Sparse-Dense, 00:01:43/00:02:11
```

1 New source (S_i,G) packet arrives down Shared tree.

PIM SM SPT-Switchover

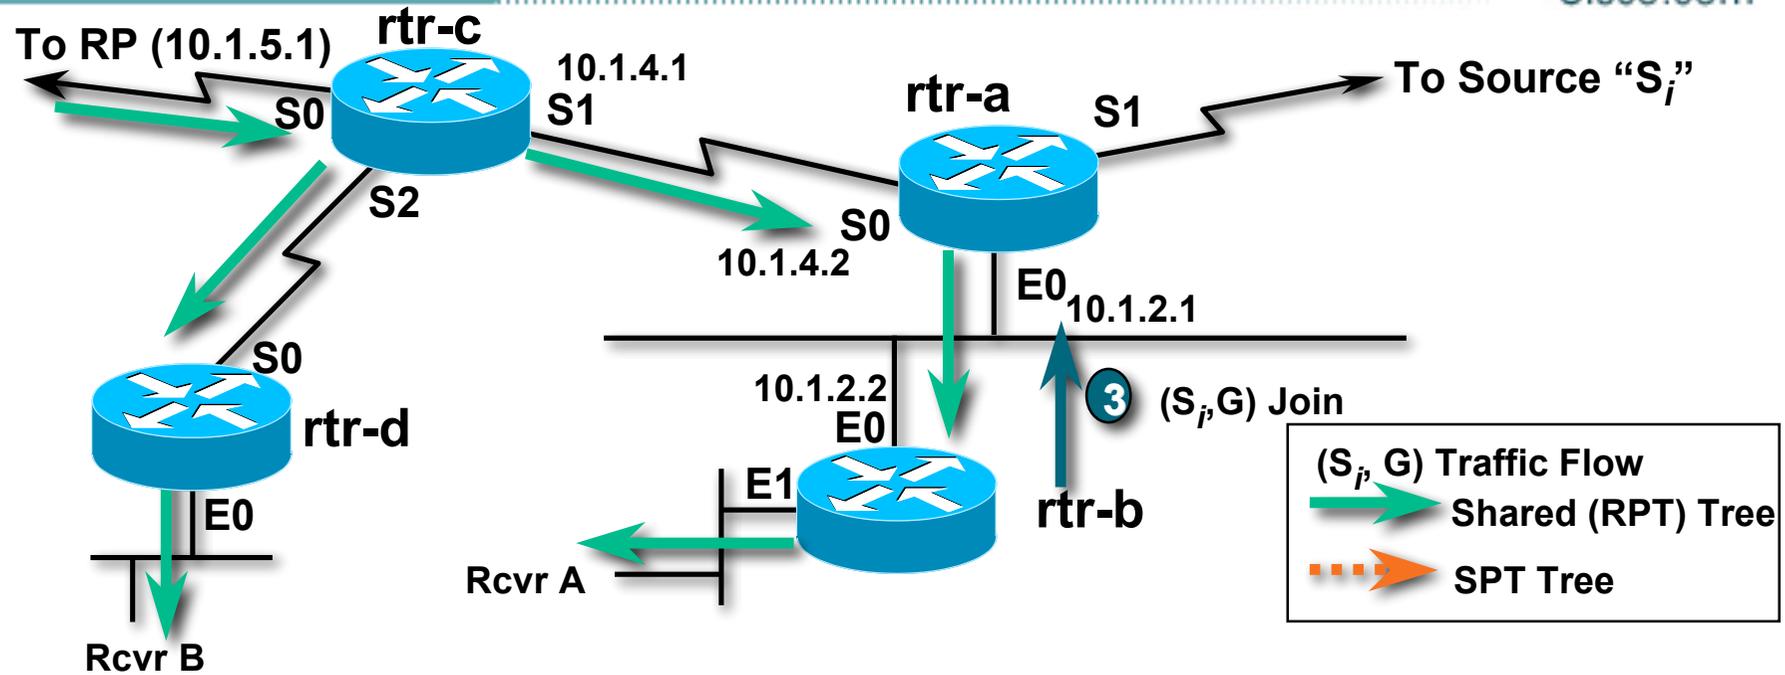


```
(*, 224.1.1.1), 00:01:43/00:02:13, RP 10.1.5.1, flags: SCJ
Incoming interface: Ethernet0, RPF nbr 10.1.2.1,
Outgoing interface list:
  Ethernet1, Forward/Sparse-Dense, 00:01:43/00:02:11

(171.68.37.121/32, 224.1.1.1), 00:00:28/00:02:51, flags: CJ
Incoming interface: Ethernet0, RPF nbr 10.1.2.1
Outgoing interface list:
  Ethernet1, Forward/Sparse-Dense, 00:00:28/00:02:32
```

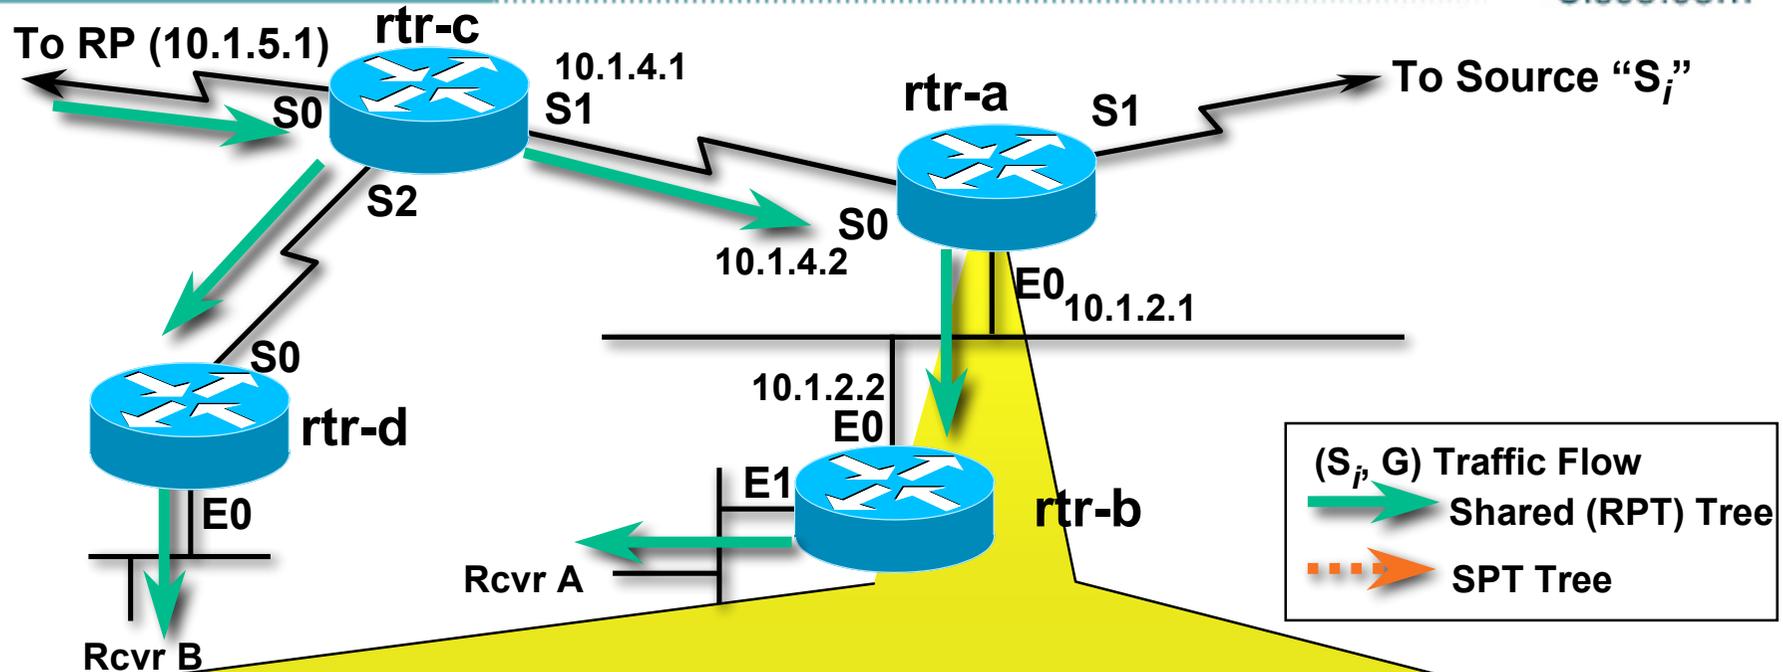
2 Create (S_i,G) state.

PIM SM SPT-Switchover



③ Send (S_i, G) Join towards S_i.

PIM SM SPT-Switchover

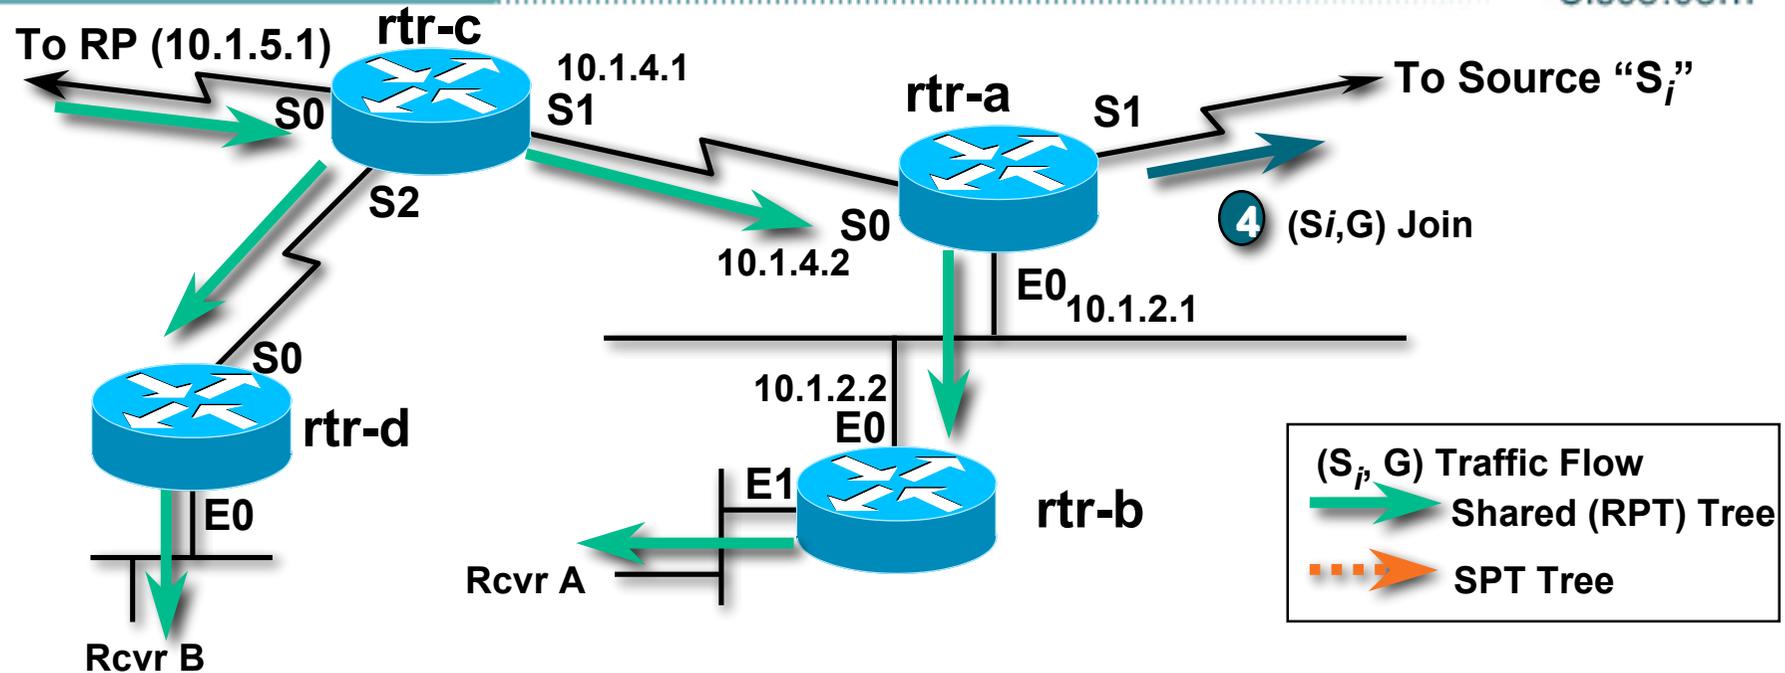


```
(*, 224.1.1.1), 00:01:43/00:02:13, RP 10.1.5.1, flags: S
Incoming interface: Serial0, RPF nbr 10.1.4.1,
Outgoing interface list:
  Ethernet0, Forward/Sparse-Dense, 00:01:43/00:02:11

(171.68.37.121/32, 224.1.1.1), 00:13:28/00:02:53, flags:
Incoming interface: Serial1, RPF nbr 10.1.9.2
Outgoing interface list:
  Ethernet0, Forward/Sparse-Dense, 00:13:25/00:02:30
```

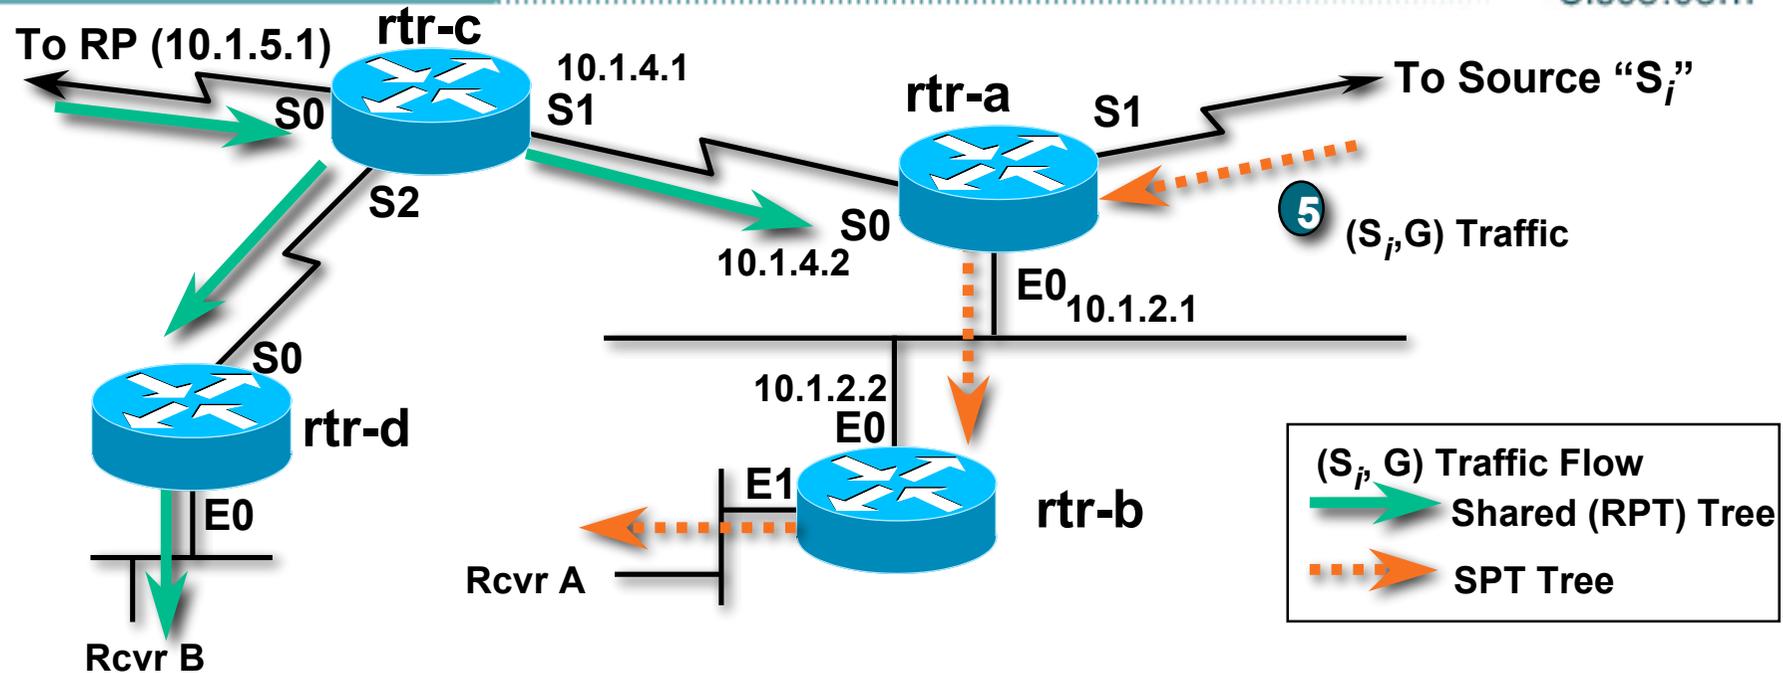
New state in "rtr-a"

PIM SM SPT-Switchover



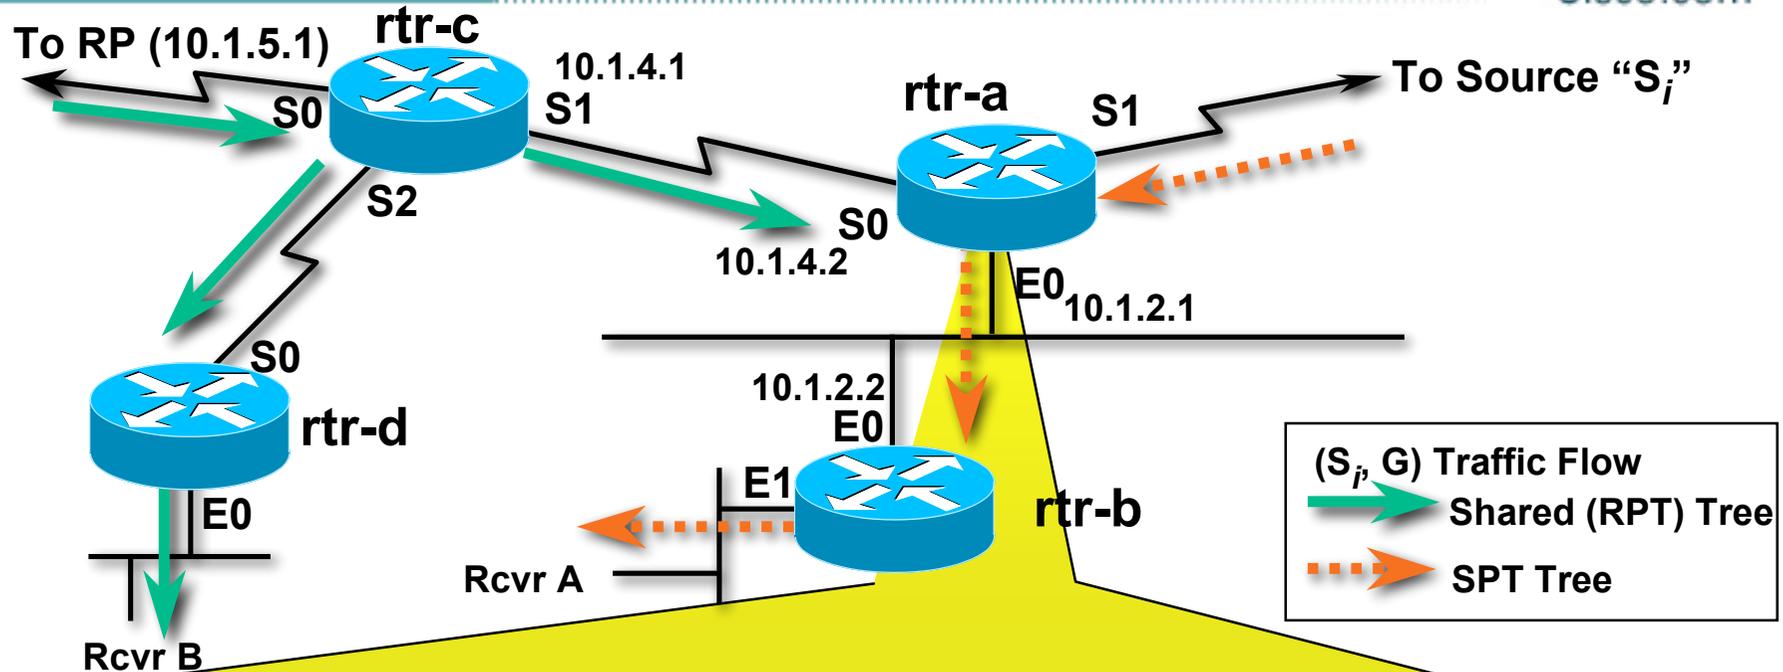
④ "rtr-a" triggers (S_i,G) Join toward S_i.

PIM SM SPT-Switchover



- ④ “rtr-a” forwards (S_i, G) Join toward S_i.
- ⑤ (S_i, G) traffic begins flowing down SPT tree.

PIM SM SPT-Switchover

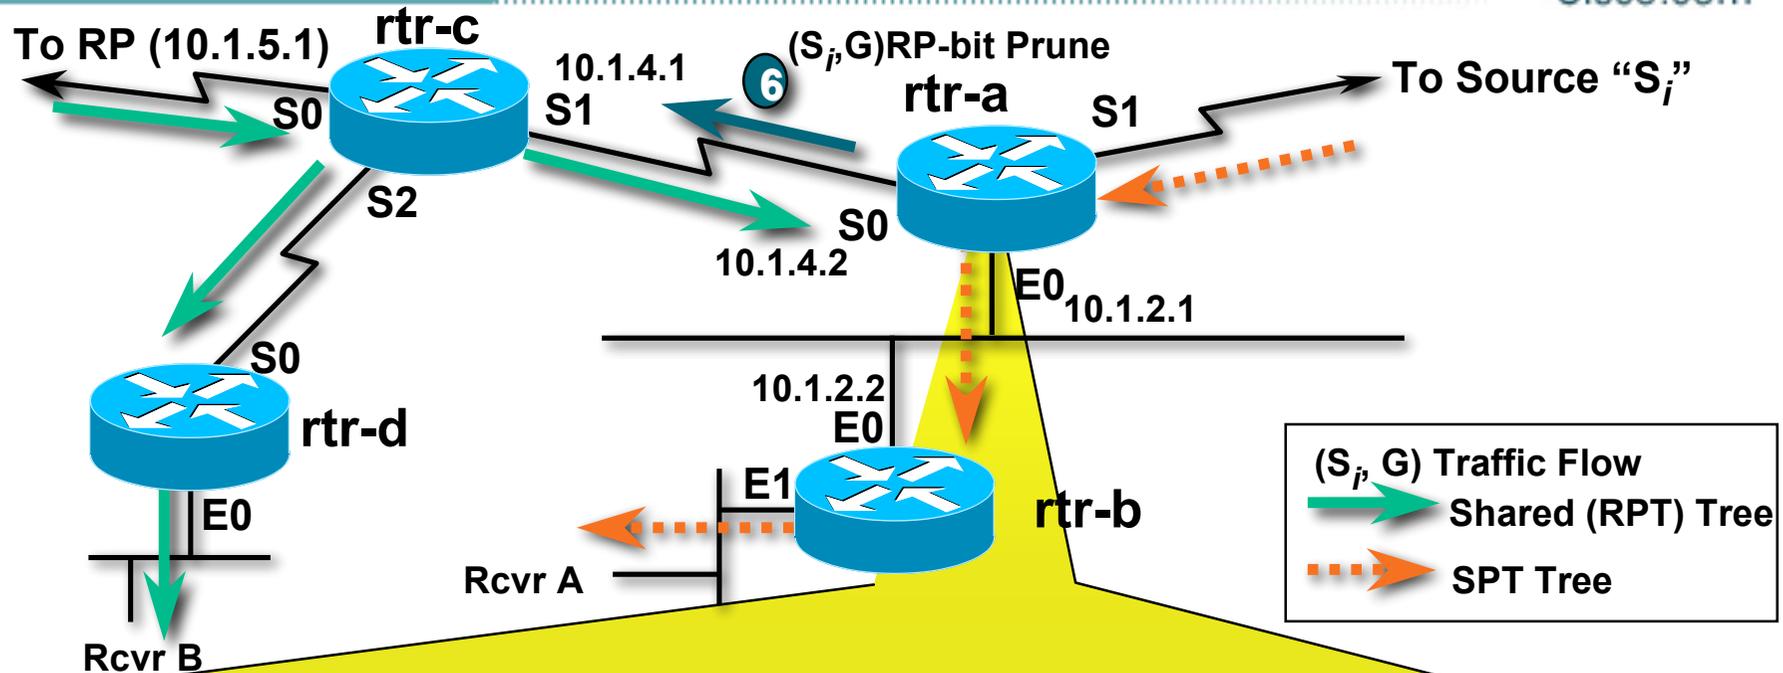


```
(*, 224.1.1.1), 00:01:43/00:02:13, RP 10.1.5.1, flags: S
  Incoming interface: Serial0, RPF nbr 10.1.4.1,
  Outgoing interface list:
    Ethernet0, Forward/Sparse-Dense, 00:01:43/00:02:11

(171.68.37.121/32, 224.1.1.1), 00:13:28/00:02:53, flags:T ←
  Incoming interface: Serial1, RPF nbr 10.1.9.2
  Outgoing interface list:
    Ethernet0, Forward/Sparse-Dense, 00:13:25/00:02:30
```

← "T" Flag set by arriving traffic on SPT

PIM SM SPT-Switchover



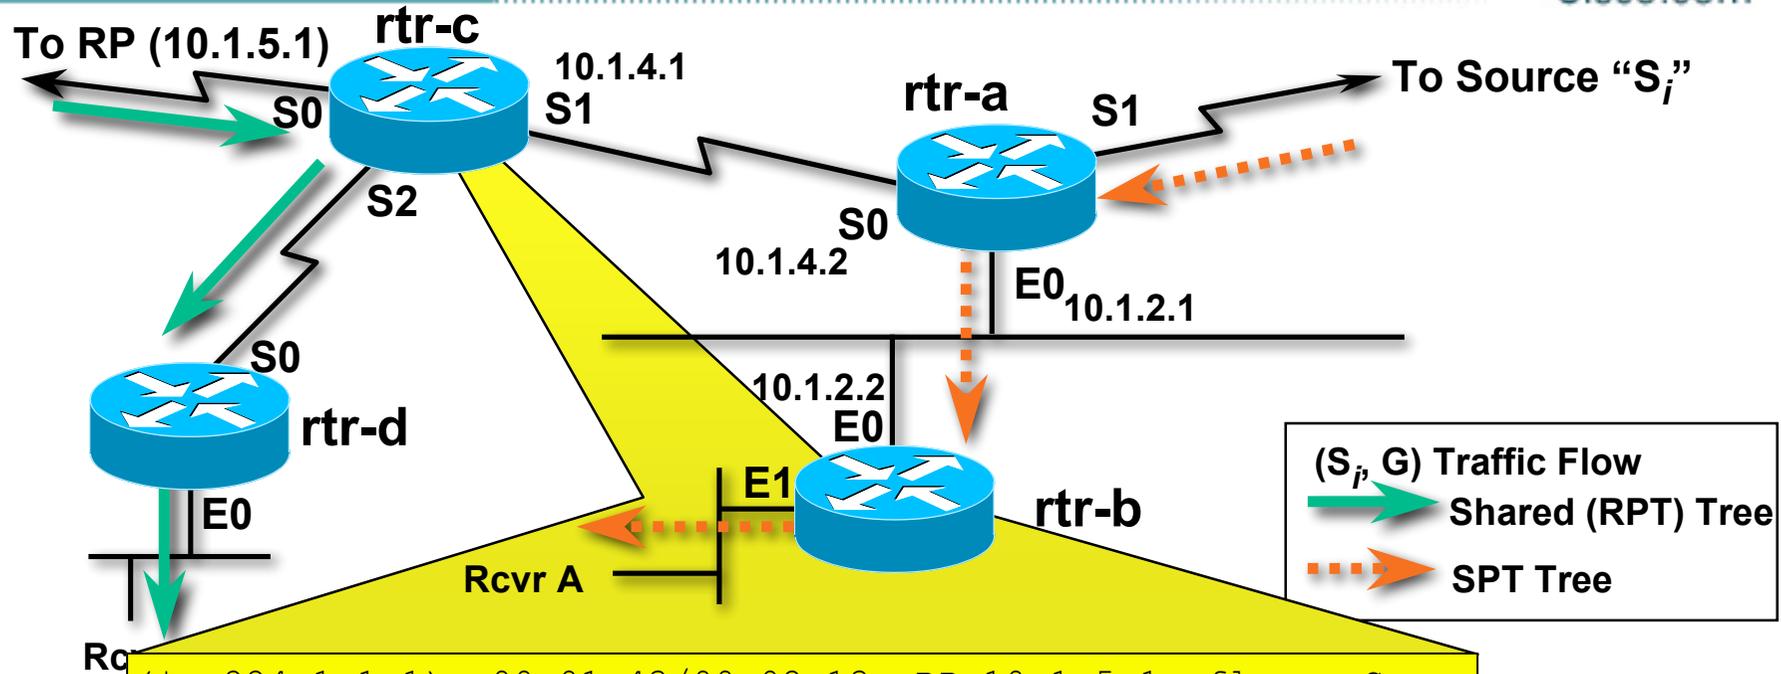
```
(*, 224.1.1.1), 00:01:43/00:02:13, RP 10.1.5.1, flags: S
  Incoming interface: Serial0, RPF nbr 10.1.4.1,
  Outgoing interface list:
    Ethernet0, Forward/Sparse-Dense, 00:01:43/00:02:11

(171.68.37.121/32, 224.1.1.1), 00:13:28/00:02:53, flags:T
  Incoming interface: Serial1, RPF nbr 10.1.9.2
  Outgoing interface list:
    Ethernet0, Forward/Sparse-Dense, 00:13:25/00:02:30
```

Note RPF Info does not match. This indicates SPT & RPT diverge.

6 Once “T” flag is set, rtr-a triggers (S_i, G) RP-bit Prunes toward RP.

PIM SM SPT-Switchover



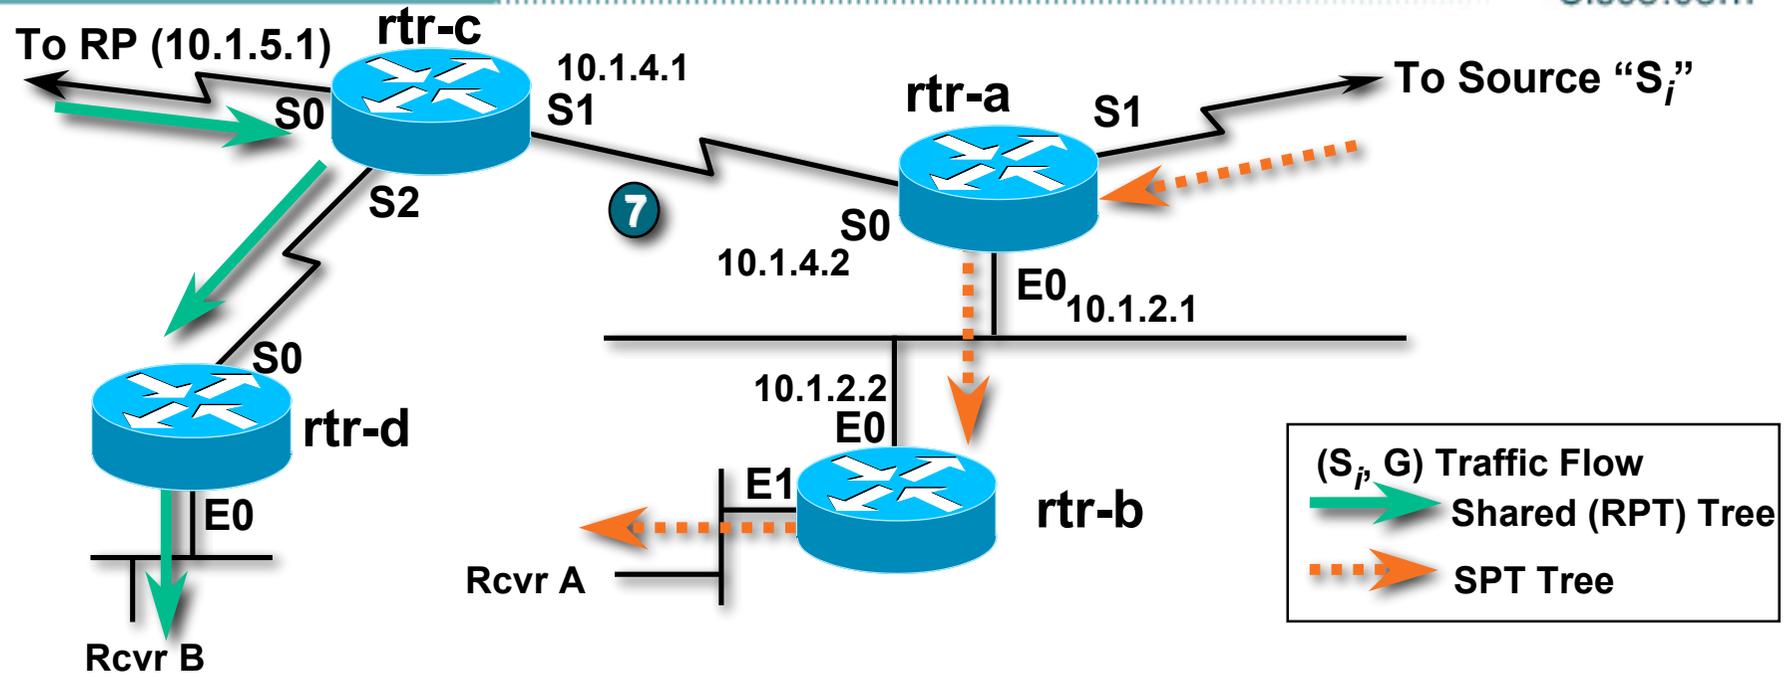
```

Rcvr A
(*, 224.1.1.1), 00:01:43/00:02:13, RP 10.1.5.1, flags: S
  Incoming interface: Serial0, RPF nbr 10.1.5.1,
  Outgoing interface list:
    Serial1, Forward/Sparse-Dense, 00:01:43/00:02:11
    Serial2, Forward/Sparse-Dense, 00:00:32/00:02:28

(171.68.37.121/32, 224.1.1.1), 00:13:28/00:02:53, flags: R
  Incoming interface: Serial0, RPF nbr 10.1.5.1
  Outgoing interface list:
    Serial2, Forward/Sparse-Dense, 00:00:32/00:02:28
    
```

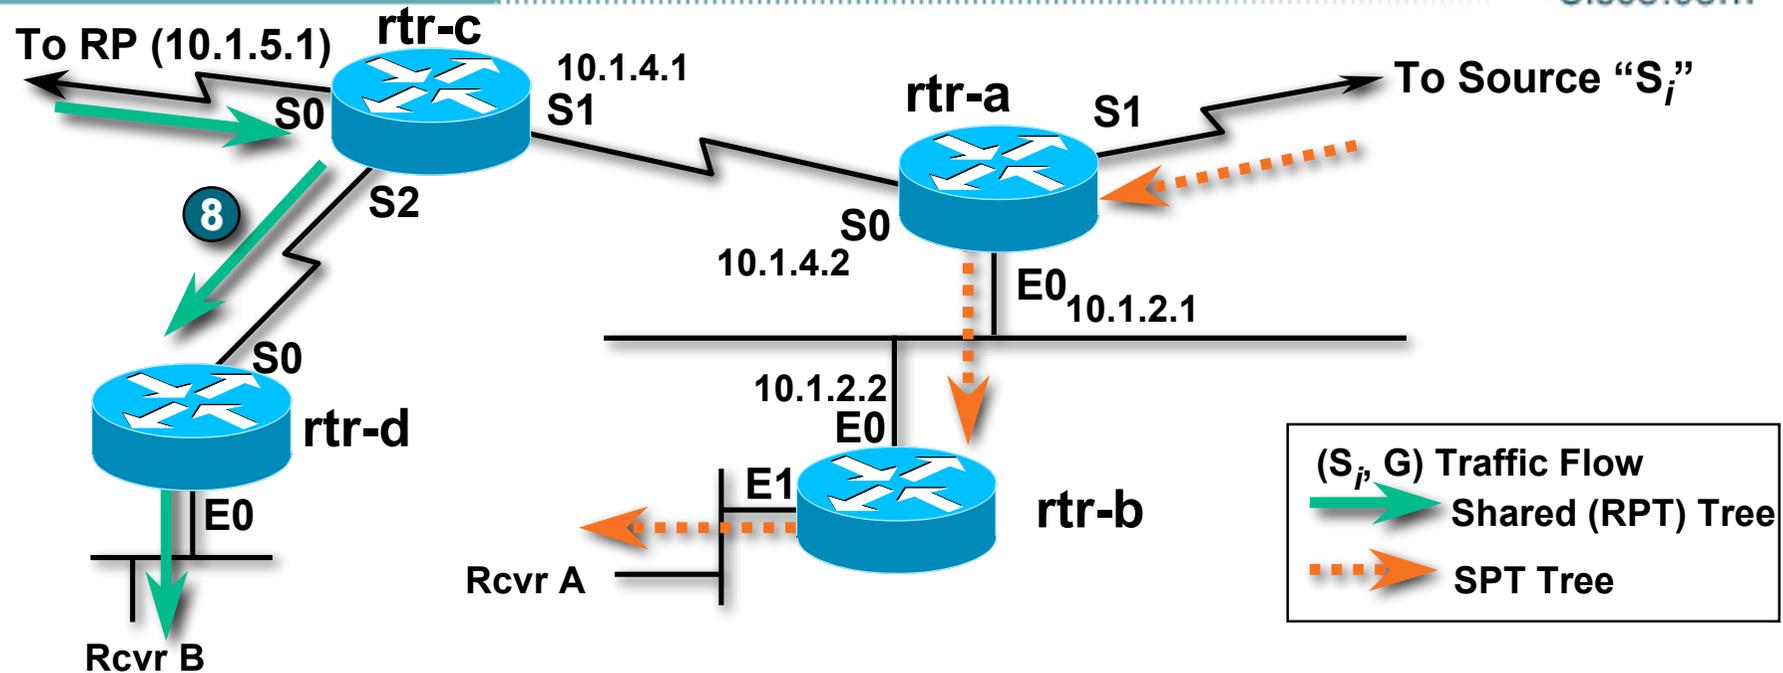
State in "rtr-c" after receiving the (S_i, G) RP-bit Prune

PIM SM SPT-Switchover



7 Unnecessary (S_i, G) traffic is pruned from the Shared tree.

PIM SM SPT-Switchover



- ⑦ Unnecessary (S_i, G) traffic is pruned from the Shared tree.
- ⑧ (S_i, G) traffic still flows via other branches of the Shared tree.

Agenda

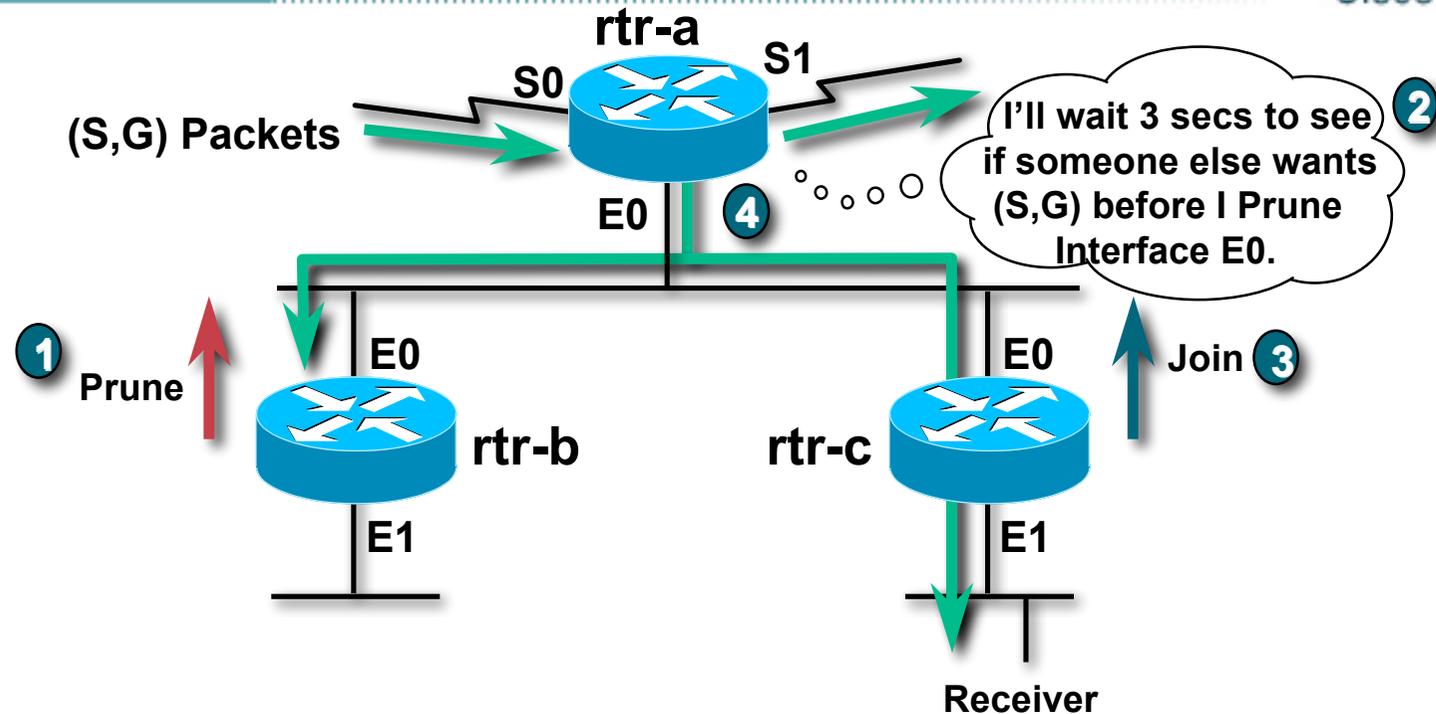
- PIM Packet Formats
- PIM Neighbor Discovery
- PIM State
- PIM SM Joining
- PIM SM Registering
- PIM SM SPT-Switchover
- **PIM SM Pruning**
- PIM Asserts

PIM SM Pruning

- **IGMP group times out / last host sends Leave**
- **Interface removed from all (*,G) & (S,G) entries**
 - **IF OIL of the (*,G) becomes empty (Null);
THEN trigger (*,G) Prune up shared tree toward RP**
 - **IF OIL of the (S,G) becomes empty (Null);
THEN trigger (S,G) Prune up SPT toward Source**

PIM Prune Delay on Multiaccess Networks

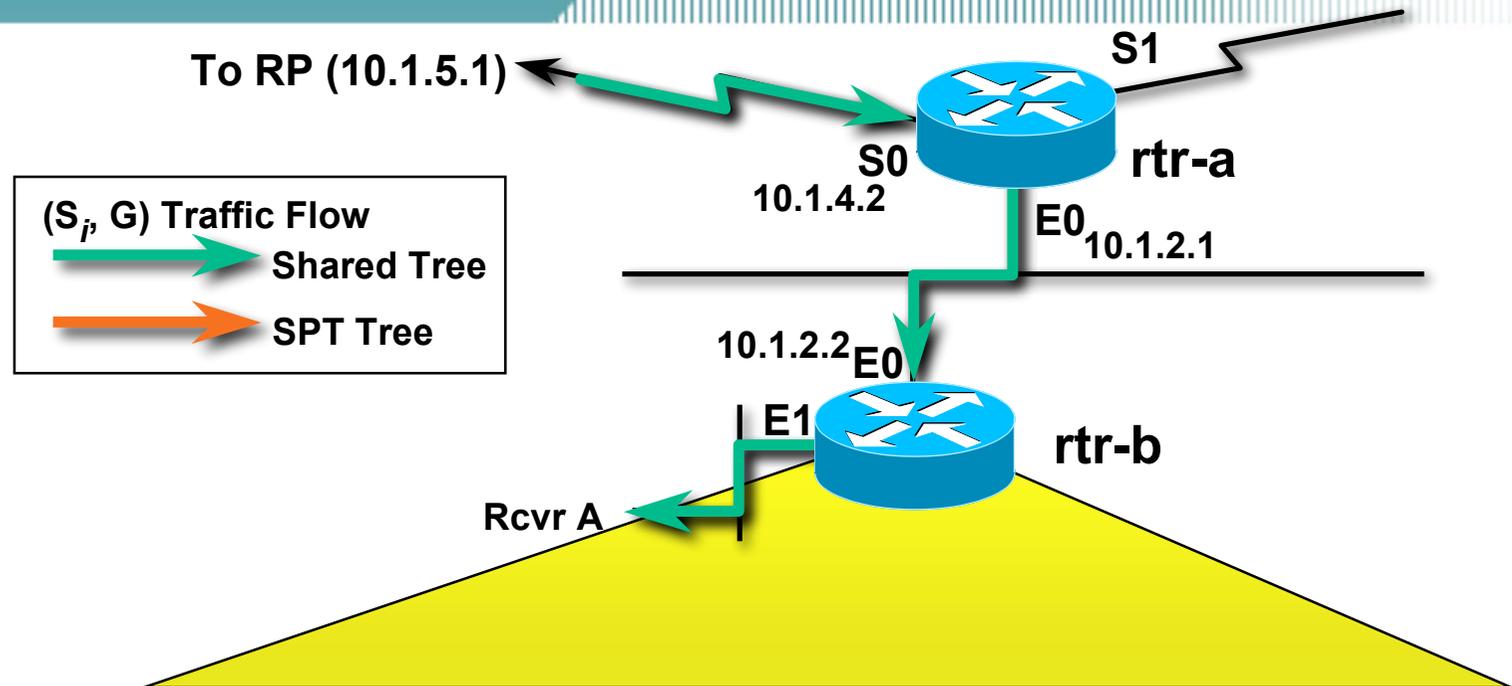
Cisco.com



- 1 “rtr-b” is a leaf node w/o receivers. Sends Prune for (S,G).
- 2 “rtr-a” schedules a Prune for (S,G) to occur in 3 seconds.
- 3 “rtr-c” hears Prune from “rtr-b”. Overrides with a Join.
- 4 “rtr-a” hears Join and cancels Prune for (S,G).

PIM SM Pruning

Shared Tree Case

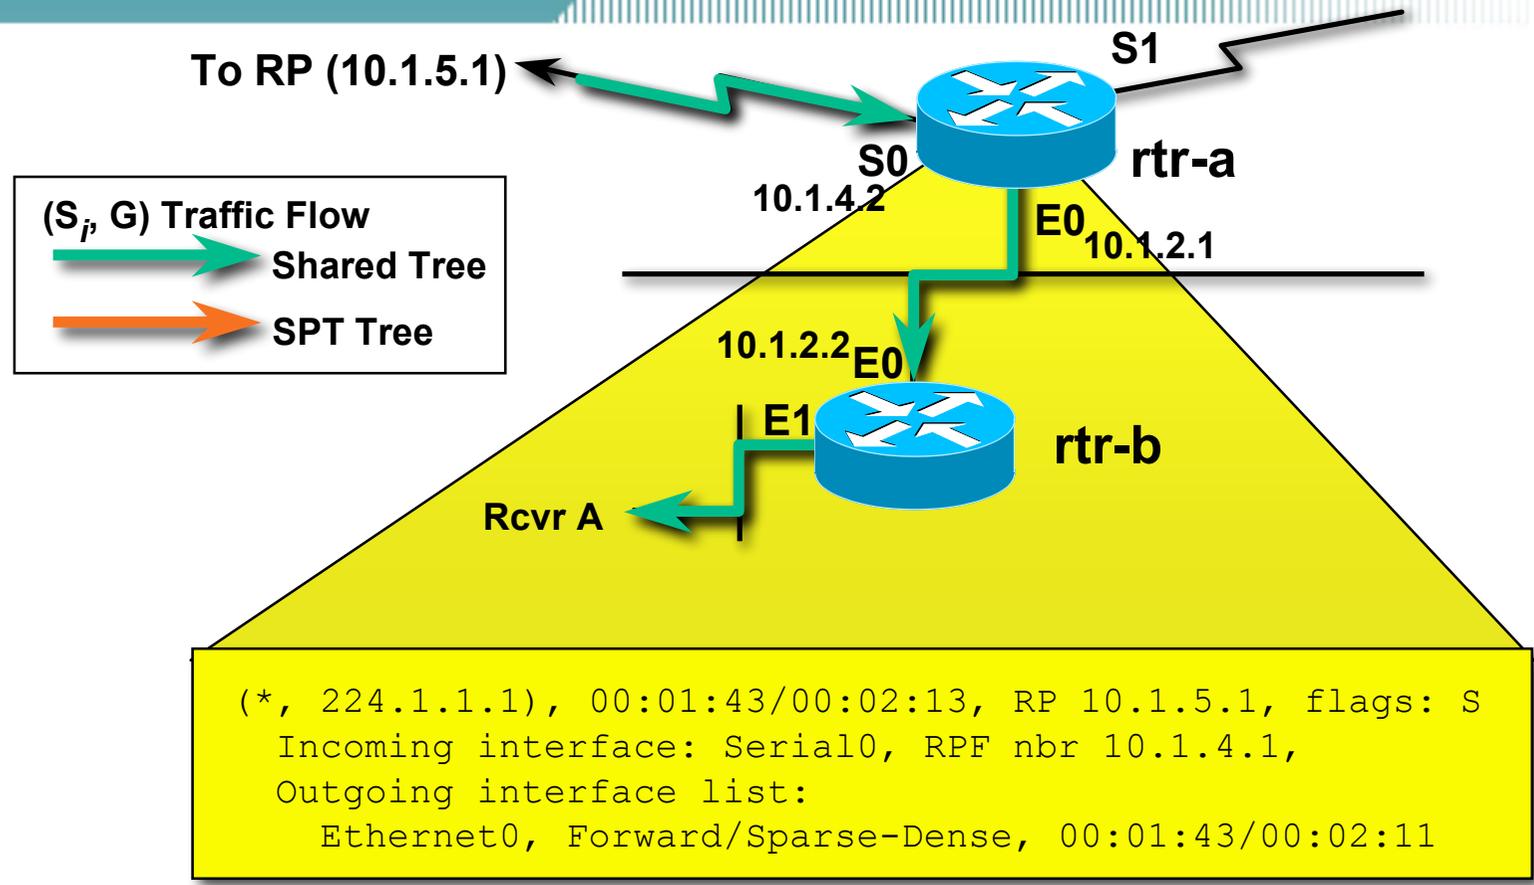


```
(*, 224.1.1.1), 00:01:43/00:02:13, RP 10.1.5.1, flags: SC
Incoming interface: Ethernet0, RPF nbr 10.1.2.1,
Outgoing interface list:
Ethernet1, Forward/Sparse-Dense, 00:01:43/00:02:11
```

State in "rtr-b" before Pruning

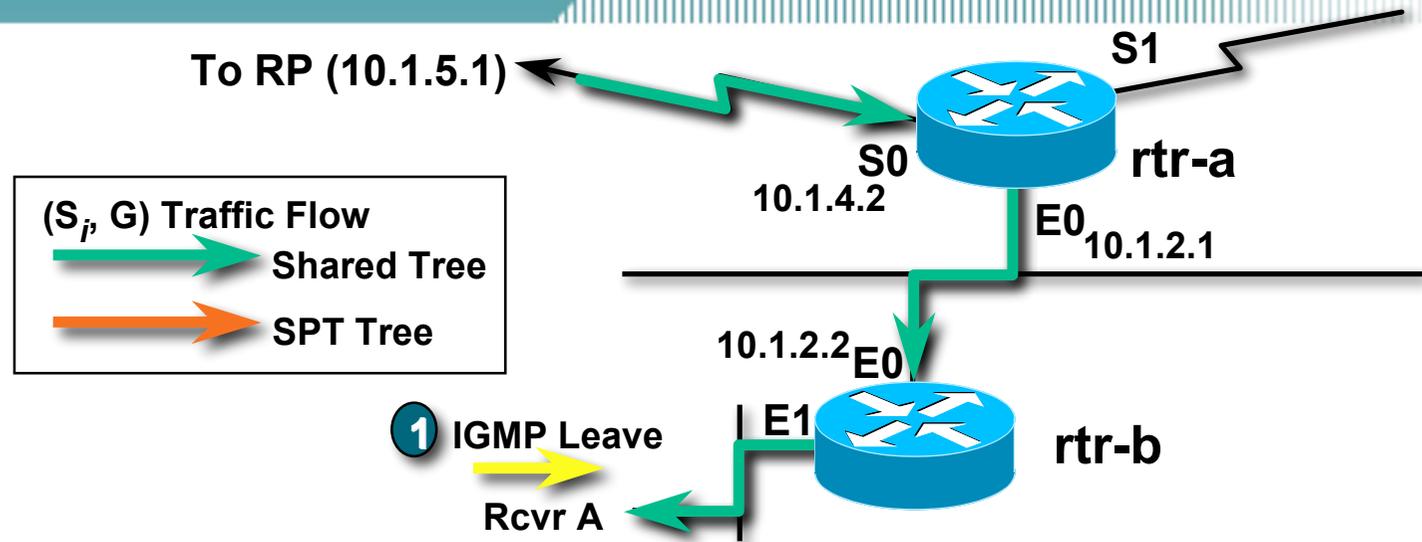
PIM SM Pruning

Shared Tree Case



PIM SM Pruning

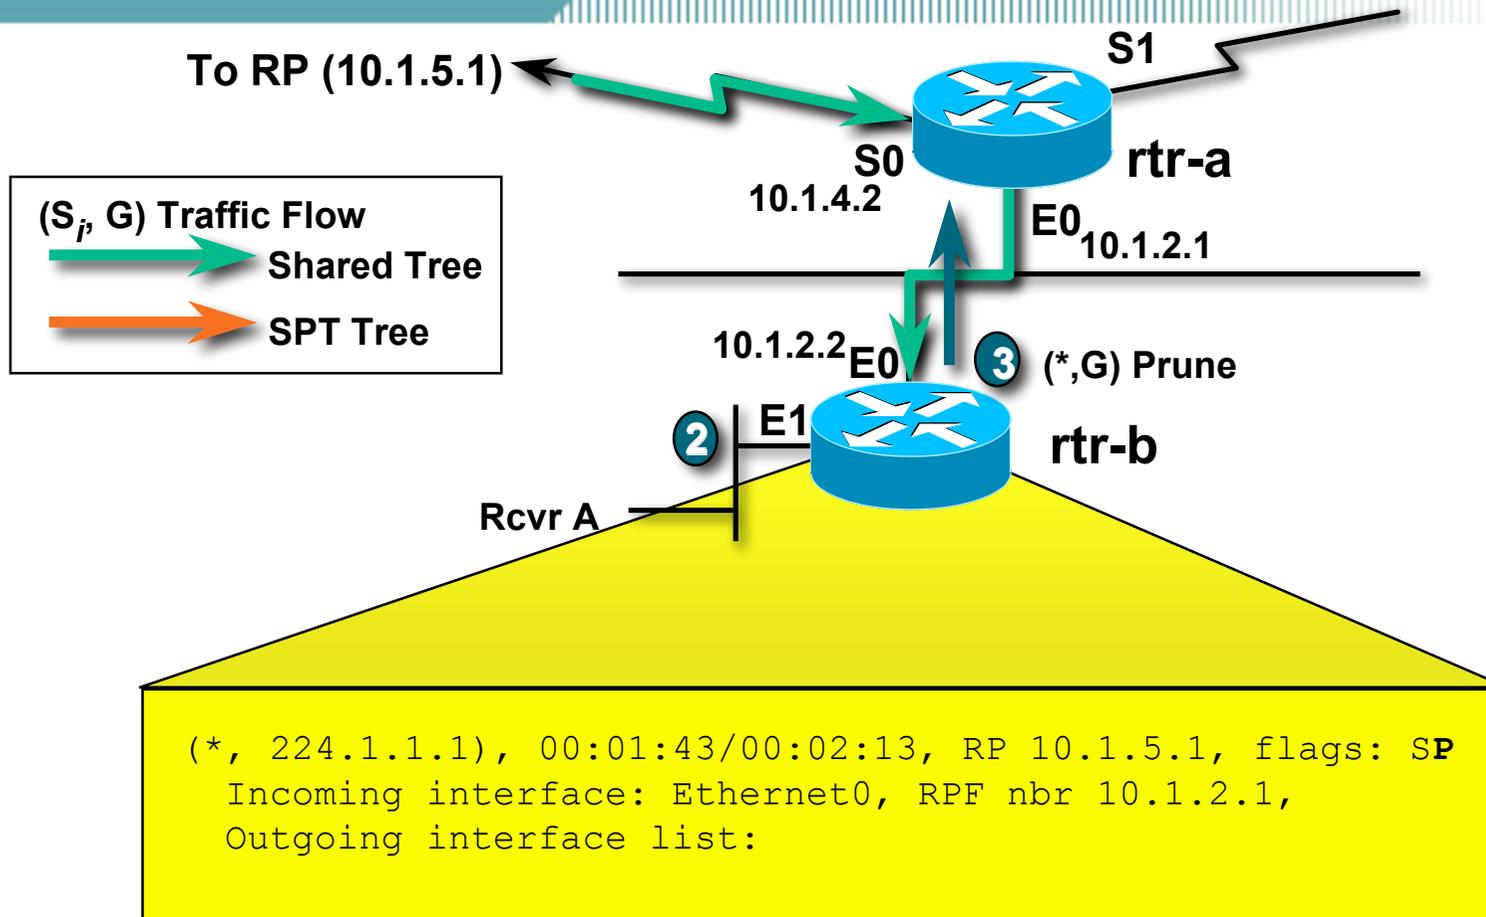
Shared Tree Case



1 “rtr-b” is a Leaf router. Last host “Rcvr A”, leaves group G.

PIM SM Pruning

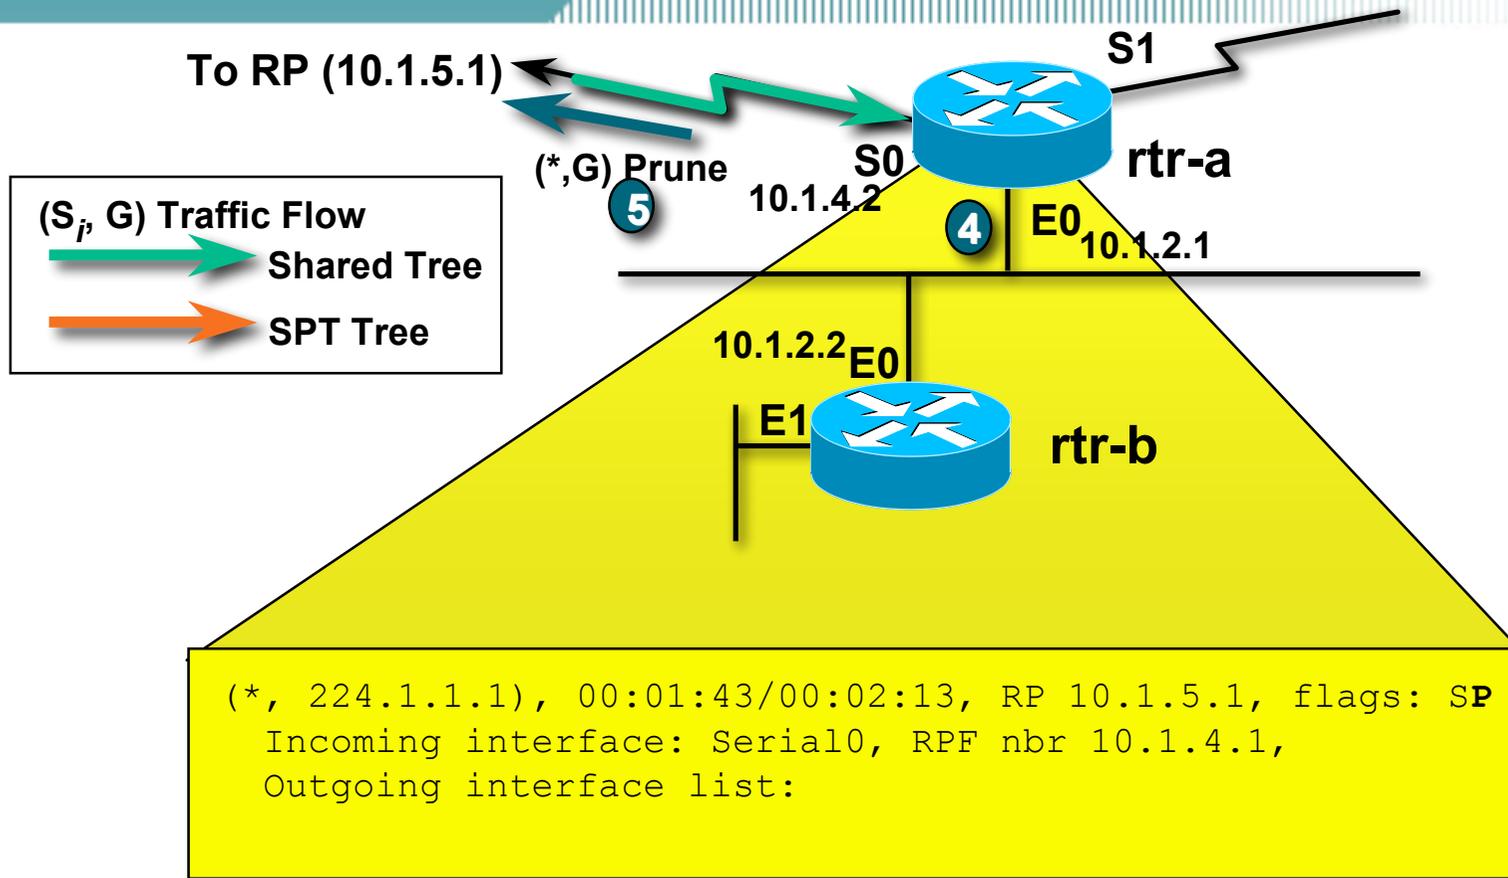
Shared Tree Case



- 2** “rtr-b” removes E1 from (*,G) and any (S_i,G) “oilists”.
- 3** “rtr-b” (*,G) “oilist” now empty; triggers (*,G) Prune toward RP.

PIM SM Pruning

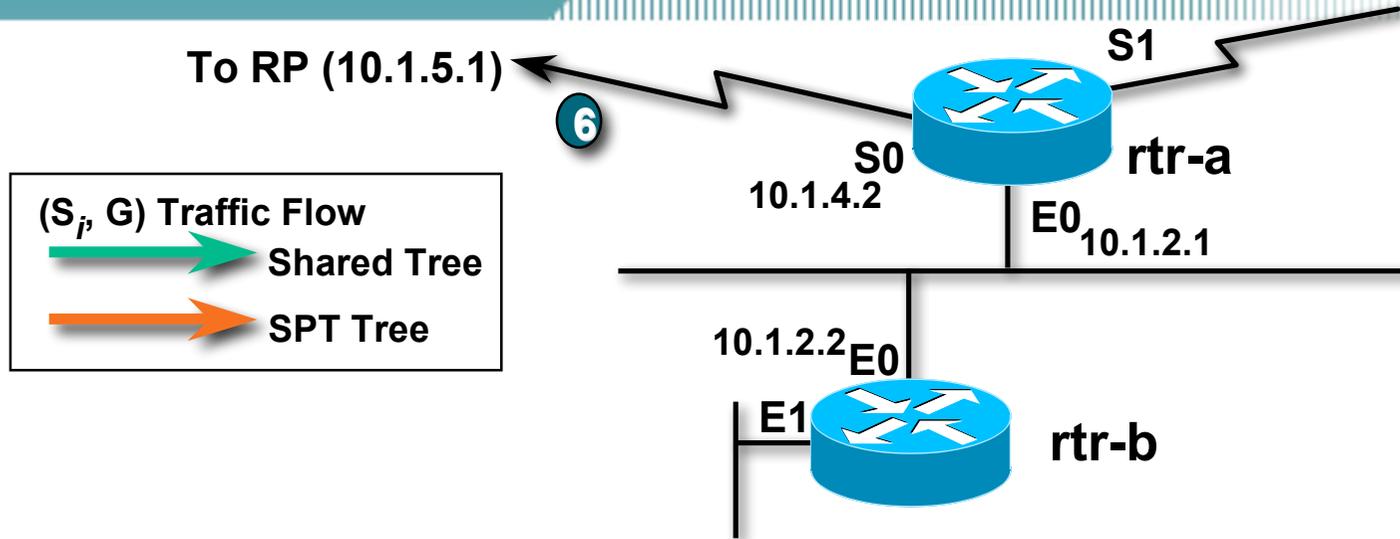
Shared Tree Case



- ④ “rtr-a” receives Prune; removes E0 from (*,G) “oilist”.
(After the 3 second Multi-access Network Prune delay.)
- ⑤ “rtr-a” (*,G) “oilist” now empty; triggers (*,G) Prune toward RP.

PIM SM Pruning

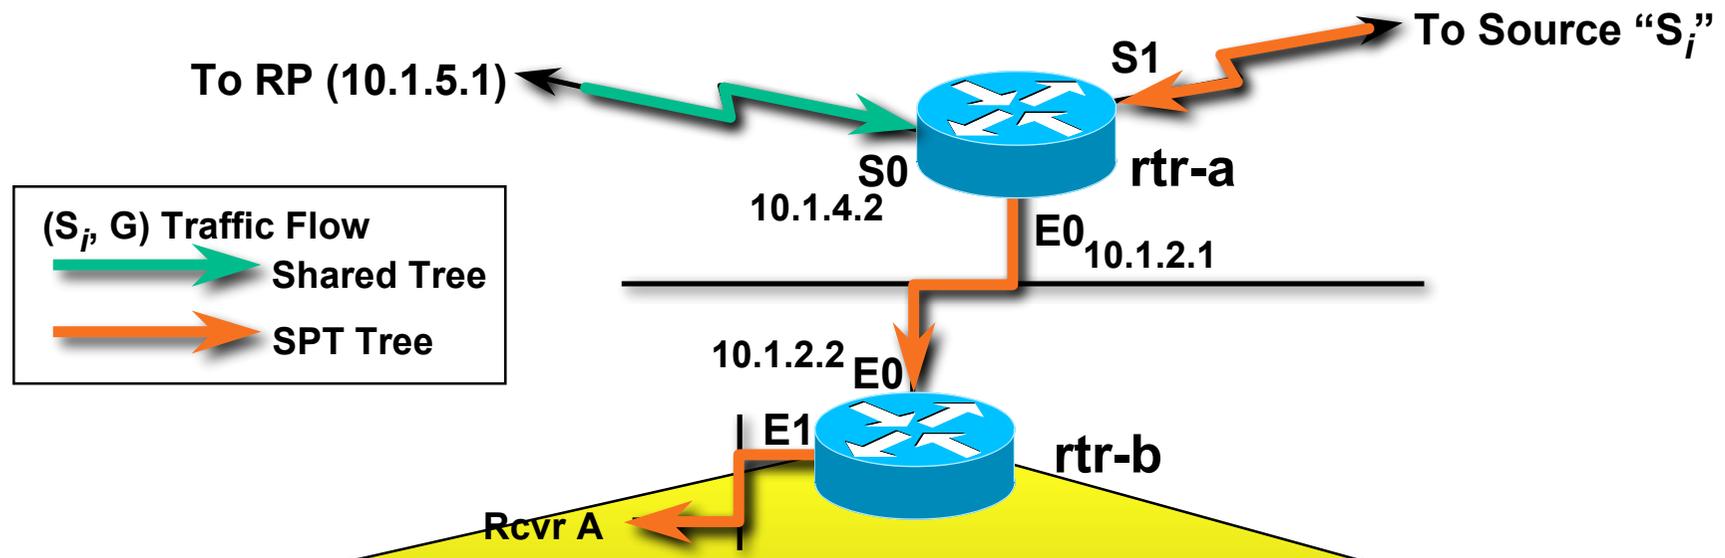
Shared Tree Case



6 Pruning continues back toward RP.

PIM SM Pruning

Source (SPT) Case



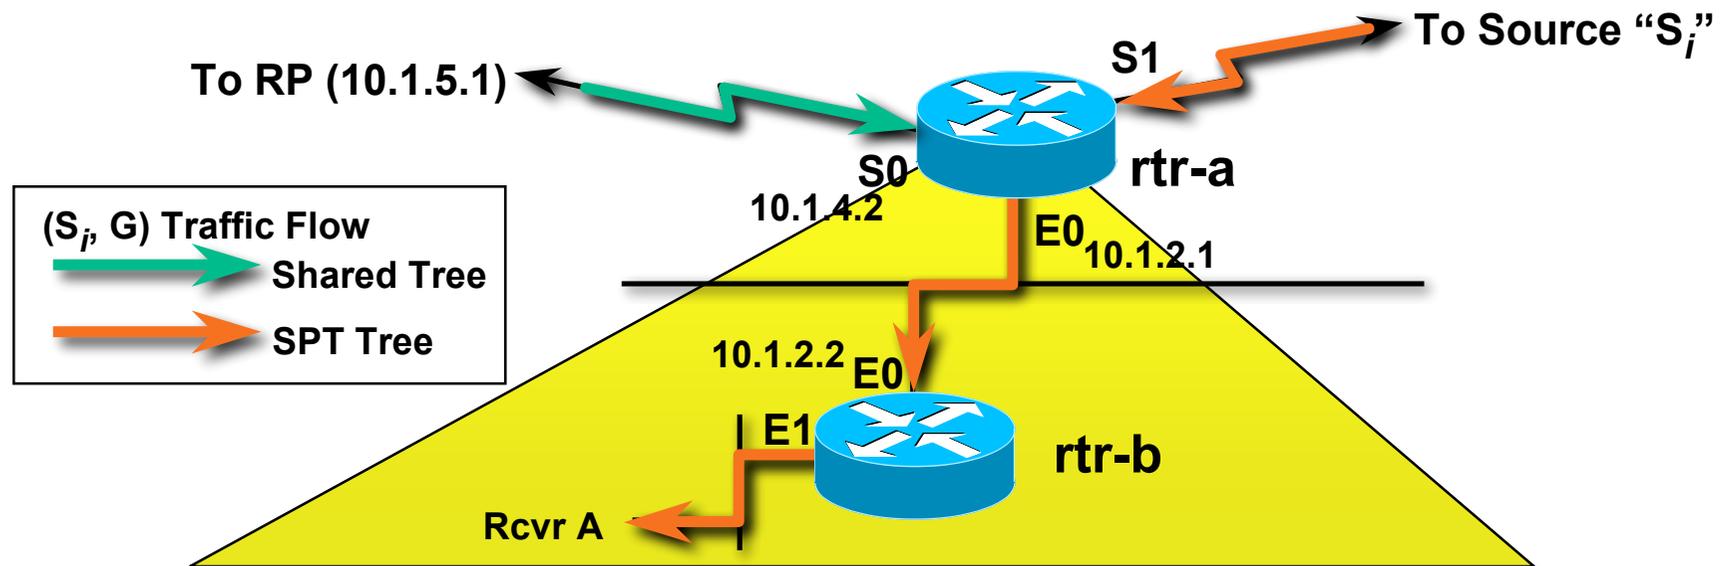
```
(*, 224.1.1.1), 00:01:43/00:02:59, RP 10.1.5.1, flags: SC
Incoming interface: Ethernet0, RPF nbr 10.1.2.1,
Outgoing interface list:
  Ethernet1, Forward/Sparse-Dense, 00:01:43/00:02:11

(171.68.37.121/32, 224.1.1.1), 00:01:05/00:01:55, flags: CJT
Incoming interface: Ethernet0, RPF nbr 10.1.2.1
Outgoing interface list:
  Ethernet1, Forward/Sparse-Dense, 00:01:05/00:02:55
```

State in "rtr-b" before Pruning

PIM SM Pruning

Source (SPT) Case



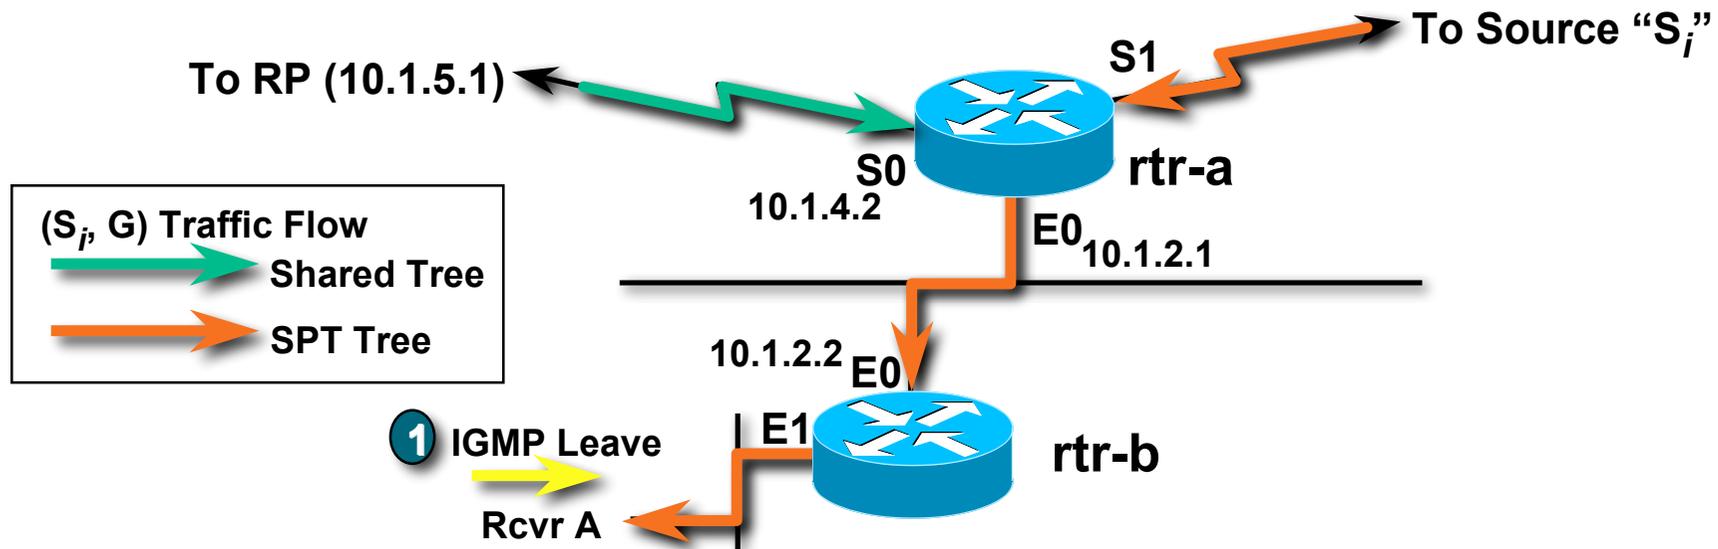
```
(*, 224.1.1.1), 00:01:43/00:02:59, RP 10.1.5.1, flags: S
Incoming interface: Serial0, RPF nbr 10.1.4.1,
Outgoing interface list:
  Ethernet0, Forward/Sparse-Dense, 00:01:43/00:02:11

(171.68.37.121/32, 224.1.1.1), 00:01:05/00:01:55, flags: T
Incoming interface: Serial1, RPF nbr 10.1.9.2
Outgoing interface list:
  Ethernet0, Forward/Sparse-Dense, 00:01:05/00:02:55
```

State in "rtr-a" before Pruning

PIM SM Pruning

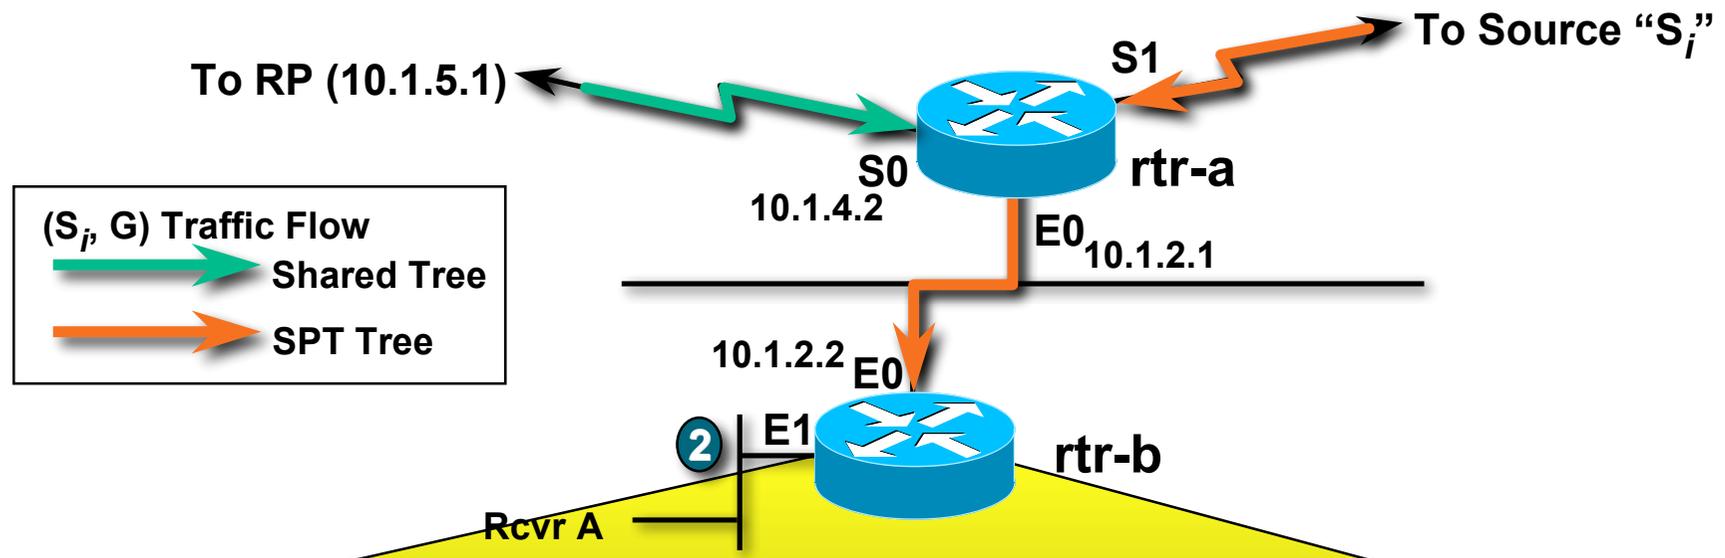
Source (SPT) Case



1 "rtr-b" is a Leaf router. Last host "Rcvr A", leaves group G.

PIM SM Pruning

Source (SPT) Case



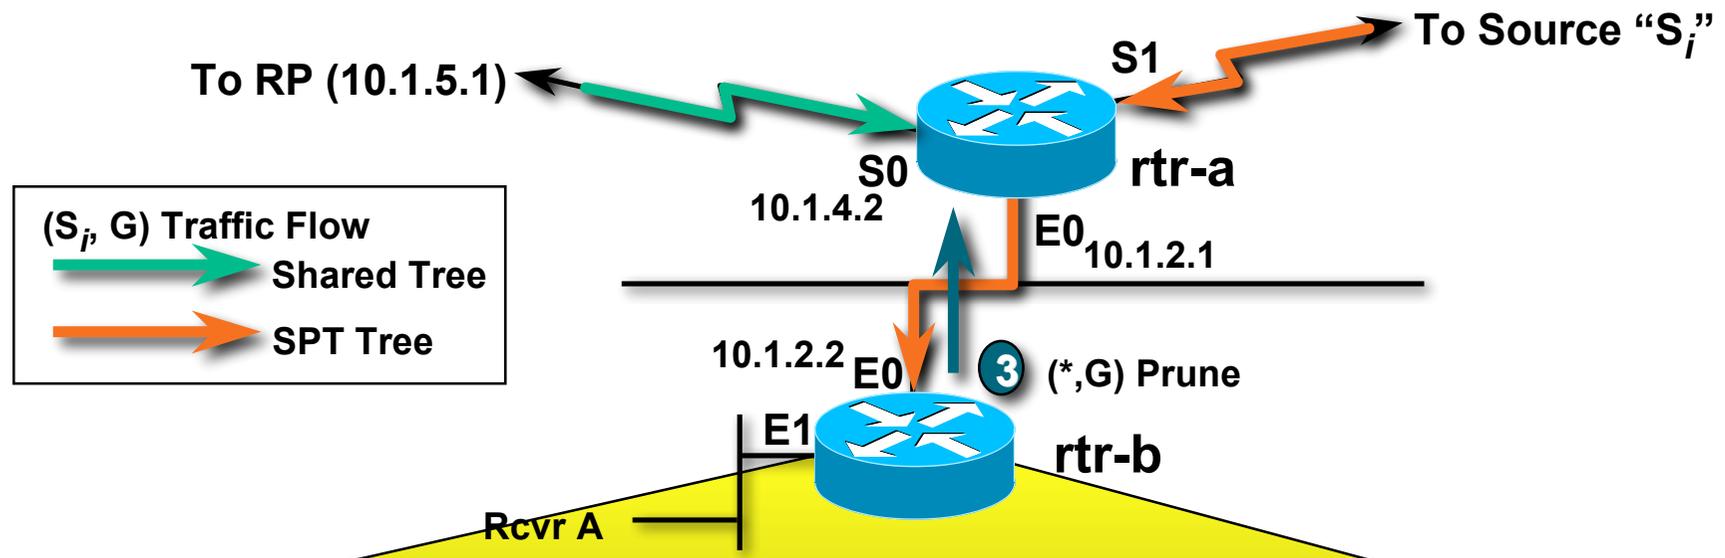
```
(*, 224.1.1.1), 00:01:43/00:02:59, RP 10.1.5.1, flags: SP
Incoming interface: Ethernet0, RPF nbr 10.1.2.1,
Outgoing interface list:

(171.68.37.121/32, 224.1.1.1), 00:01:05/00:01:55, flags: CJPT
Incoming interface: Ethernet0, RPF nbr 10.1.2.1
Outgoing interface list:
```

② "rtr-b" removes E1 from (*,G) and all (S,G) OIL's.

PIM SM Pruning

Source (SPT) Case



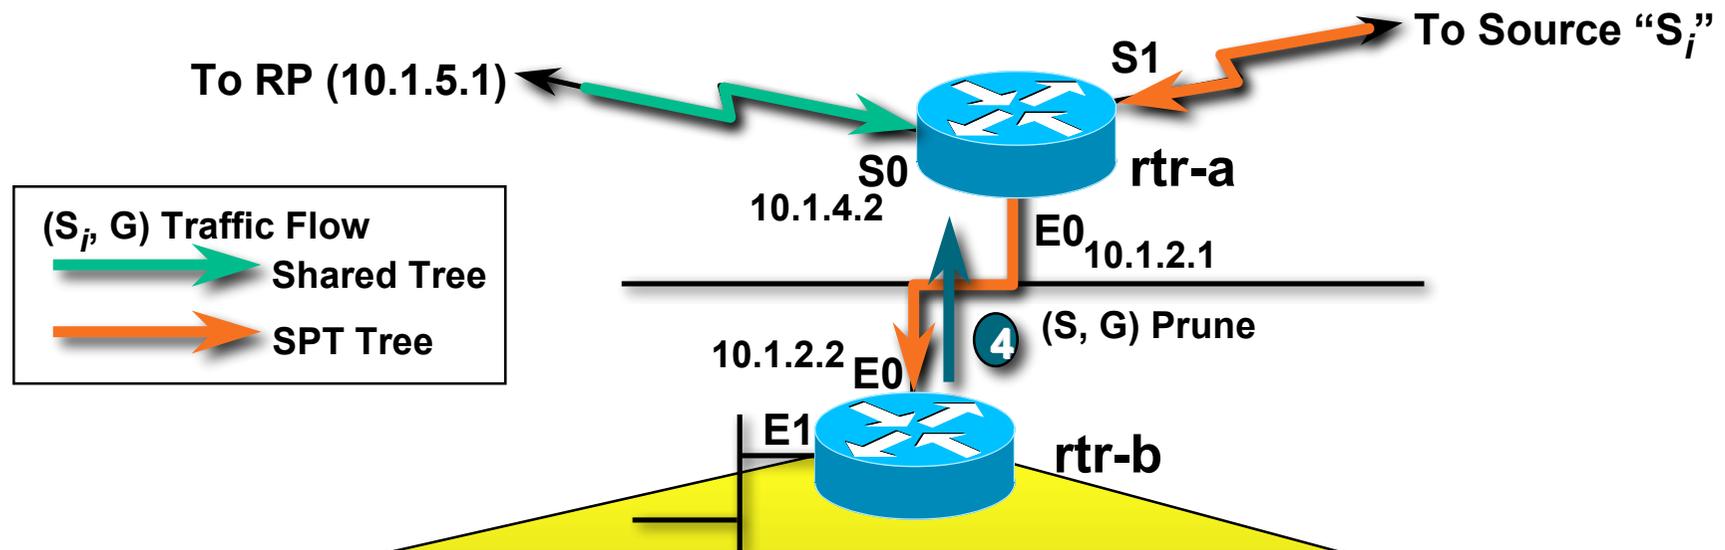
```
(*, 224.1.1.1), 00:01:43/00:02:59, RP 10.1.5.1, flags: SP
Incoming interface: Ethernet0, RPF nbr 10.1.2.1,
Outgoing interface list:

(171.68.37.121/32, 224.1.1.1), 00:01:05/00:01:55, flags: CJPT
Incoming interface: Ethernet0, RPF nbr 10.1.2.1
Outgoing interface list:
```

3 “rtr-b” (*,G) OIL now empty; triggers (*,G) Prune toward RP.

PIM SM Pruning

Source (SPT) Case



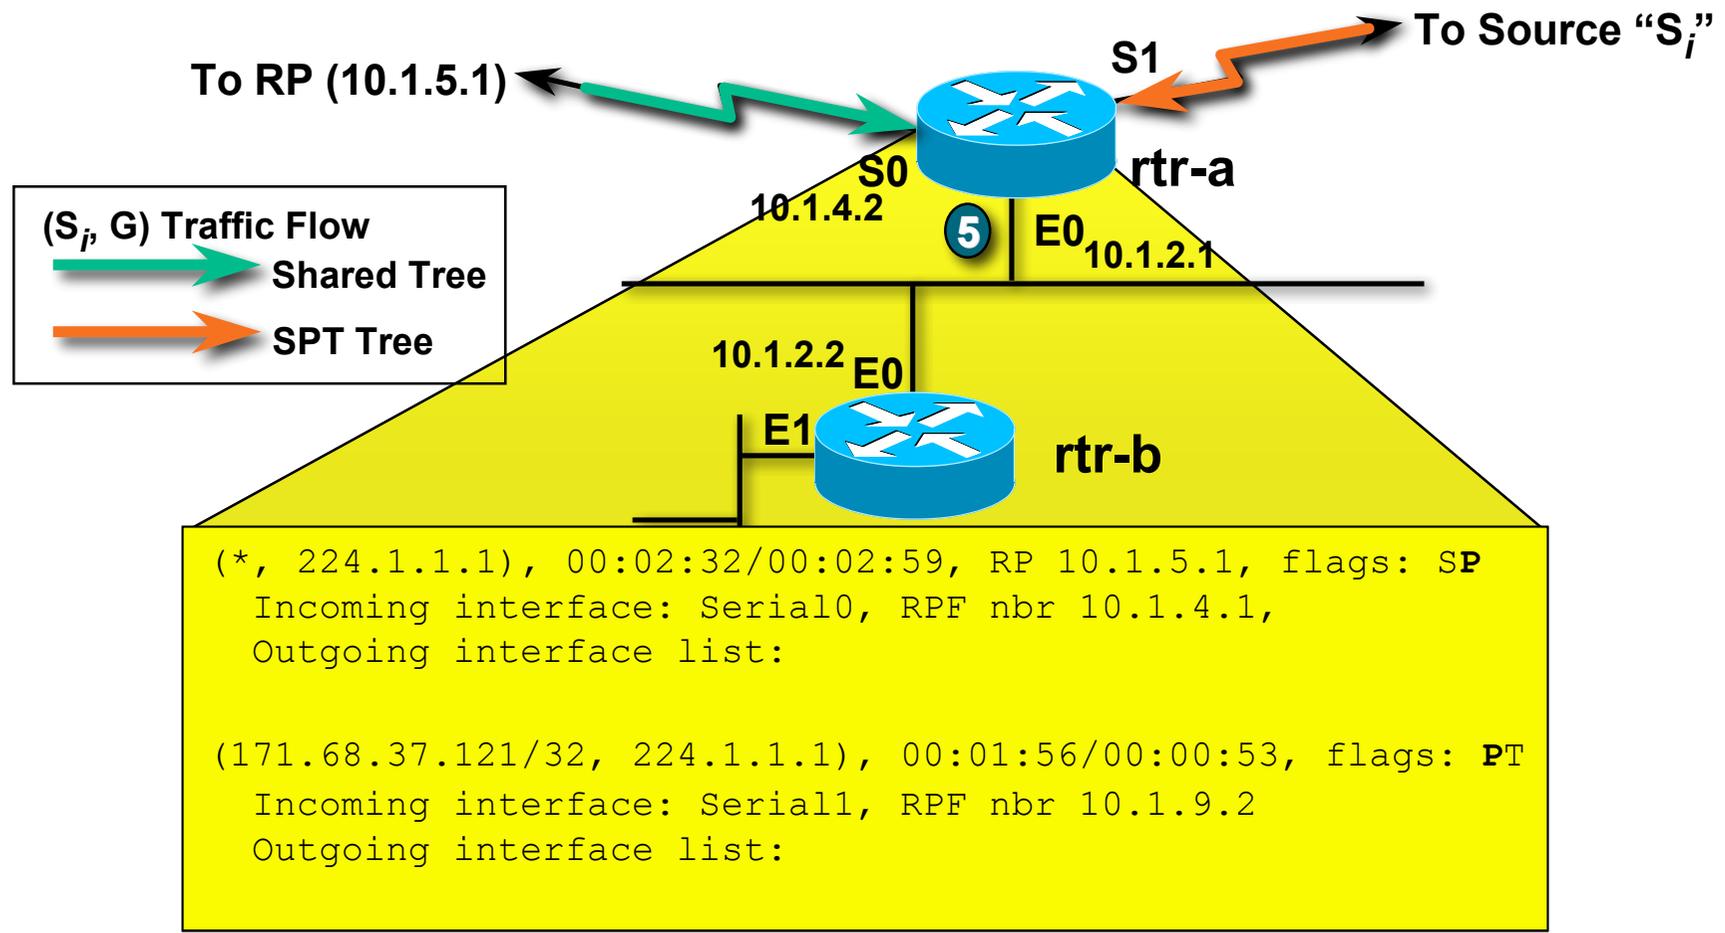
```
(*, 224.1.1.1), 00:01:43/00:02:59, RP 10.1.5.1, flags: SP
Incoming interface: Ethernet0, RPF nbr 10.1.2.1,
Outgoing interface list:

(171.68.37.121/32, 224.1.1.1), 00:01:05/00:01:55, flags: CJPT
Incoming interface: Ethernet0, RPF nbr 10.1.2.1
Outgoing interface list:
```

4 "rtr-b" (S,G) OIL also now empty; triggers (S, G) Prune towards S_i.

PIM SM Pruning

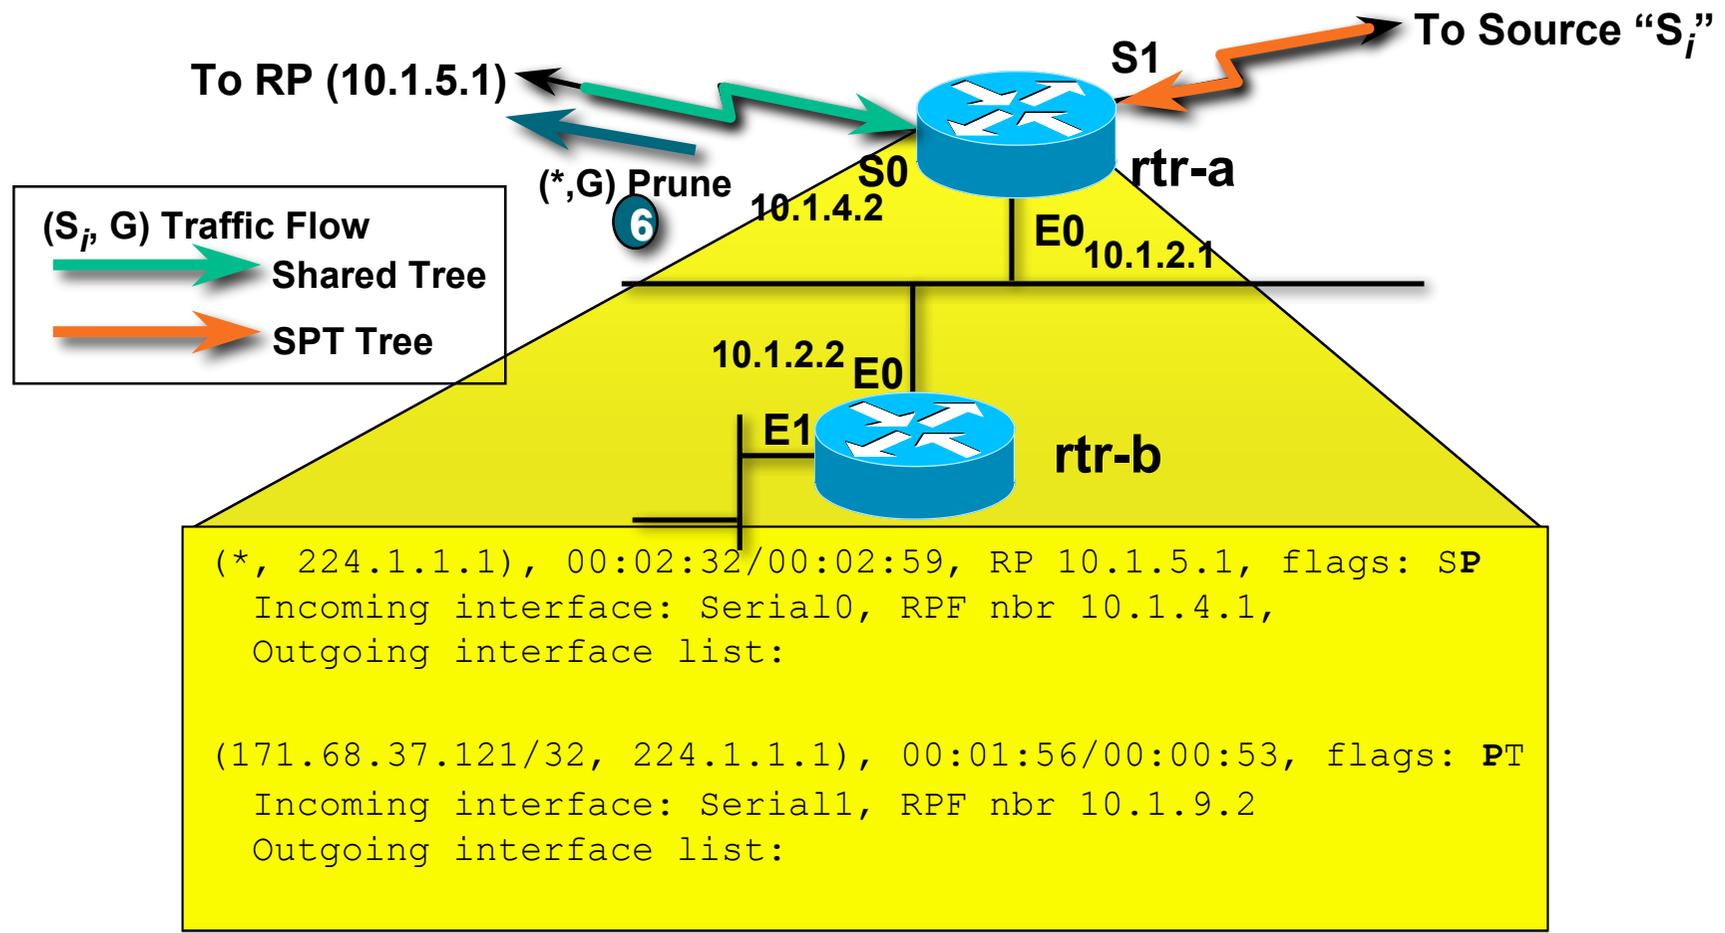
Source (SPT) Case



5 “rtr-a” receives (*, G) Prune; removes E0 from (*,G) & (S,G) OIL’s
 (After the 3 second Multi-access Network Prune delay.)

PIM SM Pruning

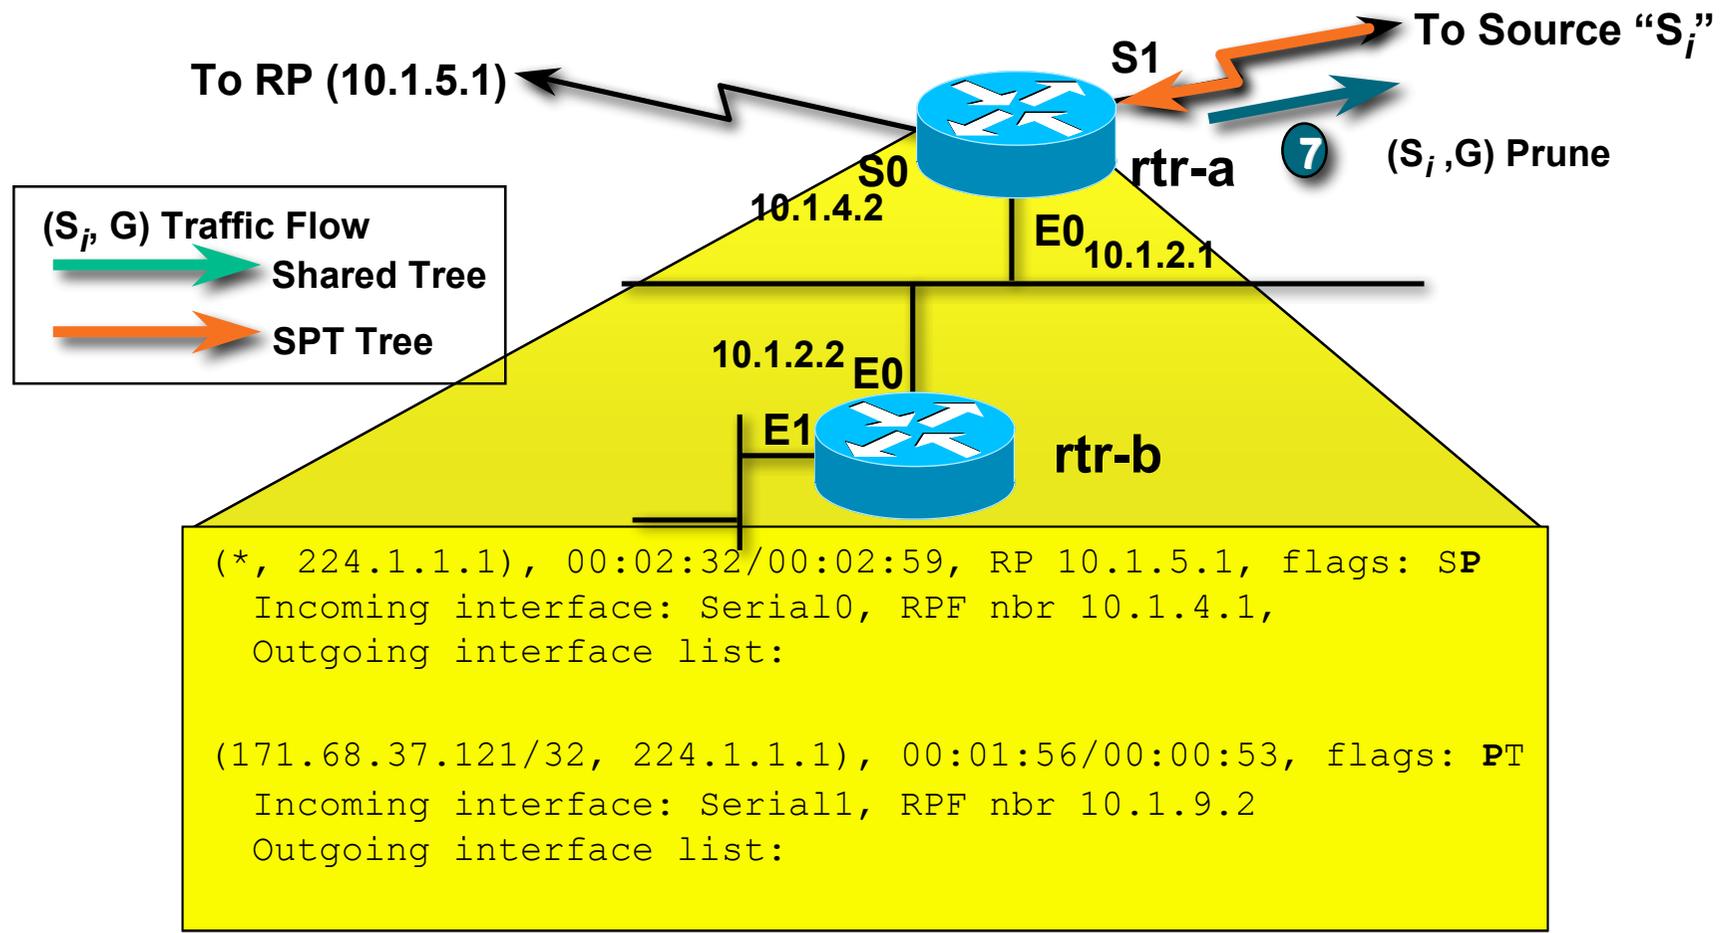
Source (SPT) Case



6 "rtr-a" (*,G) OIL now empty; triggers (*,G) Prune toward RP.

PIM SM Pruning

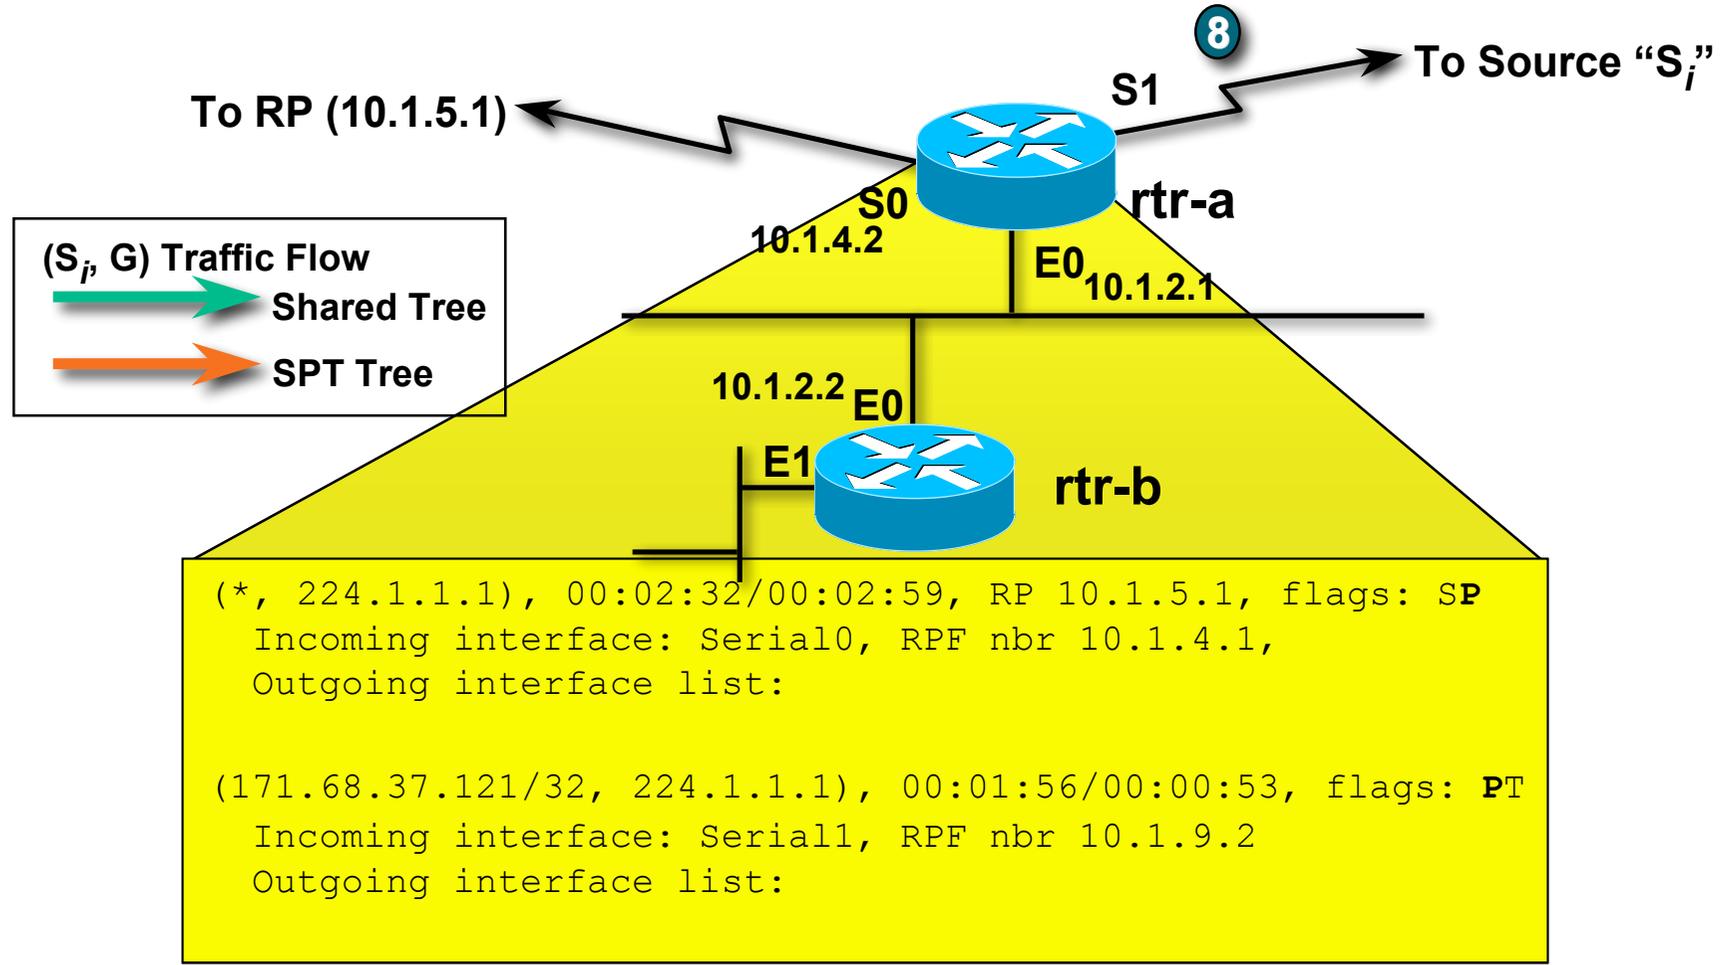
Source (SPT) Case



7 “rtr-a” (S,G) OIL also now empty; triggers (S,G) Prune towards S_i.

PIM SM Pruning

Source (SPT) Case

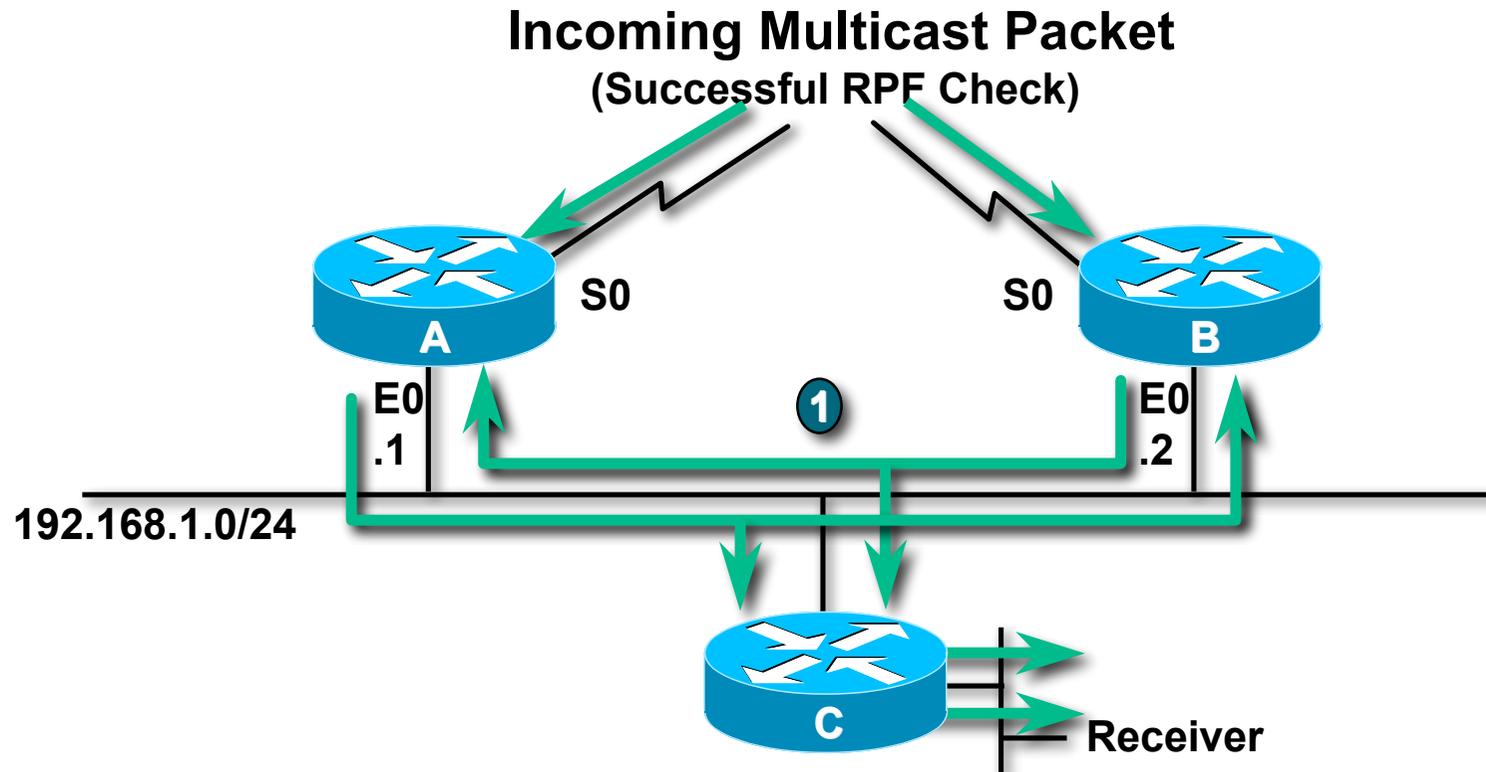


8 (S_i,G) traffic ceases flowing down SPT.

Agenda

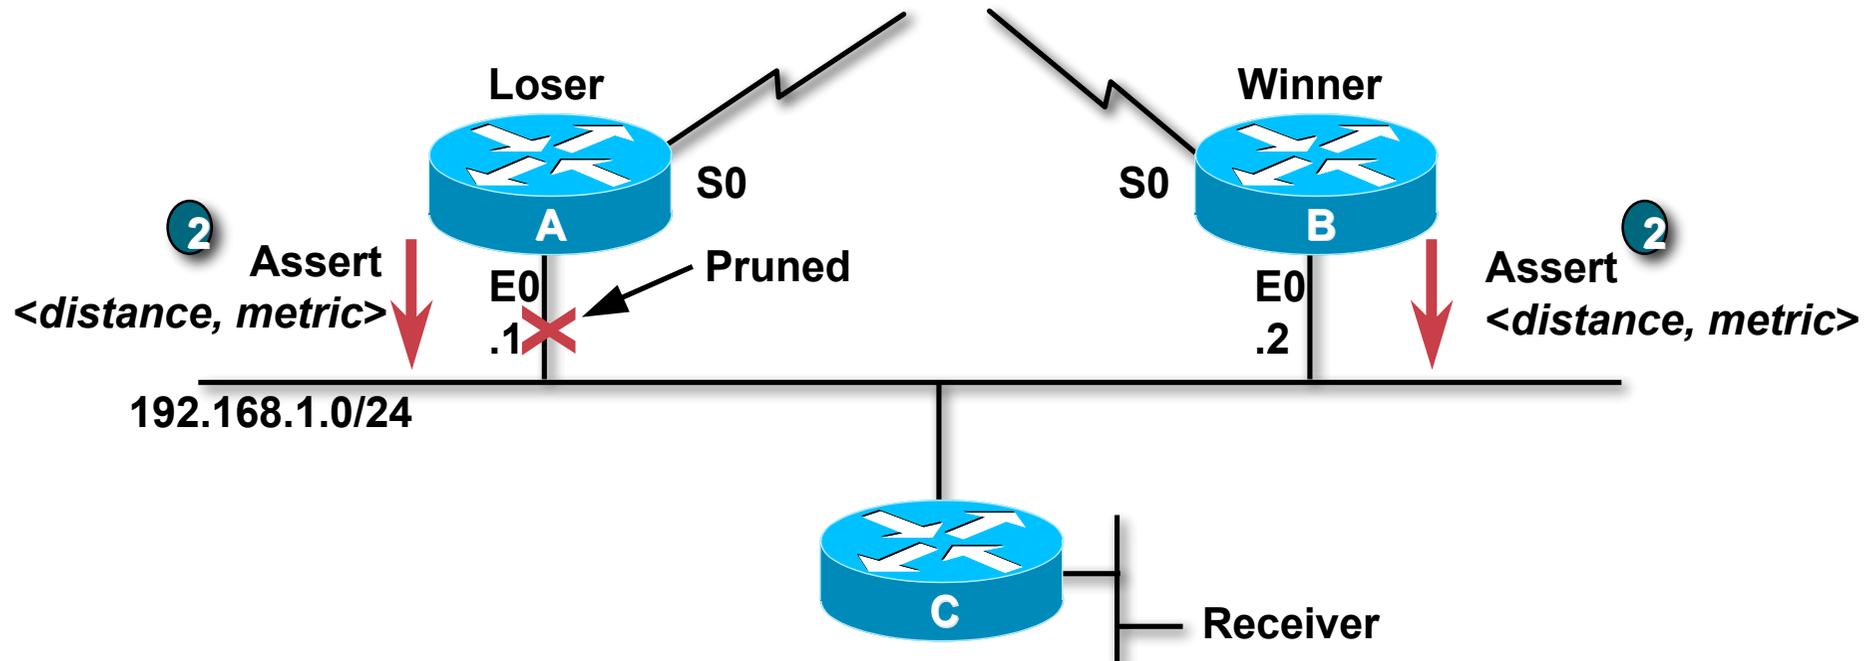
- **PIM Packet Formats**
- **PIM Neighbor Discovery**
- **PIM State**
- **PIM SM Joining**
- **PIM SM Registering**
- **PIM SM SPT-Switchover**
- **PIM SM Pruning**
- **PIM Asserts**

PIM Assert Mechanism



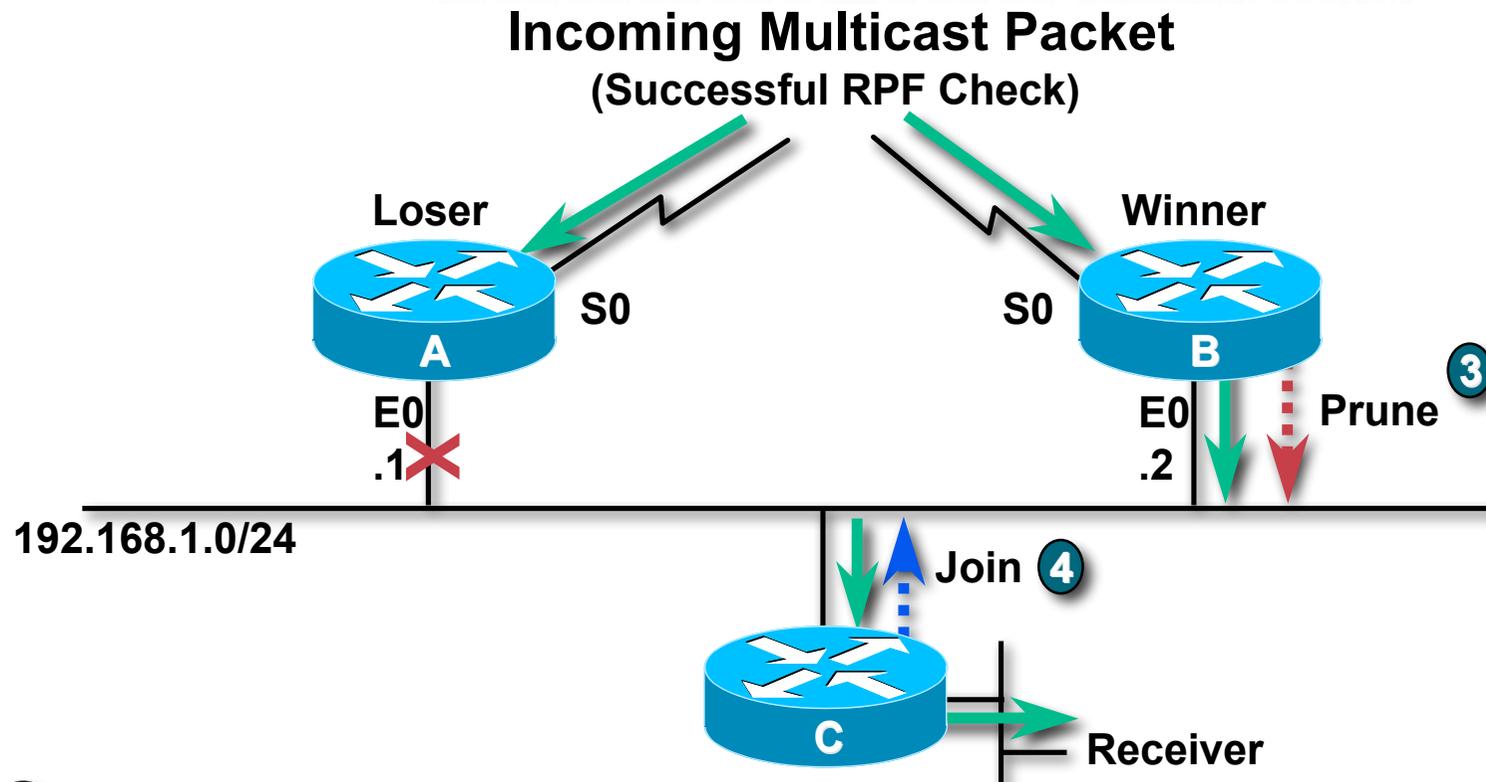
- 1 Routers A & B **receive** packet on an interface in their “*oilist*”!!
 - Only one router should continue sending to avoid duplicate packets.

PIM Assert Mechanism



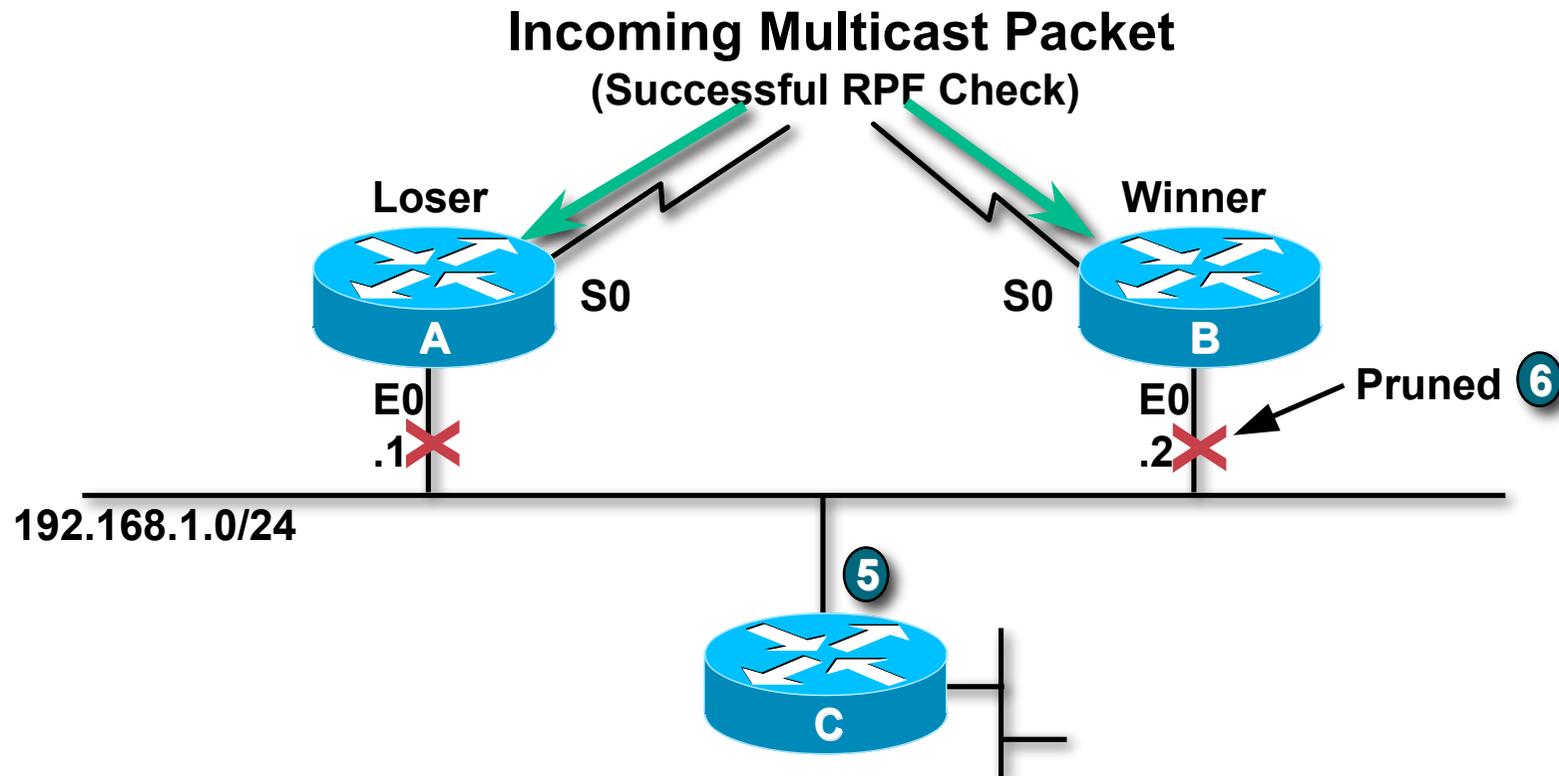
- 2 Routers send "PIM Assert" messages
- Compare *distance* and *metric* values
 - Router with best route to source wins
 - If *metric* & *distance* equal, highest IP adr wins
 - Losing router stops sending (prunes interface)

PIM Assert Mechanism



- 3** If there are no directly connected members on the interface, the winning router sends a Prune and waits 3 seconds for a Join override.
 - This will shutoff traffic if it is not needed somewhere downstream.
- 4** Router C *does* need traffic. Sends join to override.

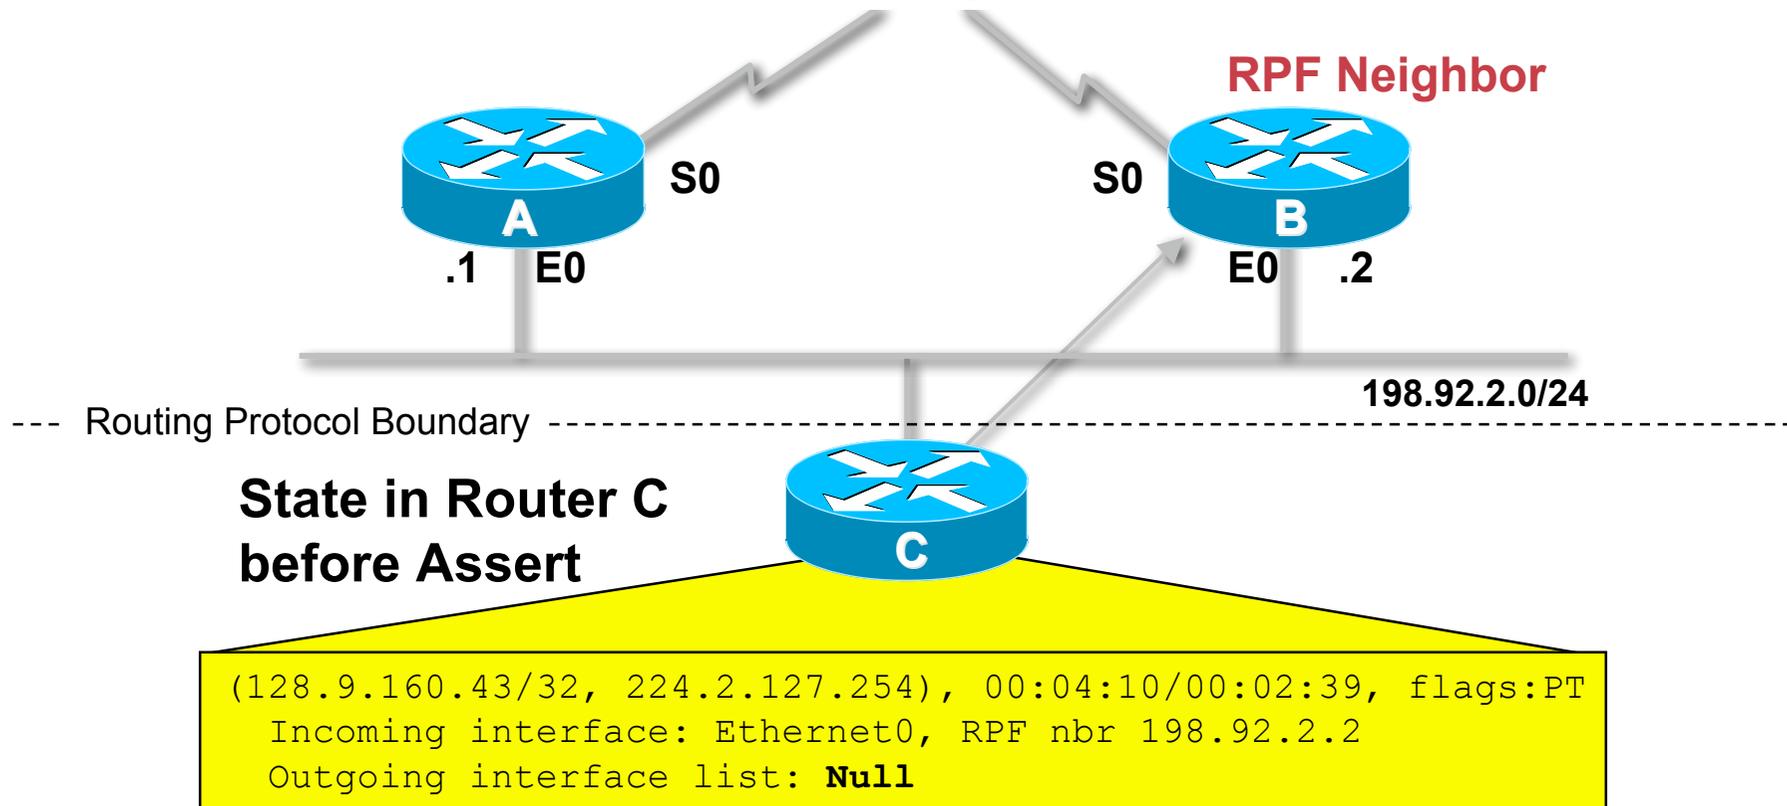
PIM Assert Mechanism



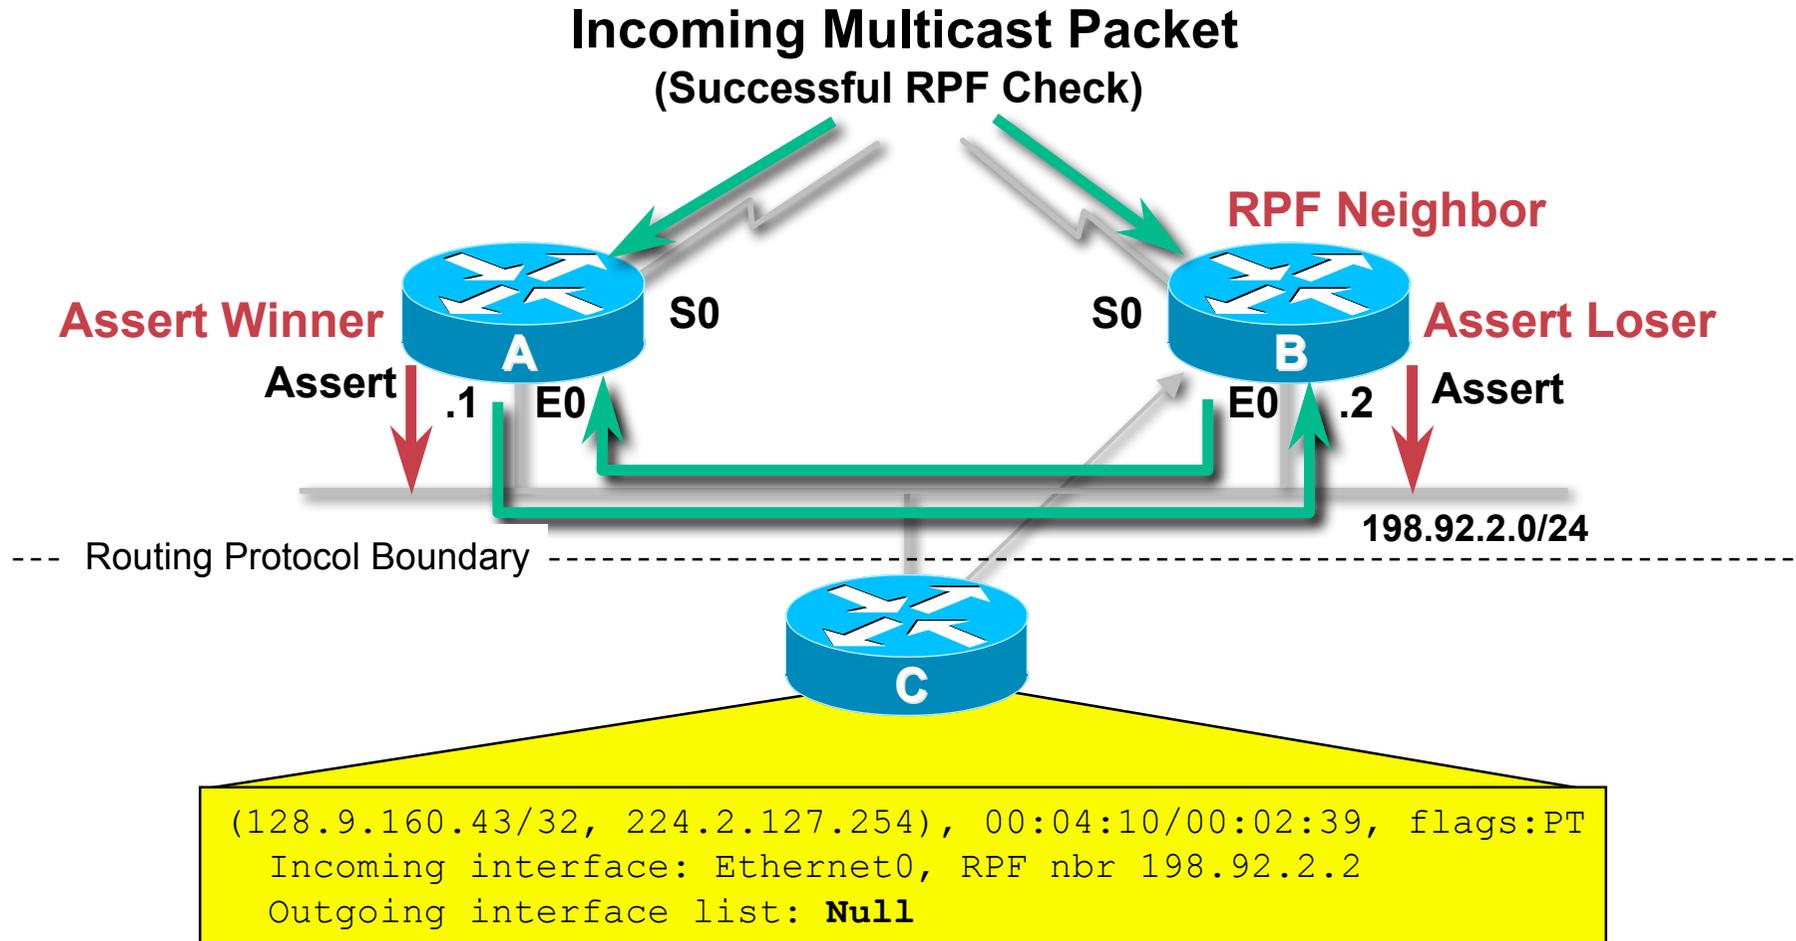
- 5 If router *C* *doesn't* need the traffic, no Join override is sent.
- 6 Router B prunes its interface after 3 second prune delay.
– This stops the flow of traffic onto the transit LAN.

PIM Assert Mechanism

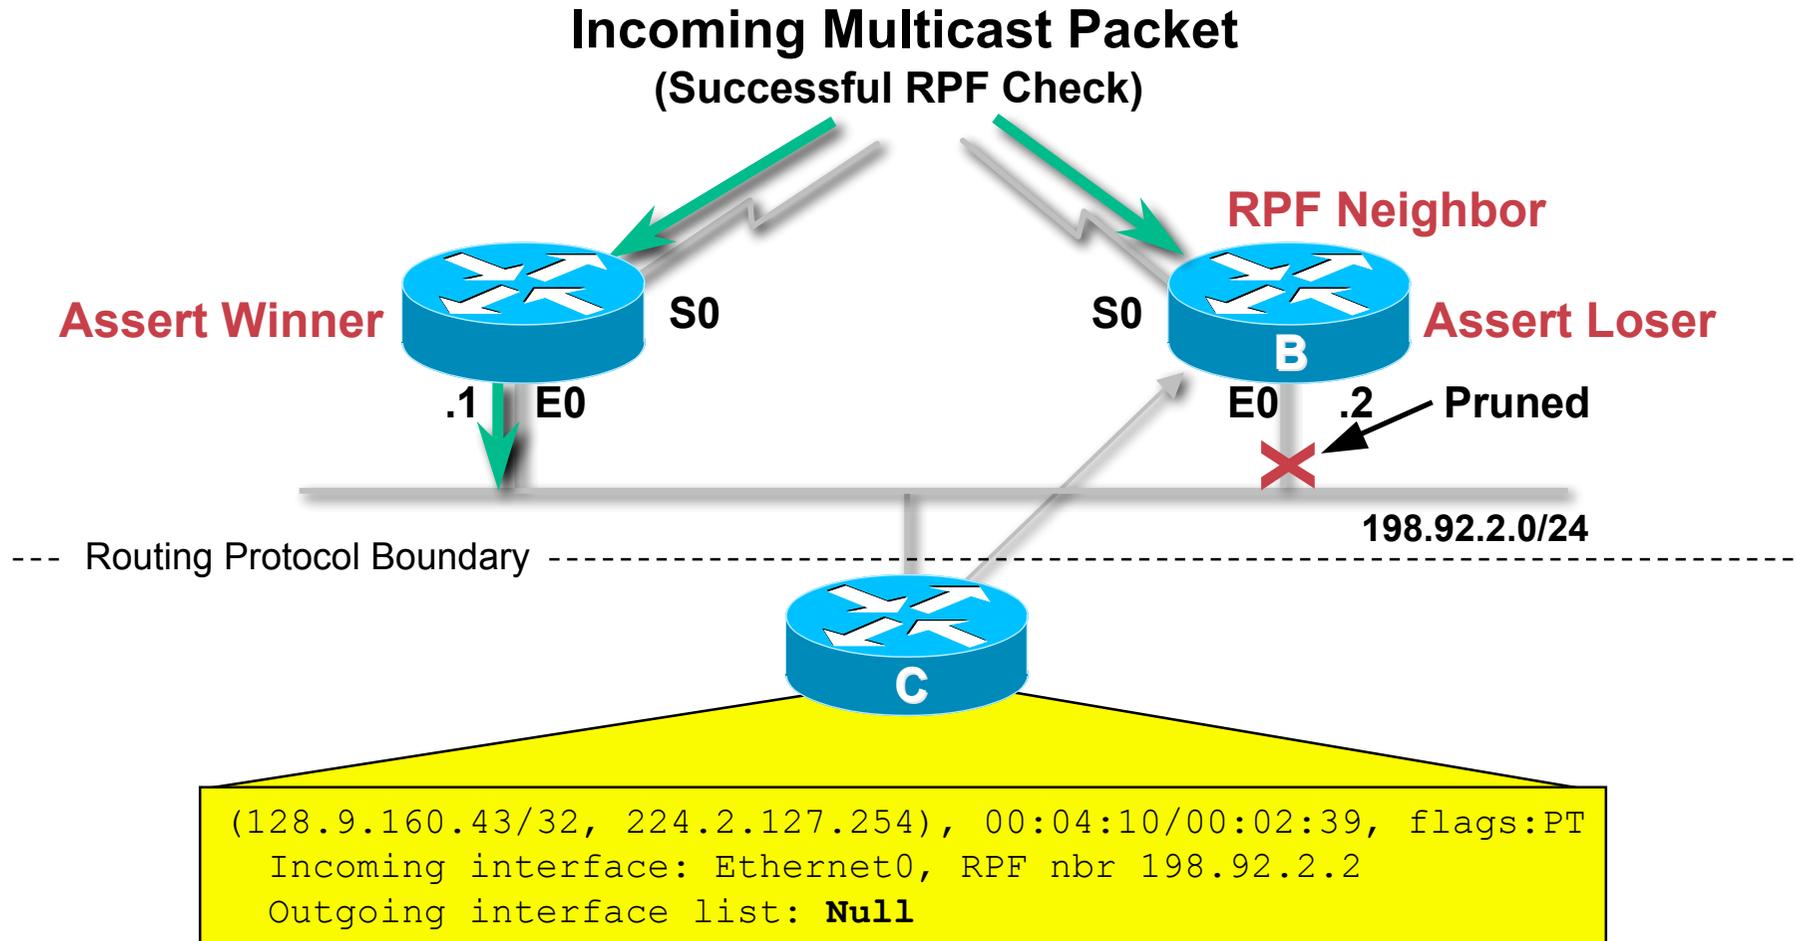
- Downstream routers must listen for the assert winner to know which router to send prunes and grafts



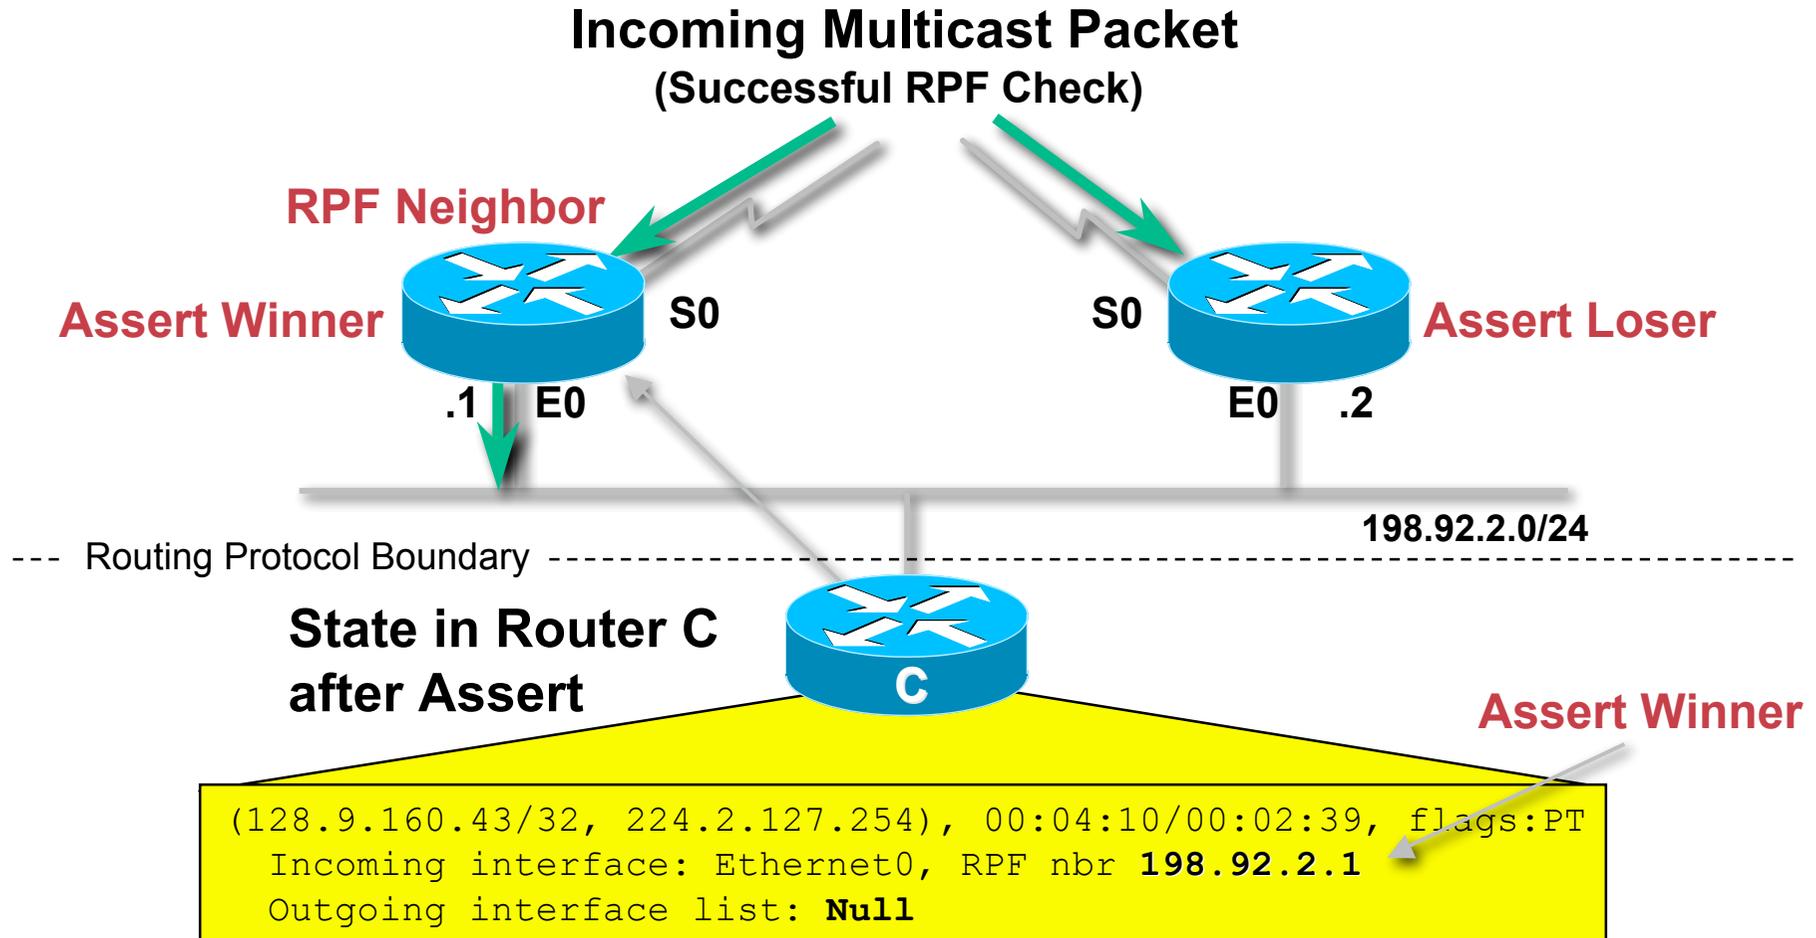
PIM Assert Mechanism



PIM Assert Mechanism



PIM Assert Mechanism



CISCO SYSTEMS

