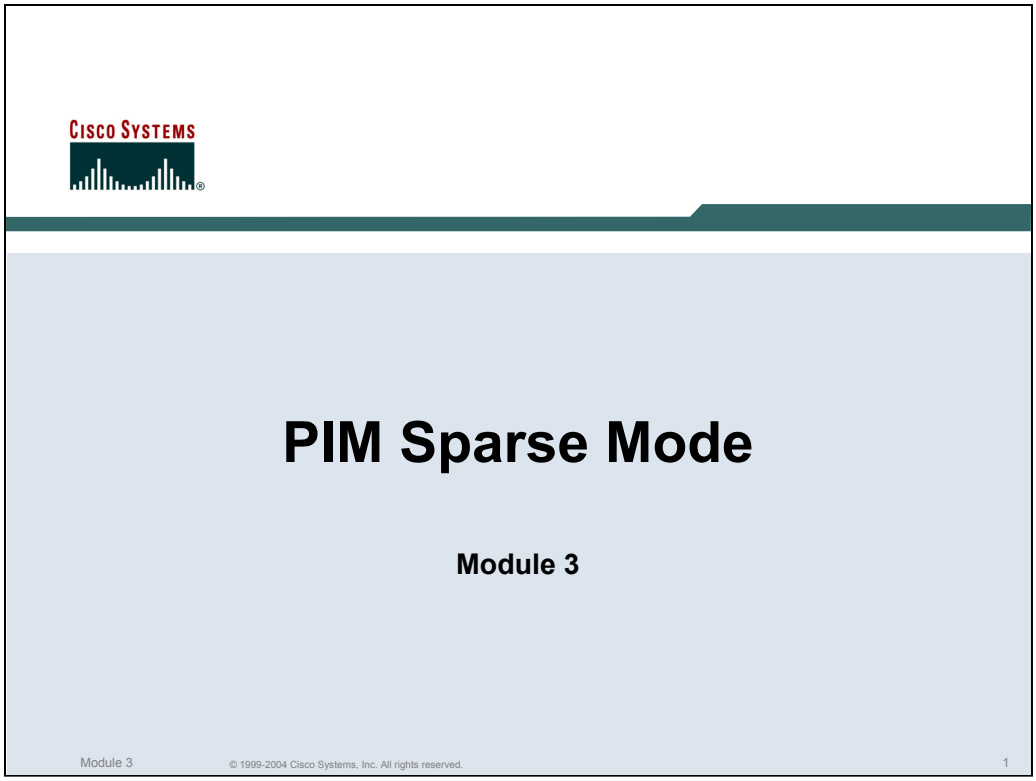
Cisco Systems
PIM Sparse Mode
Module 3
Module 3
© 1999-2004 Cisco Systems, Inc. All rights reserved.
1

Module Agenda

Cisco.com

- PIM Neighbor Discovery
- PIM State
- PIM SM Joining
- PIM SM Registering
- PIM SM SPT-Switchover
- PIM SM Pruning

PIM Neighbor Discovery

PIM Neighbor Discovery

Cisco.com

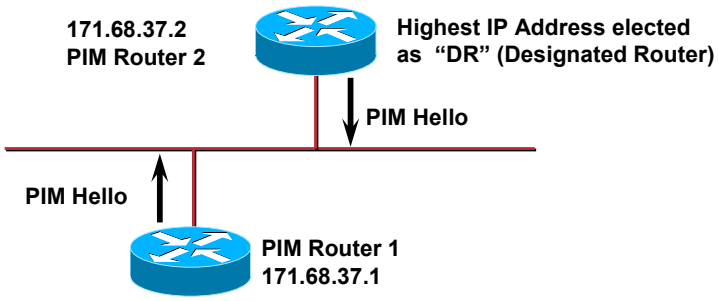

- **PIMv2 Hellos are periodically multicast to the "All-PIM-Routers" (224.0.0.13) group address. (Default = 30 seconds)**
- **If the "DR" times-out, a new "DR" is elected.**
- **The "DR" is responsible for sending all Joins and Register messages for any receivers or senders on the network.**

Module 3 4

- **PIM Neighbor Discovery**
  - PIM Hellos are sent periodically to discover the existence of other PIM routers on the network and to elect the Designated Router.
  - For Multi-Access networks (e.g. Ethernet), the PIM Hello messages are multicast to the "All-PIM-Routers" (224.0.0.13) multicast group address.
- **Designated Router (DR)**
  - For multi-access networks, a Designated Router (DR) is elected. In PIM Sparse mode networks, the DR is responsible for sending Joins to the RP for members on the multi-access network and for sending Registers to the RP for sources on the multi-access network. For Dense mode, the DR has no meaning. The exception to this is when IGMPv1 is in use. In this case, the DR also functions as the IGMP Querier for the Multi-Access network.
- **Designated Router (DR) Election**
  - To elect the DR, each PIM node on a multi-access network examines the received PIM Hello messages from its neighbors and compares the IP Address of its interface with the IP Address of its PIM Neighbors. The PIM Neighbor with the highest IP Address is elected the DR.
  - If no PIM Hellos have been received from the elected DR after some period (configurable), the DR Election mechanism is run again to elect a new DR.

PIM Neighbor Discovery

Cisco.com

```
wan-gw8>show ip pim neighbor
PIM Neighbor Table
Neighbor          Interface          Uptime/Expires      Ver    Mode
Address                                                         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
```

Module 3 5

- **"show ip pim neighbor" command output**
  - Neighbor Address - the IP address of the PIM Neighbor
  - Interface - the interface where the PIM Hello of this neighbor was received.
  - Uptime - the period of time that this PIM Neighbor has been active.
  - Expires - the period of time after which this PIM Neighbor will no longer be considered as active. (Reset by the receipt of a another PIM Query.)
  - Mode - PIM mode (Sparse, Dense, Sparse/Dense) that the PIM Neighbor is using.
  - "(DR)" - Indicates that this PIM Neighbor is the Designated Router for the network.

PIM State

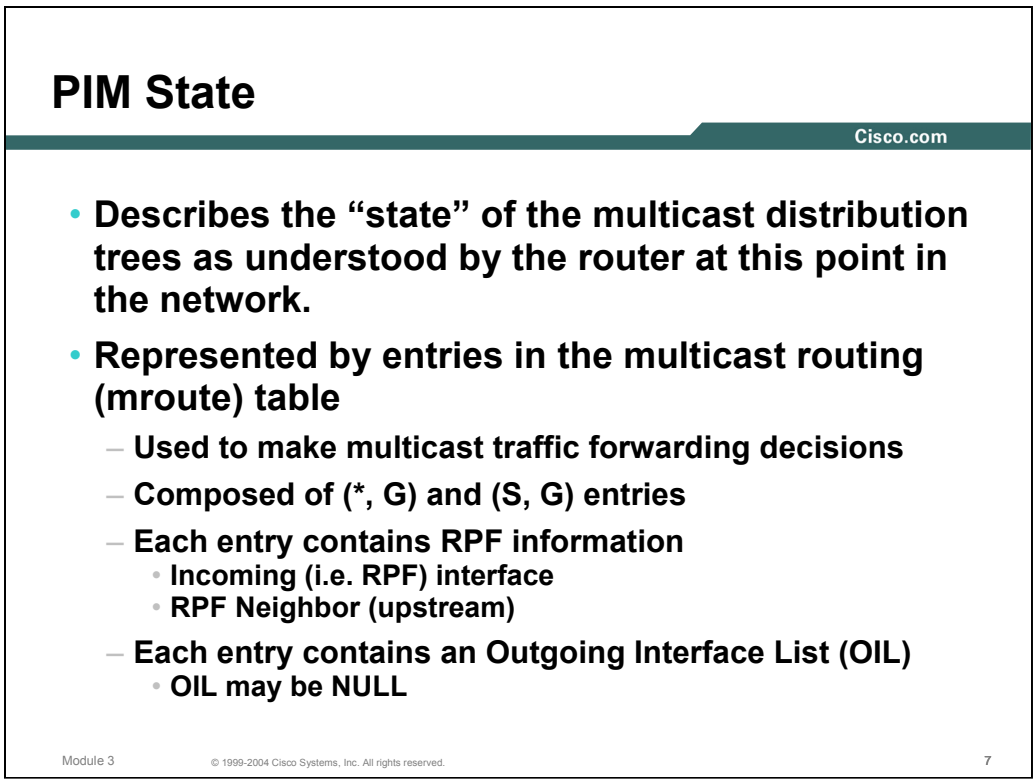

- **PIM State**
  - In general, Multicast State basically describes the multicast distribution tree as it is understood by the router at this point in the network.
  - However to be completely correct, "Multicast State" describes the multicast traffic "forwarding" state that is used by the router to forward multicast traffic.
- **Multicast Routing (mroute) Table**
  - Multicast "state" is stored in the multicast routing (mroute) table and which can be displayed using the show ip mroute command.
  - Entries in the mroute table are composed of (*, G) and (S, G) entries each of which contain:
    - RPF Information consisting of an Incoming (or RPF) interface and the IP address of the RPF (i.e. upstream) neighbor router in the direction of the source. (In the case of PIM-SM, this information in a (*, G) entry points toward the RP. PIM-SM will be discussed in a later module.)
    - Outgoing Interface List (OIL) which contains a list of interfaces that the multicast traffic is to be forwarded. (Multicast traffic must arrive on the Incoming interface before it will be forwarded out this interfaces. If multicast traffic does not arrive on the Incoming interface, it is simply discarded.)

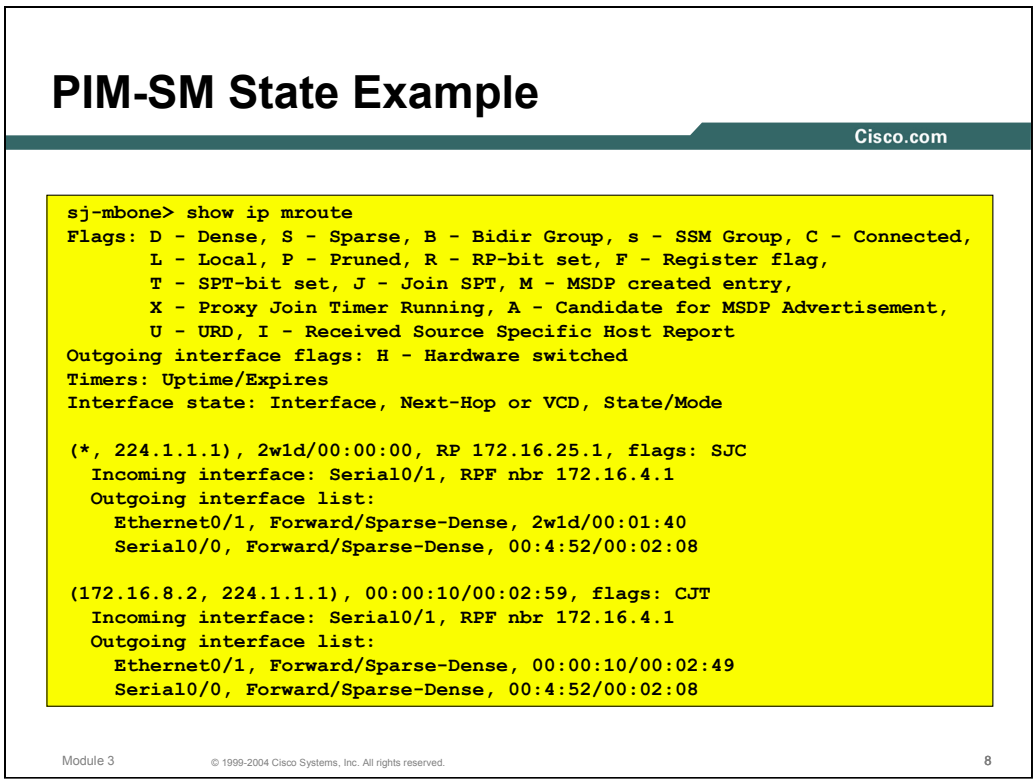

- **PIM-SM State Example**
  - (*, G) Entry - The (*, 224.1.1.1) entry shown in sample output of the show ip mroute command is the (*, G) entry. If there is no matching entry for a particular (S, G) entry, this entry is used to forward traffic down the Shared Tree.
    - The Expires countdown timer in the first line of the (*, G) entry which shows when the entry will expire and be deleted. This entry will remain at roughly 3 minutes as long as there is an interface in the Outgoing Interface list.
    - The Incoming interface information is used to RPF check arriving (*, G) multicast traffic and is computed in the direction of the RP (in this case, 10.1.5.1.).
    - The Outgoing Interface list which reflects the interfaces where (*,G) Joins have been received or where directly connected members of group “G” reside. Traffic flowing down the Shared Tree are forwarded out these interfaces. The Expires countdown timers on these interfaces are reset to 3 minutes by the receipt of periodic (*, G) Joins. If the count ever reaches zero, the entry in the OIL is deleted.
  - (S, G) Entry - The (128.9.160.43/32, 224.1.1.1) entry is an example of an (S, G) entry in the mroute table. This entry is used to forward any multicast traffic sent by source 128.9.160.43 to group 224.1.1.1. Notice the following:
    - The Expires countdown timer in the first line of the (S, G) entry which shows when the entry will expire and be deleted. This entry is reset to 3 minutes whenever an (S, G) multicast packet is forwarded.
    - The Incoming interface information is used to RPF check arriving (S, G) multicast traffic. If a packet does not arrive via this interface, the packet is discarded.
    - The Outgoing Interface list which reflects the interfaces where (S,G) packets are to be forwarded.

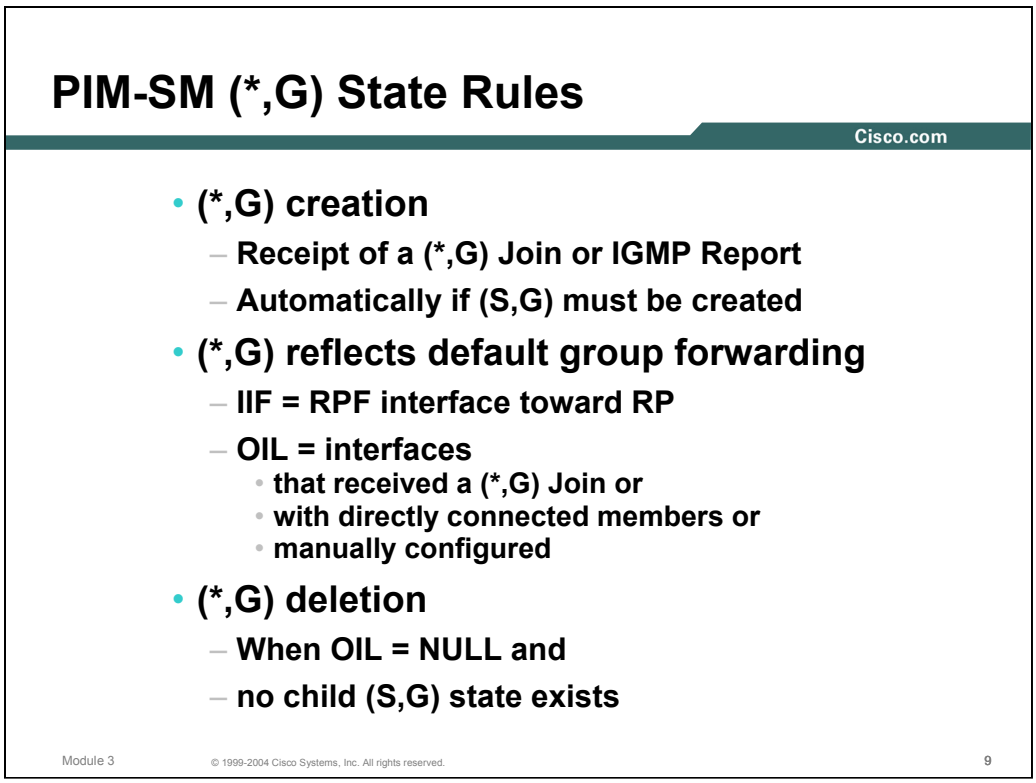

- **PIM-SM (*,G) State Rules**
  - A (*, G) entry is created when a (*, G) Join or an IGMP Report is received
    - The later condition can be simulated by manually configuring the interface to join the group.
  - (*, G) entries are also automatically created whenever an (S, G) entry for the group must be created.
    - The (*, G) entry is created first and then the (S, G) entry. The reason for this will become clear shortly.
  - The IIF reflects the RPF interface/neighbor in the direction of the RP.
  - The OIL of a PIM-SM (*, G) entry reflects interfaces that:
    - Have received a (*, G) Join or
    - Where a directly connected member has joined the group
    - The interface was manually configured to join the group. (Note: This may be accomplished using the ip igmp static-group <group> command.)
  - (*, G) entries are deleted when its Expires timer counts down to zero. This will only occur when:
    - The OIL is Null and
    - No child (S, G) entry exists

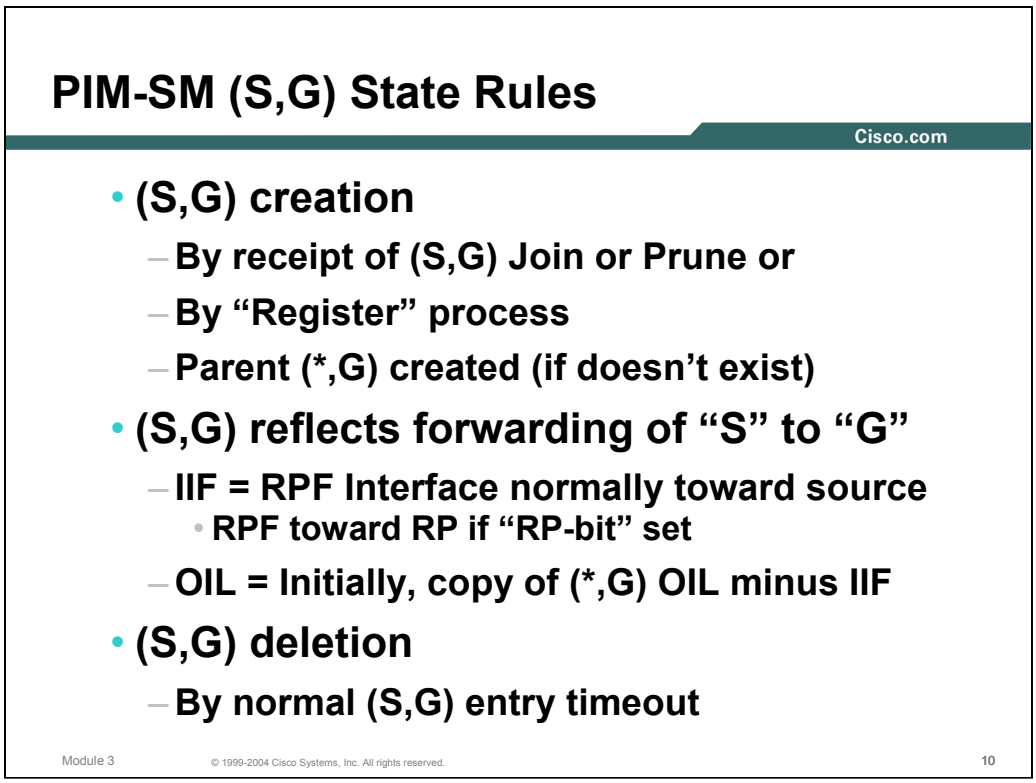

- **PIM-SM (S, G) Rules**
  - In PIM-SM, (S, G) state is created as a result of:
    - The receipt of an (S, G) Join or Prune or
    - The PIM-SM Register process which is triggered by a first-hop router receiving a packet from a directly connected source.
  - When an (S, G) entry must be created, the following steps occur:
    - If a corresponding (*, G) entry does not exist, it is created first.
    - The RPF Information is computed for the source "S". This information is stored in the (S, G) entry as the Incoming interface and the RPF neighbor (i.e. the PIM neighbor in the direction of the source).
    - The exception to this rule is if the RP-bit is set in the (S, G) entry, the RPF interface is pointed up the Shared Tree. This mechanism allows duplicate (S, G) traffic to be blocked from flowing down the Shared Tree after a downstream router has switched to the Shortest Path Tree. (More on this later.)
    - The OIL of the (S, G) entry is populated with a copy of the OIL from the parent (*, G) entry less the Incoming interface. (The Incoming interface must not appear in the OIL otherwise a multicast route loop could occur.)
  - In PIM-SM, (S, G) entries are deleted when their Expires timer counts down to zero. The Expires timer is reset whenever an (S, G) packet is received and forwarded.

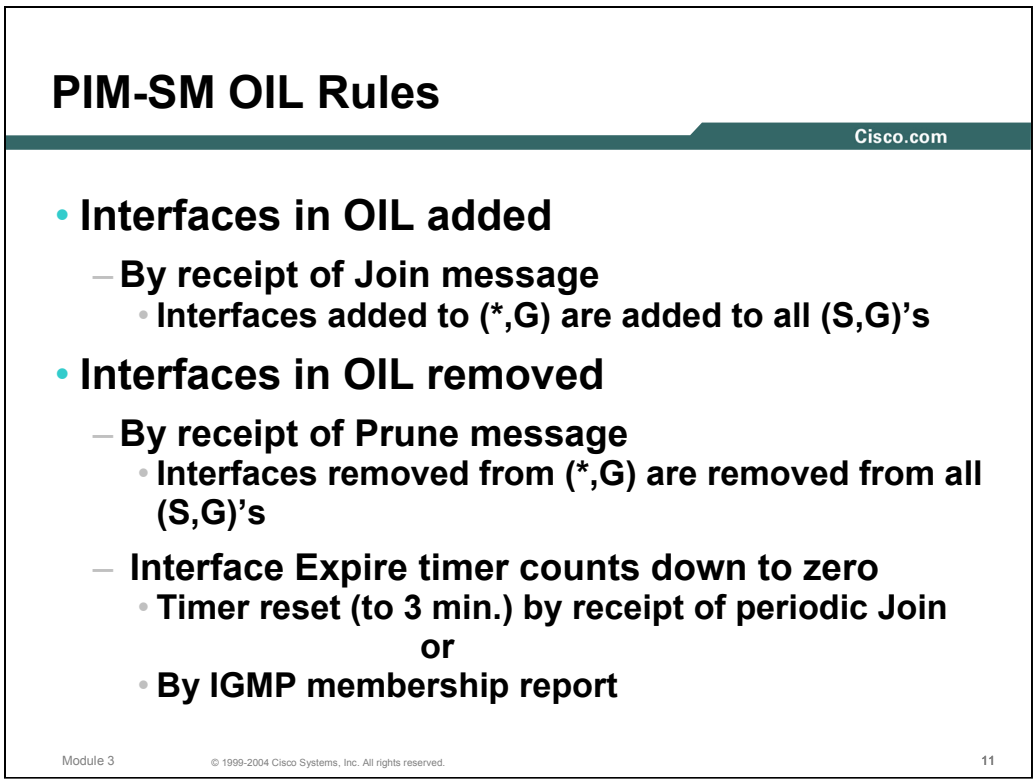

- **PIM-SM Outgoing Interface List Rules**
  - Adding an interface
    - Interfaces are added to an (S, G) OIL when a (S, G) Join message is received on an interface.
    - Interfaces are added to the (*, G) OIL when a (*, G) Join message is received on an interface.
    - Anytime an interface is added to the (*, G) OIL, the interface is added to the OIL of all associated (S, G) OIL's. (Note: A check is always made to prevent the IIF from appearing in the OIL.)
  - Removing an interface
    - Interfaces are removed from the OIL of a (*, G) or (S, G) entry if the interface's Expires timer counts down to zero.
    - Note: The interface Expires timer is reset to 3 minutes by the receipt of periodic Join messages sent by downstream routers once per minute or by an IGMP Report sent by a directly connected member on the interface.
    - Interfaces are removed from the OIL if an Prune message is received (and it is not overridden by another router if the interface is a multi-access network).
    - Interfaces removed from a (*, G) OIL, are removed from the OIL of all associated (S, G) OIL's.

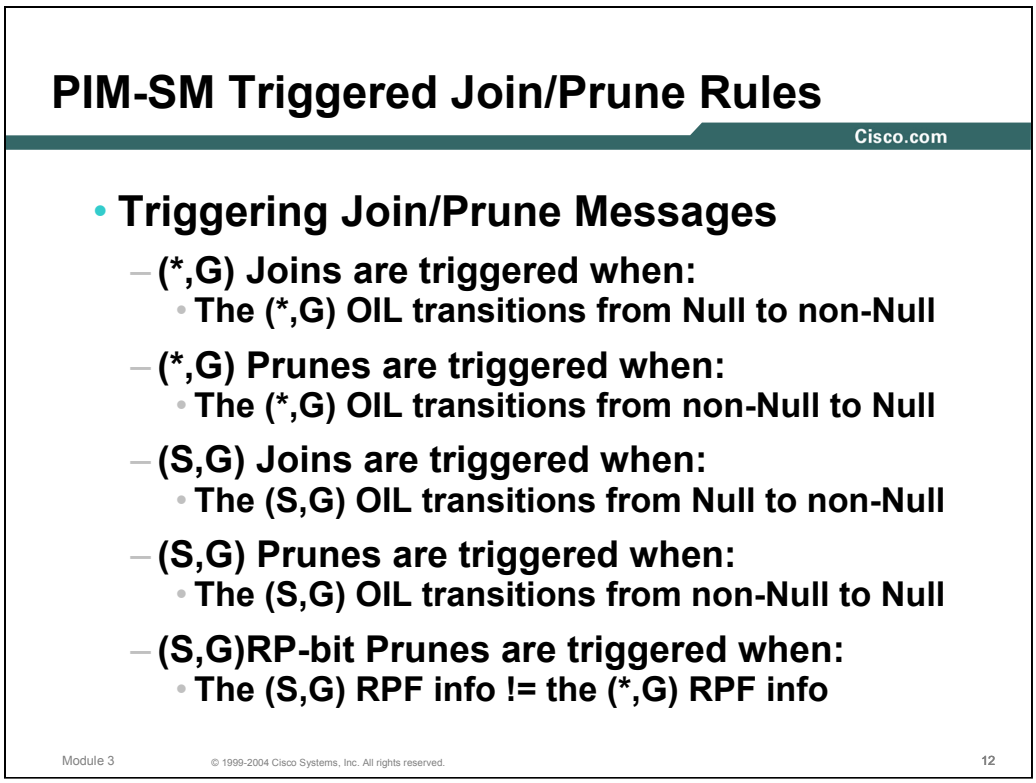

- **PIM-SM Outgoing Interface List Rules**
  - Triggering Join/Prune Messages
    - (*,G) Joins are triggered whenever the (*,G) OIL is empty (Null) and an interface is added making the OIL non-Null.
    - (*,G) Prunes are triggered whenever the last interface is removed from the (*,G) OIL.
    - (S,G) Joins are triggered whenever the (S,G) OIL is empty (Null) and an interface is added making the OIL non-Null.
    - (S,G) Prunes are triggered whenever the last interface is removed from the (S,G) OIL.
    - Note: Prior to 12.1(E) and 12.2, there was an optimization that attempted to minimize the sending of (S,G) Prunes. Instead of sending the (S,G) Prune immediately when the last interface is removed, the state is just allowed to time out. However, if (S,G) traffic is still flowing, then the arrival of the next (S,G) packet would cause the prune to be sent.
    - (S,G)RP-bit Prunes are sent whenever the (S,G) RPF information (incoming interface and RPF-neighbor) is not the same as the (*,G) RPF information. This indicates that the SPT and the Shared-Tree diverge at this point and that (S,G) traffic should be pruned from the Shared-Tree.

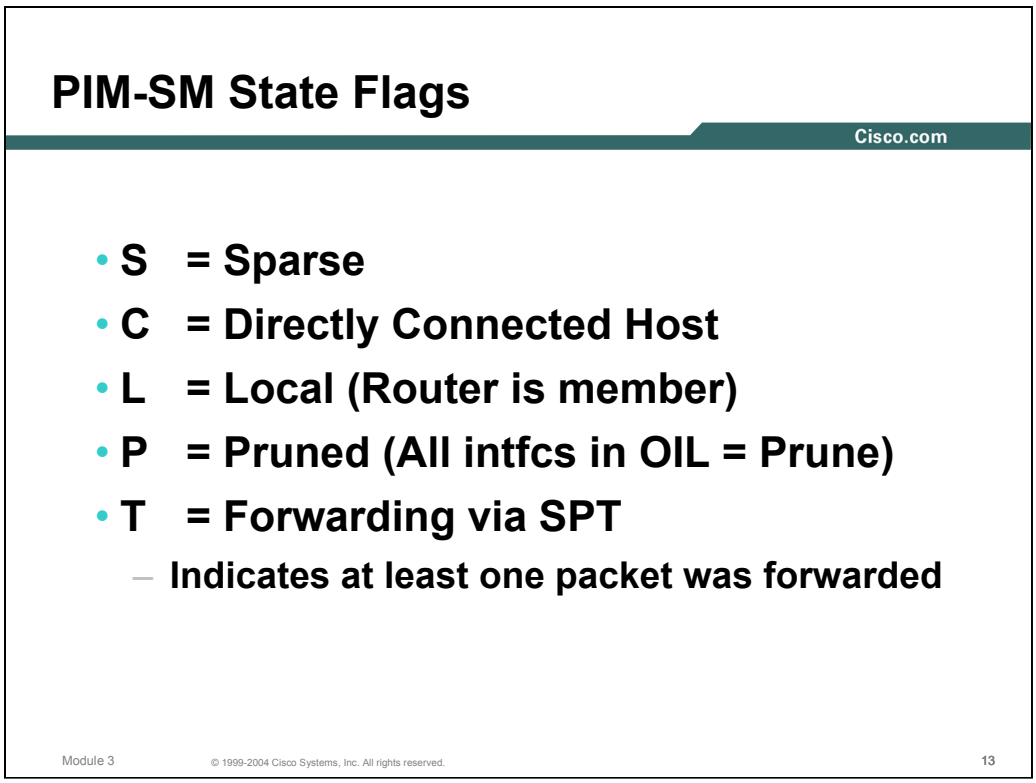

- **PIM-SM State Flags**
  - "S" Flag ((*, G) entries only)
    - Indicates the group is operating in Sparse mode. (Appears only on (*, G) entries.)
  - "C" Flag
    - Indicates that there is a member of the group directly connected to the router.
  - "L" Flag
    - Indicates the router itself is a member of this group and is receiving the traffic. (This would be the case for the Auto-RP Discovery group 224.0.1.40 which all Cisco routers join automatically.)
  - "P" Flag
    - Set whenever all interfaces in the outgoing interface list of an entry are Pruned (or the list is Null). This general means that the router will send Prune messages to the RPF neighbor to try to shutoff this traffic.)
  - "T" Flag ((S, G) entries only)
    - Indicates that at least one packet was received via the SPT

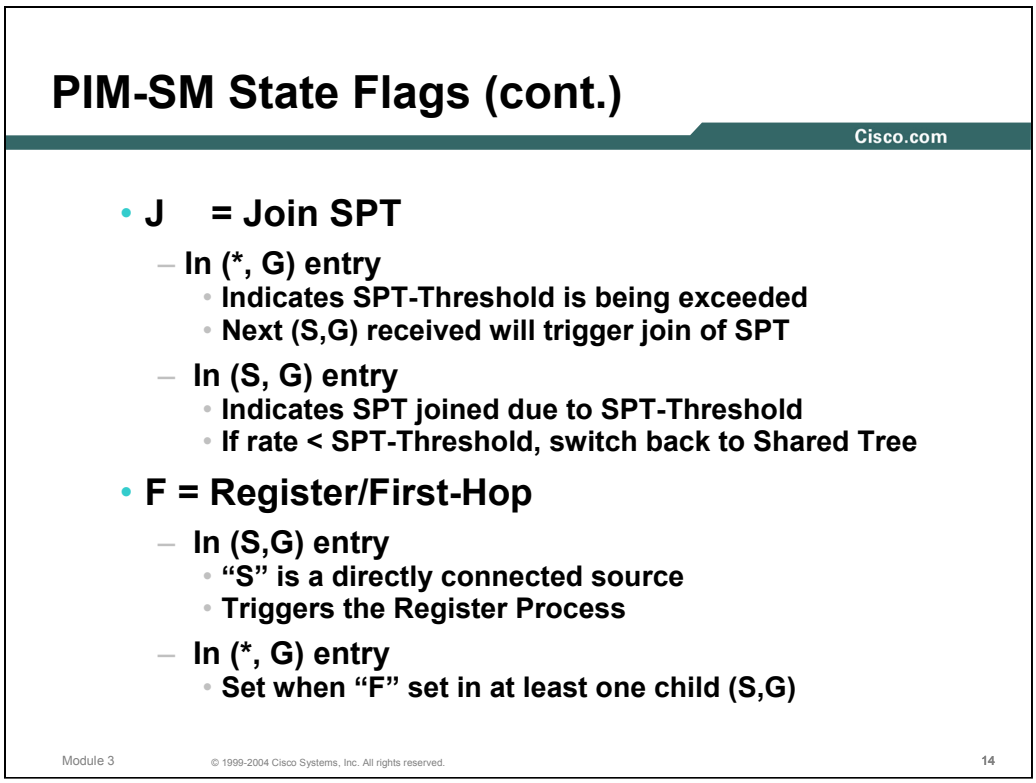

- **PIM-SM State Flags**
  - "J" Flag (Join SPT)
    - When this flag is set in a (*, G) entry, it indicates that the rate of traffic flowing down the Shared Tree is above the SPT-Threshold and will cause a switch to the SPT for the next packet received down the shared tree. (More on this later.)
    - When this flag is set in an (S, G) entry, it indicates that the (S, G) entry (and hence the SPT) was created as a result of the SPT-Threshold being exceeded. If the rate of this (S, G) traffic drops back below the SPT, the router will attempt to switch this traffic flow back to the Shared Tree.
  - "F" Flag (Register/First-Hop)
    - This flag is set on an (S, G) entry when source "S" is directly connected to the router. This indicates that this router is a "first-hop" router and triggers it to send Register messages to the RP to inform the RP of this active source.
    - This flag can also be set for arriving (S, G) entries created at a border router such as a router that borders on a DVMRP or other dense mode cloud. This causes the router to perform a proxy-register operation and send (S, G) Register messages to the RP on behalf of the downstream DVMRP routers. This proxy-register operation follows the same rules as for directly connected sources.
    - The "F" flag is also set on a (*, G) entry if any associated (S, G) entries have the "F" flag set.

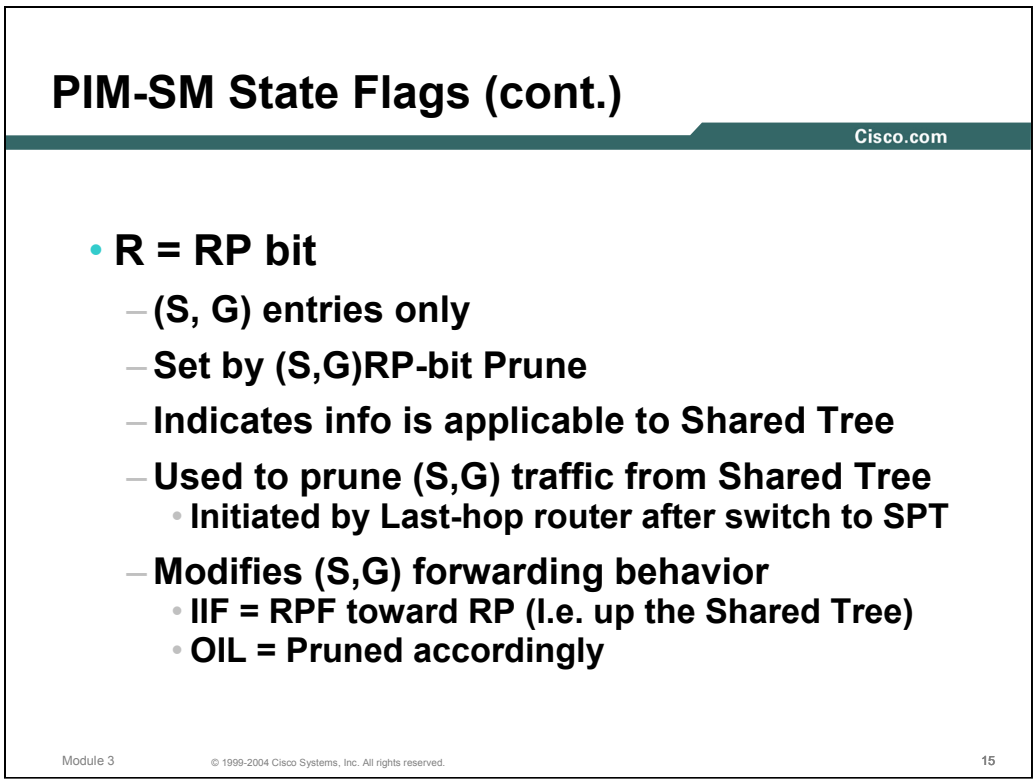

- **PIM-SM Flags**
  - "R" Flag (RP-Bit)
    - This flag is set on (S, G) entries only and indicates that the (S, G) forwarding information in the entry is applicable to (S, G) traffic flowing down the Shared Tree.
    - The "R" flag is set on an (S, G) entry by the receipt of an (S, G)RP-bit Prune message. These messages are sent by downstream routers on the Shared Tree that are requesting that this specific (S, G) traffic flow be pruned off of the Shared Tree. This is done to eliminate duplicate (S, G) traffic after a downstream router has switched to the (S, G) Shortest-Path Tree.
    - Whenever the "R" flag is set on an (S, G) entry, the RPF information must be changed to point toward the RP instead of pointing at source "S". This is done because the (S, G) entry is now applicable to (S, G) traffic arriving down the Shared Tree. As a result, the RPF information must point up the Shared Tree in order for arriving (S, G) packets to RPF correctly. (This should be made clear later.)

PIM SM Joining

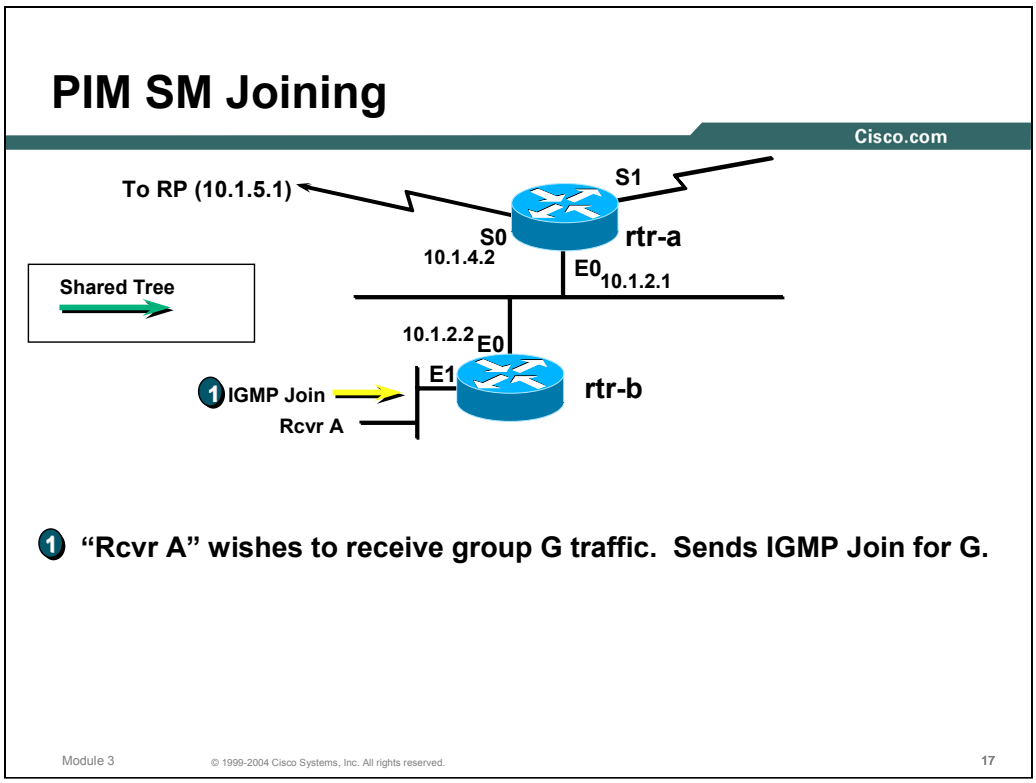

- **PIM SM Joining Example**
  - 1) Receiver "A" wishes to receive group "G" multicast traffic and therefore sends an IGMP Host Membership message (sometimes loosely referred to as an IGMP Join) which is received by "rtr-b".
  - "rtr-b" has no existing (*, G) state for group "G" and therefore creates an entry. (See next slide.)

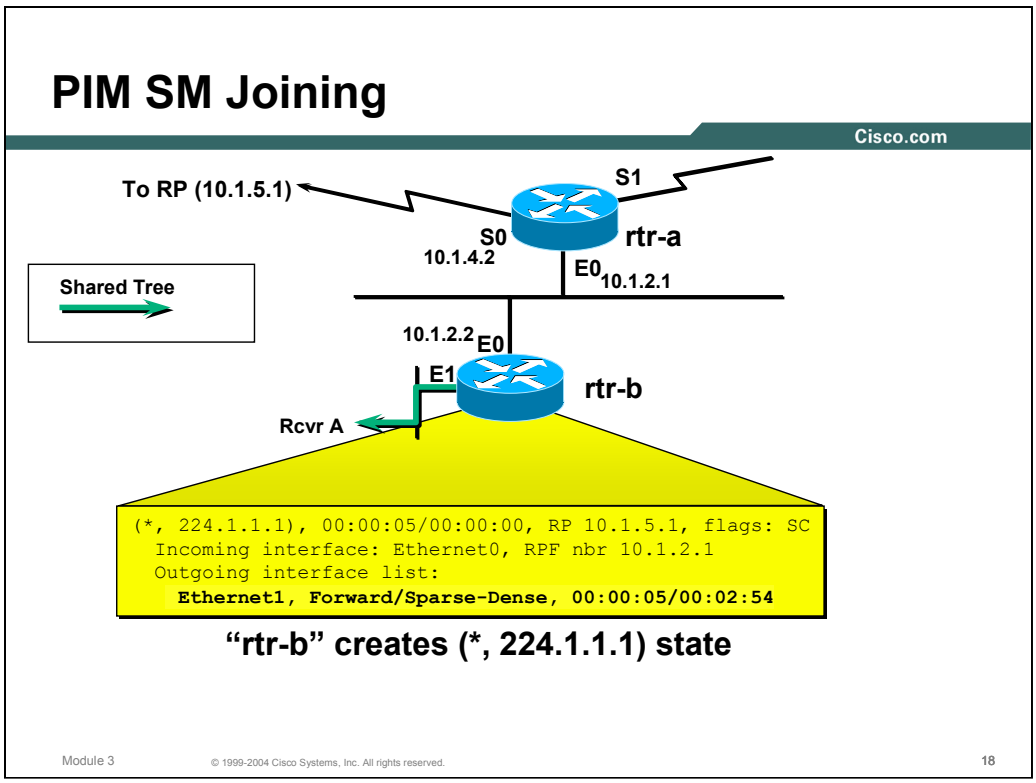

- **State in "rtr-b" after Joining (*, 224.1.1.1)**
  - (*, 224.1.1.1)
    - indicates the (*, G) entry.
  - 00:00:05/00:02:54
    - indicates that the entry has existed for 5 seconds and will expire in 2 minutes and 54 seconds.
  - RP 10.1.5.1
    - is the IP Address of the Rendezvous Point for Group 224.1.1.1
  - flags: SC
    - indicates that this is a Sparse mode group (S) and that there is a member of this group directly connected (C) to the router.
  - Incoming interface: Ethernet0, RPF nbr 10.1.2.1
    - indicates the Incoming interface (up the Shared Tree toward RP) and
    - the RPF neighbor's IP address (in the direction of the RP) is 10.1.2.1
  - Outgoing interface list:
    - lists the interfaces that are in the outgoing interface list for this entry.
  - Ethernet1, Forward/Sparse-Dense, 00:00:05/00:02:54
    - indicates Ethernet 1 is in the oilist; it's in the "Forward" state; Sparse-Dense mode and that it has been in the list for 5 seconds and will expire in 2 minutes and 54 seconds if no further (*, G) Join or IGMP Report is received on this interface.

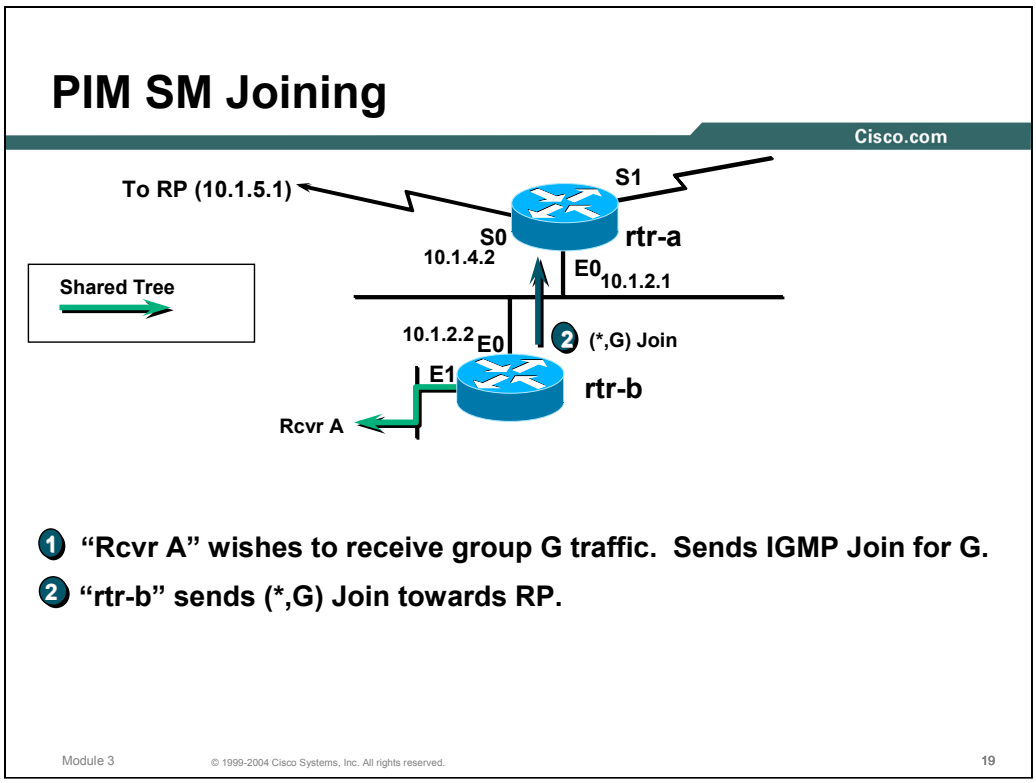

- **PIM SM Joining Example**
  - 2) Because the OIL of the (*, G) transitioned from Null to non-Null (when "rtr-b" added Ethernet 1 to the OIL of the newly created entry), a PIM (*, G) Join is sent to rtr-b's up-stream PIM neighbor (rtr-a) in the direction of the RP.
  - When "rtr-a" receives the (*, G) Join it creates (*, G) state. (See next slide.)

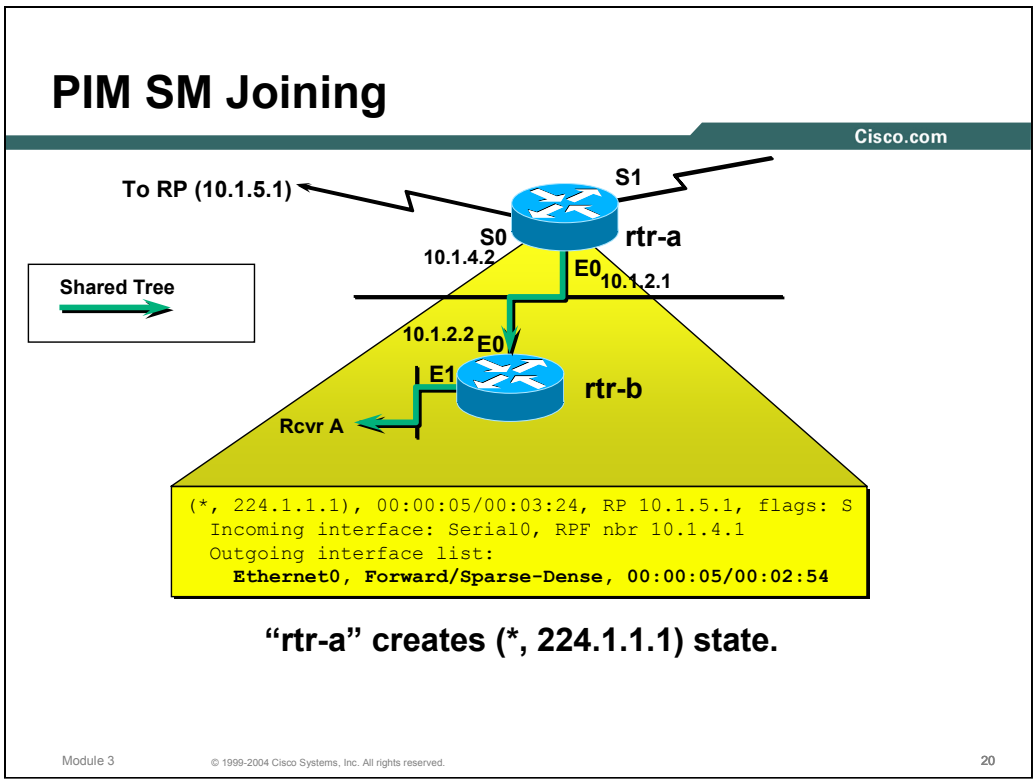

- **State in "rtr-a" after Joining (*, 224.1.1.1)**
  - (*, 224.1.1.1)
    - indicates the (*, G) entry.
  - 00:00:05/00:02:54
    - indicates that the entry has existed for 5 seconds and will expire in 2 minutes and 54 seconds.
  - RP 10.1.5.1
    - is the IP Address of the Rendezvous Point for Group 224.1.1.1
  - flags: S
    - indicates that this is a Sparse mode group (S).
  - Incoming interface: Serial0, RPF nbr 10.1.4.1
    - indicates the Incoming interface (up the Shared Tree toward RP) and
    - the RPF neighbor's IP address (in the direction of the RP) is 10.1.4.1
  - Outgoing interface list:
    - lists the interfaces that are in the outgoing interface list for this entry.
  - Ethernet0, Forward/Sparse-Dense, 00:00:05/00:02:54
    - indicates Ethernet 0 is in the oilist; it's in the "Forward" state; Sparse-Dense mode and that it has been in the list for 5 seconds and will expire in 2 minutes and 54 seconds if no further (*, G) Join or IGMP Report is received on this interface.

PIM SM Joining

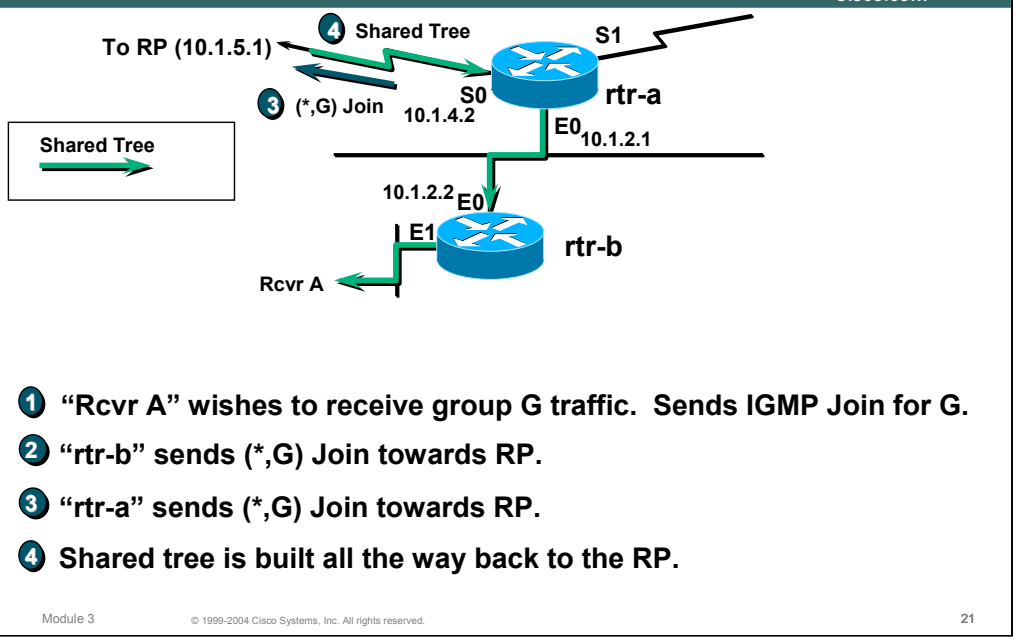

- **PIM SM Joining Example**
  - 3) Because the OIL of the (*, G) transitioned from Null to non-Null (when "rtr-a" added Ethernet 0 to the OIL of the newly created entry), a PIM (*, G) Join is sent to rtr-a's up-stream PIM neighbor in the direction of the RP.
  - When the upstream router receives the (*, G) Join it too creates (*, G) state and creates a branch of the Shared Tree.
  - 4) This process continues all the way back to the RP (or until a router is reached that is already on the Shared Tree and therefore already has a (*, G) entry.)

PIM SM Registering

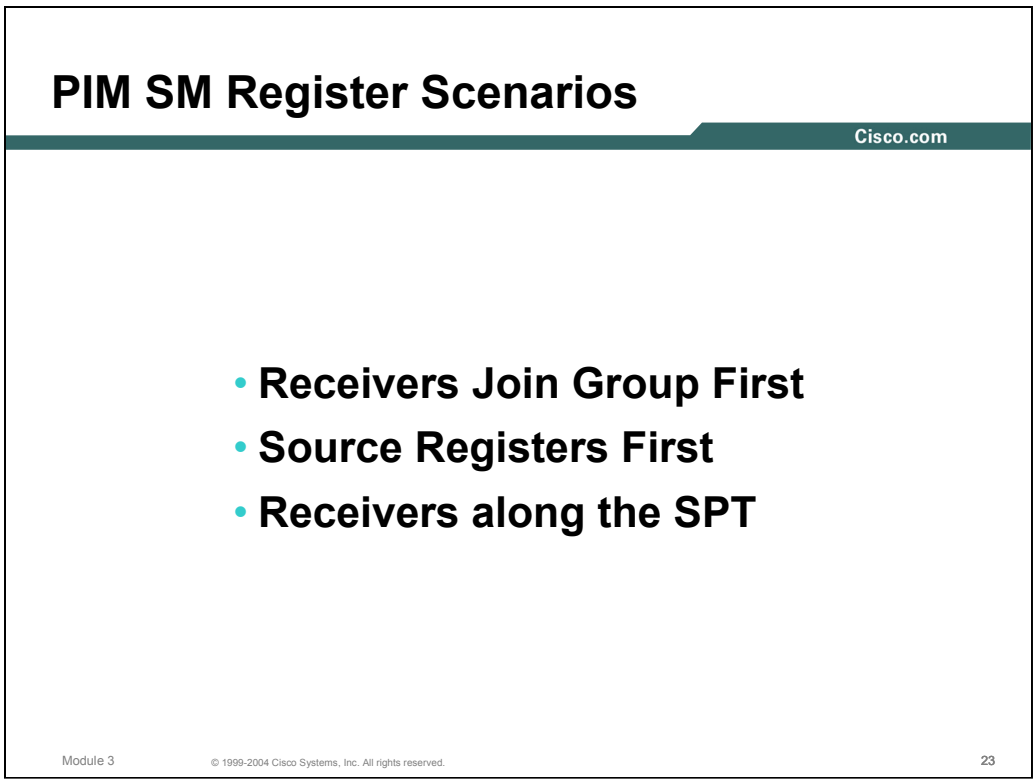

- **PIM SM Register Examples**
  - Depending on whether there are any existing Receivers for group “G” on the Shared Tree (RPT), the RP hands the Register process a little different.

    In the following examples we will consider the Register process for the cases when:
    - Receivers join group “G” first;
    - The Source Registers first.
    - Receivers along the SPT.

PIM SM Registering – Receivers Join Group First

PIM SM Registering
Receiver Joins Group First

Cisco.com

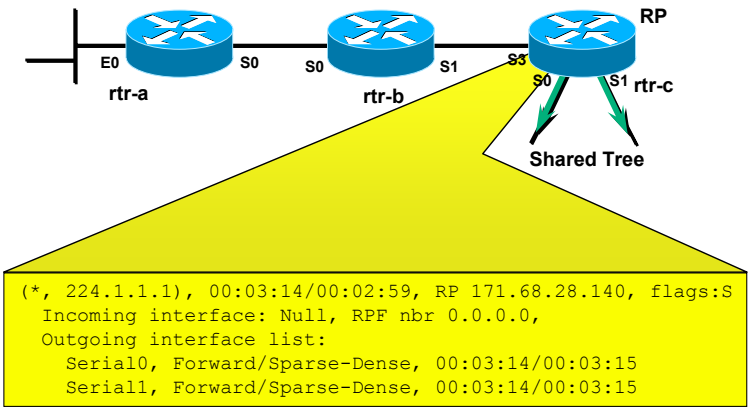

State in "RP" before any source registers
(with receivers on Shared Tree)

- **State in RP before Registering (Rcvr's on Shared Tree)**
  - Pay particular attention to the following in the (*, G) entry:
    - The "Incoming interface:" is NULL and the "RPF nbr" is 0.0.0.0. This indicates that this router is the RP.
    - The "Outgoing interface list:" contains Serial0 and Serial1 which are assumed to be the only two active branches of the Shared Tree (RPT).

PIM SM Registering
Receiver Joins Group First

Cisco.com

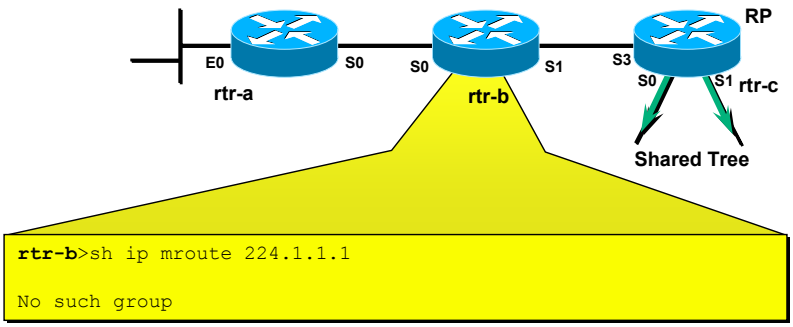

State in "rtr-b" before any source registers
(with receivers on Shared Tree)

- **State in "rtr-b" before source registers**
  - Note that there is no group state information for this Group yet.

PIM SM Registering
Receiver Joins Group First

Cisco.com

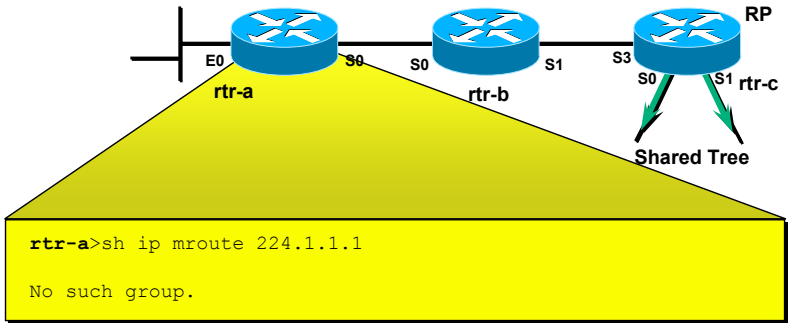

State in "rtr-a" before any source registers
(with receivers on Shared Tree)

- **State in 1st-hop router (rtr-a) before source registers**
  - Note that there is no group state information for this Group yet.

PIM SM Registering

Receiver Joins Group First

Cisco.com

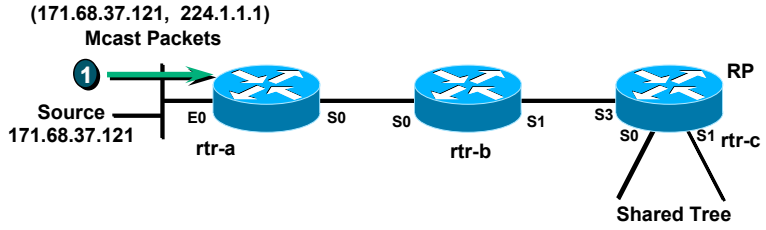

1 "Source" begins sending group G traffic.

- **Receivers Join Group First Example**
  - 1) Source "S" begins sending traffic to group "G".
  - 2) 1st-hop router ("rtr-a") creates (*, G) and (S, G) state; encapsulates the multicast packets in PIM Register message(s) and unicasts it(them) to the RP.
  - 3) The RP ("rtr-c") de-encapsulates the packets and sees that the packet is for group "G" for which it already has (*, G) state. It then forwards the packets down the Shared Tree.

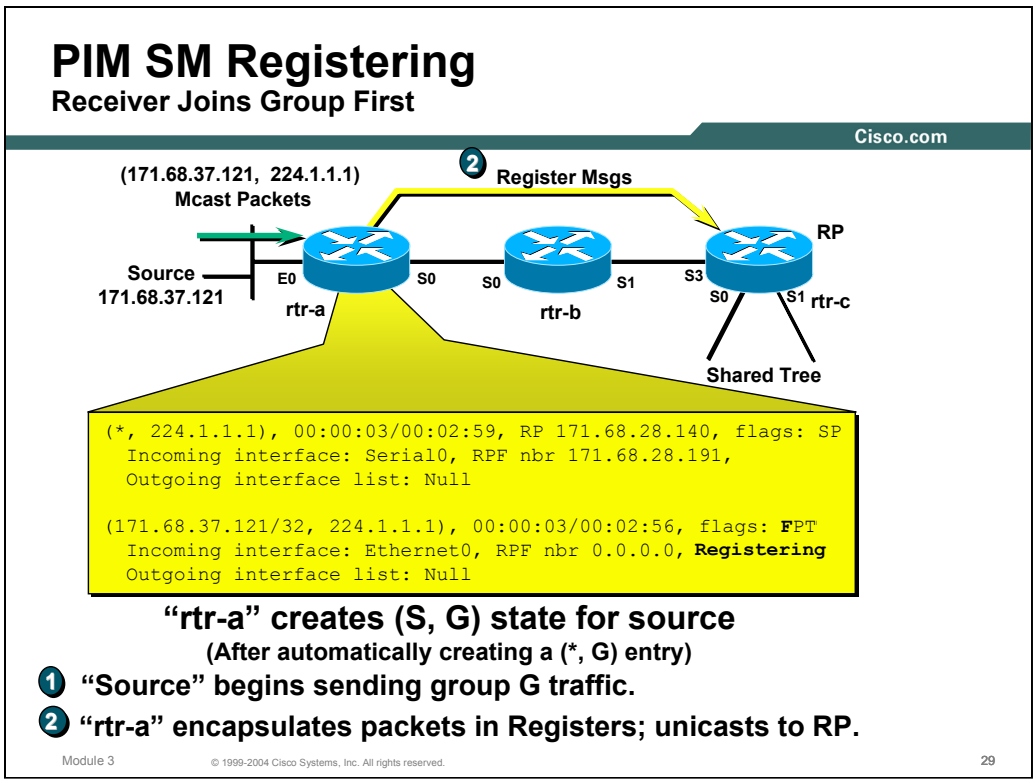

- **1st-hop router (rtr-a) creates (S, G) state**
  - A (*, G) entry must be created before the (S, G) entry can be created. Note that:
    - The RPF information for this entry points up the Shared Tree via Serial0 with the RPF neighbor of 171.68.28.191. (Serial 0 of "rtr-b".)
    - Because in this example no members have joined the group (the sender is only sending), the OIL of the (*, G) entry is Null.
    - The "P" flag (Pruned) is set since the OIL is Null.
  - The (S, G) entry is then created. Pay particular attention to the following:
    - The RPF information for this entry points towards the source via Ethernet0. The RPF neighbor is 0.0.0.0 because the source is directly connected.
    - The (S, G) OIL receives a copy of the (*, G) OIL. (Which is Null.)
    - The "F" flags are set in the (S, G) entry which indicates that this is a directly connected Source.
    - The "Registering" flag is set in the (S, G) entry which indicates that we are still sending Register messages to the RP for this Source.
    - The "P" flag (Pruned) is set since the OIL is Null.
  - 2) The 1st-hop router encapsulates the multicast packets in PIM Register message(s) and unicasts them to the RP.

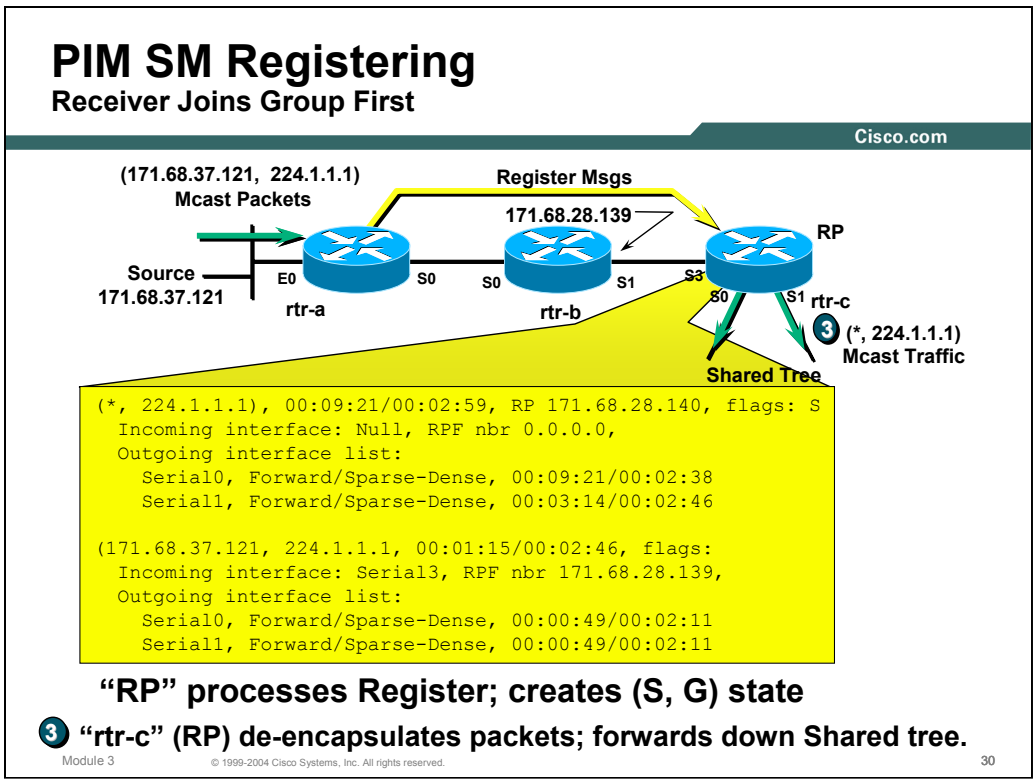

- **The RP creates (S, G) state**
  - As a result of the Register message that was received from "rtr-a", the RP creates (S, G) state as follows:
    - The RPF information is calculated using the source address contained in the multicast packet encapsulated inside of the register message. This results in an IIF of Serial3 and an RPF neighbor of 171.68.28.139.
    - Next, the OIL of the parent (*, G) entry is copied into the OIL of the new (S,G) entry. (An additional check is made to insure that the IIF does not appear in the OIL. If it does, it is removed to prevent a route loop.)
  - Now the router is ready to forward the (S, G) packet that was encapsulated in the Register message using the newly created (S, G) state. (Note that traffic is always forwarded using the matching (S, G) entry if one exists. Otherwise, the (*, G) entry is used.) This is accomplished as follows:
    - Because this packet was received inside of a Register message, the RPF check is skipped.
    - Next, the router forwards a copy of the packet out all interfaces in the (S, G) OIL. In this case a copy is sent out Serial0 and Serial1 which corresponds to the two branches of the Shared Tree.
    - The "T" flag is not yet set in the (S, G) entry. However, when the first (S, G) packet is received natively (via the Incoming interface) and forwarded using this entry, the "T" flag will be set.

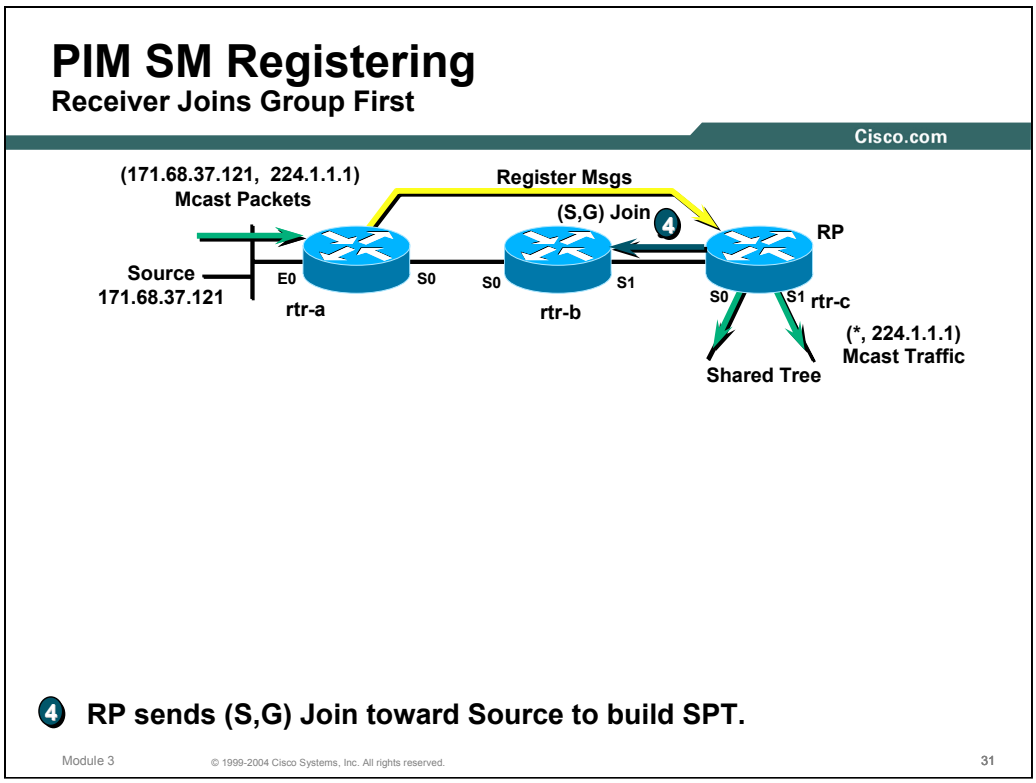

- **Receivers Join Group First Example (cont.)**
  - 4) Because RP has existing (*, G) state (I.e. Receivers already waiting on the Shared Tree), it sends an (S, G) Join toward source “S” to build a Shortest-Path Tree (SPT) from source “S” to the RP.

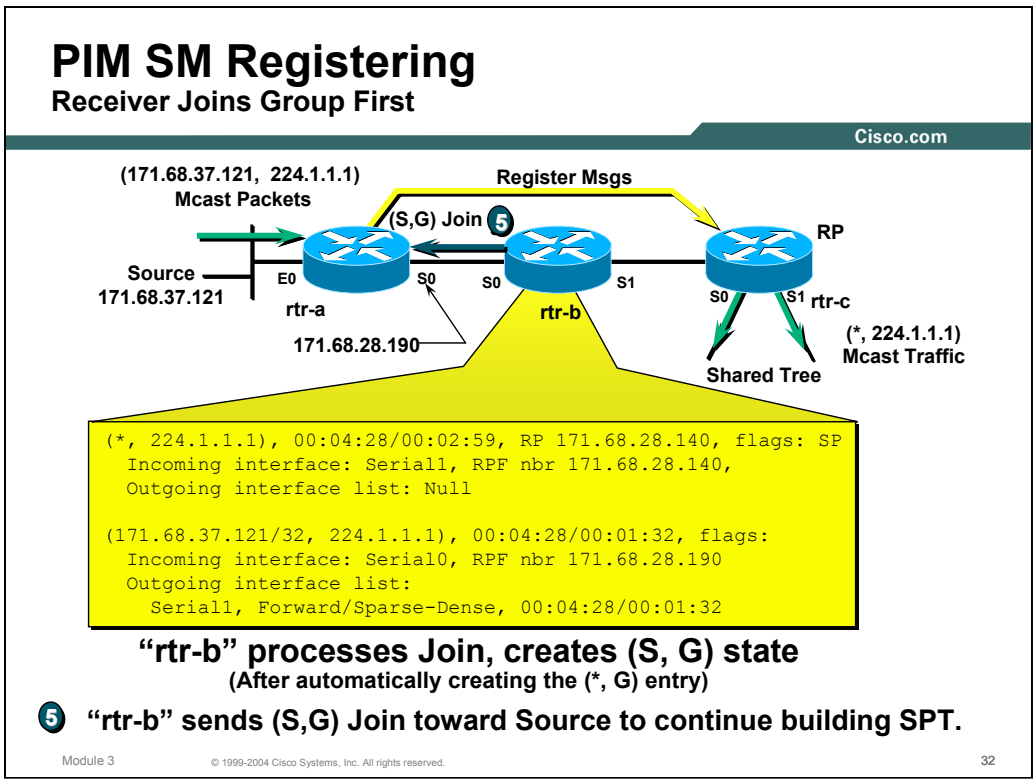

- **"rtr-b" processes the (S, G) Join and creates state**
  - A (*, G) entry must be created before the (S, G) entry can be created. Note that:
    - The RPF information for the (*, G) entry points up the Shared Tree via Serial1 with the RPF neighbor of 171.68.28.140. (Serial 3 of the RP.)
    - Because in this example no members have joined the group, the OIL of the (*, G) entry is Null.
    - The "P" flag (Pruned) is set since the OIL is Null.
  - The (S, G) entry is then created. Pay particular attention to the following:
    - The RPF information for this entry points towards the source via Serial0. The RPF neighbor is 171.68.28.190. (Serial 0 of "rtr-a".)
    - The (S, G) OIL initially receives a copy of the (*, G) OIL. (Which is Null.)
    - Interface Serial1 (which is the interface that received the (S, G) Join) is added to the (S, G) OIL.
    - The "T" flag is not yet set in the (S, G) entry. However, when the first (S, G) packet is forwarded using this entry, the flag will be "T" set.
  - 5) Because the OIL of the (S, G) transitioned from Null to non-Null (when "rtr-b" added Serial1 to the OIL of the newly created entry), a PIM (S, G) Join is sent to rtr-a's to continue the process of joining the SPT.

PIM SM Registering
Receiver Joins Group First

Cisco.com

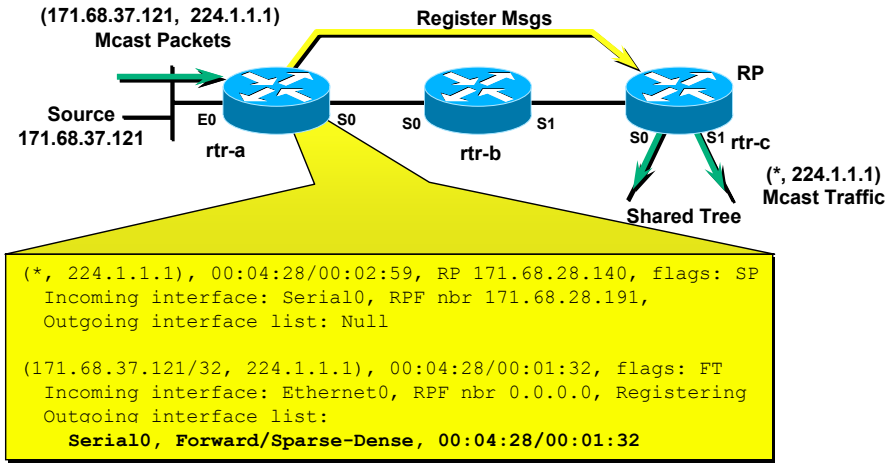

"rtr-a" processes the (S, G) Join; adds Serial 0 to OIL

Module 3 33

- **"rtr-a" processes the (S, G) Join**
  - Because an (S, G) entry already existed, "rtr-a" simply added the interface on which it received the (S, G) join to the OIL. This results in the following:
    - Serial0 is now listed in the "Outgoing interface list" (OIL) since the RP joined the SPT via this interface.
    - The "P" flag (Pruned) is cleared since the OIL is no longer Null.

PIM SM Registering

Receiver Joins Group First

Cisco.com

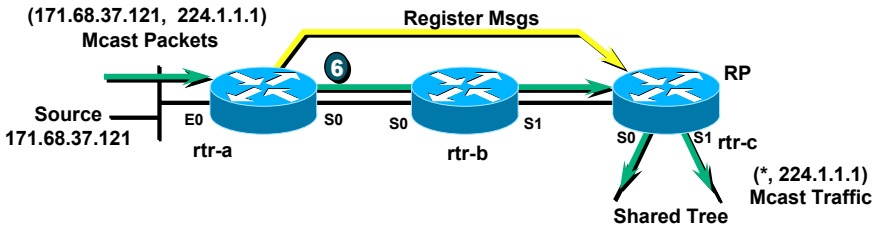

6 **RP begins receiving (S,G) traffic down SPT.**

- **A branch of the (S,G) SPT has been built to the RP.**
  - 6) Now that the SPT has been built from source “S” to the RP, traffic begins flowing down the Shortest-Path Tree (SPT). At this point, the RP is receiving the (S, G) traffic “natively” down the SPT. (This causes the “T” flags to be set in the (S, G) entries along this path including in the RP.)

PIM SM Registering

Receiver Joins Group First

Cisco.com

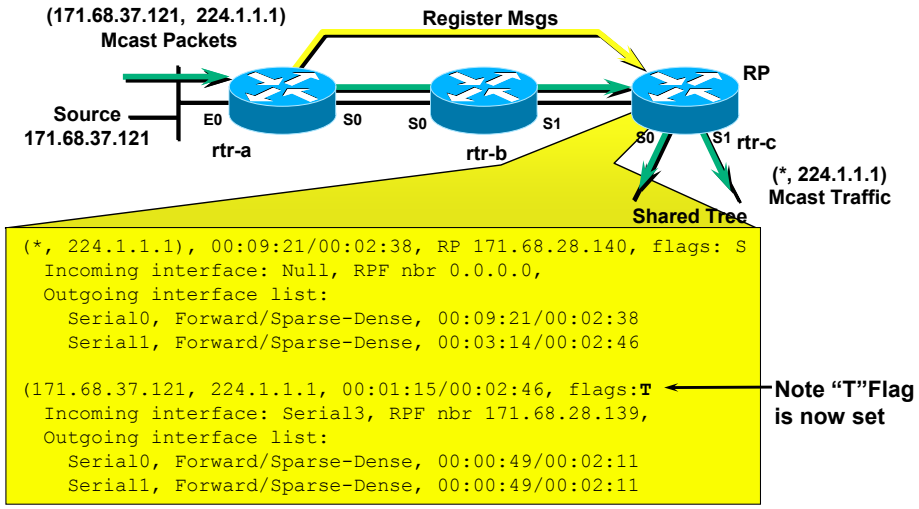


```
(*, 224.1.1.1), 00:09:21/00:02:38, RP 171.68.28.140, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0,
  Outgoing interface list:
    Serial0, Forward/Sparse-Dense, 00:09:21/00:02:38
    Serial1, Forward/Sparse-Dense, 00:03:14/00:02:46

(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.)

- **A branch of the (S,G) SPT has been built to the RP.**
  - The (S, G) traffic arriving via the SPT is forwarded down the Shared Tree via the (S, G) entry. This causes the "T" flag to be set in the (S, G) entry.

PIM SM Registering

Receiver Joins Group First

Cisco.com

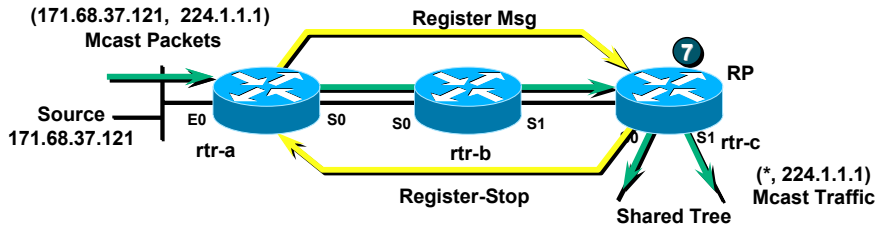

7 **Once "T" Flag is set, next "Register" causes RP to send back a "Register-Stop" to "rtr-a".**

- **A branch of the (S,G) SPT has been built to the RP.**
  - 7) Once the "T" flag is set in the (S,G) entry, the next (S,G) "Register" message causes the RP to send an (S, G) "Register-Stop" back to the 1st-hop router to inform it that the encapsulated group "G" Register messages from source "S" are no longer necessary.

PIM SM Registering
Receiver Joins Group First

Cisco.com

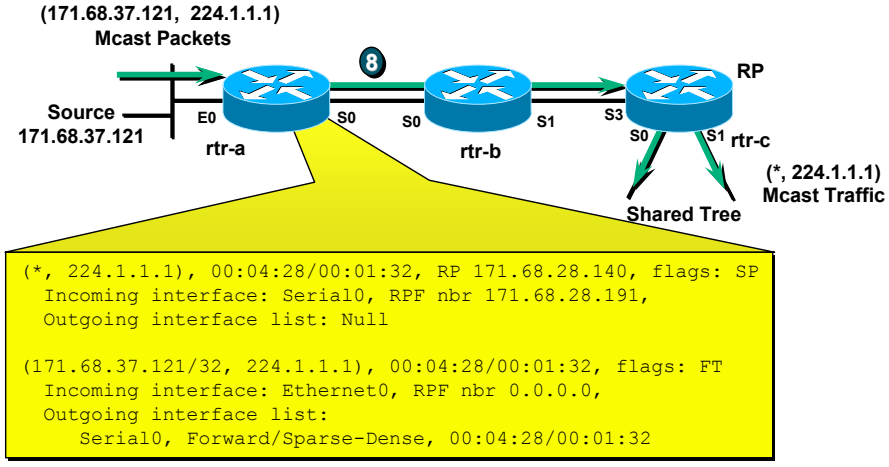


```
(*, 224.1.1.1), 00:04:28/00:01:32, RP 171.68.28.140, flags: SP
  Incoming interface: Serial0, RPF nbr 171.68.28.191,
  Outgoing interface list: Null

(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.**

- **"rtr-a" stops sending Register messages**
  - When the 1st-hop router (rtr-a) receives the (S, G) Register-Stop message it ceases sending encapsulated Register messages for (S, G) traffic.
    - Notice that the "Registering" flag on the second line of the (S, G) entry is no longer being displayed indicating that "rtr-a" is not sending Registers.
    - This is the final state in "rtr-a" after the Registration process.
  - 8) (S, G) traffic is now only flowing down the Shortest-Path Tree (SPT).

PIM SM Registering
Receiver Joins Group First

Cisco.com

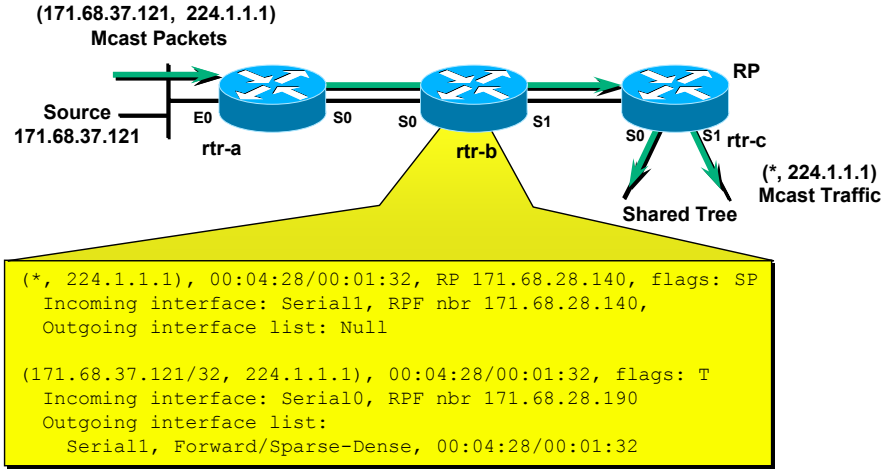


```
(*, 224.1.1.1), 00:04:28/00:01:32, RP 171.68.28.140, flags: SP
  Incoming interface: Serial1, RPF nbr 171.68.28.140,
  Outgoing interface list: Null

(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"

- **Final state in "rtr-b" after the Registration process**
  - Pay particular attention to the following in the (S, G) entry:
    - The "T" flag is now set indicating that (S, G) traffic is flowing along this path.
  - The (*, G) entry still has a Null OIL and the "P" flag is still set.
    - This is because there are no members that have joined the Shared Tree.

PIM SM Registering
Receiver Joins Group First

Cisco.com

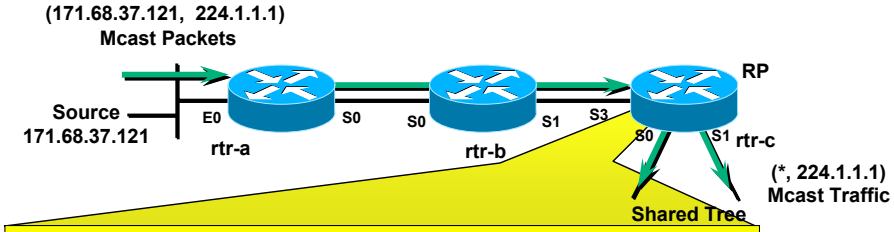


```
(*, 224.1.1.1), 00:09:21/00:02:38, RP 171.68.28.140, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0,
  Outgoing interface list:
    Serial0, Forward/Sparse-Dense, 00:09:21/00:02:38
    Serial1, Forward/Sparse-Dense, 00:03:14/00:02:46

(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
```

Final state in the "RP"
(with receivers on Shared Tree)

Module 3 39

- **Final state in the RP after the Registration process**
  - Pay particular attention to the following in the newly created (S, G) entry:
    - The "T" flag is now set indicating that (S, G) traffic is flowing along this path.

PIM SM Registering – Source Registers First

PIM SM Registering
Source Registers First

Cisco.com

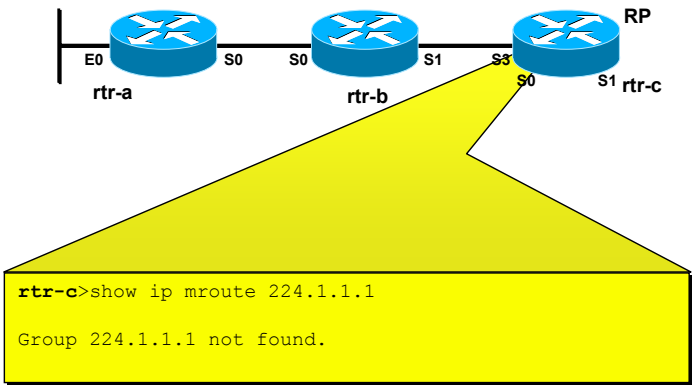

State in "RP" before Registering
(without receivers on Shared Tree)

- **State in RP before Registering (w/o Rcvr's on Shared Tree)**
  - Notice that no state for group "G" exists since there are no Receivers on the Shared Tree yet.

PIM SM Registering
Source Registers First

Cisco.com

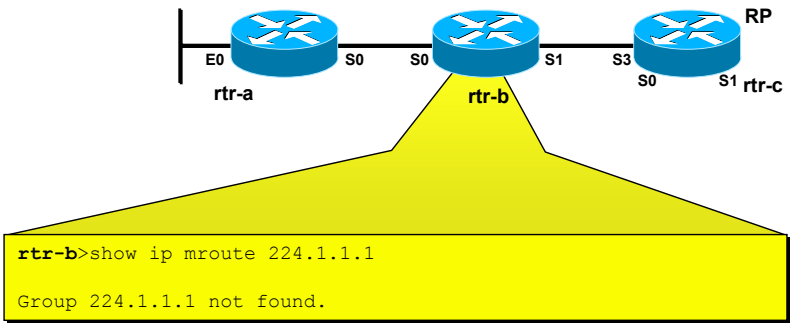

State in "rtr-b" before any source registers
(with receivers on Shared Tree)

- **State in "rtr-b" before source registers**
  - Note that there is no group state information for this Group yet.

PIM SM Registering
Source Registers First

Cisco.com

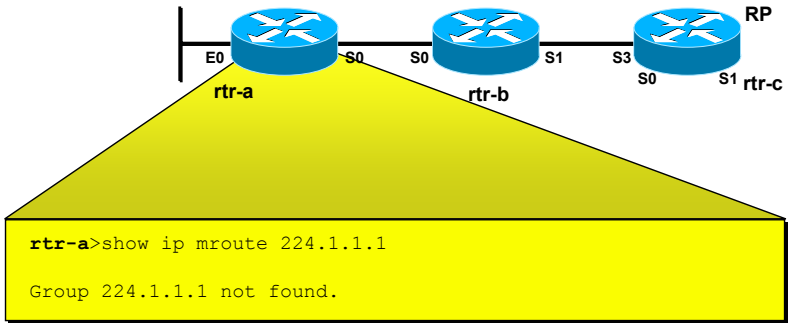

State in "rtr-a" before any source registers
(with receivers on Shared Tree)

- **State in 1st-hop router (rtr-a) before source registers**
  - Note that there is no group state information for this Group yet.

PIM SM Registering
Source Registers First

Cisco.com

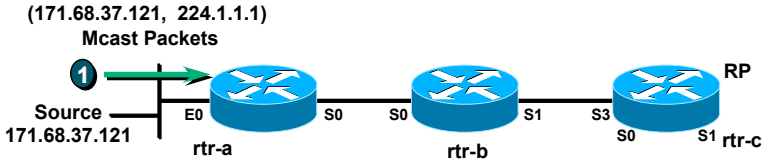

1 **"Source" begins sending group G traffic.**

- **Source Registers First Example**
  - 1) Source "S" begins sending traffic to group "G".
  - 2) 1st-hop router ("rtr-a") creates (*, G) and (S, G) state; encapsulates the multicast packets in PIM Register message(s) and unicasts it(them) to the RP.
  - 3) The RP ("rtr-c") de-encapsulates the (S, G) packet and creates (*, G) and (S, G) state. Since no one has joined the Shared Tree yet, the OIL's of these entries will be NULL.. Because the OIL of the (S, G) entry (just created) is NULL, the packet is discarded.

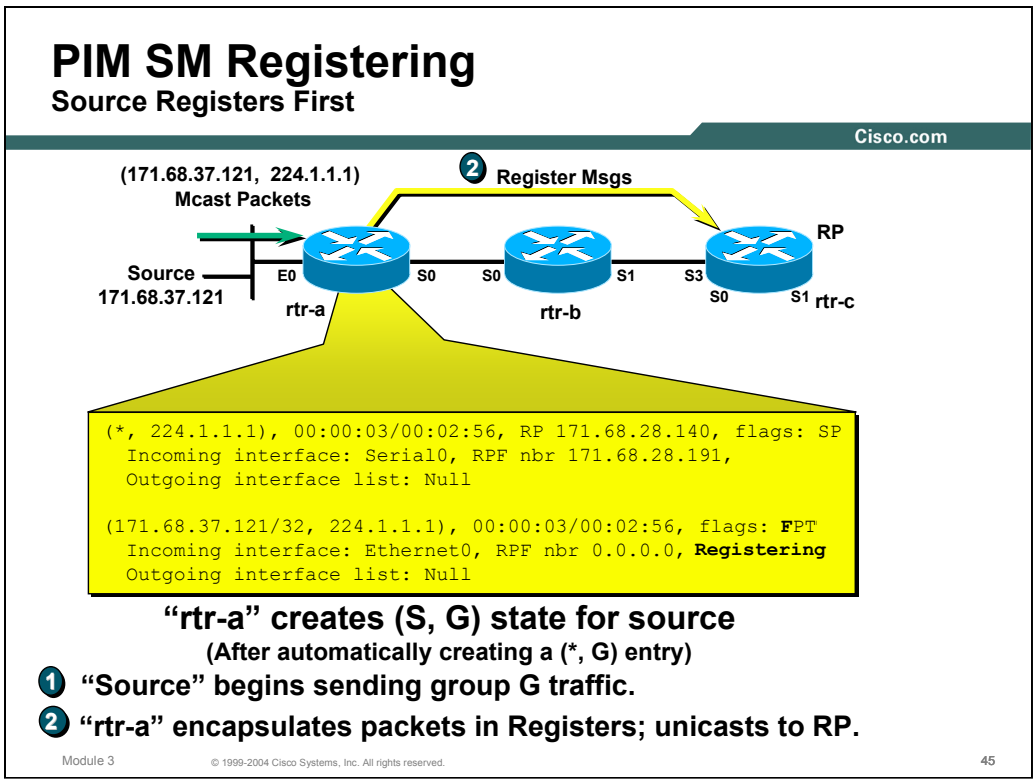

- **1st-hop router (rtr-a) creates (S, G) state**
  - A (*, G) entry must be created before the (S, G) entry can be created. Note that:
    - The RPF information for this entry points up the Shared Tree via Serial0 with the RPF neighbor of 171.68.28.191. (Serial 0 of "rtr-b".)
    - Because in this example no members have joined the group (the sender is only sending), the OIL of the (*, G) entry is Null.
    - The "P" flag (Pruned) is set since the OIL is Null.
  - The (S, G) entry is then created. Pay particular attention to the following:
    - The RPF information for this entry points towards the source via Ethernet0. The RPF neighbor is 0.0.0.0 because the source is directly connected.
    - The (S, G) OIL receives a copy of the (*, G) OIL. (Which is Null.)
    - The "F" flags are set in the (S, G) entry which indicates that this is a directly connected Source.
    - The "Registering" flag is set in the (S, G) entry which indicates that we are still sending Register messages to the RP for this Source.
    - The "P" flag (Pruned) is set since the OIL is Null.
  - 2) The 1st-hop router encapsulates the multicast packets in PIM Register message(s) and unicasts them to the RP.

PIM SM Registering
Source Registers First

Cisco.com

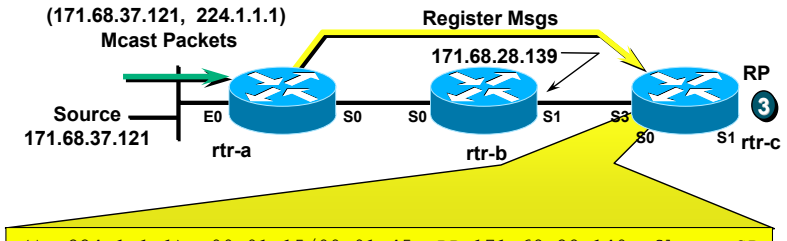


```
(*, 224.1.1.1), 00:01:15/00:01:45, RP 171.68.28.140, flags: SP
  Incoming interface: Null, RPF nbr 0.0.0.0,
  Outgoing interface list: Null

(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
```

"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.

- **The RP creates (S, G) state**
  - As a result of the Register message that was received from "rtr-a", the RP creates (*, G) and (S, G) state. However, because no previous (*, G) state existed, it must be created before the (S,G) entry can be created.
    - This (*, G) entry is created as shown above. Notice that the (*, G) OIL is NULL. This is because the RP has not yet received any (*, G) Joins for this group. (Remember, in this example, the source registers first.)
  - Next, the (S, G) entry can be created and is accomplished as follows:
    - The RPF information is calculated using the source address contained in the multicast packet encapsulated inside of the register message. This results in an IIF of Serial3 and an RPF neighbor of 171.68.28.139.
    - Next, the OIL of the parent (*, G) entry is copied into the OIL of the new (S,G) entry. Since the OIL of the (*, G) entry is NULL, this results in a NULL (S, G) OIL.
  - Now the router is ready to forward the (S, G) packet that was encapsulated in the Register message using the newly created (S, G) state. This is accomplished as follows:
    - Because this packet was received inside of a Register message, the RPF check is skipped.
    - Next, the router forwards a copy of the packet out all interfaces in the matching (S, G) OIL. However, because the (S, G) OIL is NULL (i.e. there are no branches of the Shared Tree), the packet is simply discarded.

PIM SM Registering
Source Registers First

Cisco.com

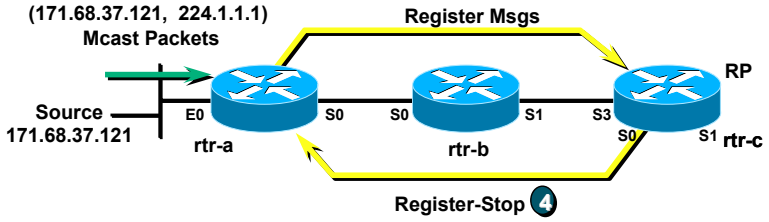

4 RP sends "Register-Stop" to "rtr-a".

- **Source Registers First Example**
  - 4) Since the RP has no (*, G) state and hence no receivers on the Shared Tree, it does not need the (S, G) traffic. Therefore the RP sends an (S, G) "Register-Stop" message to the 1st-hop router so it will stop sending Register messages.

PIM SM Registering
Source Registers First

Cisco.com

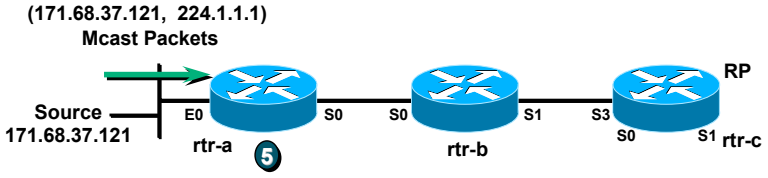

5 **"rtr-a" stops encapsulating traffic in Register Messages; drops packets from Source.**

- **Source Registers First Example**
  - 5)The 1st-hop router receives the (S, G) Register-Stop message and stops sending Register messages for (S, G) traffic.

    Note: Eventually, the original (S, G) entry will time out (approx. 3 min.) and be deleted. The Register process will start over again when the 1st-hop router receives the next multicast packet from directly connected source "S". The RP will again respond with a Register-Stop which will prevent the (S,G) traffic from flowing to the RP until it is needed.

PIM SM Registering
Source Registers First

Cisco.com

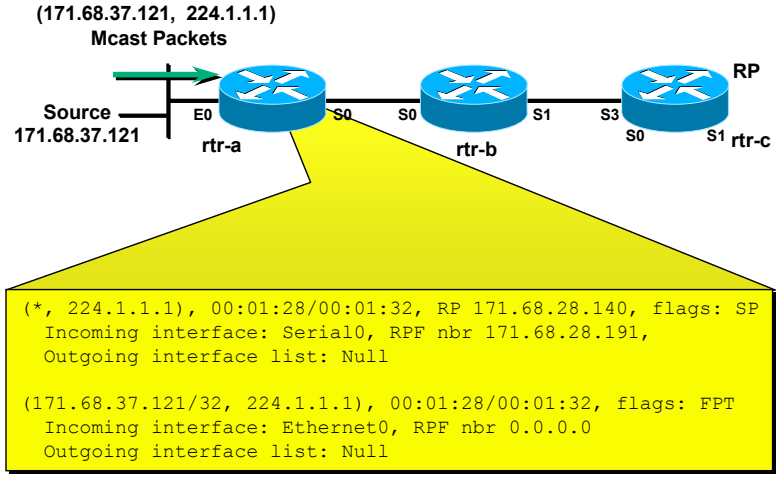

State in "rtr-a" after Registering
(without receivers on Shared Tree)

- **State in 1st-hop router after Registering (w/o Rcvr's on Shared Tree)**
  - Pay particular attention to the following in the (S, G) entry:
    - The "Registering" flag is now cleared.
    - The "Outgoing interface list" is still Null since the RP did not join the SPT.
    - The "P" flag (Pruned) is still set since the oilist is still Null.
    - The "00:01:32" Expiration time value will count down to zero at which time the (S, G) entry will be deleted. (The Register process will begin all over again when the next multicast packet is received from source "S".)

PIM SM Registering
Source Registers First

Cisco.com

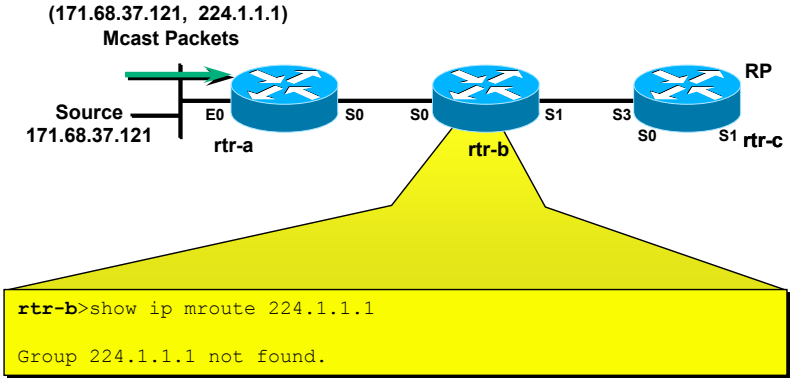

State in "rtr-b" after "rtr-a" Registers
(without receivers on Shared Tree)

- **State in "rtr-b" after Registering (w/o Rcvr's on Shared Tree)**
  - Notice that no state exists in "rtr-b" at this point in time.

PIM SM Registering
Source Registers First

Cisco.com

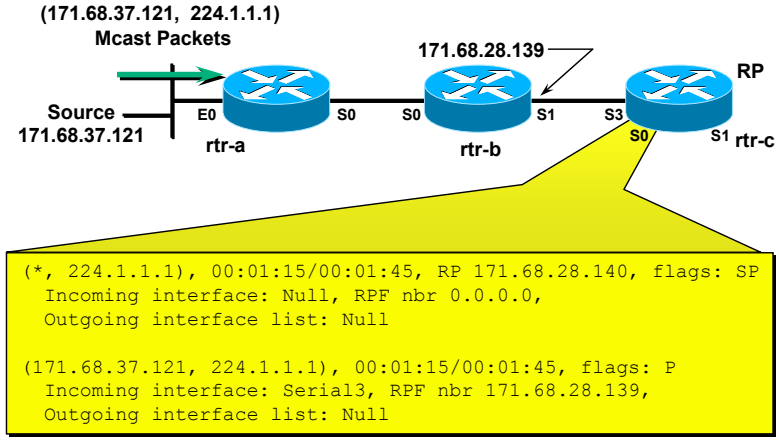

State in "RP" after "rtr-a" Registers
(without receivers on Shared Tree)

- **State in RP after Registering (w/o Rcvr's on Shared Tree)**
  - Pay particular attention to the following in the newly created (S, G) entry:
    - The "RPF nbr" is the IP Address of "rtr-b".
    - The "Incoming interface:" is Serial3 which is the RPF interface towards source "S" via "rtr-b".
    - The "Outgoing interface list:" is Null since the (*,G) OIL is also Null. (Indicates there are no Receivers on the Shared Tree yet.)
    - The "P" flag (Pruned) is set since the OIL is Null.
  - The (S,G) state will remain in the RP as long as the source is still actively sending. This is accomplished by fact that the first-hop route will continue sending periodic Register messages to the RP as long as the first-hop router is receiving traffic from the source.

PIM SM Registering
Source Registers First

Cisco.com

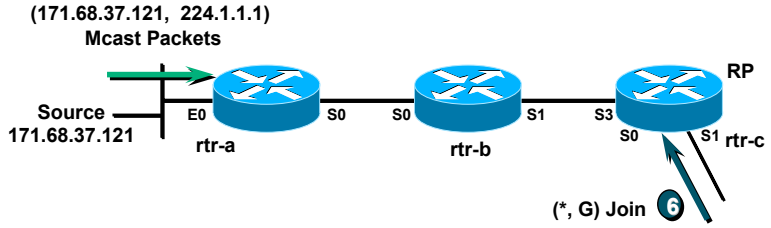

Receivers begin joining the Shared Tree

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

- **Source Registers First Example**
  - 6)The RP now begins receiving (*, G) Joins from Last-hop routers with Receivers that wish to join the Shared Tree.

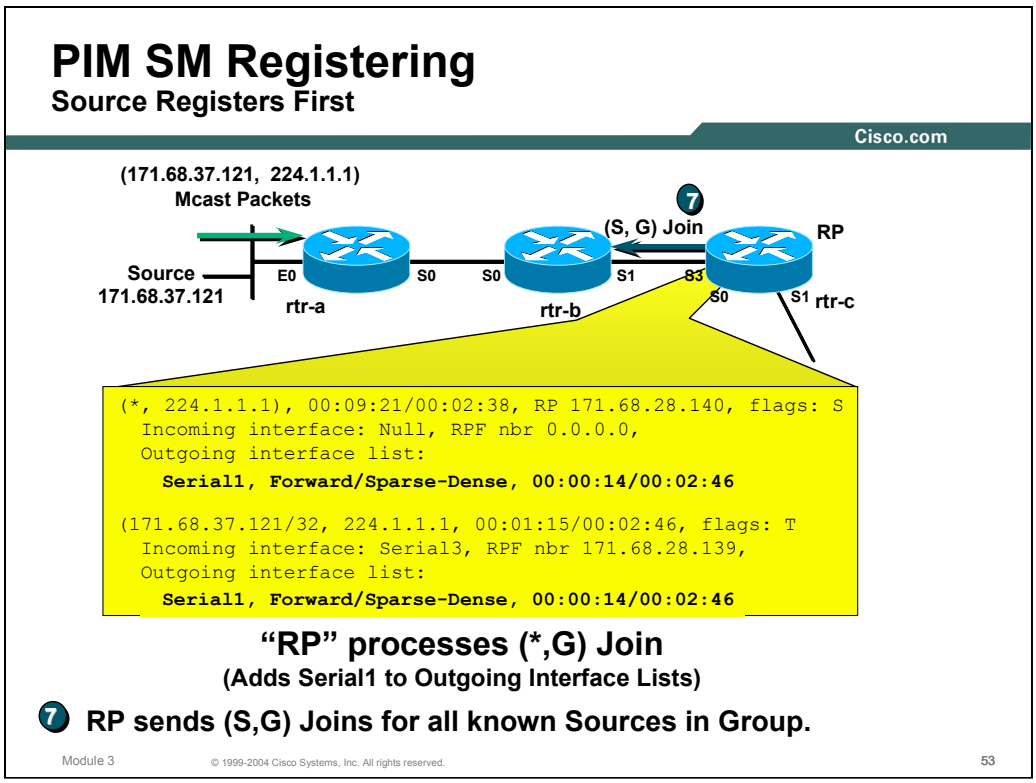

- **The RP process the (*, G) Join**
  - In the (*, G) entry:
    - Serial1 has been added to the (*, G) entry since a (*,G) Join was received on this interface which is the only active branch of the Shared Tree (RPT).
  - In the (S, G) entry:
    - Serial1 has also been added to the (S, G) OIL because the OIL's of all (S,G) entries are always kept in sync with their parent (*, G).
    - Note: When the (S, G) OIL's are synchronized with the OIL of their parent (*,G) OIL, a check is made to insure that the IIF of the (S, G) does not appear in the OIL of the (S, G). This could result in a route loop.
  - 7) The transitioning of the (*, G) OIL from Null to non-Null triggers the RP to scan its list of (S, G) entries for group "G" and send (S, G) Joins towards all sources. (This will cause a SPT to be built from each active source back to the RP which will eventually start the flow of (S, G) traffic to the RP.)

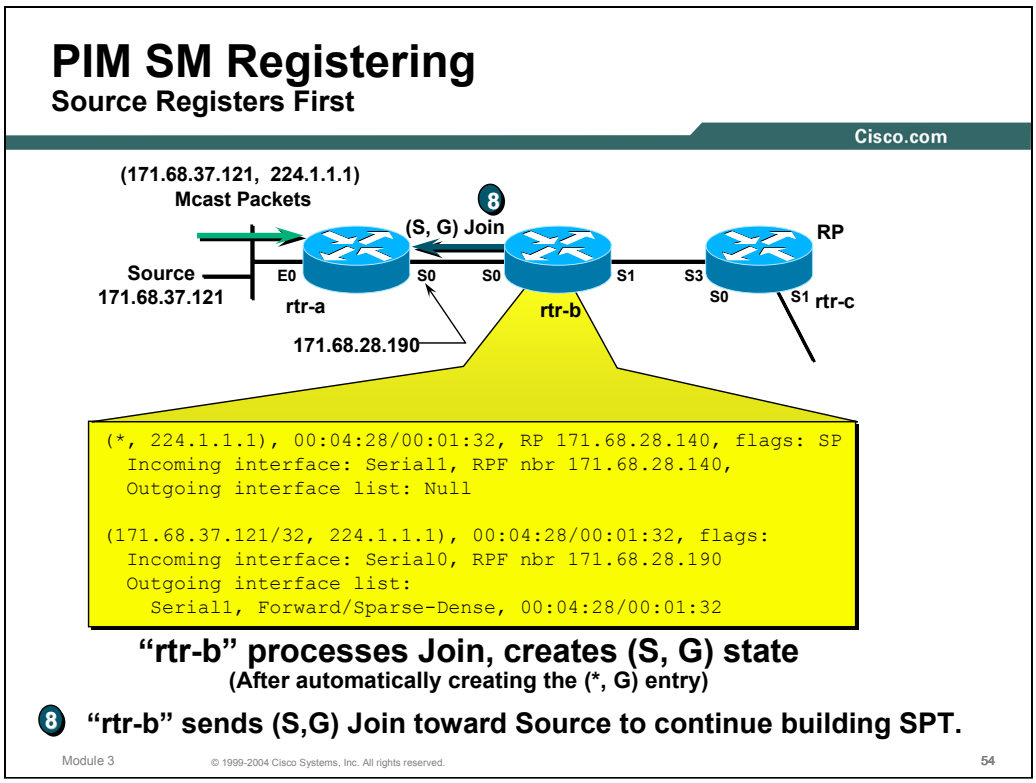

- **"rtr-b" processes the (S, G) Join and creates state**
  - A (*, G) entry must be created before the (S, G) entry can be created. Note that:
    - The RPF information for the (*, G) entry points up the Shared Tree via Serial1 with the RPF neighbor of 171.68.28.140. (Serial 3 of the RP.)
    - Because in this example no members have joined the group, the OIL of the (*, G) entry is Null.
    - The "P" flag (Pruned) is set since the OIL is Null.
  - The (S, G) entry is then created. Pay particular attention to the following:
    - The RPF information for this entry points towards the source via Serial0. The RPF neighbor is 171.68.28.190. (Serial 0 of "rtr-a".)
    - The (S, G) OIL initially receives a copy of the (*, G) OIL. (Which is Null.)
    - Interface Serial1 (which is the interface that received the (S, G) Join) is added to the (S, G) OIL.
    - The "T" flag is not yet set in the (S, G) entry. However, when the first (S, G) packet is forwarded using this entry, the flag will be "T" set.
  - 8) Because the OIL of the (S, G) transitioned from Null to non-Null (when "rtr-b" added Serial1 to the OIL of the newly created entry), a PIM (S, G) Join is sent to rtr-a's to continue the process of joining the SPT.

PIM SM Registering

Source Registers First

Cisco.com

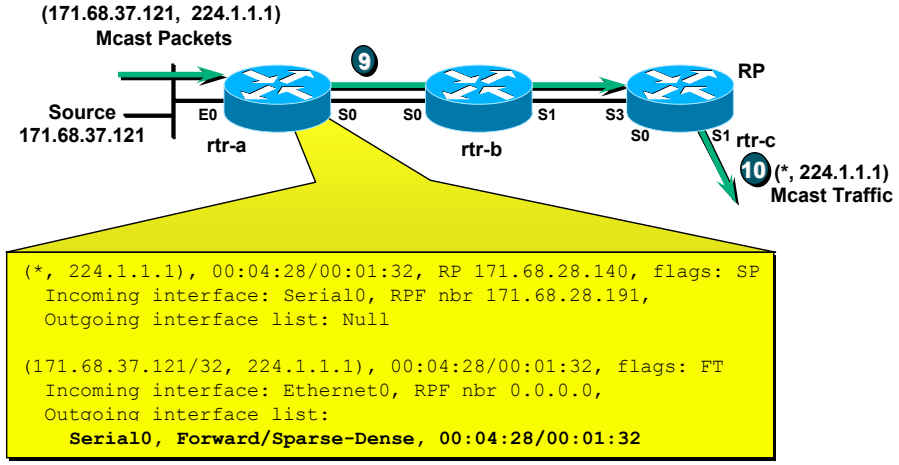


```
(*, 224.1.1.1), 00:04:28/00:01:32, RP 171.68.28.140, flags: SP
  Incoming interface: Serial0, RPF nbr 171.68.28.191,
  Outgoing interface list: Null

(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.**

- **1st-hop router (rtr-a) processes the (S, G) Join**
  - The (S, G) Join is processed as follows:
    - Serial0 is added to the "Outgoing interface list" (OIL). (This is the interface on which the (S, G) Join arrived.)
    - The "P" flag (Pruned) is cleared since the OIL is no longer Null.
  - 9) As a result of Serial0 being added to the (S, G) OIL, traffic begins to flow down the SPT from the source to the RP.
  - 10) The RP then forwards all incoming (S, G) traffic to the Receivers down the Shared Tree.

PIM SM Registering
Source Registers First

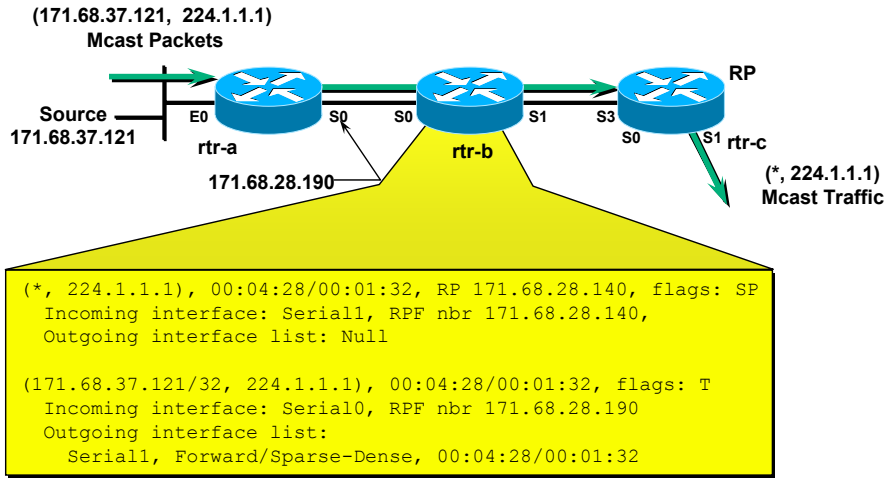

Final state in "rtr-b" after Receivers Join

- **State in "rtr-b" after Receivers Join**
  - Pay particular attention to the following:
    - Both (*, G) and (S, G) state was created as a result of the (S, G) Join received from the RP.
    - The "P" flag set in the (*, G) entry since there are no receivers on the Shared Tree at this point in the network.
    - The "T" flag is set in the (S, G) entry indicating that traffic is flowing down the Shortest-Path Tree.
    - The "RPF nbr" is the IP Address of "rtr-a".
    - Serial0 is the "Incoming interface" of the (S, G) entry since this is the RPF interface for source "S" via "rtr-a".
    - Serial1 is listed in the "Outgoing interface list" of the (S, G) entry since the RP joined the SPT via this interface.

PIM SM Registering
Source Registers First

Cisco.com

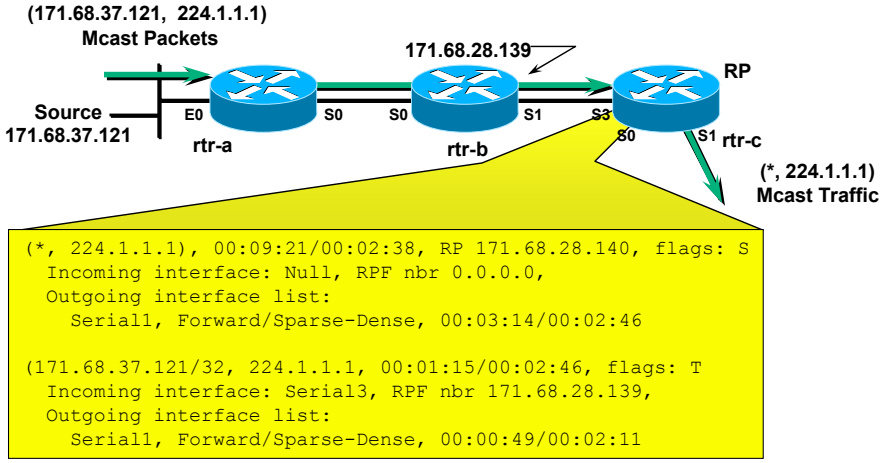

Final state in "RP" after Receivers Join

- **State in RP after Receivers Join**
  - In the (*, G) entry:
    - Serial1 has been added to the (*, G) entry since a (*,G) Join was received on this interface which is the only active branch of the Shared Tree (RPT).
  - In the (S, G) entry:
    - Serial1 has also been added to the (S, G) OIL because the OIL's of all (S,G) entries are always kept in sync with their parent (*, G).
    - Note: When the (S, G) OIL's are synchronized with the OIL of their parent (*, G) OIL, a check is made to insure that the IIF of the (S, G) does not appear in the OIL of the (S, G). This could result in a route loop.

PIM SM Registering – Receivers Along the SPT

PIM SM Registering
Receivers along the SPT

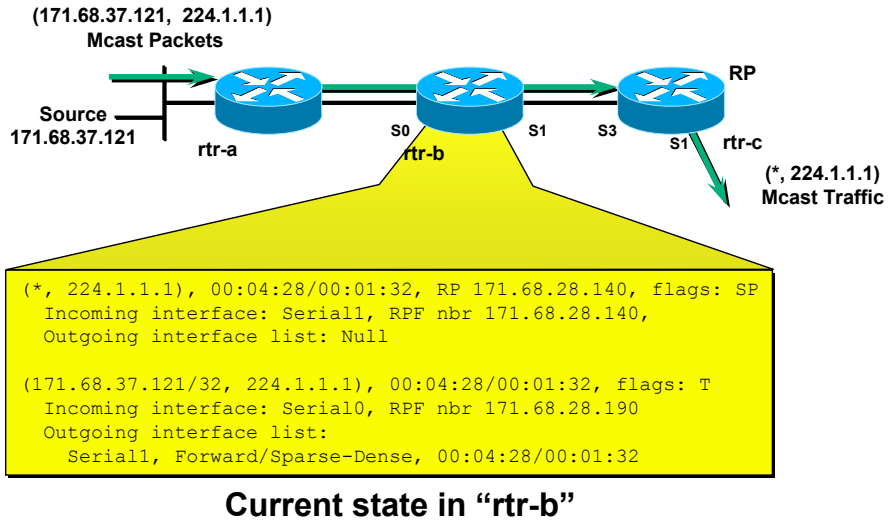

Current state in "rtr-b"

- **State in "rtr-b" with traffic flowing on the SPT**
  - Pay particular attention to the following:
    - Both (*, G) and (S, G) state was created as a result of the (S, G) Join received from the RP.
    - The "P" flag set in the (*, G) entry since there are no receivers on the Shared Tree at this point in the network.
    - The "T" flag is set in the (S, G) entry indicating that traffic is flowing down the Shortest-Path Tree.
    - The "RPF nbr" is the IP Address of "rtr-a".
    - Serial0 is the "Incoming interface" of the (S, G) entry since this is the RPF interface for source "S" via "rtr-a".
    - Serial1 is listed in the "Outgoing interface list" of the (S, G) entry since the RP joined the SPT via this interface.

PIM SM Registering
Receivers along the SPT

Cisco.com

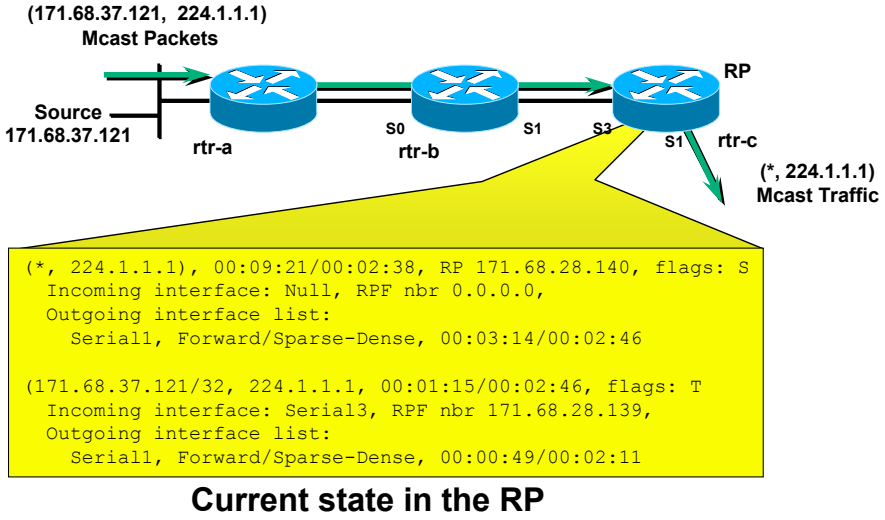

Current state in the RP

- **State in the RP with traffic flowing on the SPT**
  - Pay particular attention to the following:
    - The (*, G) entry only has Serial1 in its outgoing interface list.
    - In the (S, G) entry, Serial0 is the “Incoming interface” since this is the RPF interface for source “S” via “rtr-b”.
    - Serial1 is listed in the “Outgoing interface list” of the (S, G) entry because the OIL of the (S, G) entry is always kept in sync with the (*, G) OIL.

PIM SM Registering
Receivers along the SPT

Cisco.com

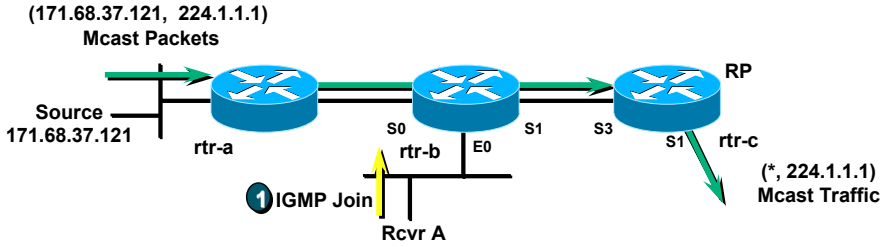

1 **"Rcvr A" wishes to receive group G traffic. Sends IGMP Join for G.**

- **Receivers along the SPT**
  - Step 1: A host directly connected to "rtr-b", Receiver "A", joins multicast group 224.1.1.1 by sending an IGMP Report.

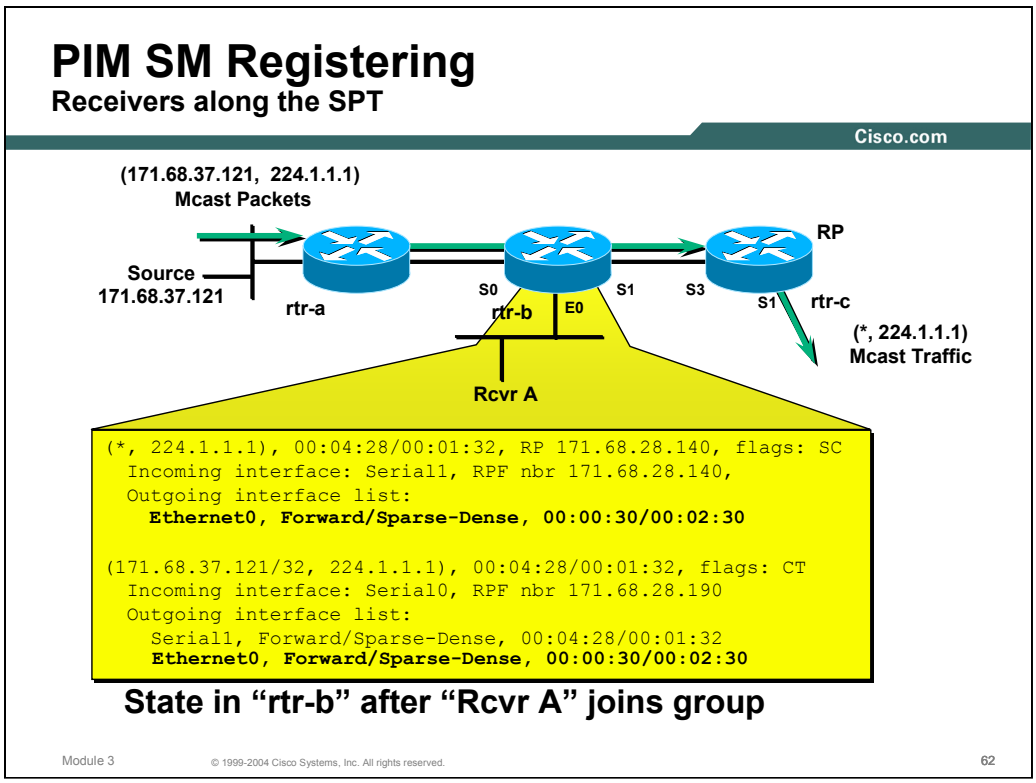

- **Receivers along the SPT**
  - As a result of the IGMP Report sent by Receiver A for group 224.1.1.1, “rtr-b” updates its state for this group as follows.
    - Ethernet0 is added to the OIL of the (*, G) entry. This is done to permit any (*, 224.1.1.1) traffic flowing down the Shared Tree to be forwarded to Receiver “A”.
    - Next, the OIL’s of all child (S, G) entries are synchronized with the OIL change just made to the OIL of the (*, G). This results in Ethernet0 being added to the OIL of the (171.68.37.121/32, 224.1.1.1) entry. This permits traffic from this source to be “picked off” as it flows along the SPT through “rtr-b” on its way to the RP. (Note that this traffic does not flow to the RP and then back out the same interface to reach “rtr-b”. This is a common misperception.)

PIM SM Registering
Receivers along the SPT

Cisco.com

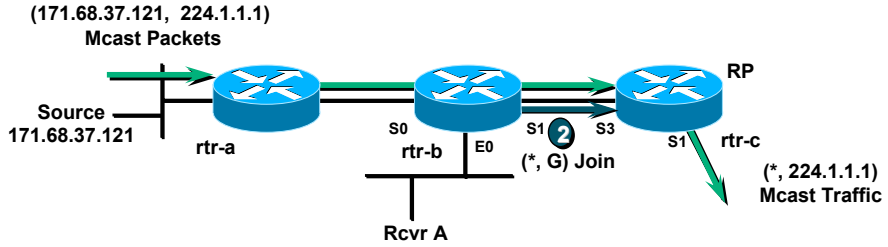

2 "rtr-b" triggers a (*,G) Join to join the Shared Tree

- **Receivers along the SPT**
  - Step 2: Because the OIL of the (*, G) entry in "rtr-b" transitioned from NULL to non-Null (Ethernet0 is now in the (*, G) OIL), a (*, G) Join message is triggered. This message is sent up the Shared Tree so that "rtr-b" will be placed on a branch of the Shared Tree.

PIM SM Registering
Receivers along the SPT

Cisco.com

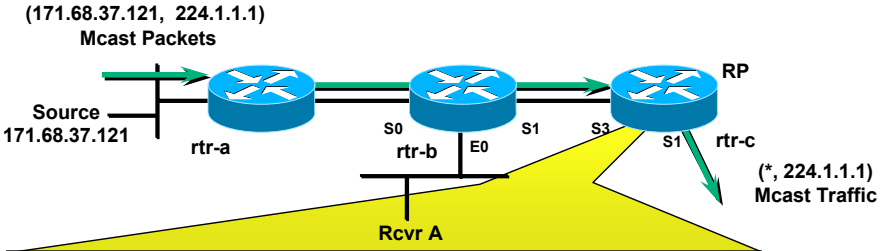


```
(*, 224.1.1.1), 00:09:21/00:02:38, RP 171.68.28.140, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0,
  Outgoing interface list:
    Serial1, Forward/Sparse-Dense, 00:03:14/00:02:46
    Serial3, Forward/Sparse-Dense, 00:00:10/00:02:50

(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

Module 3 64

- **Receivers along the SPT**
  - When the RP receives the (*, G) Join sent by "rtr-b", it adds Serial3 to the (*, G) OIL.
  - Next, the RP synchronizes the OIL's of all (S, G) entries by adding Serial3 to each (S, G) OIL. However in this case, Serial3 is the Incoming interface for the (171.68.37.121/32, 224.1.1.1) entry and is therefore not added to the OIL. (If it were, Serial3 would appear in both the incoming and outgoing interface list which could cause a route loop.)

PIM SM Registering
Receivers along the SPT

Cisco.com

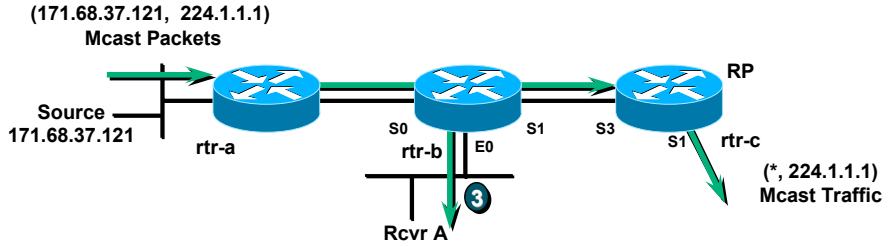

3 **Group G traffic begins to flow to "Rcvr A".**

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

- **Receivers along the SPT**
  - Step 3: Traffic from source 171.68.37.121 is now being "picked off" by "rtr-b" and forwarded out Ethernet0 as the traffic flows down the SPT to the RP.
  - Again, it is important to note that this source traffic does not flow to the RP and then turn around and come back out on the same interface that it arrived. (Refer to the state in the RP shown on the previous page.)

PIM SM SPT-Switchover

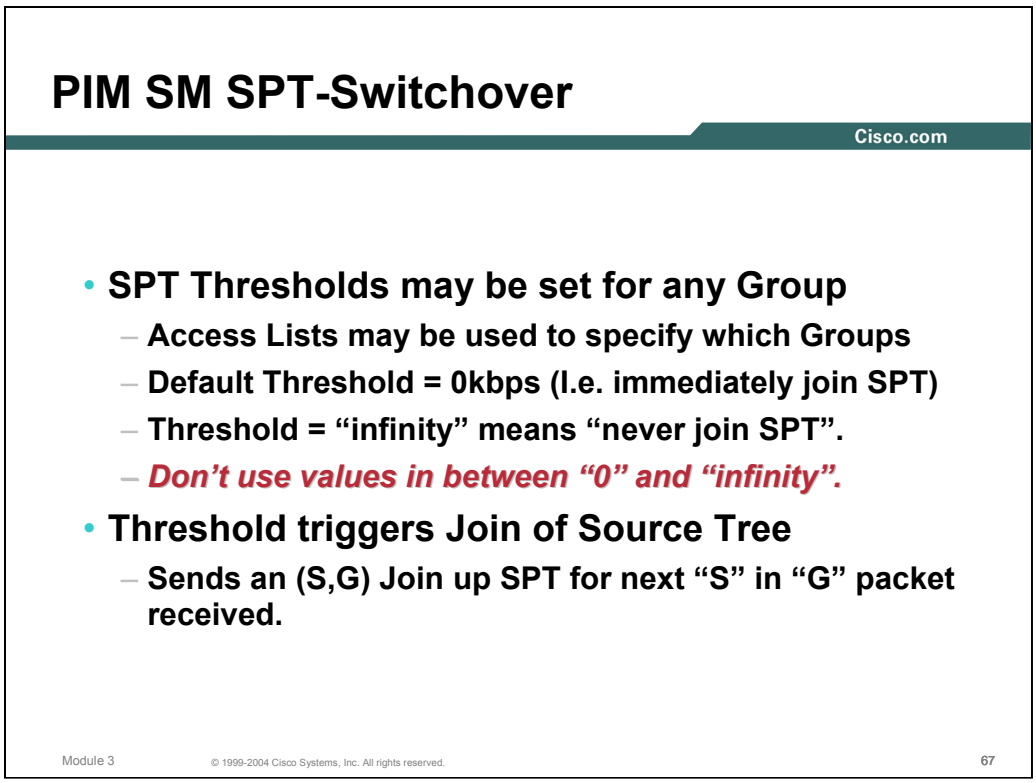

- **SPT Thresholds**
  - In PIM Sparse mode, SPT Thresholds may be configured to control when to switch to the Shortest-Path Tree (SPT).
  - SPT Thresholds are specified in Kbps and can be used with Access List to specify to which Group(s) the Threshold applies.
  - The default SPT-Threshold is 0Kbps. This means that any and all sources are immediately switched to the Shortest-Path Tree.
  - If an SPT-Threshold of “Infinity” is specified for a group, the sources will not be switched to the Shortest-Path Tree (SPT) and will remain on the Shared Tree.
- **Exceeding the Threshold**
  - When the Group’s SPT-Threshold is exceed in a last-hop router, the next received packet for the group will cause an (S, G) Join to be sent toward the source of the packet. This builds a Shortest-Path Tree from the source “S” to the last-hop router.
- **PROS**
  - By switching to the Shortest-Path Tree (SPT), the most optimal (usually) path is used to deliver the multicast traffic. Depending on the location of the source in relation to the RP, this switch to the SPT can reduce network latency substantially.
- **CONS**
  - In networks with large numbers of senders (remember most multicast applications such as IP/TV Client, send RTCP multicast packets in the background and are therefore senders), an increased amount of state must be kept in the routers. In some cases, an Infinity threshold may be used to force certain groups to remain on the Shared Tree when latency is not an issue.

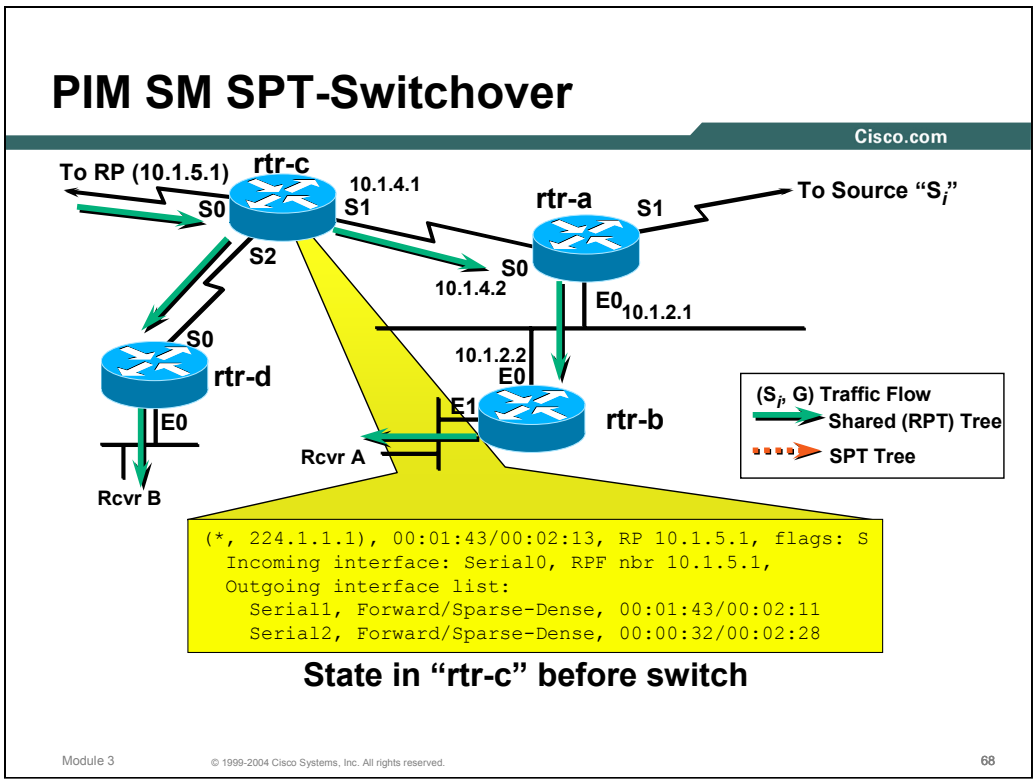

- **PIM-SM SPT-Switchover Example**
  - Receivers A & B have joined multicast group 224.1.1.1 which has resulted in traffic flowing down the Shared Tree as shown by the solid arrows.
  - The state is "rtr-c" prior to the switchover is as follows:
    - The IIF of the (*, G) entry points toward the RP via Serial0.
    - The OIL of the (*, G) entry contains Serial1 and Serial2 as a result of (*, G) Joins that were sent up the Shared Tree by "rtr-a" and "rtr-d", respectively.

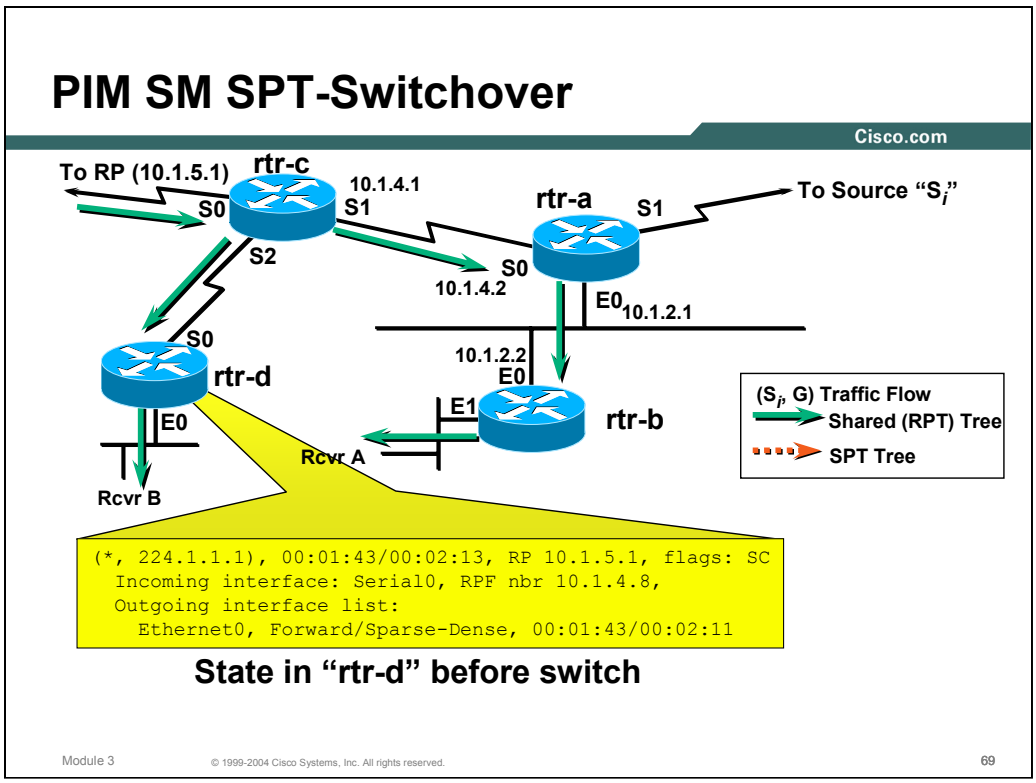

- **PIM-SM SPT-Switchover Example**
  - The state is "rtr-d" prior to the switchover is as follows:
    - The IIF of the (*, G) entry points toward the RP via Serial0.
    - The OIL of the (*, G) entry contains Ethernet0 as a result of the IGMP Reports for group 224.1.1.1 that are sent by Receiver "B".

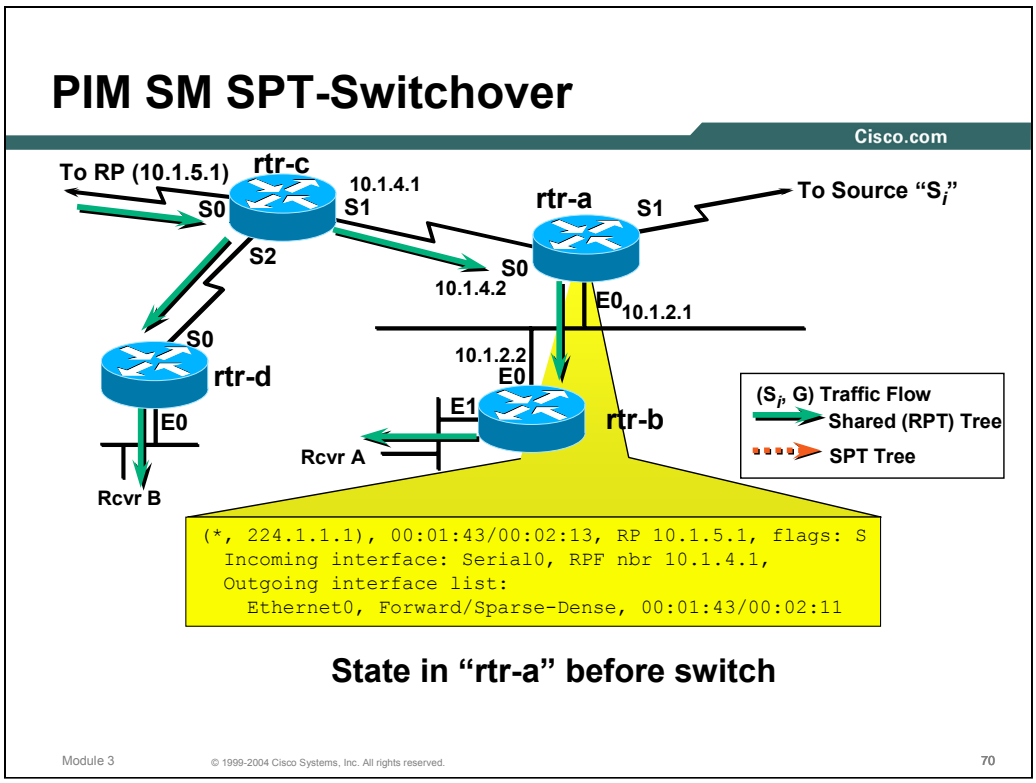

- **PIM-SM SPT-Switchover Example**
  - The state is "rtr-a" prior to the switchover is as follows:
    - The IIF of the (*, G) entry points toward the RP via Serial0.
    - The OIL of the (*, G) entry contains Ethernet0 as a result of (*, G) Joins that were sent up the Shared Tree by "rtr-b".

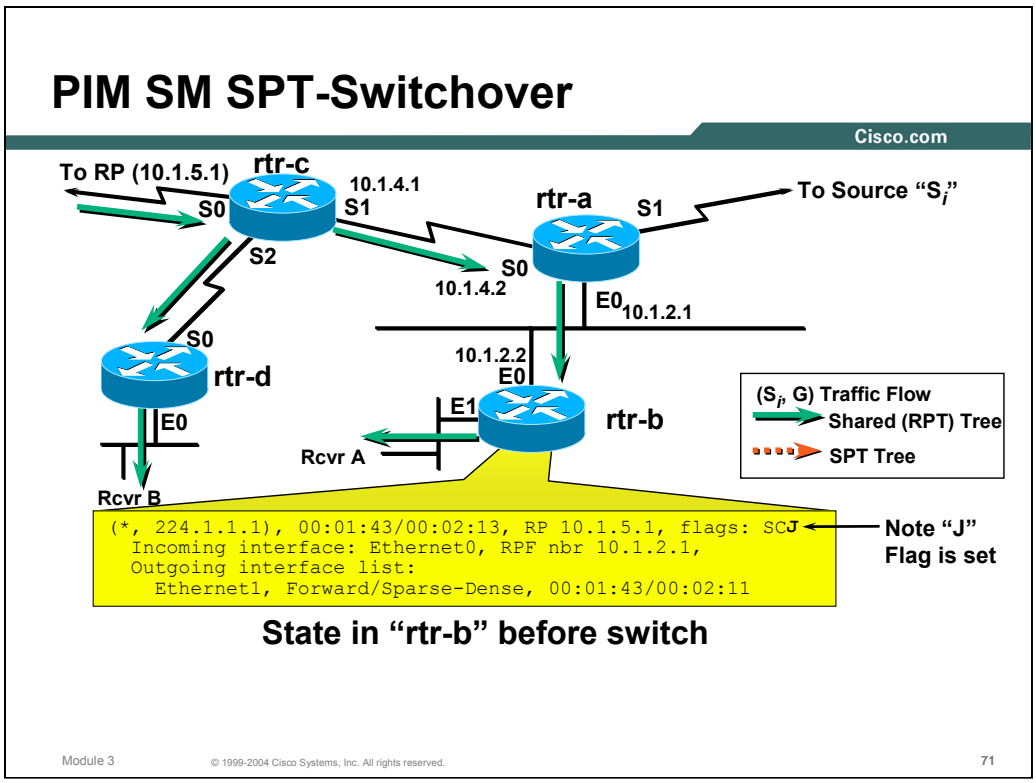

- **PIM-SM SPT-Switchover Example**
  - The state is "rtr-b" prior to the switchover is as follows:
    - The IIF of the (*, G) entry points toward the RP via Ethernet0.
    - The OIL of the (*, G) entry contains Ethernet1 as a result of the IGMP Reports for group 224.1.1.1 that are sent by Receiver "A".
    - Note that the "J" flag is set on the (*,G) entry. This indicates that the SPT Threshold is zero. (To be precise, when the SPT Threshold is zero, it is being exceeded at all times and is therefore always set.)

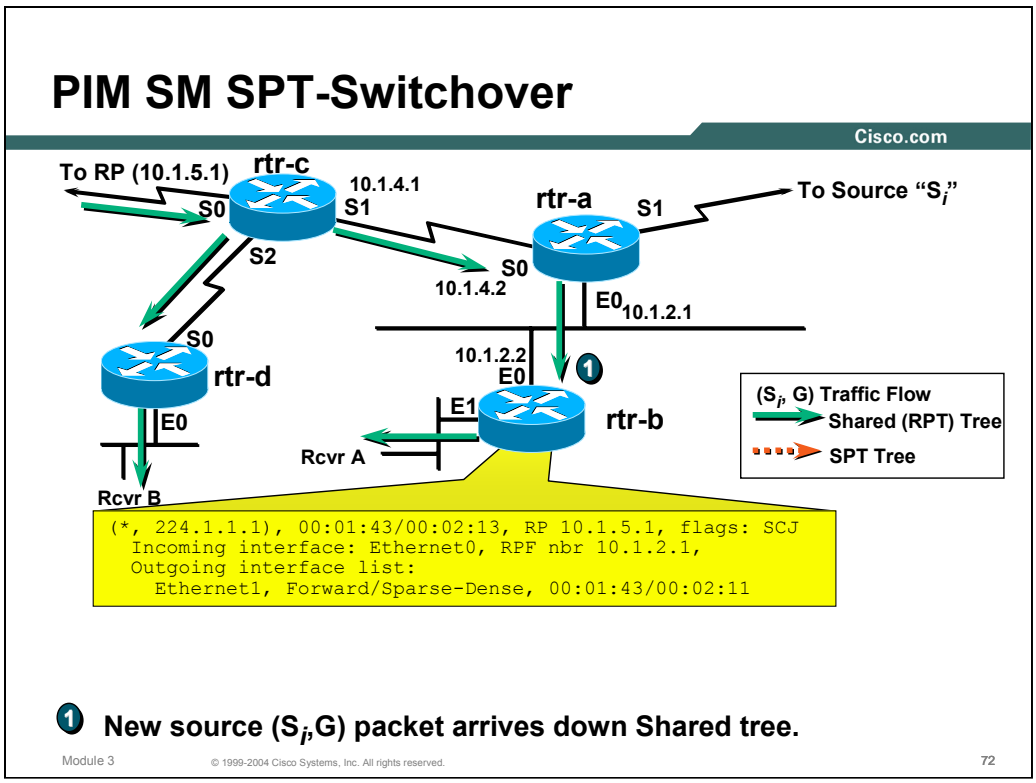

- **PIM-SM SPT-Switchover Example**
  - Step 1: The very next packet to arrive via the Shared Tree happens to be from source (Si, G). Because there is a member directly connected to this router (denoted by the "C" flag) and the traffic rate is above the SPT-Threshold (denoted by the "J" flag), "rtr-b" initiates a switch to the SPT for (Si, G) traffic.

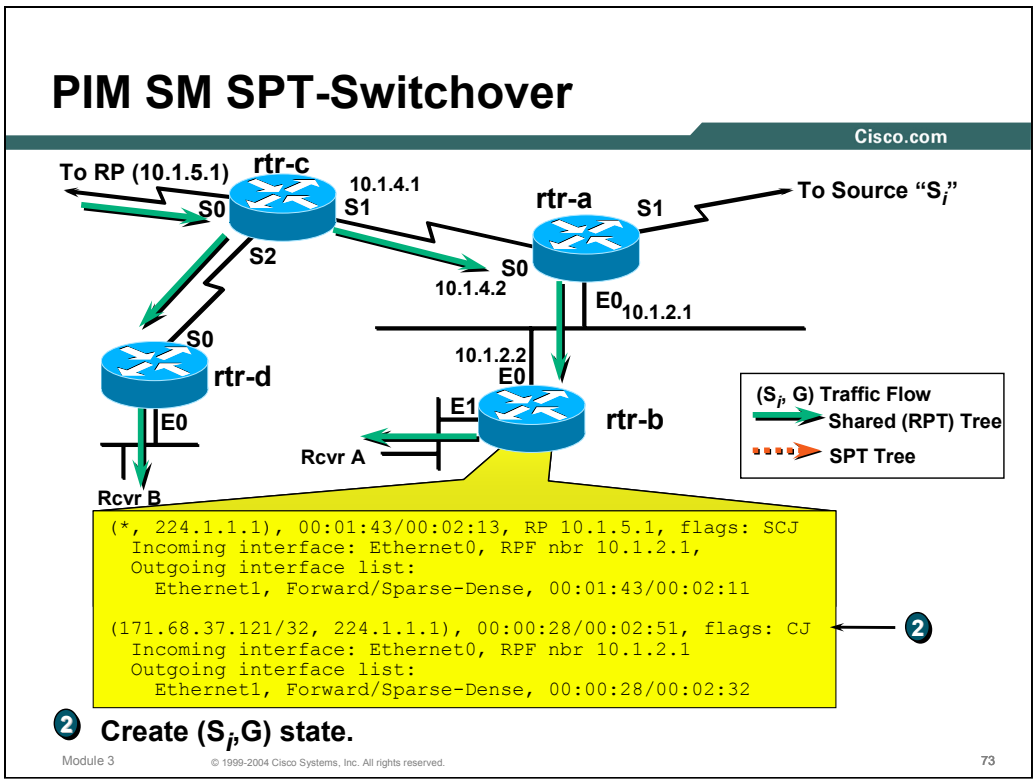

- **PIM-SM SPT-Switchover Example**
  - Step 2: Next, (Si, G) state is created for source "Si" sending to group "G" shown above is created as follows:
    - The RPF information is calculated in the direction of source "Si". This results in an Incoming interface of Ethernet0, and an RPF neighbor address of "10.1.2.1.". (Note: That the RPF information for the (S, G) entry is the same as the (*, G) entry. This indicates that the Shared Tree and the SPT are following the same path at this point.)
    - The OIL for the (S, G) entry is constructed by copying the OIL from the (*, G) entry and then removing the IIF from this list to prevent a possible route loop. This results in an (S, G) OIL containing only Ethernet1.

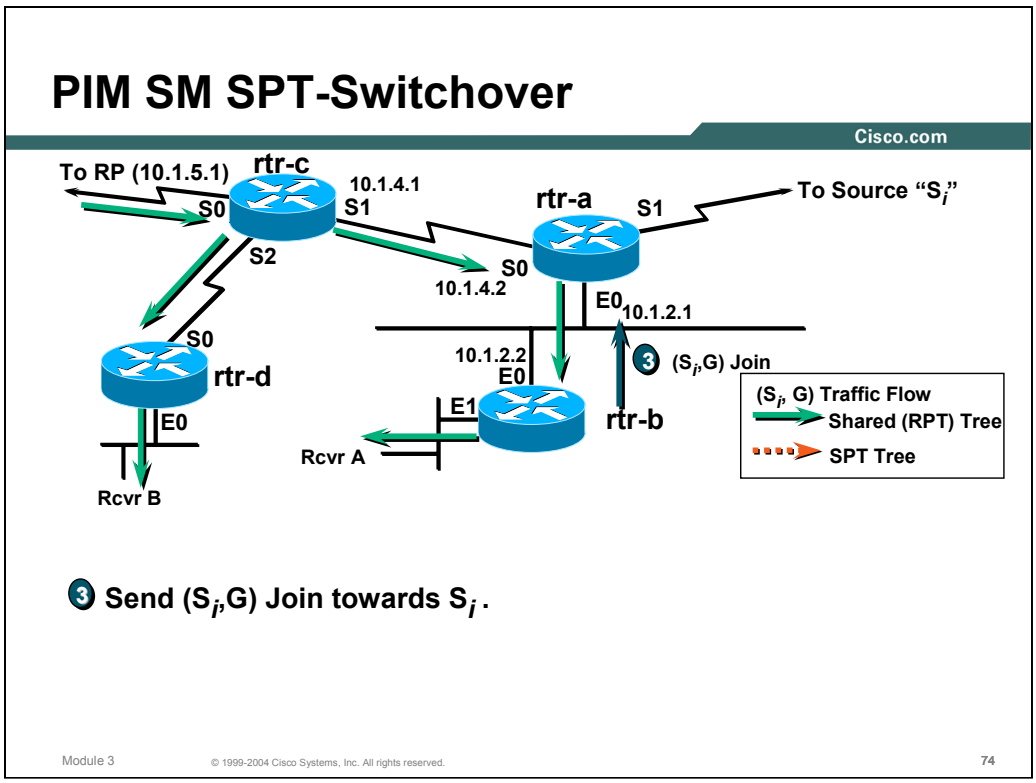

- **PIM-SM SPT-Switchover Example**
  - Step 3: Once state has been created for (Si, G), an (S, G) Join is sent toward source "Si" to build a branch of the SPT to "rtr-b". These (Si, G) Joins will continue to be sent periodically (once a minute) as long as the (Si, G) entry is not Pruned (i.e. does not have a Null OIL).

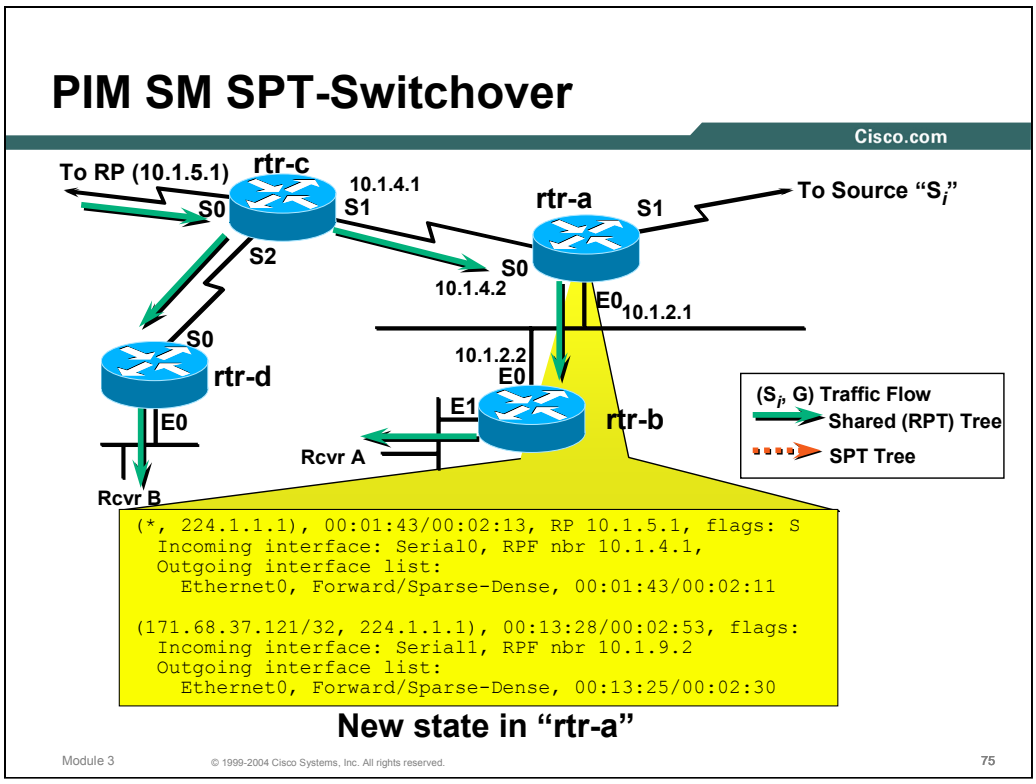

- **PIM-SM SPT-Switchover Example**
  - When the (Si, G) Join is received by "rtr-a", the (171.68.37.121/32, 224.1.1.1) entry shown above is created as follows:
    - The RPF information is calculated in the direction of source "Si". This results in an Incoming interface of Serial1, and an RPF neighbor address of "10.1.9.1.". (Note: That the RPF information for the (S, G) entry is not the same as the (*, G) entry. This indicates that the paths of the Shared Tree and the SPT diverge at this point.)
    - The OIL for the (S, G) entry is constructed by copying the OIL from the (*, G) entry and then removing the IIF from this list to prevent a possible route loop. This results in an (S, G) OIL containing only Ethernet0.

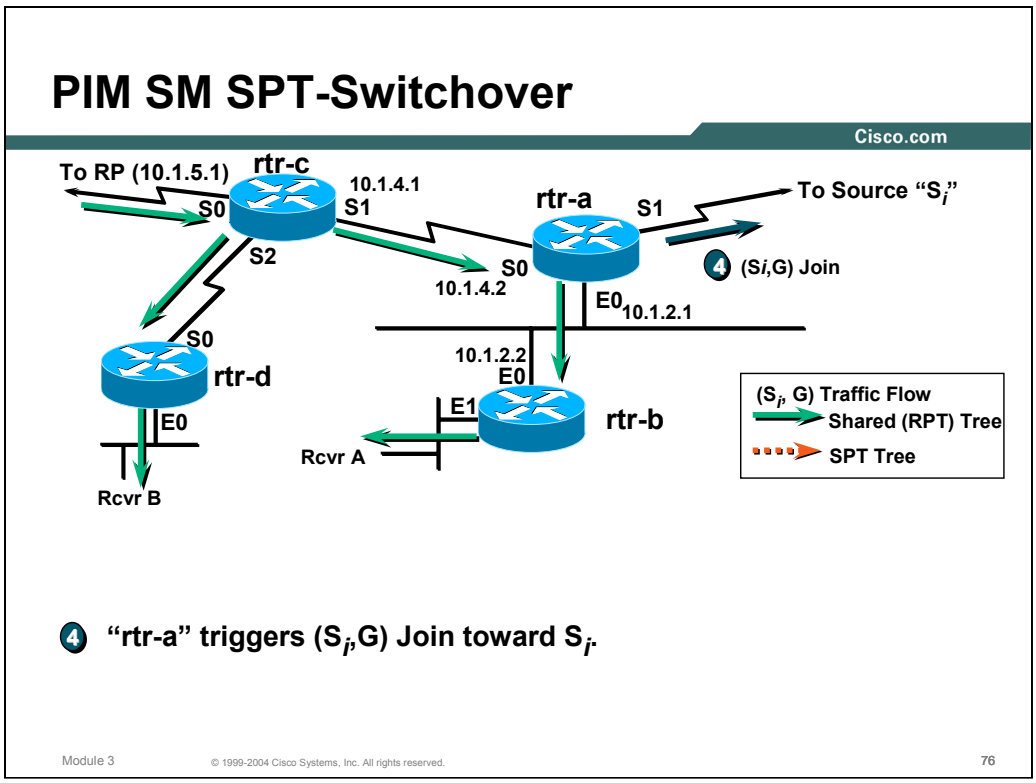

- **PIM-SM SPT-Switchover Example**
  - Step 4: When the (Si, G) state is created at "rtr-a", an (Si, G) Join is sent toward source "Si". These (Si, G) Joins will continue to be sent periodically (once a minute) as long as the (Si, G) entry is not Pruned (i.e. does not have a Null OIL).

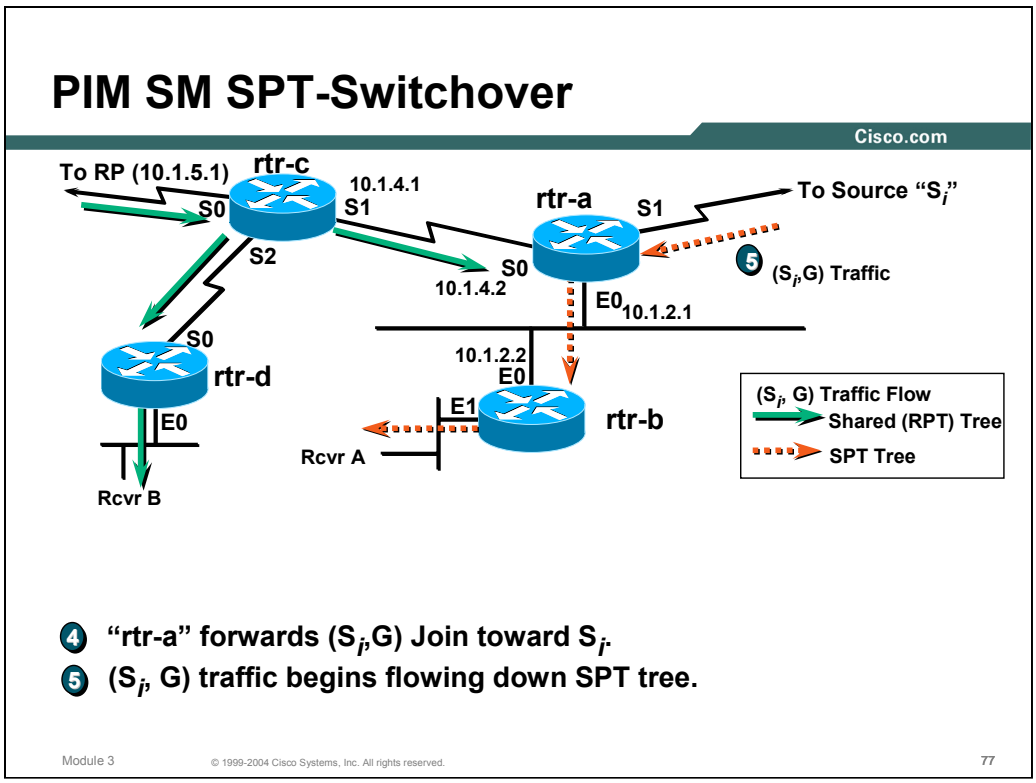

- **PIM-SM SPT-Switchover Example**
  - Step 5: When the (Si, G) Joins reach the first-hop router directly connected to source "Si", a complete branch of the SPT has been built (shown by the dashed arrows). This permits (Si, G) traffic to flow via the SPT to "rtr-b" and receiver "A".

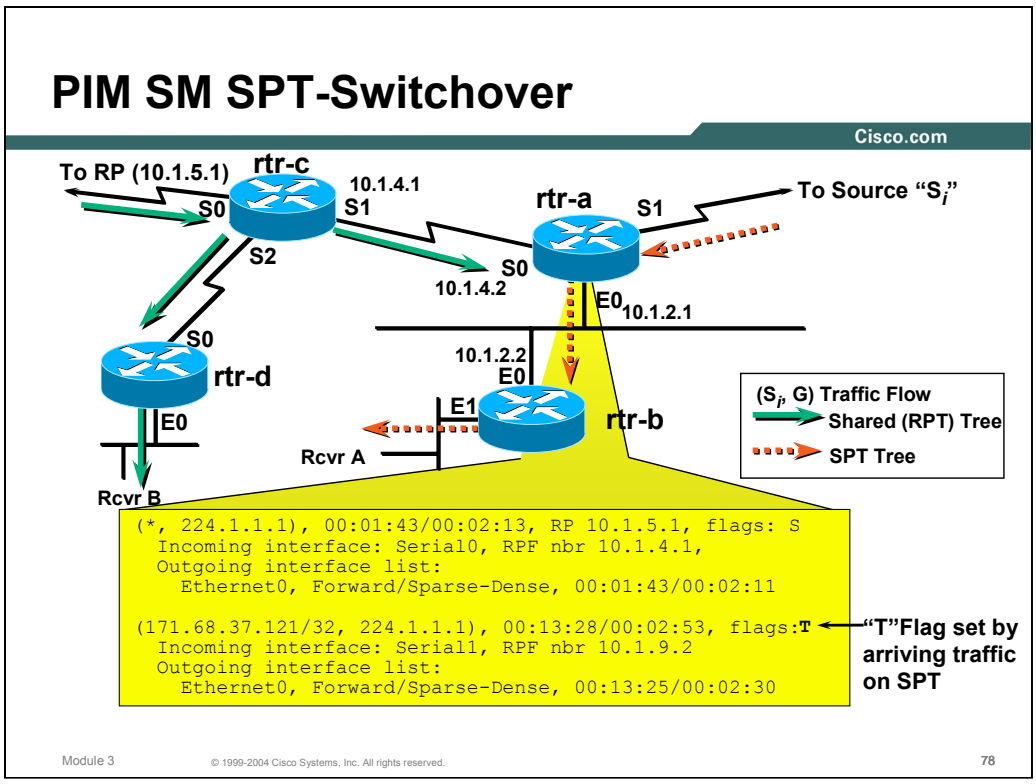

- **PIM-SM SPT-Switchover Example**
  - Traffic arriving via the SPT causes the "T" flag to be set in the (Si, G) entry in the mroute table.

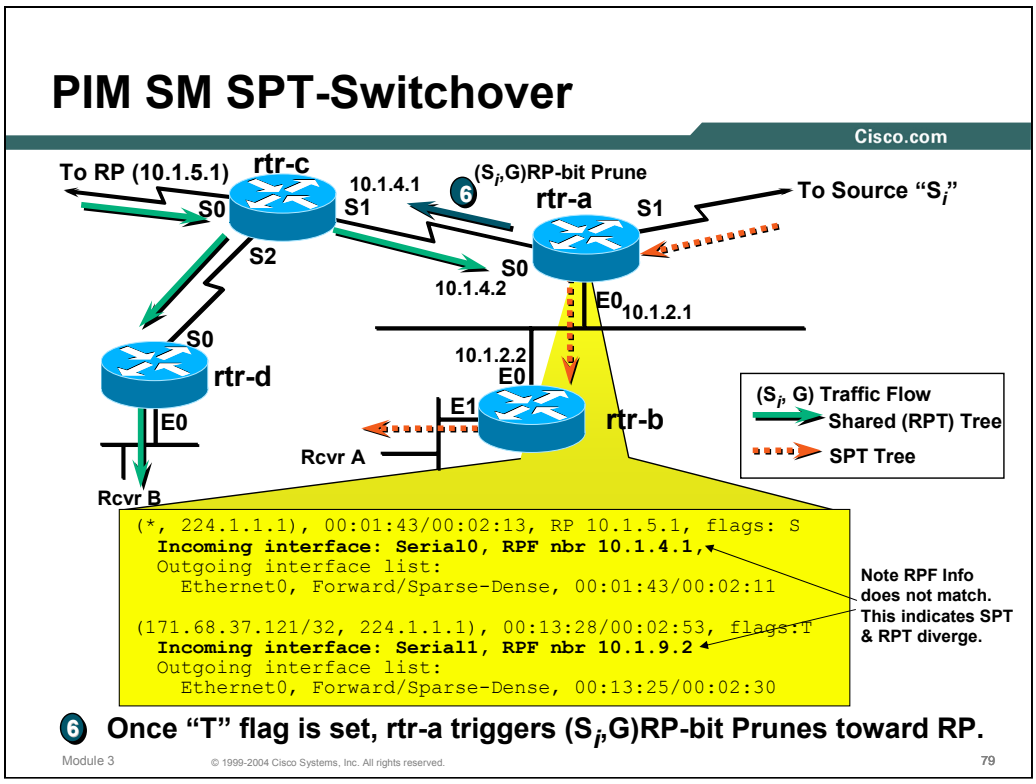

- **PIM-SM SPT-Switchover Example**
  - Step 6: Now that the "T" flag is set in the (S,G) and because the paths of the Shared Tree and the SPT diverge at "rtr-a", (note the difference in RPF information), this causes "rtr-a" to begin sending (Si, G)RP-bit Prune messages up the Shared Tree to stop the flow of redundant (Si, G) traffic down the Shared Tree.

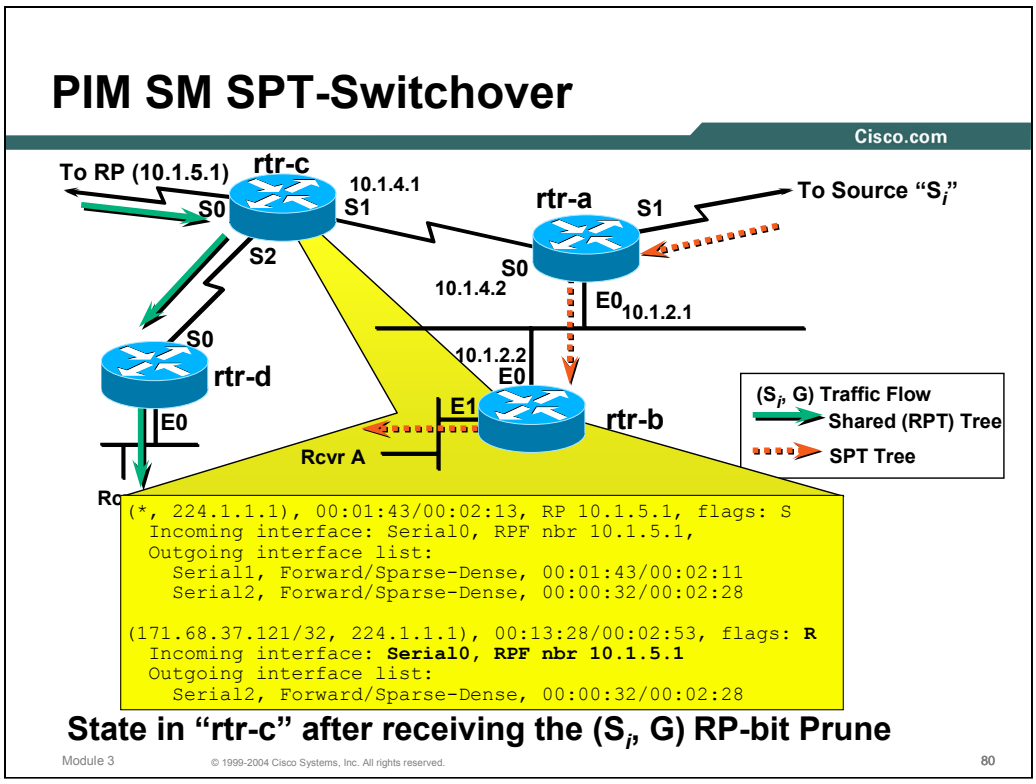

- **PIM-SM SPT-Switchover Example**
  - When the (Si, G)RP-bit Prune reaches "rtr-c", the (171.68.37.121/32, 224.1.1.1) entry shown above is created as follows:
    - Because this (S, G) entry was created as a result of the receipt of an (S,G)RP-bit Prune, the "R" bit is set to denote that this forwarding state is applicable to traffic flowing down the Shared Tree and not the Source Tree.
    - Because the "R" bit is set, the RPF information is calculated in the direction of the RP instead of source "Si". (Remember, this entry is applicable to (Si,G) traffic flowing down the Shared Tree and therefore the RPF information must point up the Shared Tree.) This results in an Incoming interface of Serial0, and an RPF neighbor address of "10.1.5.1.".
    - The OIL for the (S, G) entry is constructed by copying the OIL from the (*, G) entry minus the interface that the (Si, G)RP-bit Prune was received. Next, the IIF is removed from the OIL to prevent a possible route loop. These steps results in an (S, G) OIL containing only Serial2.
  - At this point, (Si, G) traffic flowing down the Shared Tree will be forwarded using the (Si, G) entry. (Si, G) traffic arriving at "rtr-a" will RPF correctly because the RPF information in the (Si, G) entry is pointing up the Shared Tree (as a result of the "R" bit) and will then be forwarded out all interfaces in the (Si, G) OIL. In this case, only Serial2 remains in the (Si, G) OIL and therefore (Si, G) traffic will be sent to "rtr-d" but not "rtr-a". This successfully prunes the redundant (Si, G) traffic from the branch of the Shared Tree between "rtr-c" and "rtr-a".

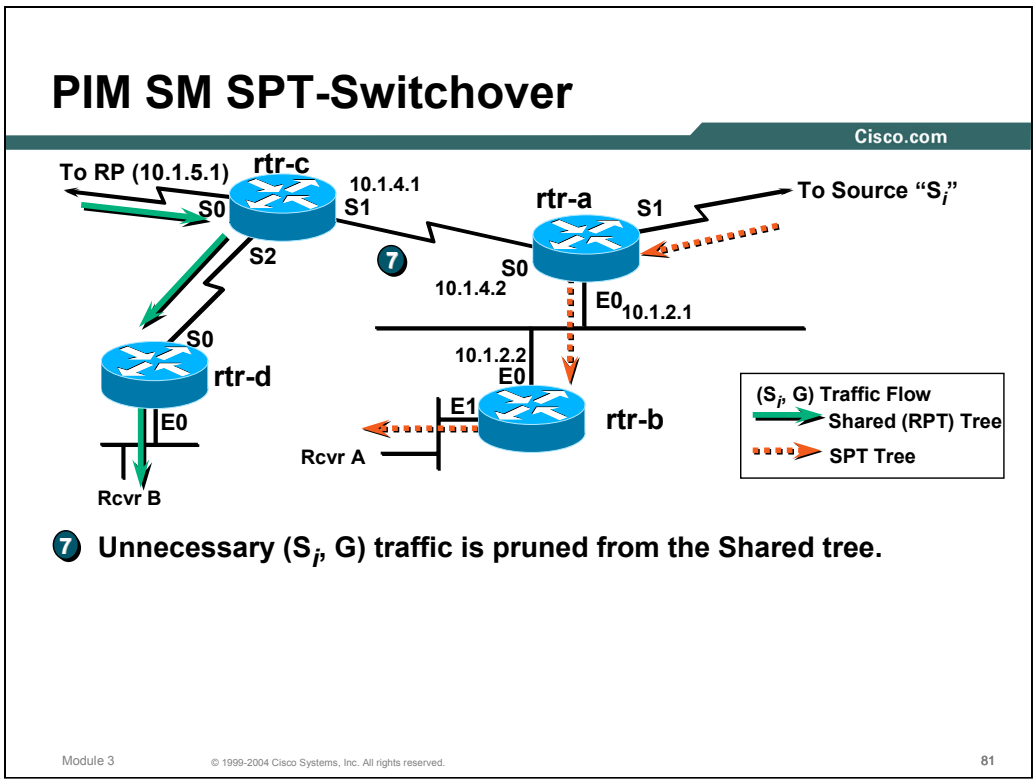

- **PIM-SM SPT-Switchover Example**
  - Step 7: At this point, the redundant (Si, G) traffic is pruned from the Shared Tree branch from "rtr-c" to "rtr-a". (Si, G) traffic is reaching receiver "A" via the SPT through "rtr-a" and "rtr-b".

PIM SM SPT-Switchover

Cisco.com

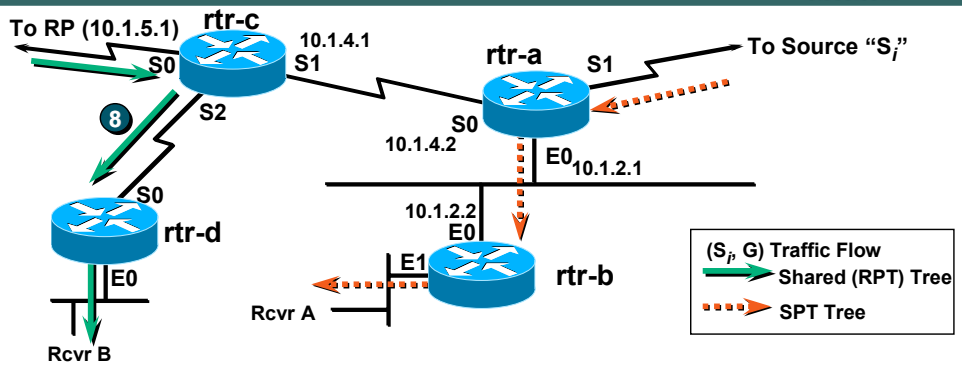

7. **Unnecessary (S_i, G) traffic is pruned from the Shared tree.**
8. **(S_i, G) traffic still flows via other branches of the Shared tree.**

- **PIM-SM SPT-Switchover Example**
  - Step 8: (Si, G) traffic is still reaching receiver "B" via a branch of the Shared Tree through "rtr-c" and "rtr-d". This is because the (Si, G) state in "rtr-c" still has Serial2 in its OIL.

PIM SM Pruning

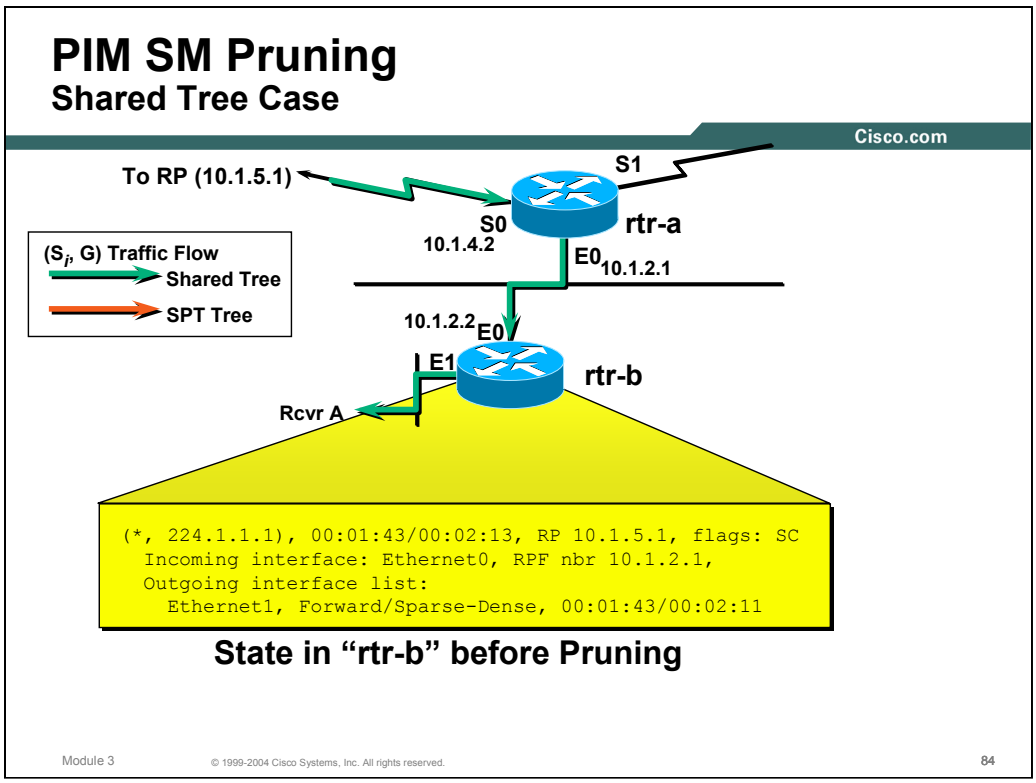

- **State in "rtr-b" before Pruning**
  - Pay particular attention to the following:
    - Traffic is flowing down the Shared Tree. (Denoted by the existance of only the (*, G) entry.)
    - The "Incoming interface" is Ethernet0.
    - The "Outgoing interface list" contains Ethernet1.
    - The "C" flag is set in the (*, G) which denotes that there is a locally connected host for this group. (Rcvr A)

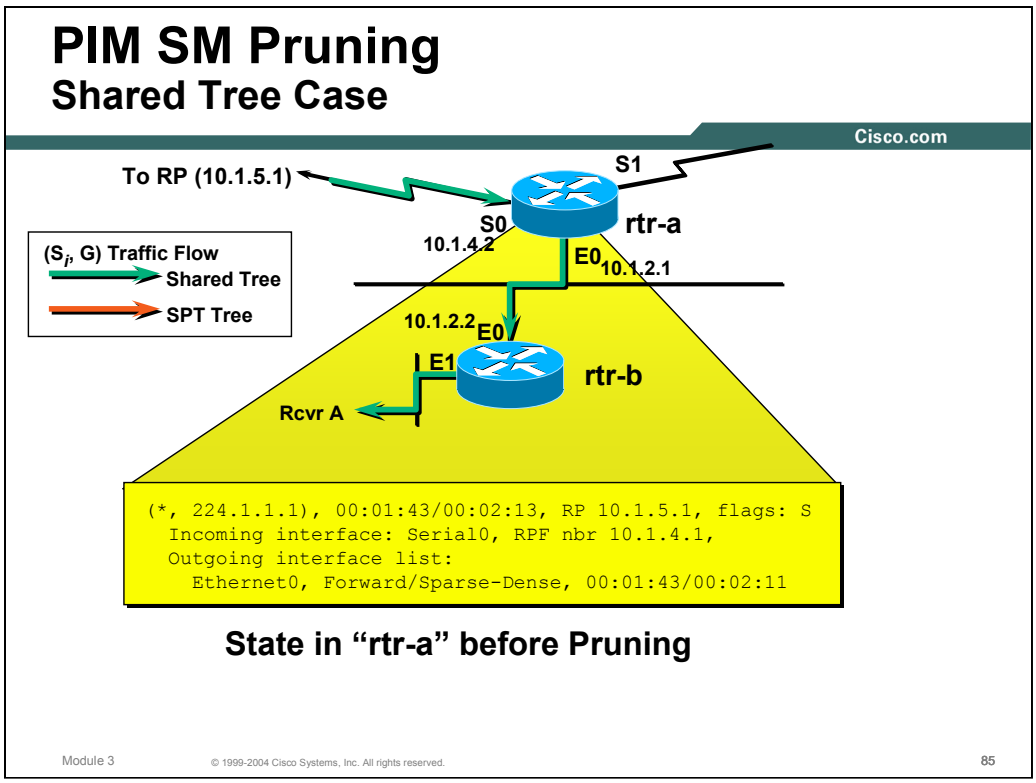

- **State in "rtr-a" before Pruning — RPT Case**
  - Pay particular attention to the following:
    - Traffic is flowing down the Shared Tree. (Denoted by the existance of only the (*, G) entry.)
    - The "Incoming interface" is Serial0.
    - The "Outgoing interface list" contains Ethernet0.

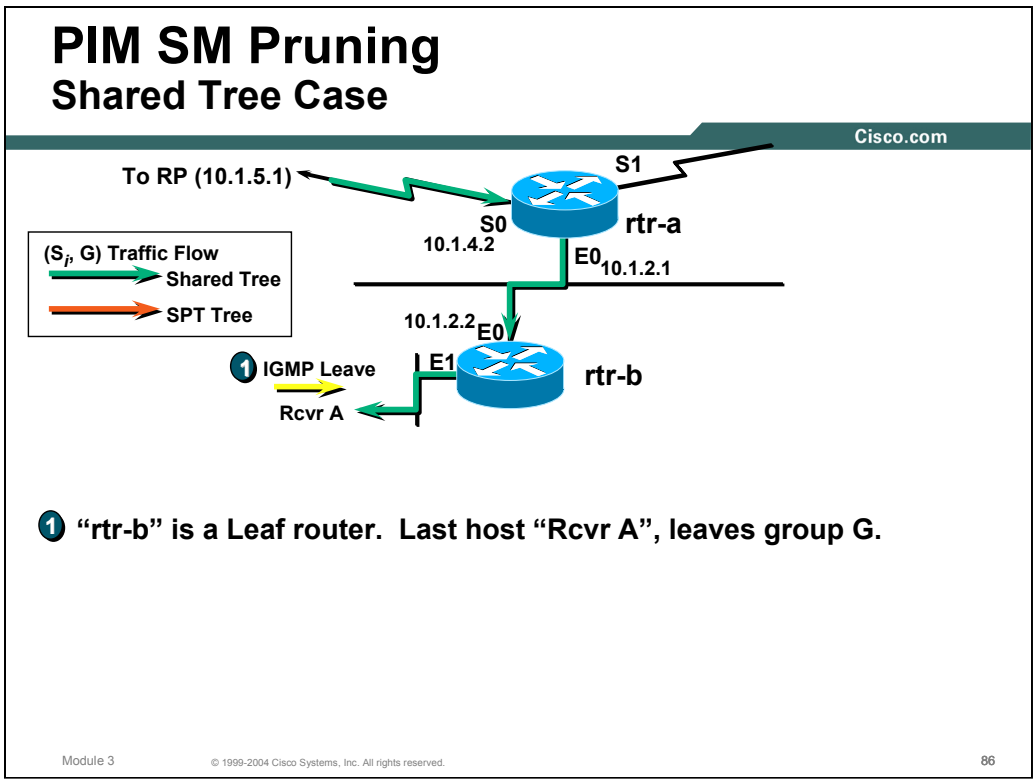

- **PIM SM Pruning Example — RPT Case**
  - 1) The last-hop or Leaf router (rtr-b) receives an IGMP Group Leave message from "Rcvr A" for group "G".

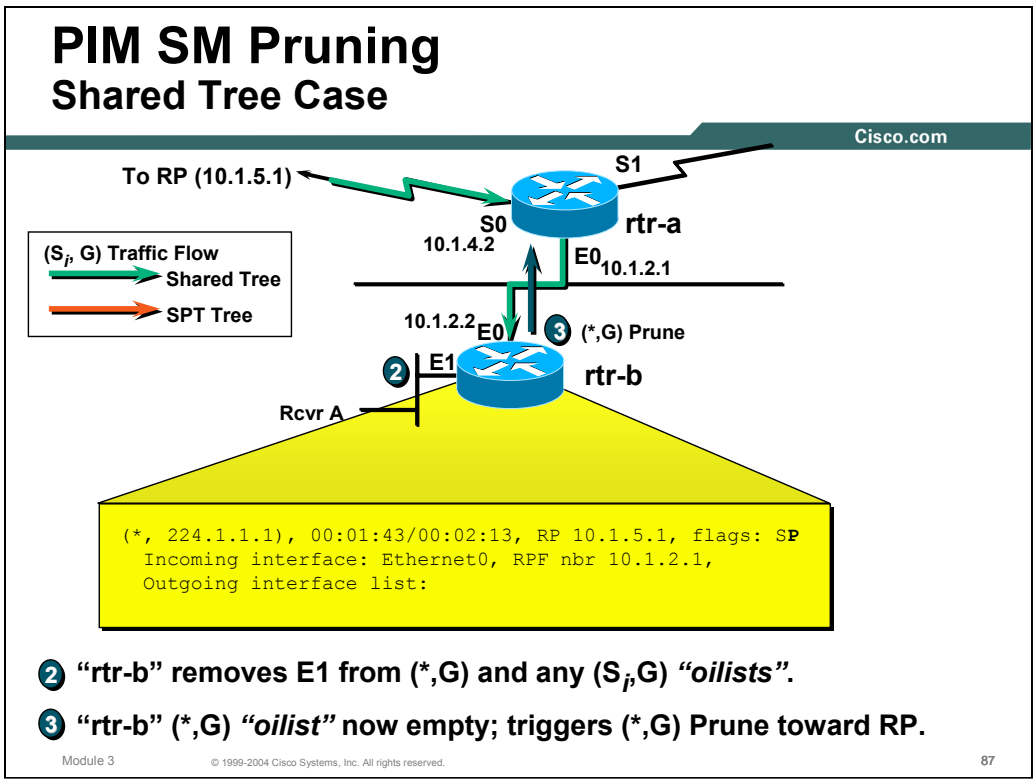

- **PIM SM Pruning Example — RPT Case**
  - 2) After performing the normal IGMP Leave processing and finding that "Rcvr A" was the last host to leave. This causes interface "E1" to be removed from the "Outgoing interface list" of the (*, G) entry and any (Si, G) entries (in this case there are none) in the Multicast Routing Table. Because "E1" was the only interface in the (*, G) entry, it's outgoing interface list becomes Null.
  - 3) Because the (*, G) "Outgoing interface list" is now Null, a (*, G) Prune is sent up the Shared Tree (RPT) via "E0" toward the RP.

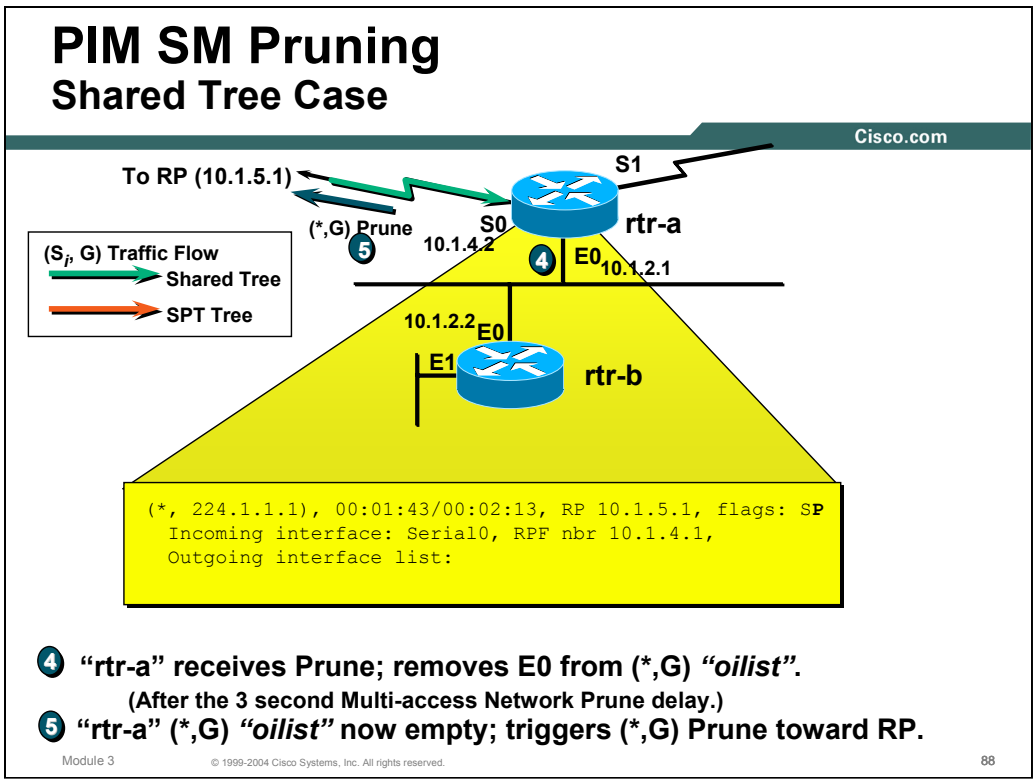

- **PIM SM Pruning Example — RPT Case (cont.)**
  - 4) The (*, G) Prune is received by "rtr-a" which causes interface "E0" to be removed from the "Outgoing interface list" of the (*, G) entry in the Multicast Routing Table.

    (Note: "rtr-a" delayed Pruning E0 from the (*, G) entry for 3 seconds since this is a Multi-Access network and it needed to wait for a possible overriding Join from another PIM neighbor. Since none was received, the interface was pruned.)
  - 5) Because the (*, G) "Outgoing interface list" is now Null, a (*, G) Prune is forwarded on up the Shared Tree (RPT) via "S0" toward the RP.

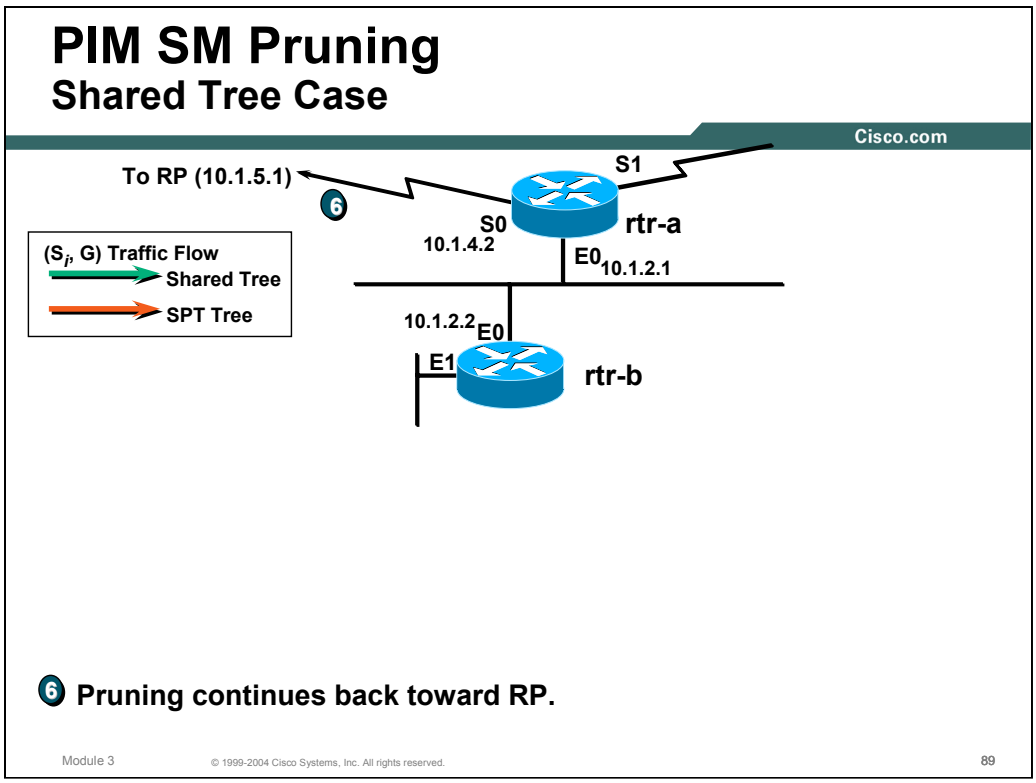

- **PIM SM Pruning Example — RPT Case (cont.)**
  - 6) This pruning continues back toward the RP or until a router is reached whose (*, G) "Outgoing interface list" doesn't go to Null as a result of the Prune.

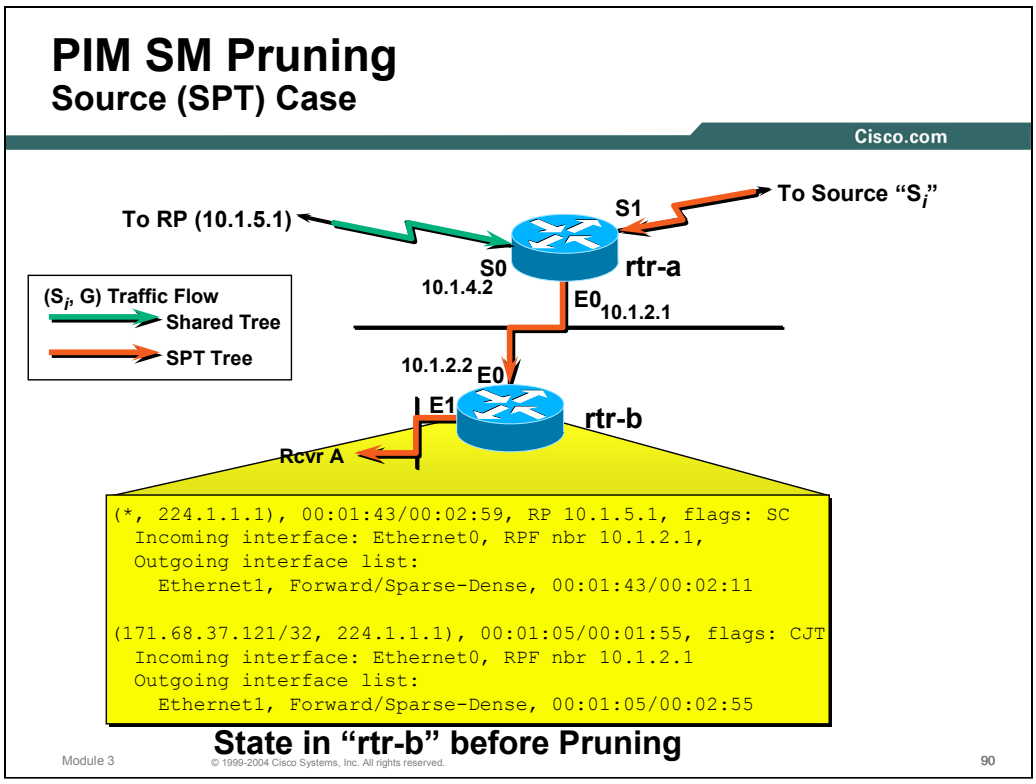

- **State in "rtr-b" before Pruning — SPT Case**
  - Pay particular attention to the following:
    - Both a (*, G) and (S, G) entries exist.
    - The "J" flag is set in the (S, G) entry. This indicates that the (S, G) state was created as a result of the SPT-Threshold being exceeded.
    - The "T" flag is set in the (S, G) entry. This indicates that (S, G) traffic is being successfully received via the Shortest-Path Tree (SPT).
    - The "Incoming interface" is the same for the (*, G) and the (S, G) entry. This indicates that Shared Tree and the Shortest-Path tree are the same at this point.

PIM SM Pruning
Source (SPT) Case

Cisco.com

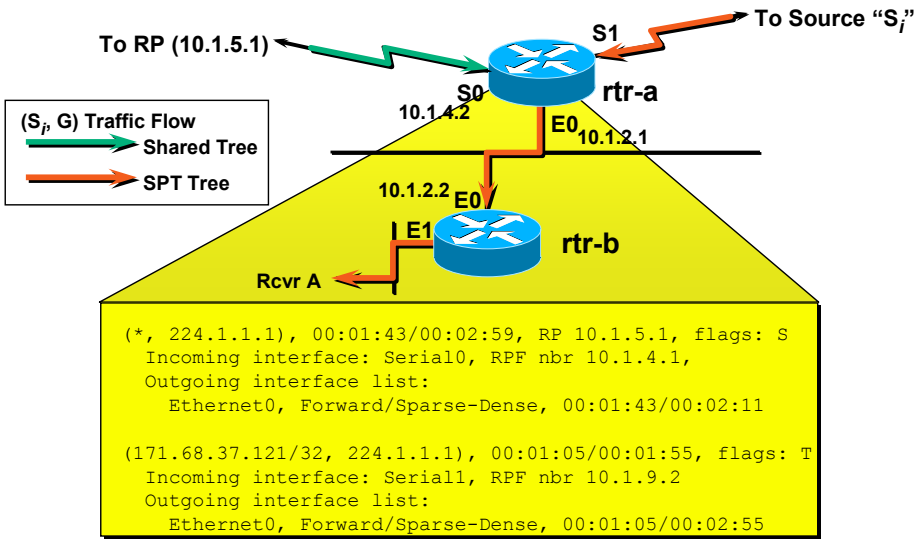

State in "rtr-a" before Pruning

- **State in "rtr-b" before Pruning — SPT Case**
  - Pay particular attention to the following:
    - Both a (*, G) and (S, G) entries exist.
    - The "T" flag is set in the (S, G) entry. This indicates that (S, G) traffic is being successfully received via the Shortest-Path Tree (SPT).
    - The "Incoming interface" is different for the (*, G) and the (S, G) entry. This indicates that Shared Tree and the Shortest-Path tree diverge at this point.

PIM SM Pruning
Source (SPT) Case

Cisco.com

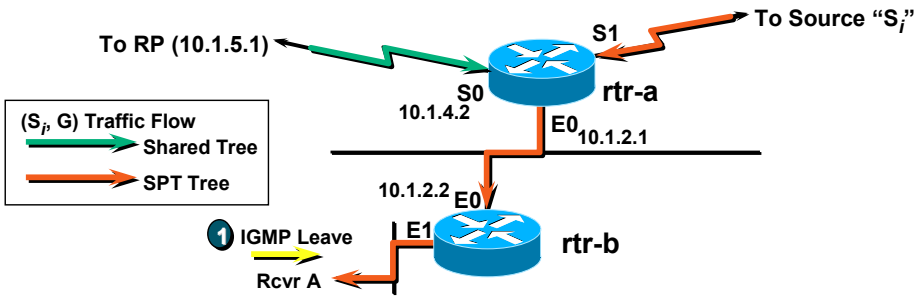

1 "rtr-b" is a Leaf router. Last host "Rcvr A", leaves group G.

- **PIM SM Pruning Example — SPT Case**
  - 1) The last-hop or Leaf router (rtr-b) receives an IGMP Group Leave message from "Rcvr A" for group "G". After performing the normal IGMP Leave processing and finding that "Rcvr A" was the last host to leave, the IGMP state for group "G" on interface "E1" is deleted.
  - 2) This causes interface "E1" to be removed from the "Outgoing interface list" of the (*, G) entry and any (Si, G) entries in the Multicast Routing Table. Because "E1" was the only interface in the (*, G) and the (Si, G) entries, their outgoing interface lists become Null.
  - 3) Because the (*, G) "Outgoing interface list" is now Null, a (*, G) Prune is sent up the Shared Tree (RPT) via "E0" toward the RP.

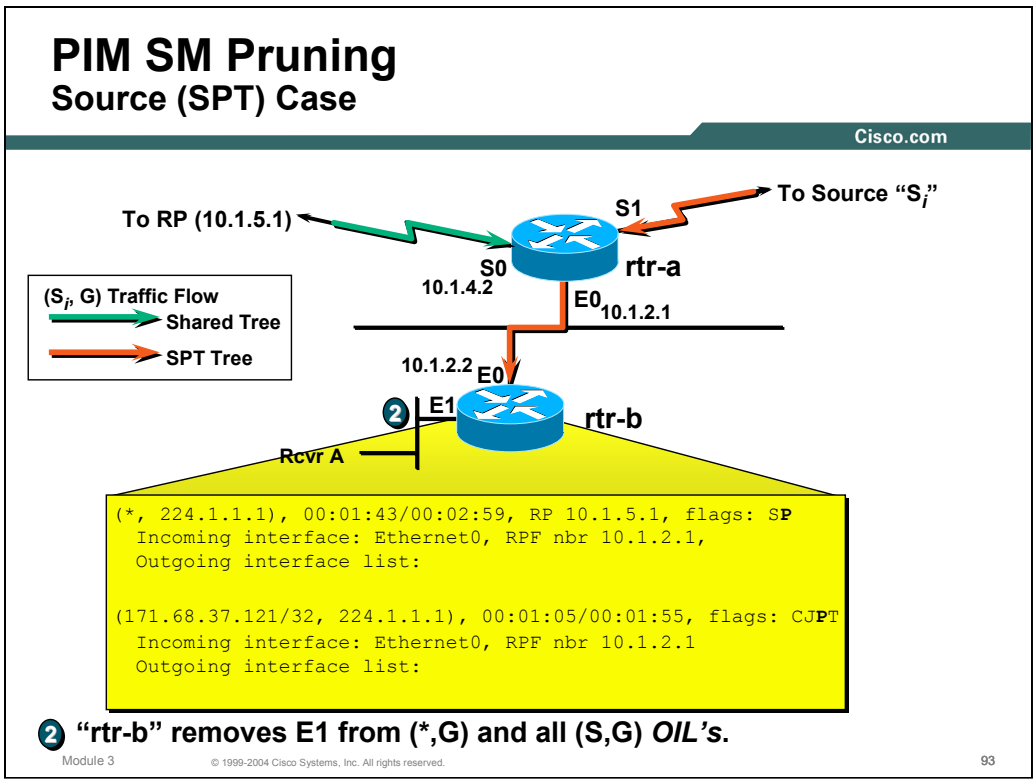

- **PIM SM Pruning Example — SPT Case**
  - 1) The last-hop or Leaf router (rtr-b) receives an IGMP Group Leave message from "Rcvr A" for group "G". After performing the normal IGMP Leave processing and finding that "Rcvr A" was the last host to leave, the IGMP state for group "G" on interface "E1" is deleted.
  - 2) This causes interface "E1" to be removed from the "Outgoing interface list" of the (*, G) entry and any (Si, G) entries in the Multicast Routing Table. Because "E1" was the only interface in the (*, G) and the (Si, G) entries, their outgoing interface lists become Null.
  - 3) Because the (*, G) "Outgoing interface list" is now Null, a (*, G) Prune is sent up the Shared Tree (RPT) via "E0" toward the RP.

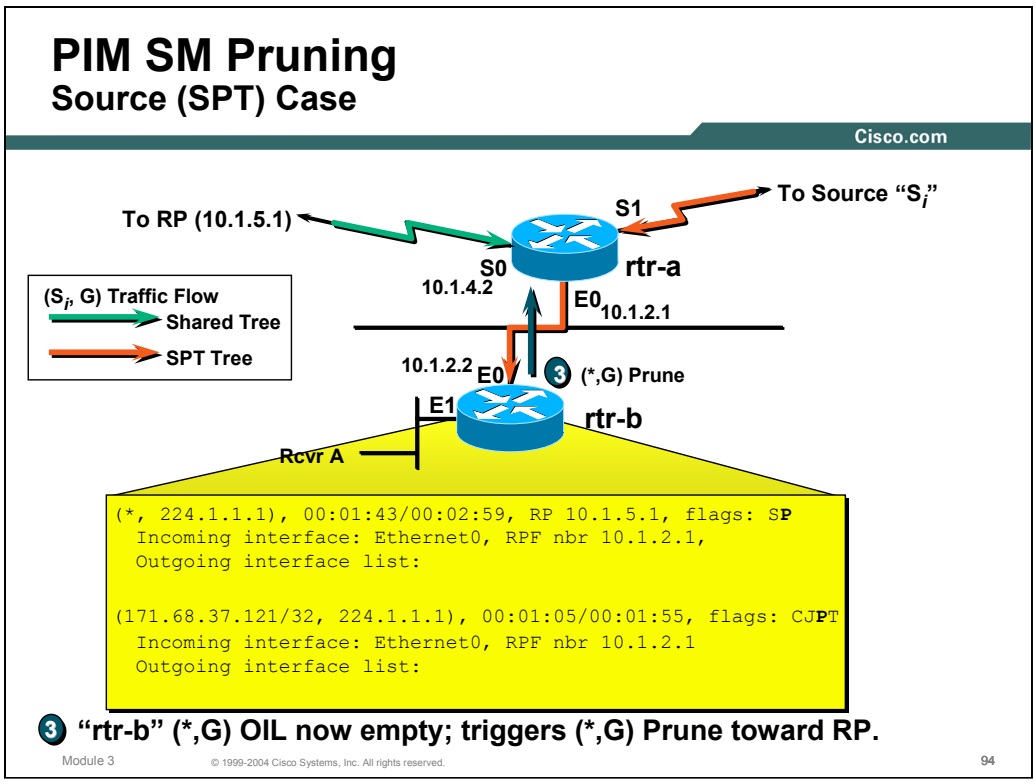

- **PIM SM Pruning Example — SPT Case**
  - 1) The last-hop or Leaf router (rtr-b) receives an IGMP Group Leave message from "Rcvr A" for group "G". After performing the normal IGMP Leave processing and finding that "Rcvr A" was the last host to leave, the IGMP state for group "G" on interface "E1" is deleted.
  - 2) This causes interface "E1" to be removed from the "Outgoing interface list" of the (*, G) entry and any (Si, G) entries in the Multicast Routing Table. Because "E1" was the only interface in the (*, G) and the (Si, G) entries, their outgoing interface lists become Null.
  - 3) Because the (*, G) "Outgoing interface list" is now Null, a (*, G) Prune is sent up the Shared Tree (RPT) via "E0" toward the RP.

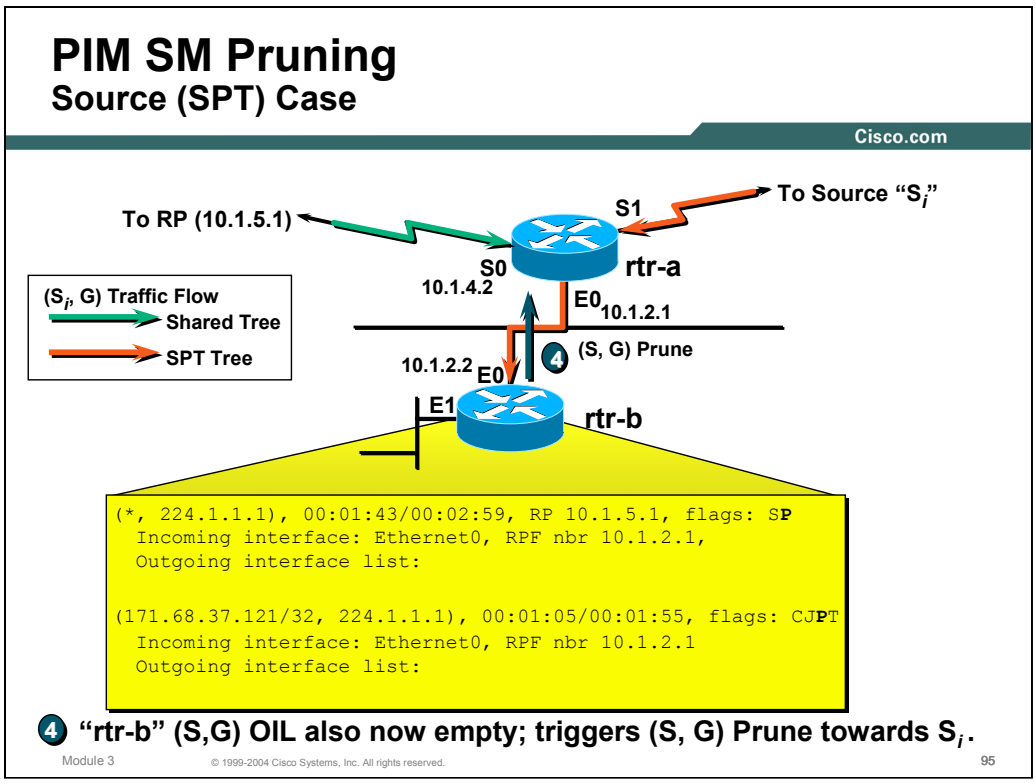

- **PIM SM Pruning Example — SPT Case (cont.)**
  - 4) Because the (Si, G) "Outgoing interface list" is now Null, "rtr-b" stops sending Periodic (Si, G) Join messages up the Shortest-Path Tree (SPT).

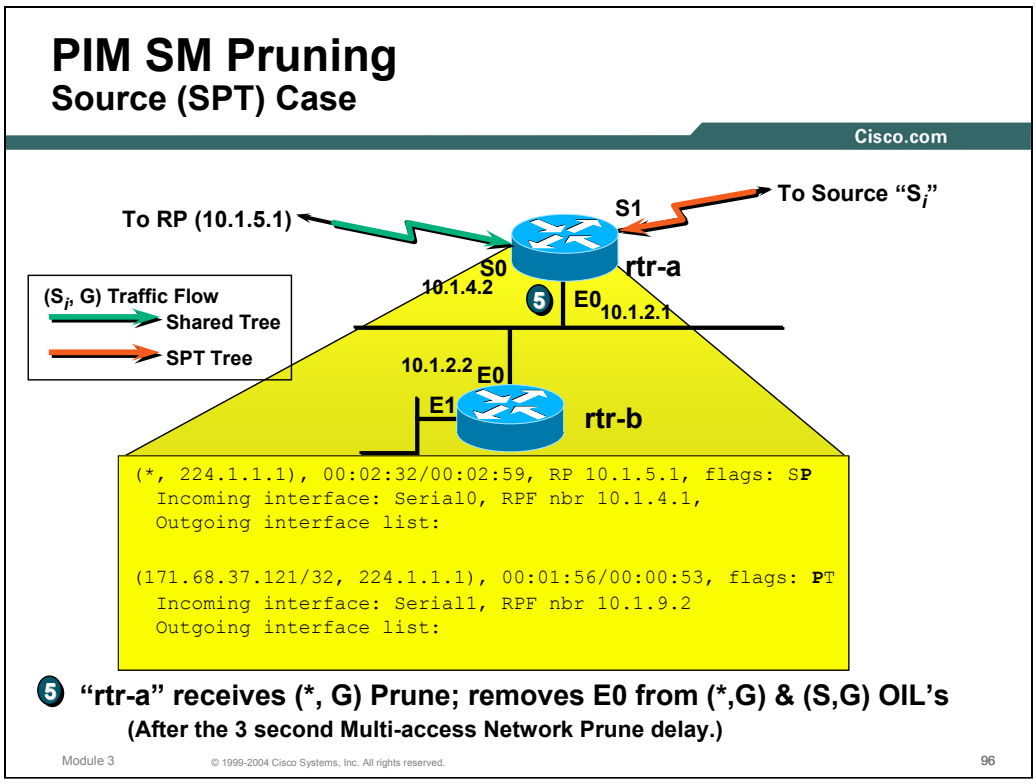

- **PIM SM Pruning Example — SPT Case (cont.)**
  - 5) The (*, G) Prune is received by "rtr-a" which causes interface "E0" to be removed from the "Outgoing interface list" of the (*, G) entry in the Multicast Routing Table.

    (Note: "rtr-a" delayed Pruning E0 from the (*, G) entry for 3 seconds since this is a Multi-Access network and it needed to wait for a possible overriding Join from another PIM neighbor. Since none was received, the interface was pruned.)

PIM SM Pruning
Source (SPT) Case

Cisco.com

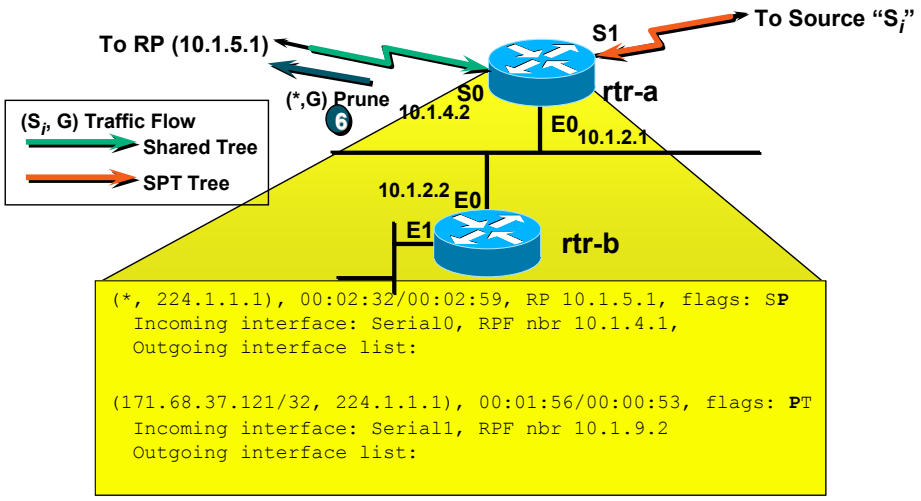

6 "rtr-a" (*,G) *OIL* now empty; triggers (*,G) Prune toward RP.

- **PIM SM Pruning Example — SPT Case (cont.)**
  - 6) Because the (*, G) "Outgoing interface list" is now Null, a (*, G) Prune is forwarded on up the Shared Tree (RPT) via "S0" toward the RP.

PIM SM Pruning
Source (SPT) Case

Cisco.com

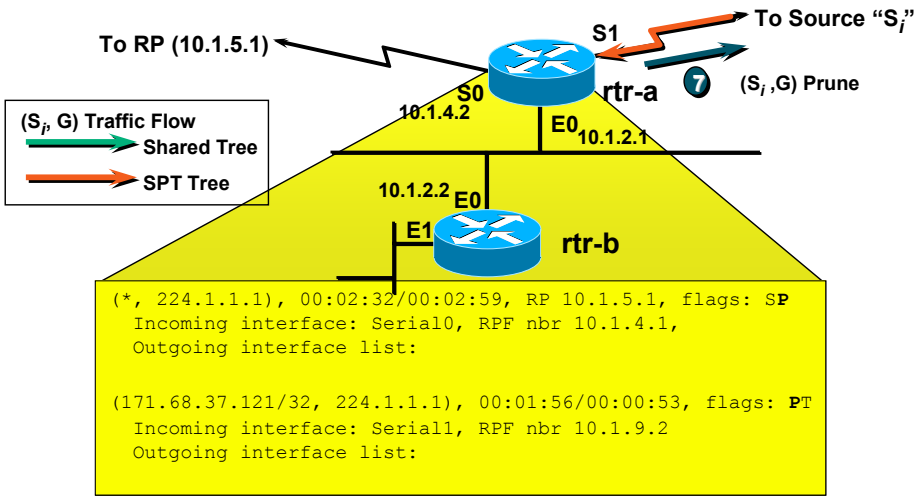

7 "rtr-a" (S,G) *OIL* also now empty; triggers (S,G) Prune towards S_i.

- **PIM SM Pruning Example — SPT Case (cont.)**
  - 7) Because the (Si, G) "Outgoing interface list" is now Null, a (Si, G) Prune is sent up the Shortest-Path Tree (SPT) towards the source "Si".

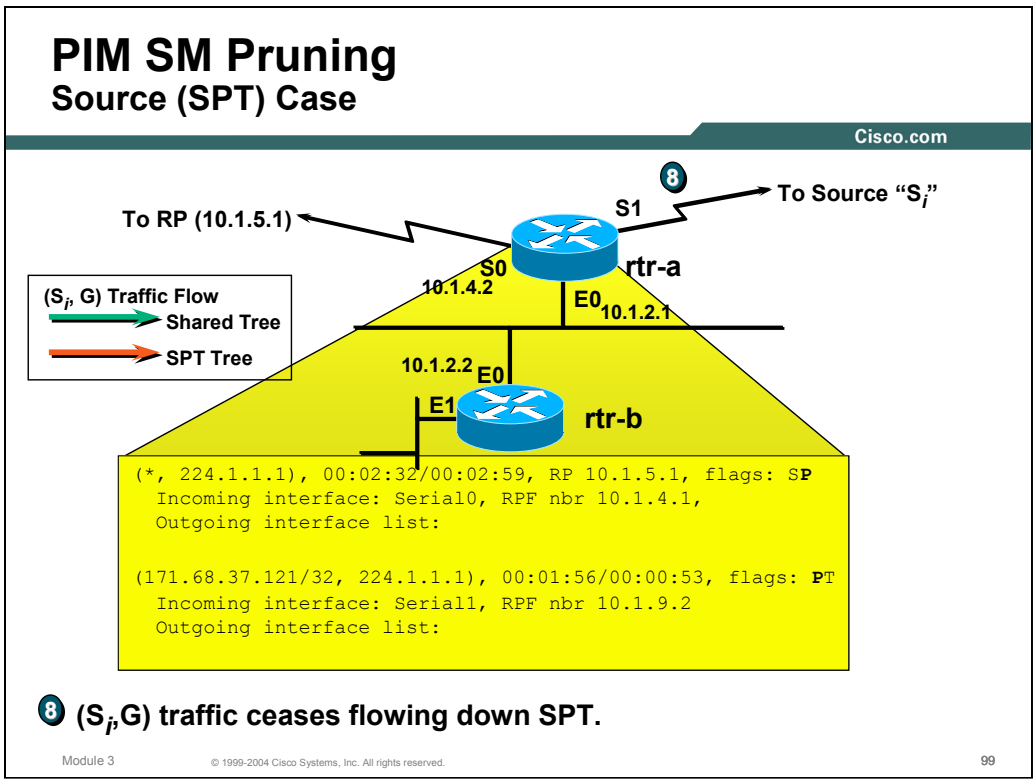

- **PIM SM Pruning Example — SPT Case (cont.)**
  - 8) Traffic stops flowing down the Shortest-Path Tree (SPT).