

# Multicast Source Discovery Protocol (MSDP)

# Agenda

- **MSDP Overview**
- **MSDP Peers**
- **MSDP SA Messages**
- **MSDP Mesh Groups**
- **MSDP State Flags**
- **MSDP Enhancements**

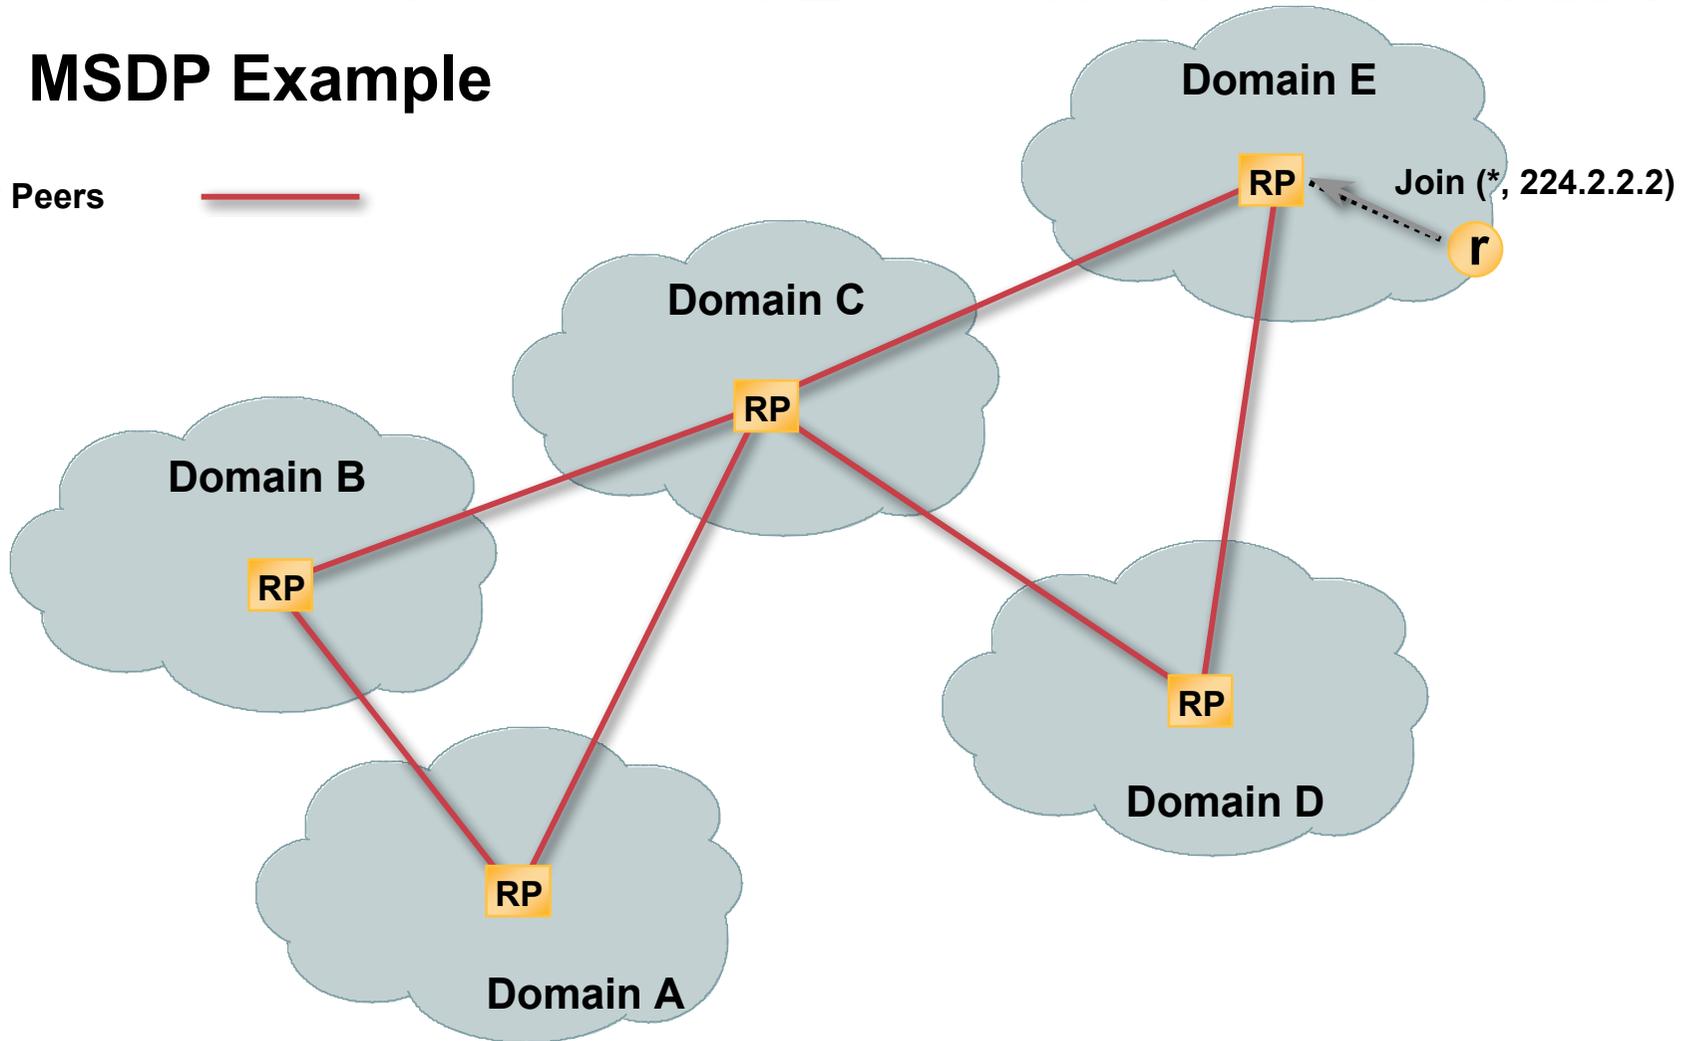
# MSDP Overview

- **Uses inter-domain source trees only.**
  - **RP's know about all sources in their domain**
    - **Sources cause a “PIM Register” to the RP**
    - **Can tell RP's in other domains of its sources**
      - **Via MSDP SA (Source Active) messages**
  - **RP's know about receivers in their domain**
    - **Receivers cause a “(\*, G) Join” to the RP**
    - **RP can join the source tree in the peer domain**
      - **Via normal PIM (S, G) joins**
      - **Only necessary if there are receivers for the group**
    - **Last-hop routers then join source tree directly.**

# MSDP Overview

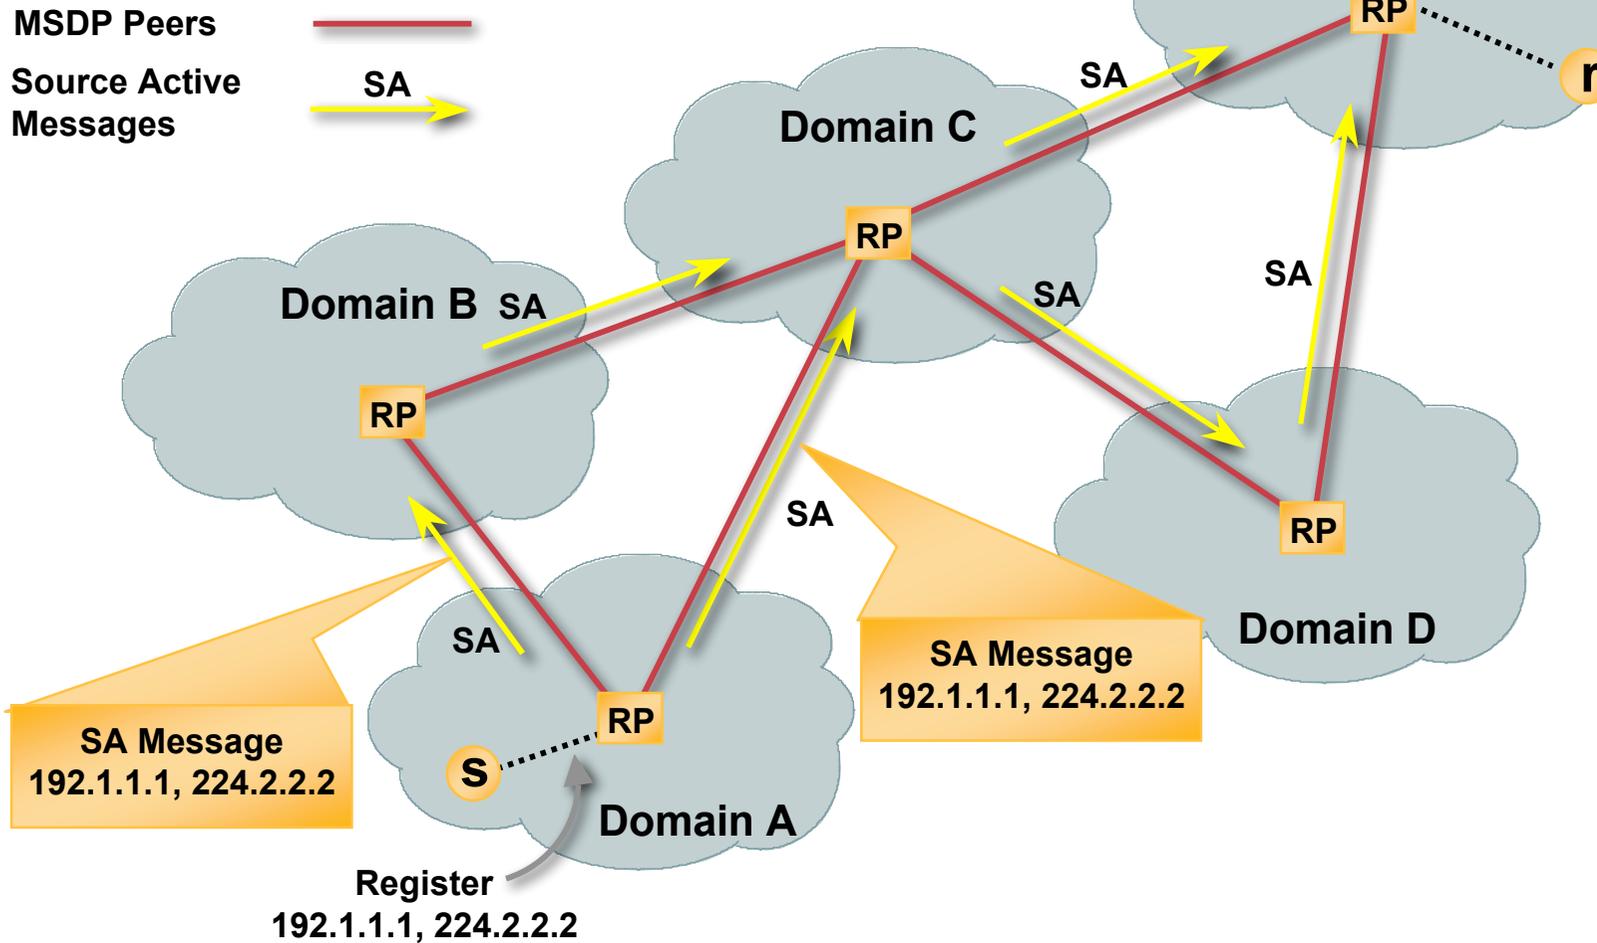
## MSDP Example

MSDP Peers



# MSDP Overview

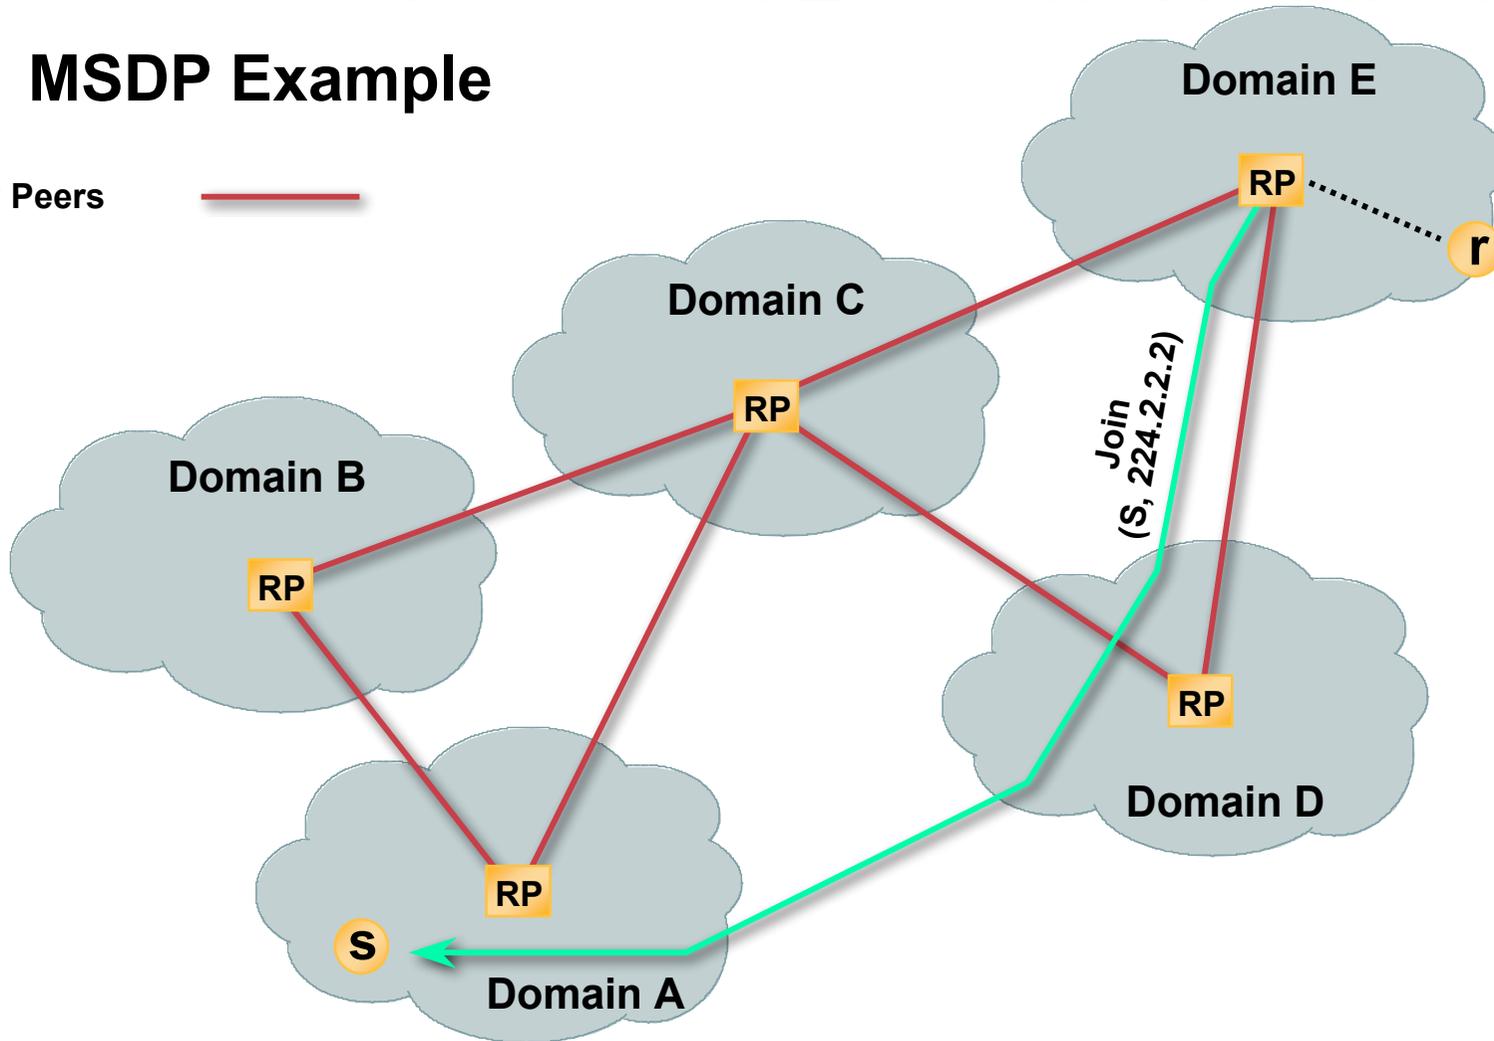
## MSDP Example



# MSDP Overview

## MSDP Example

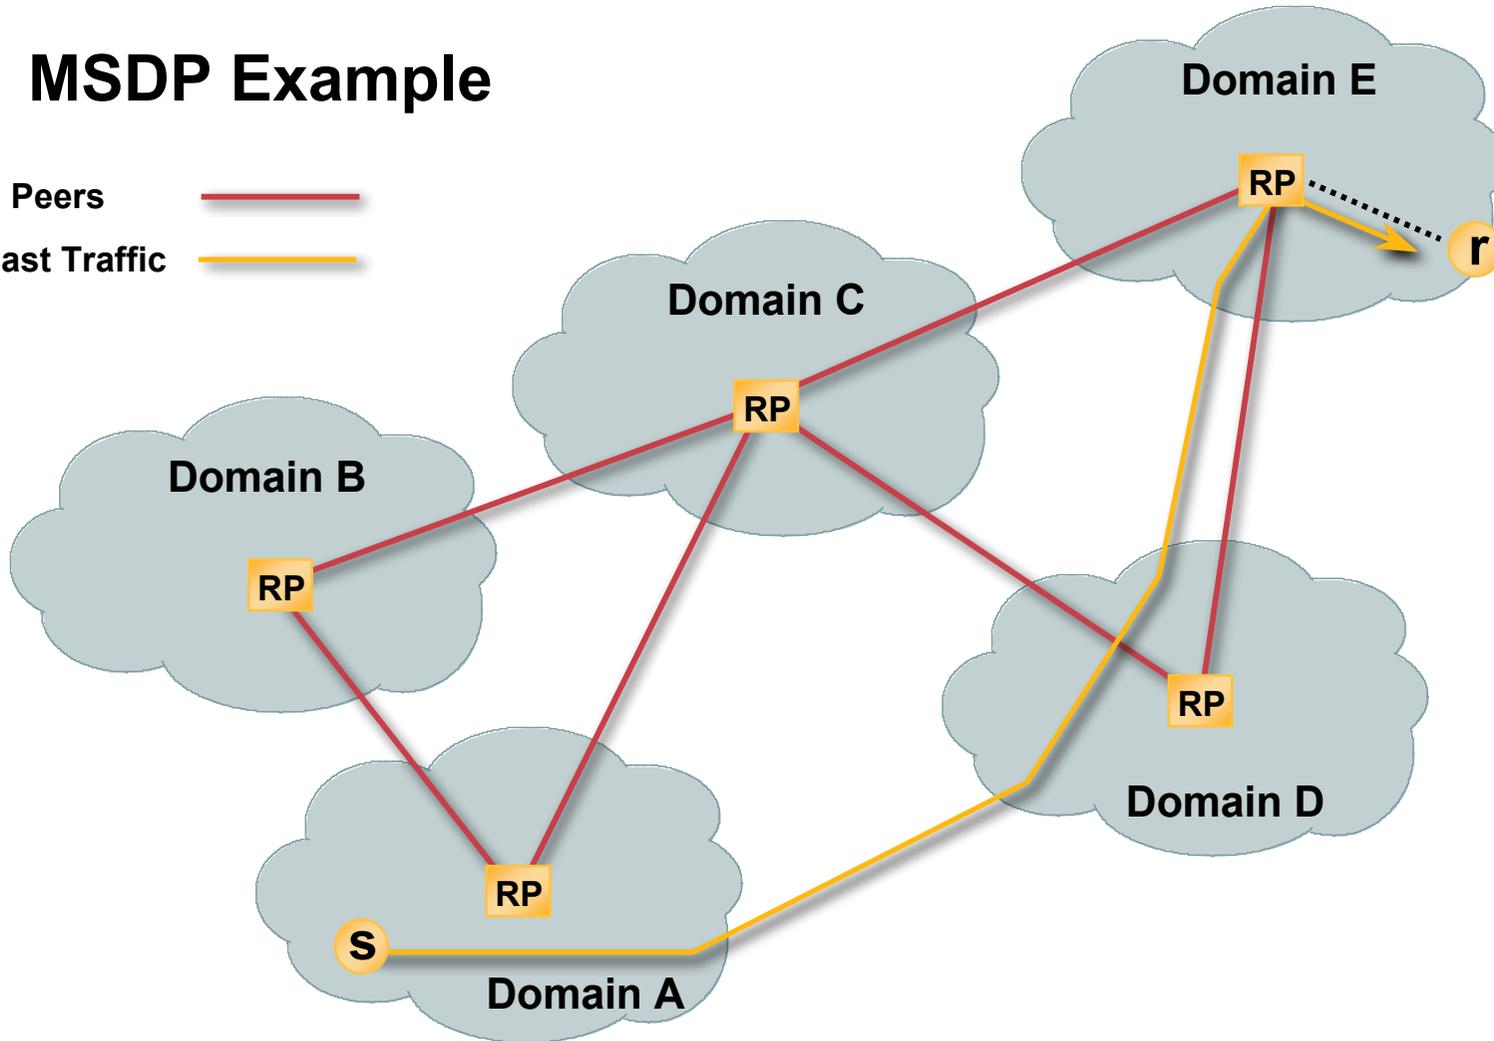
MSDP Peers



# MSDP Overview

## MSDP Example

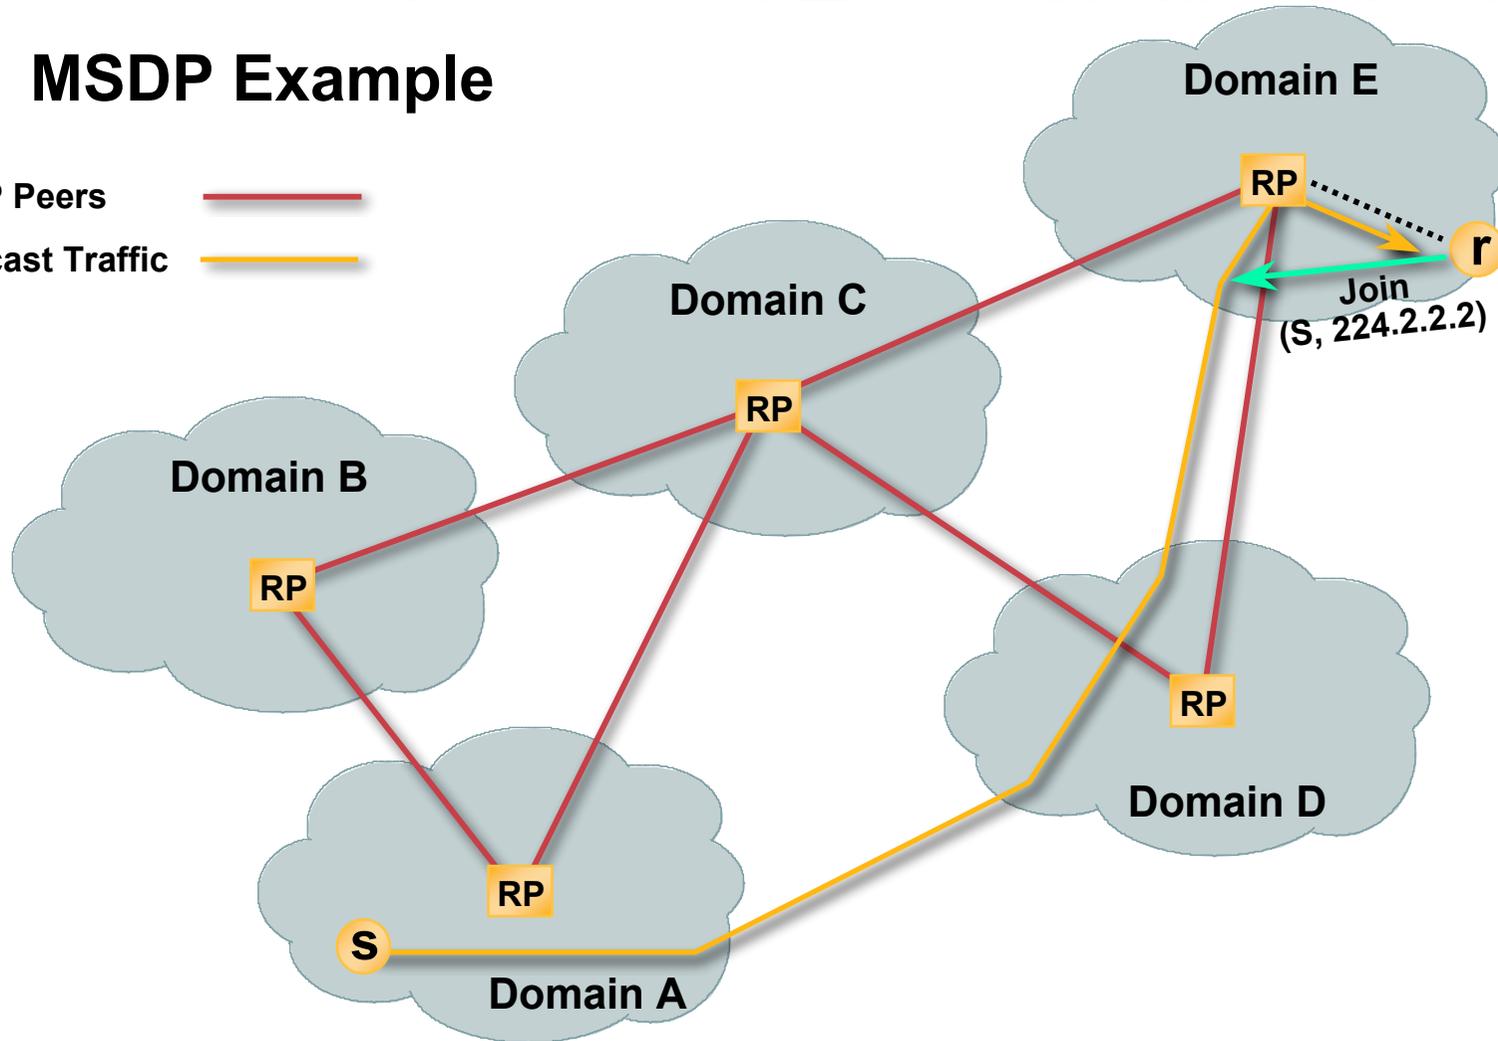
MSDP Peers   
Multicast Traffic 



# MSDP Overview

## MSDP Example

MSDP Peers   
Multicast Traffic 



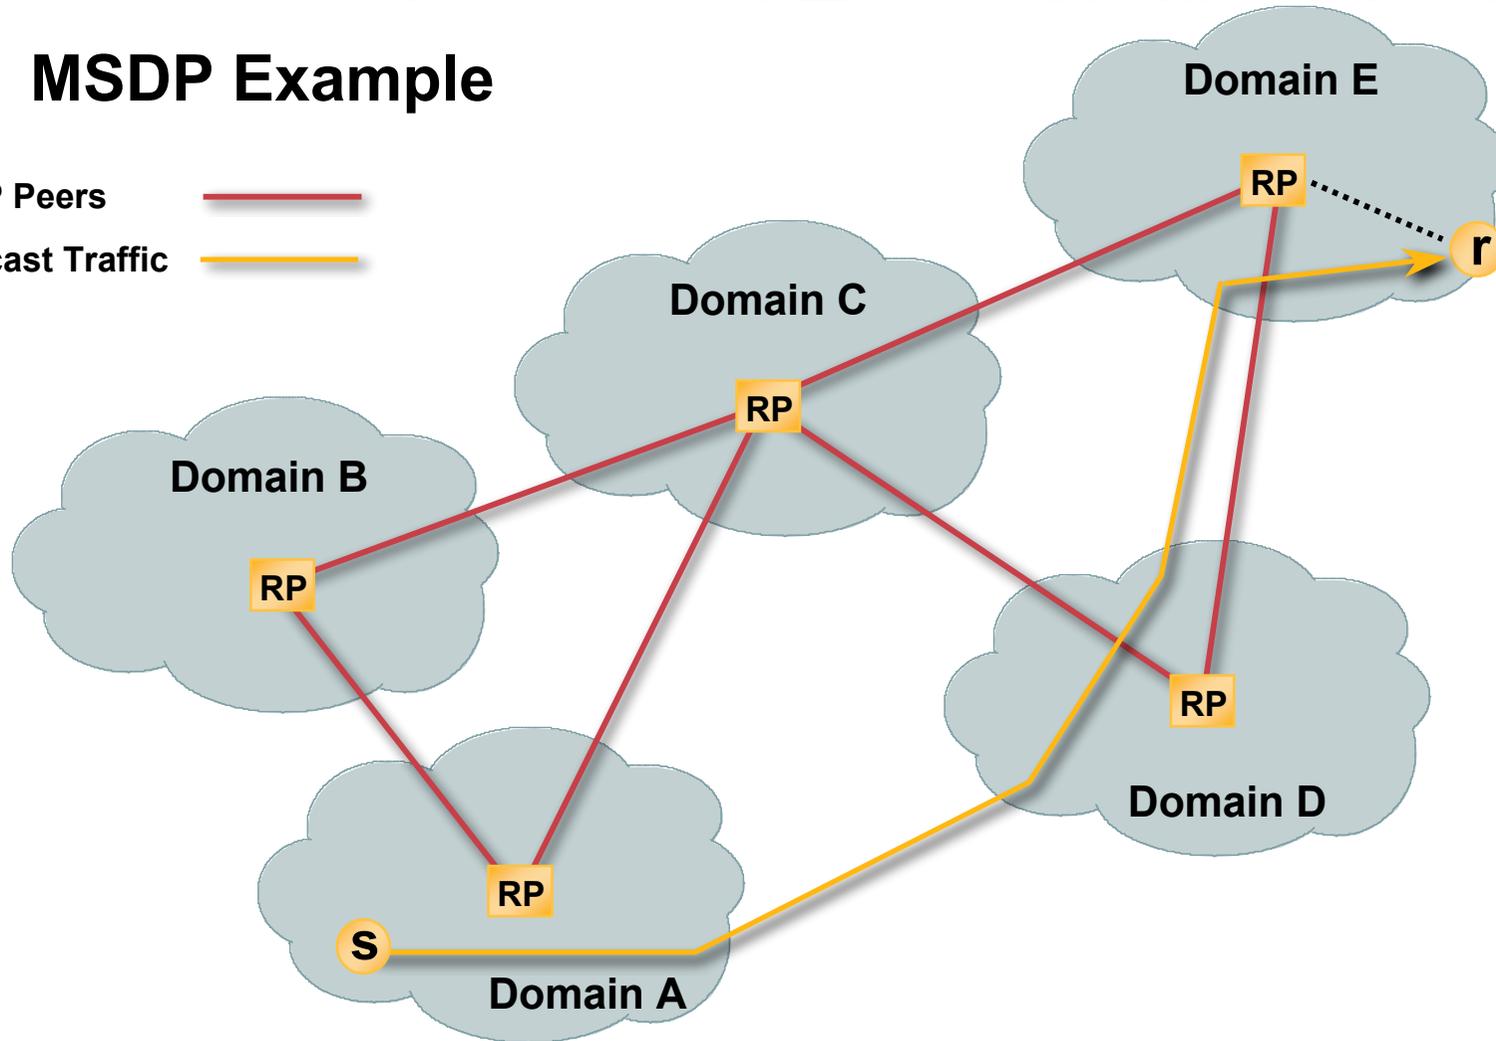
# MSDP Overview

## MSDP Example

MSDP Peers



Multicast Traffic



# Agenda

- **MSDP Overview**
- **MSDP Peers**
- **MSDP SA Messages**
- **MSDP Mesh Groups**
- **MSDP State Flags**
- **MSDP Enhancements**

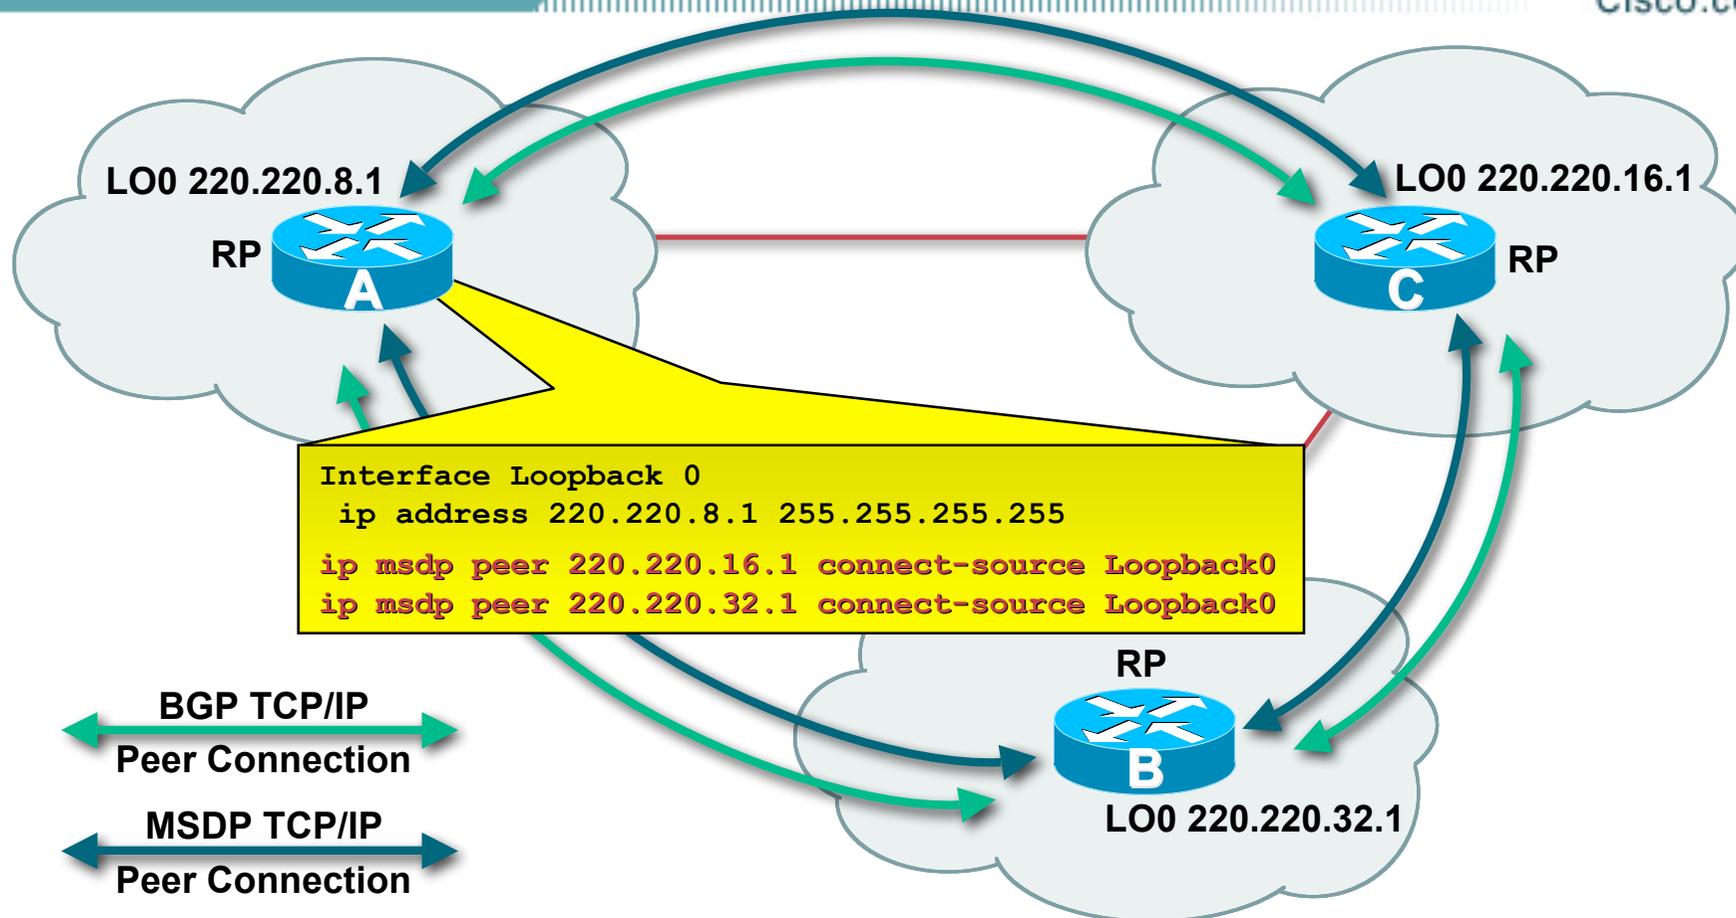
# MSDP Peers

- **MSDP Peers configured similar to BGP**
- **Peers connect using TCP port 639**
  - Lower address peer initiates connection
  - Higher address peer waits in LISTEN state
- **Peers send keepalives every 60 secs.**
- **Connection reset after 75 seconds**
  - If no MSDP packets or keepalives are received

# MSDP Peers

- **MSDP peers normally *must* run BGP!**
  - **BGP NLRI is used to RPF check SA messages.**
    - May use NLRI from M-Table, U-Table or both.
  - **RPF check prevents SA's from looping.**  
(More on that later.)
- **Exceptions:**
  - **When peering with only a single MSDP peer.**
  - **When using an MSDP Mesh-Group.**

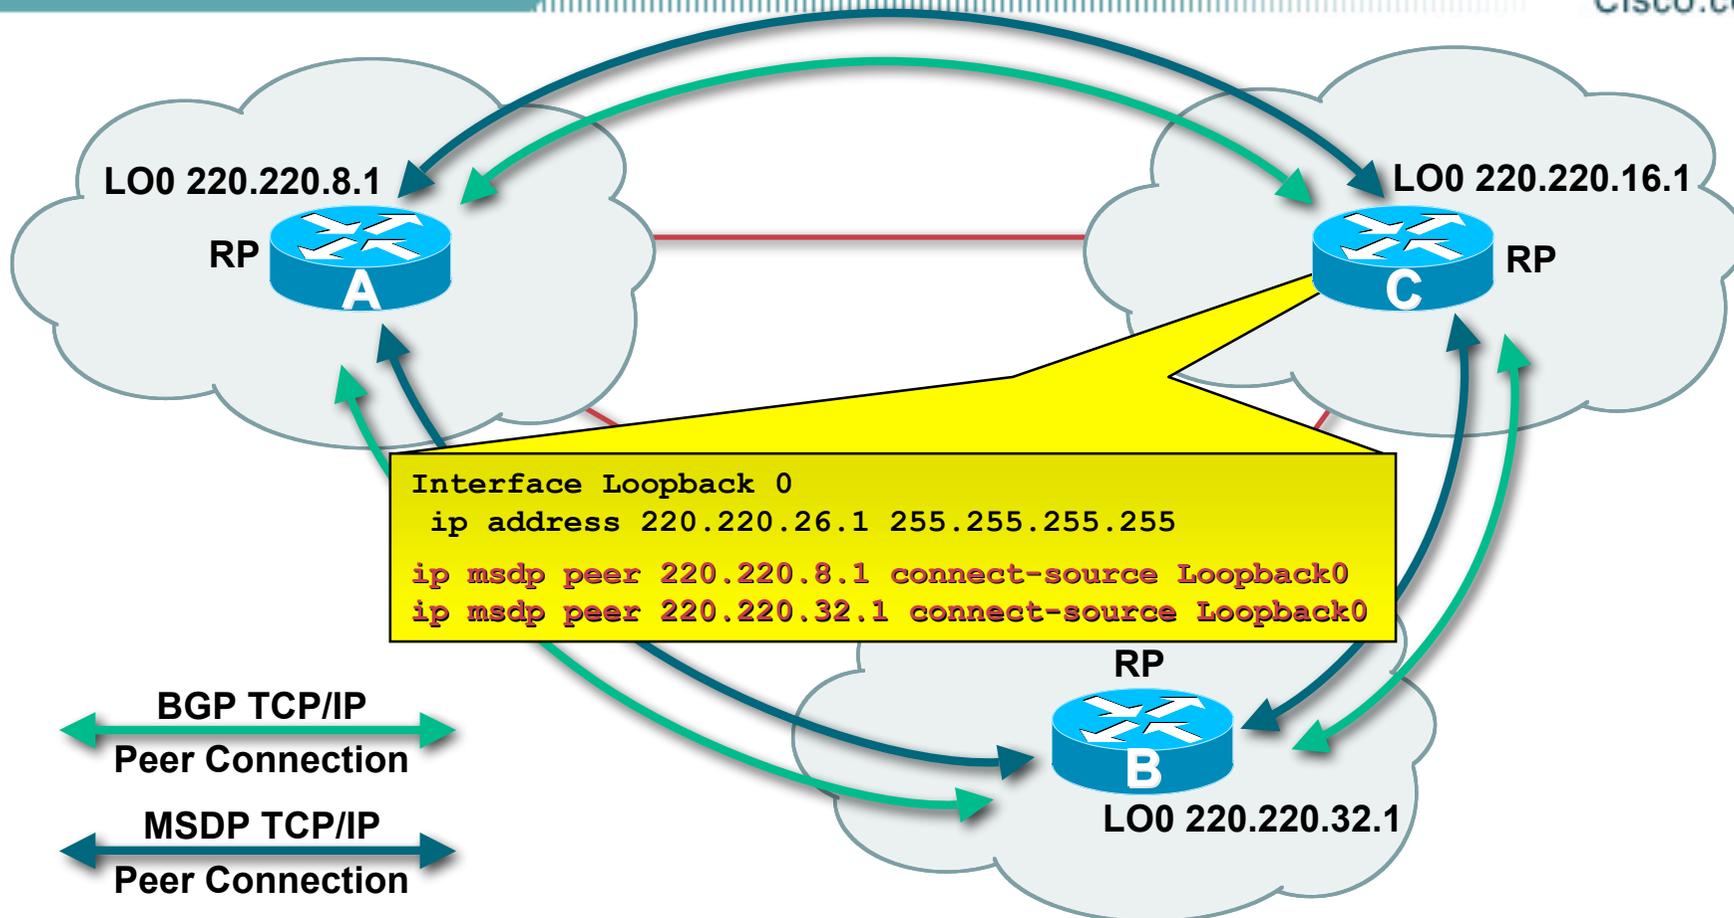
# MSDP Peers



- **MSDP peer connections are established using the MSDP “peer” configuration command**

```
ip msdp peer <ip-address> [connect-source <intfc>]
```

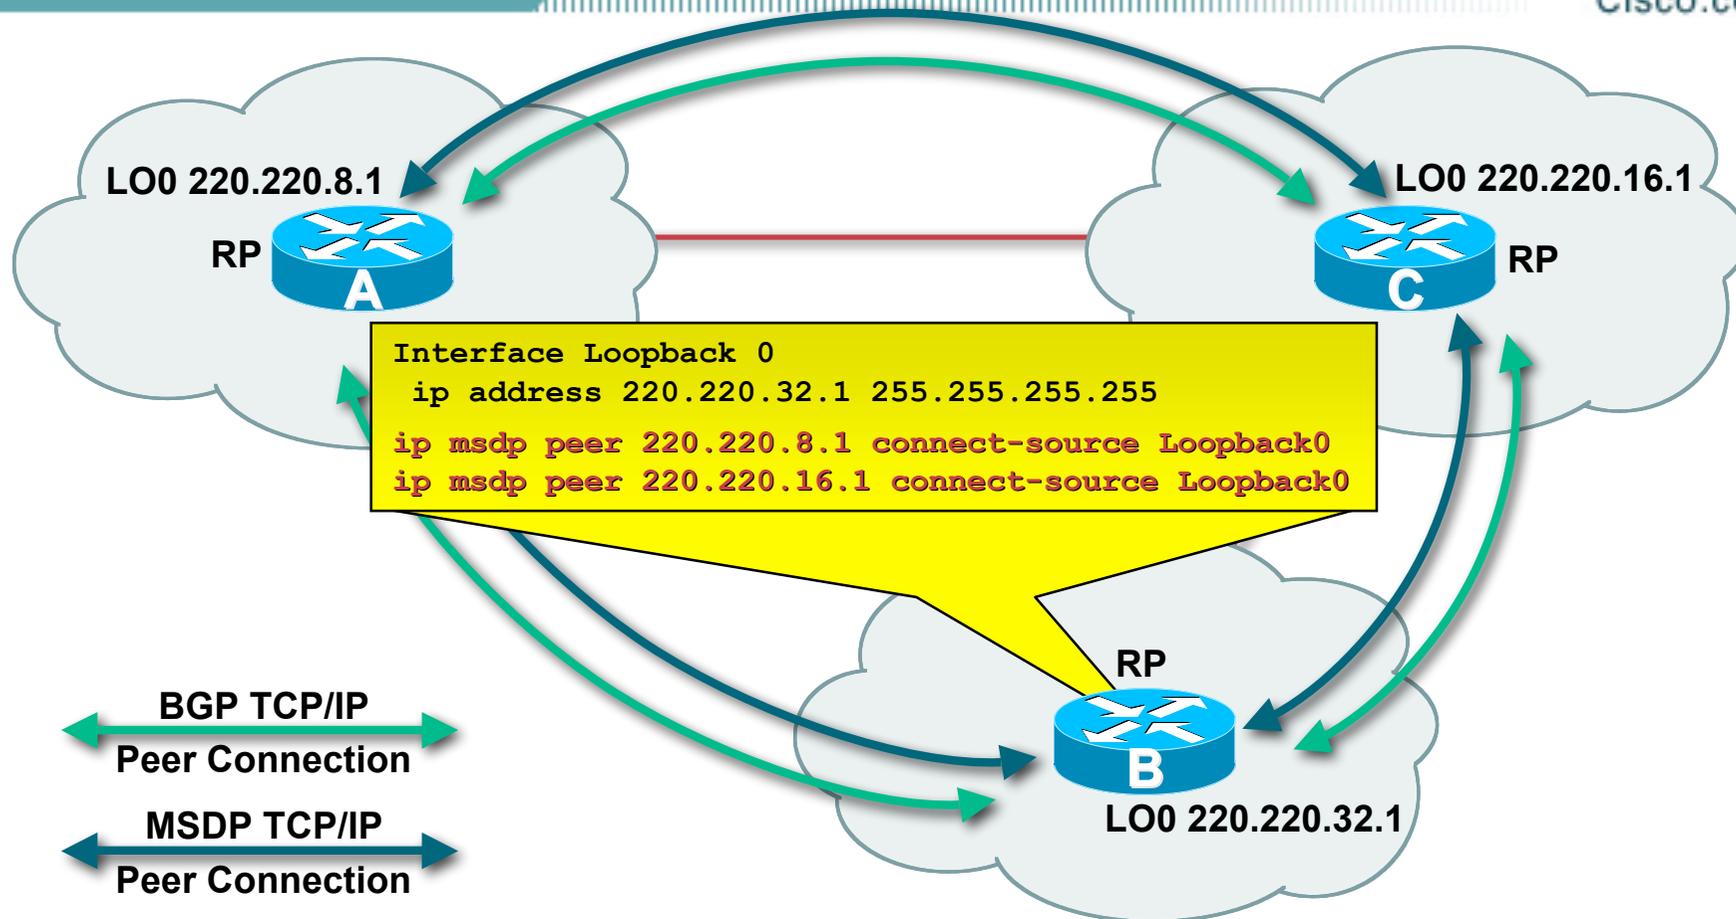
# MSDP Peers



- **MSDP peer connections are established using the MSDP “peer” configuration command**

```
ip msdp peer <ip-address> [connect-source <intfc>]
```

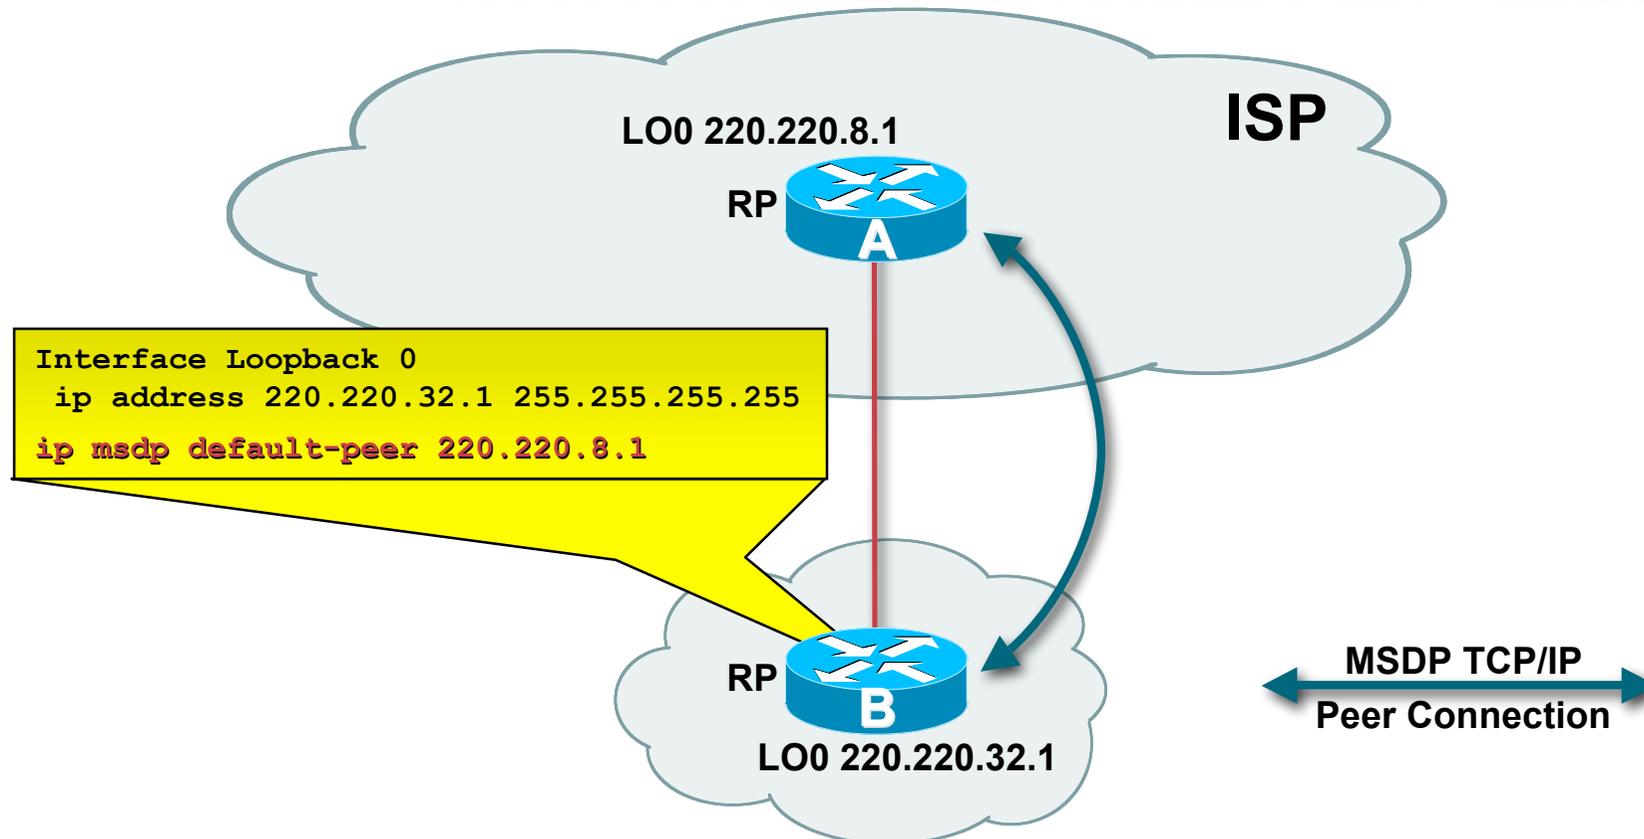
# MSDP Peers



- **MSDP peer connections are established using the MSDP “peer” configuration command**

```
ip msdp peer <ip-address> [connect-source <intfc>]
```

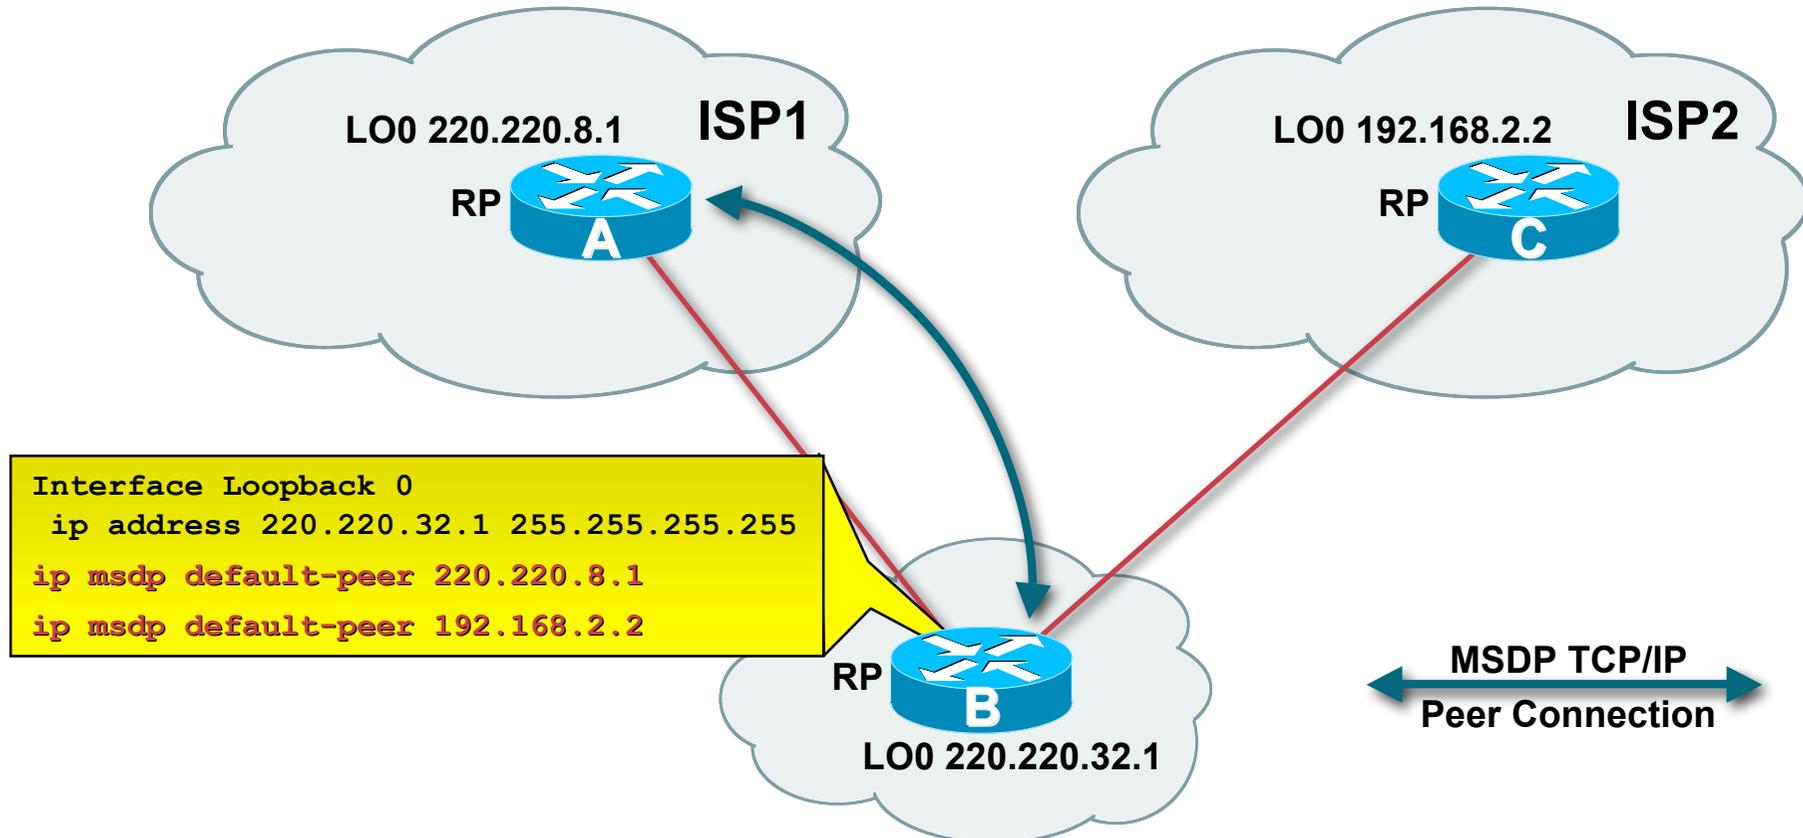
# MSDP Peers



- Stub-networks may use “default” peering without being a BGP peer by using the MSDP “default-peer” configuration command.

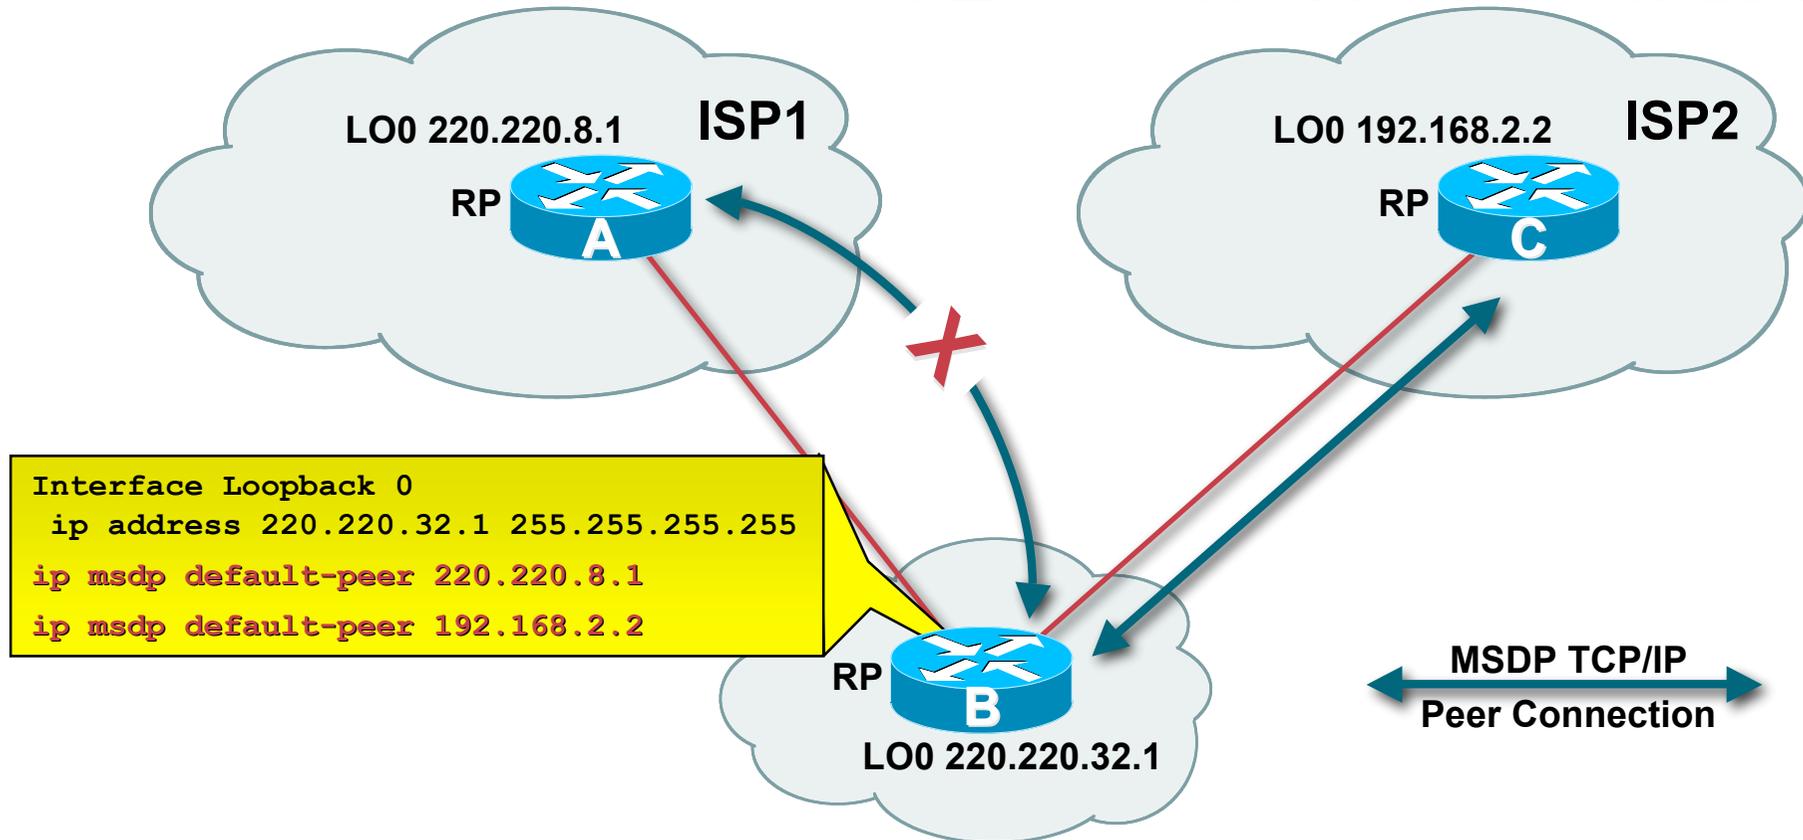
```
ip msdp default-peer <ip-address>
```

# MSDP Peers



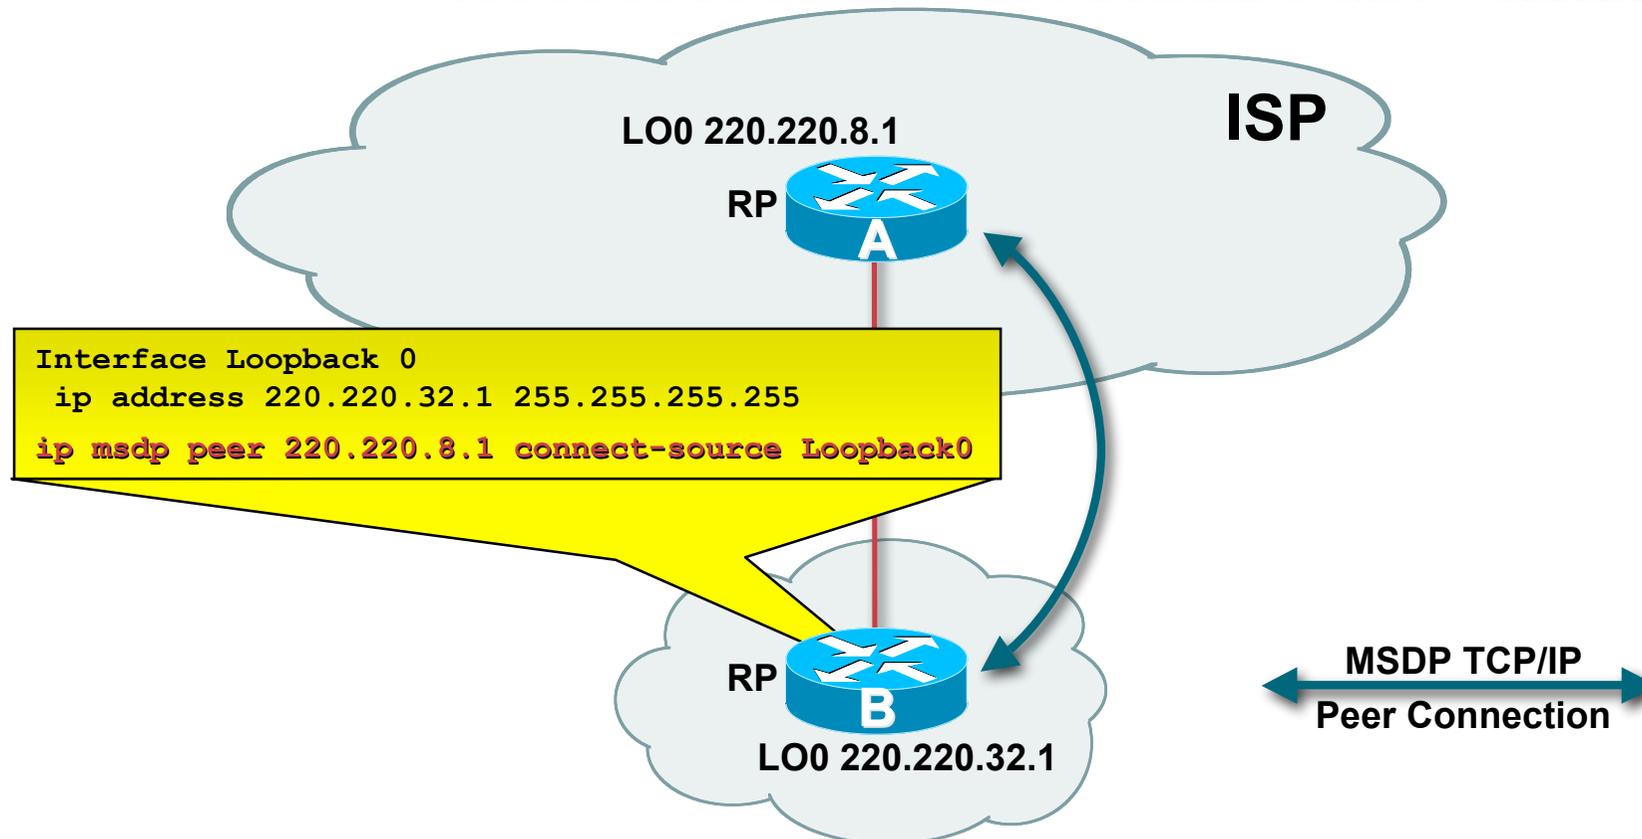
- Multiple “default-peers” may be configured in case connection to first default-peer goes down.

# MSDP Peers



- When connection to first 'default-peer' is lost, the next one in the list is tried.

# MSDP Peers



- **Stub-networks configured with only a single MSDP peer are treated in the same manner as when a single “default-peer” is configured. (i.e. BGP is not required.)**

# Agenda

- **MSDP Overview**
- **MSDP Peers**
- **MSDP SA Messages**
- **MSDP Mesh Groups**
- **MSDP State Flags**
- **MSDP Enhancements**

# SA Message Contents

- **MSDP Source Active (SA) Messages**
  - **Used to advertise active Sources in a domain**
  - **Can also carry 1st multicast packet from source**
    - **Hack for Bursty Sources (a'la SDR)**
  - **SA Message Contents:**
    - **IP Address of Originating RP**
    - **Number of (S, G)'s pairs being advertised**
    - **List of active (S, G)'s in the domain**
    - **Encapsulated Multicast packet [optional]**

# SA Messages

- **Originating SA Messages**
- **Receiving SA Messages**
- **SA Message Caching**
- **SA Input/Output Filtering**
- **SA Message RPF Checking**

# Originating SA Messages

- **Local Sources**
  - **RP's only originate SA's for local sources**
    - **Denoted by the "A" flag on an (S,G) entry on RP**
  - **A source is local if:**
    - **The RP received a "Register" for (S, G), or**
    - **The source is directly connected to RP**

# Originating SA Messages

- Use 'msdp redistribute' to control what SA's are originated.

– Think of this as '**msdp sa-originate-filter**' function

```
ip msdp redistribute [list <acl>]
                    [asn <aspath-acl>]
                    [route-map <map>]
```

- Filter by (S,G) pair using 'list <acl>'
- Filter by AS-PATH using 'asn <aspath-acl>'
- Filter based on route-map '<map>'

– Omitting all acl's stops all SA origination

Example: ip msdp redistribute

- Default: Originate SA's for all local sources
- If 'msdp redistribute' command is not configured

# Originating SA Messages

- **SA messages are triggered when any new source in the local domain goes active.**
  - **Initial multicast packet is encapsulated in an SA message.**
    - **This is an attempt at solving the bursty-source problem**

# Originating SA Messages

- **Encapsulating Initial Multicast Packets**
  - Can bypass TTL-Thresholds
    - Original TTL is inside of data portion of SA message
    - SA messages sent via Unicast with TTL = 255
- **Requires special command to control**

```
ip msdp ttl-threshold <peer-address> <ttl>
```

  - Encapsulated multicast packets with a TTL lower than <ttl> for the specific MSDP peer are not forwarded or originated.

# Originating SA Messages

- **Once a minute**
  - **Router scans mroute table**
  - **If group = sparse AND router = RP for group**
    - **For each (S,G) entry for the group:**
      - **If the 'msdp redistribute' filters permits**
      - **AND if the source is a local source**
      - **Then originate an SA message for (S,G)**

# SA Messages

- **Originating SA Messages**
- **Receiving SA Messages**
- **SA Message Caching**
- **SA Input/Output Filtering**
- **SA Message RPF Checking**

# Receiving SA Messages

- **If SA message RPF checks OK**
  - **Store in SA Cache**
  - **If new SA cache entry**
    - Immediately flood SA downstream
    - Set entry's SA-expire-timer to 6 minutes.
    - If RP for group and receivers exist
      - » Create (S,G) entry and trigger (S,G) Join
  - **If existing entry**
    - Reset entry's SA-expire-timer to 6 minutes.
      - » When timer = zero, entry has expired and is deleted.
- **Else**
  - **Discard SA**

# SA Messages

- **Originating SA Messages**
- **Receiving SA Messages**
- **SA Message Caching**
- **SA Input/Output Filtering**
- **SA Message RPF Checking**

# SA Message Cache

- **Enabling SA Caching**

- `ip msdp cache-sa-state [list <acl>]`

- **Caching is now on by default.**

- **Beginning with IOS versions 12.1(7), 12.0(14)S1.**
      - Cannot be turned off.

- **Router caches all SA messages.**

- **Cached (S, G) entries timeout after 6 minutes.**
      - If not refreshed by another (S,G) SA message.

- **Once per minute, router scans SA cache.**

- **Sends SA downstream for each entry in cache.**

# SA Message Caching

- Listing the contents of the SA Cache

```
show ip msdp sa-cache [<group-or-source>] [<asn>]
```

```
sj-mbone# show ip msdp sa-cache
MSDP Source-Active Cache - 1997 entries
(193.92.8.77, 224.2.232.0), RP 194.177.210.41, MBGP/AS 5408, 00:01:51/00:04:09
(128.119.167.221, 224.77.0.0), RP 128.119.3.241, MBGP/AS 1249, 06:40:59/00:05:12
(147.228.44.30, 233.0.0.1), RP 195.178.64.113, MBGP/AS 2852, 00:04:48/00:01:11
(128.117.16.142, 233.0.0.1), RP 204.147.128.141, MBGP/AS 145, 00:00:41/00:05:18
(132.250.95.60, 224.253.0.1), RP 138.18.100.1, MBGP/AS 668, 01:15:07/00:05:55
(128.119.40.229, 224.2.0.1), RP 128.119.3.241, MBGP/AS 1249, 06:40:59/00:05:12
(130.225.245.71, 227.37.32.1), RP 130.225.245.71, MBGP/AS 1835, 1d00h/00:05:29
(194.177.210.41, 227.37.32.1), RP 194.177.210.41, MBGP/AS 5408, 00:02:53/00:03:07
(206.190.42.106, 236.195.60.2), RP 206.190.40.61, MBGP/AS 5779, 00:07:27/00:04:04
:
:
.
```

- Clearing the contents of the SA Cache

```
clear ip msdp sa-cache [<group-address> | group-name]
```

# SA Messages

- **Originating SA Messages**
- **Receiving SA Messages**
- **SA Message Caching**
- **SA Input/Output Filtering**
- **SA Message RPF Checking**

# Filtering Incoming/Outgoing SA Messages

Cisco.com

- **SA Filter Command:**

```
ip msdp sa-filter {in|out} <peer-address> [list <acl>]
[route-map
<map>]
```

- Filters (S,G) pairs to / from peer based on specified ACL.
  - Can filter based on AS-Path by using optional route-map clause with a path-list acl.
  - You can filter flooded and originated SA's based on a specific peer, incoming and outgoing.
- **Caution: Filtering SA messages can break the Flood and Join mechanism!**

# Recommended MSDP SA Filter

```
! domain-local applications
access-list 111 deny ip any host 224.0.2.2 !
access-list 111 deny ip any host 224.0.1.3 ! Rwhod
access-list 111 deny ip any host 224.0.1.24 ! Microsoft-ds
access-list 111 deny ip any host 224.0.1.22 ! SVRLOC
access-list 111 deny ip any host 224.0.1.2 ! SGI-Dogfight
access-list 111 deny ip any host 224.0.1.35 ! SVRLOC-DA
access-list 111 deny ip any host 224.0.1.60 ! hp-device-
disc
!-- auto-rp groups
access-list 111 deny ip any host 224.0.1.39
access-list 111 deny ip any host 224.0.1.40
!-- scoped groups
access-list 111 deny ip any 239.0.0.0 0.255.255.255
!-- loopback, private addresses (RFC 1918)
access-list 111 deny ip 10.0.0.0 0.255.255.255 any
access-list 111 deny ip 127.0.0.0 0.255.255.255 any
access-list 111 deny ip 172.16.0.0 0.15.255.255 any
access-list 111 deny ip 192.168.0.0 0.0.255.255 any
access-list 111 permit ip any any
!-- Default SSM-range. Do not do MSDP in this range
access-list 111 deny ip any 232.0.0.0 0.255.255.255
access-list 111 permit ip any any
```

See “<ftp://ftp-eng.cisco.com/ipmulticast/msdp-sa-filter.txt>” for the latest updates to this list.

# SA Messages

- **Originating SA Messages**
- **Receiving SA Messages**
- **SA Message Caching**
- **SA Input/Output Filtering**
- **SA Message RPF Checking**

# SA Message RPF Checking

- **Purpose**
  - **Accept SA's via a single deterministic path**
    - Ignore all other arriving SA's
    - Necessary to prevent SA's from looping endlessly
- **Problem**
  - **Need to know MSDP topology of Internet**
    - **But, MSDP does not distribute topology data!**
- **Solution**
  - **Use BGP data to *infer* MSDP topology.**
    - **Impact:**
      - The MSDP topology must follow BGP topology.
      - An MSDP peer must *generally* also be an BGP peer.

# SA Message RPF Checking

- **RPF Check Rules depend on peering**
  - Rule 1: Sending MSDP peer = iBGP peer
  - Rule 2: Sending MSDP peer = eBGP peer
  - Rule 3: Sending MSDP peer != BGP peer
- **Exceptions:**
  - RPF check is skipped when:
    - Sending MSDP peer = Originating RP
    - Sending MSDP peer = Mesh-Group peer
    - Sending MSDP peer = only MSDP peer
      - (i.e. the 'default-peer' or the only 'msdp-peer' configured.)

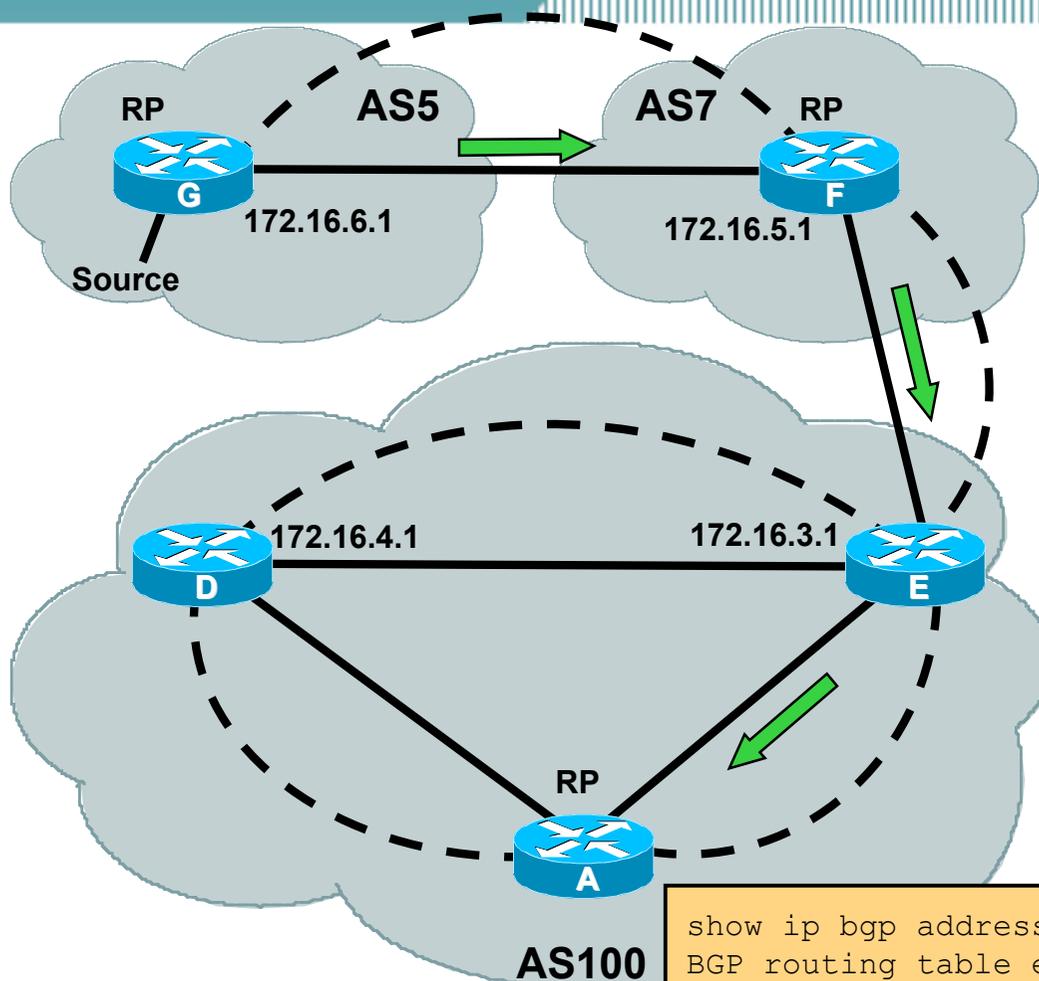
# SA Message RPF Checking

- **Determining Applicable RPF Rule**
  - Use IP address of sending MSDP peer
    - Find BGP neighbor w/matching IP address
    - IF (no match found)
      - Apply Rule 3
    - IF (matching neighbor = iBGP peer)
      - Apply Rule 1
    - ELSE {matching neighbor = eBGP peer}
      - Apply Rule 2
- ***Implication***
  - ***The MSDP peer address must be configured using the same IP address as the BGP peer!***

# RPF Check Rule 1

- **When MSDP peer = iBGP peer**
  - Find “Best Path” to RP in BGP Tables
    - Search M-Table first then U-Table.
    - If no path to Originating RP found, RPF Fails
  - Note “BGP Neighbor” that advertised path (i.e IP Address of BGP peer that sent us this path)
    - **Warning:**
      - *This is not the same as the Next-hop of the path!!!*
      - *iBGP peers normally do not set Next-hop = Self.*
      - *This is also not necessarily the same as the Router-ID!*
  - Rule 1 Test Condition:
    - MSDP Peer address = BGP Neighbor address?
      - If Yes, RPF Succeeds

# Rule1: MSDP peer = iBGP peer



iBGP peer address = 172.16.3.1  
(advertising best-path to RP)

MSDP Peer address = 172.16.3.1

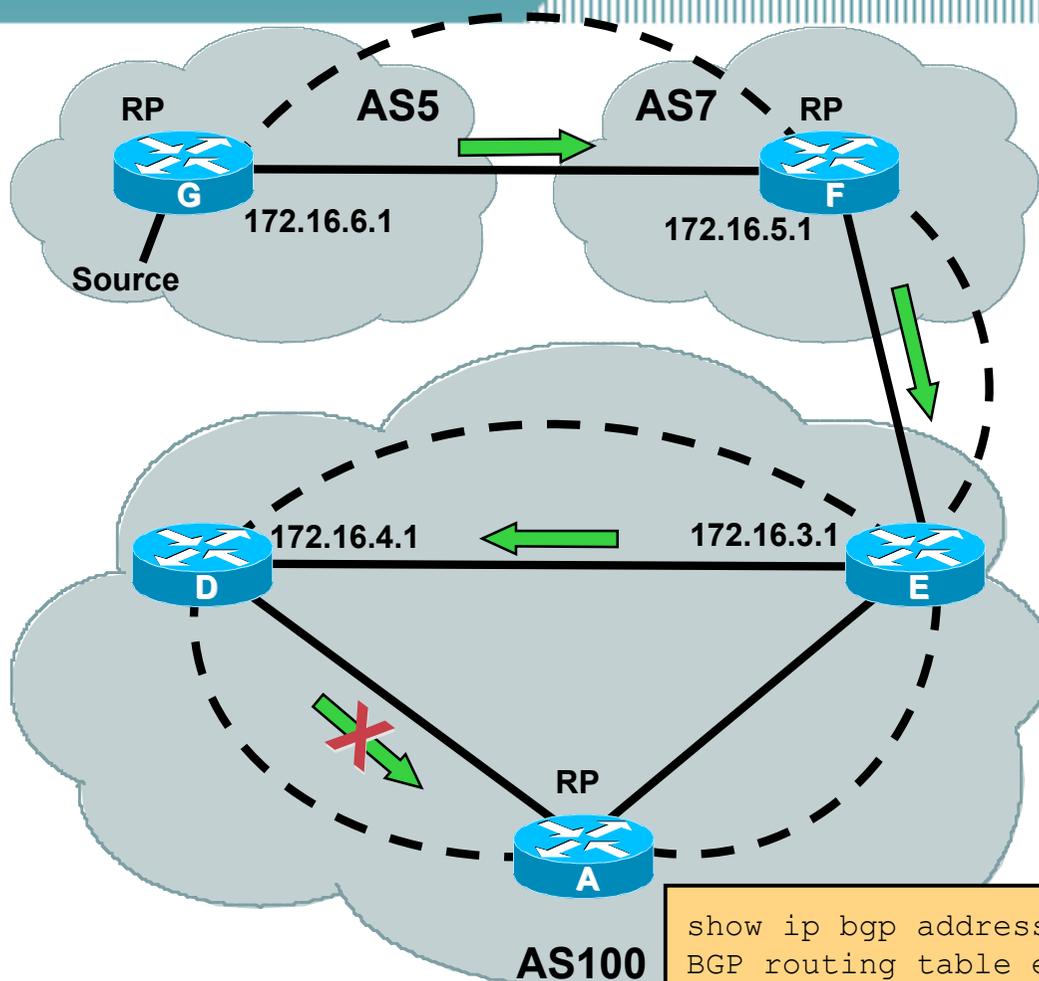
MSDP Peer address = iBGP Peer address

SA RPF Check Succeeds

```
show ip bgp address-family ipv4 multicast 172.16.6.1
BGP routing table entry for 172.16.6.0/24, version 8745118
Paths: (1 available, best #1)
 7 5, (received & used)
    172.16.5.1 (metric 68096) from 172.16.3.1 (172.16.3.1)
```

- BGP Peer
- MSDP Peer
- SA Message

# Rule1: MSDP peer = iBGP peer



iBGP Peer address = 172.16.3.1  
(advertising best-path to RP)

MSDP Peer address = 172.16.4.1

**MSDP Peer address != iBGP Peer address**

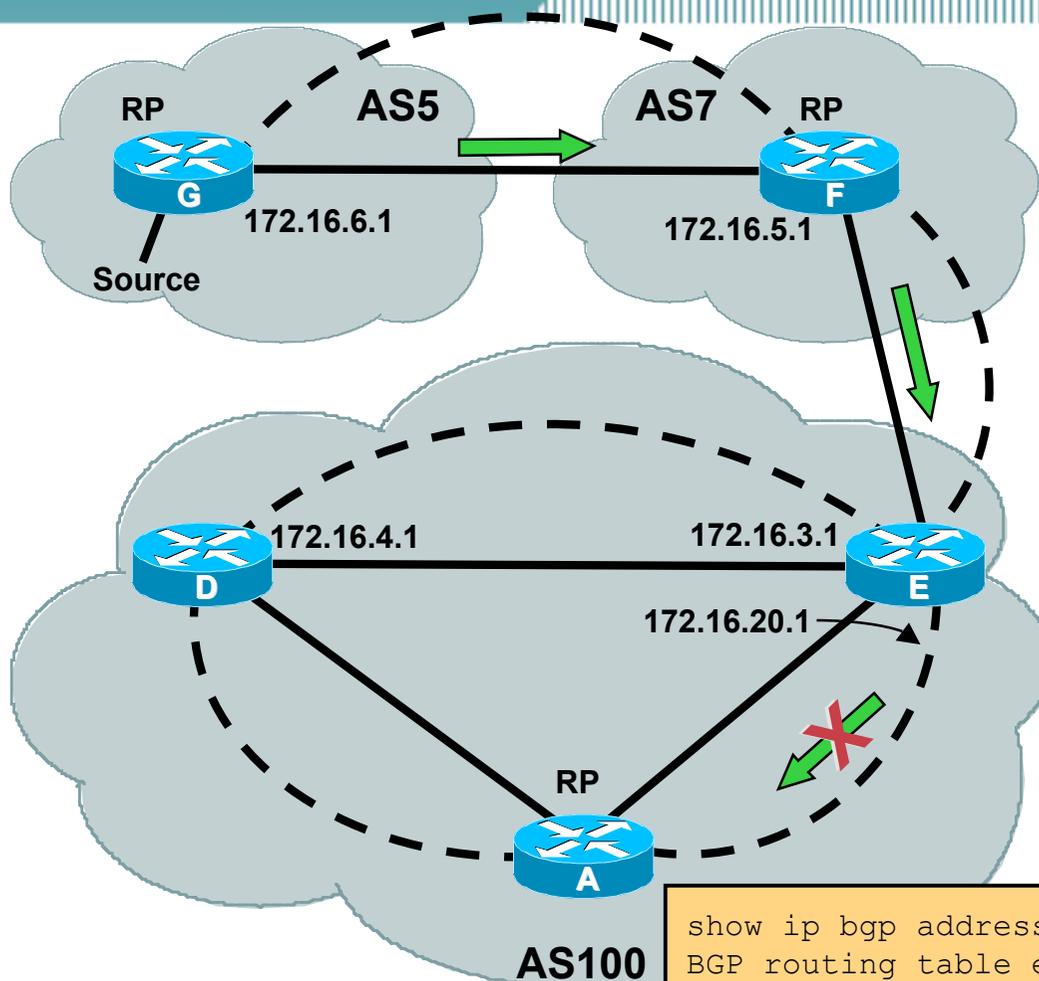
**SA RPF Check Fails**

BGP Peer ———  
MSDP Peer - - - -  
SA Message →

```
show ip bgp address-family ipv4 multicast 172.16.6.1
BGP routing table entry for 172.16.6.0/24, version 8745118
Paths: (1 available, best #1)
 7 5, (received & used)
    172.16.5.1 (metric 68096) from 172.16.3.1 (172.16.3.1)
```

# Rule1: MSDP peer = iBGP peer

Cisco.com



## Common Mistake #1:

*Failure to use same addresses for MSDP peers as iBGP peers!*

iBGP Peer address = 172.16.3.1  
(advertising best-path to RP)

MSDP Peer address = 172.16.20.1

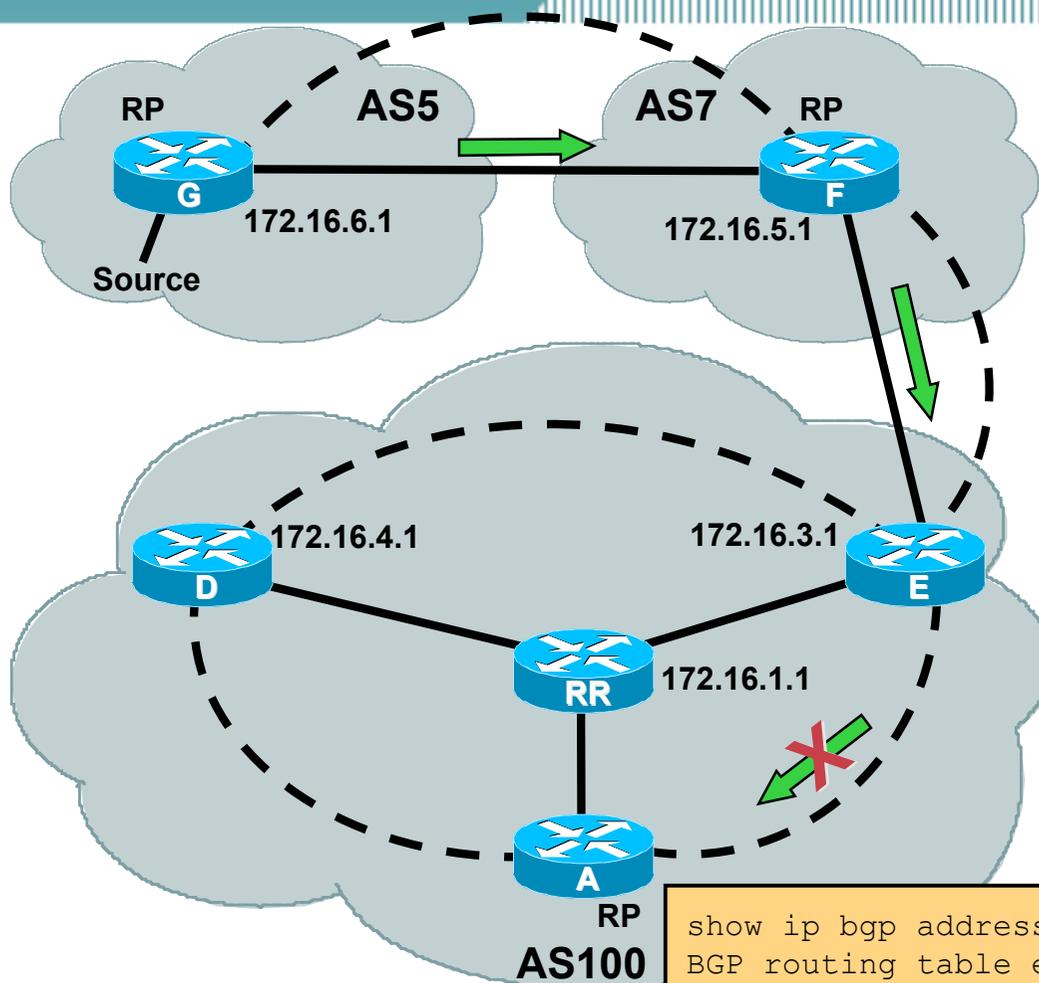
**MSDP Peer address != iBGP Peer address**

**SA RPF Check Fails**

```
show ip bgp address-family ipv4 multicast 172.16.6.1
BGP routing table entry for 172.16.6.0/24, version 8745118
Paths: (1 available, best #1)
 7 5, (received & used)
    172.16.5.1 (metric 68096) from 172.16.3.1 (172.16.3.1)
```

# Rule1: MSDP peer = iBGP peer

Cisco.com



## Common Mistake #2:

*Failure to follow iBGP topology!  
Can happen when RR's are used.*

iBGP Peer address = 172.16.1.1  
(advertising best-path to RP)

MSDP Peer address = 172.16.3.1

**MSDP Peer address != iBGP Peer address**

**SA RPF Check Fails**

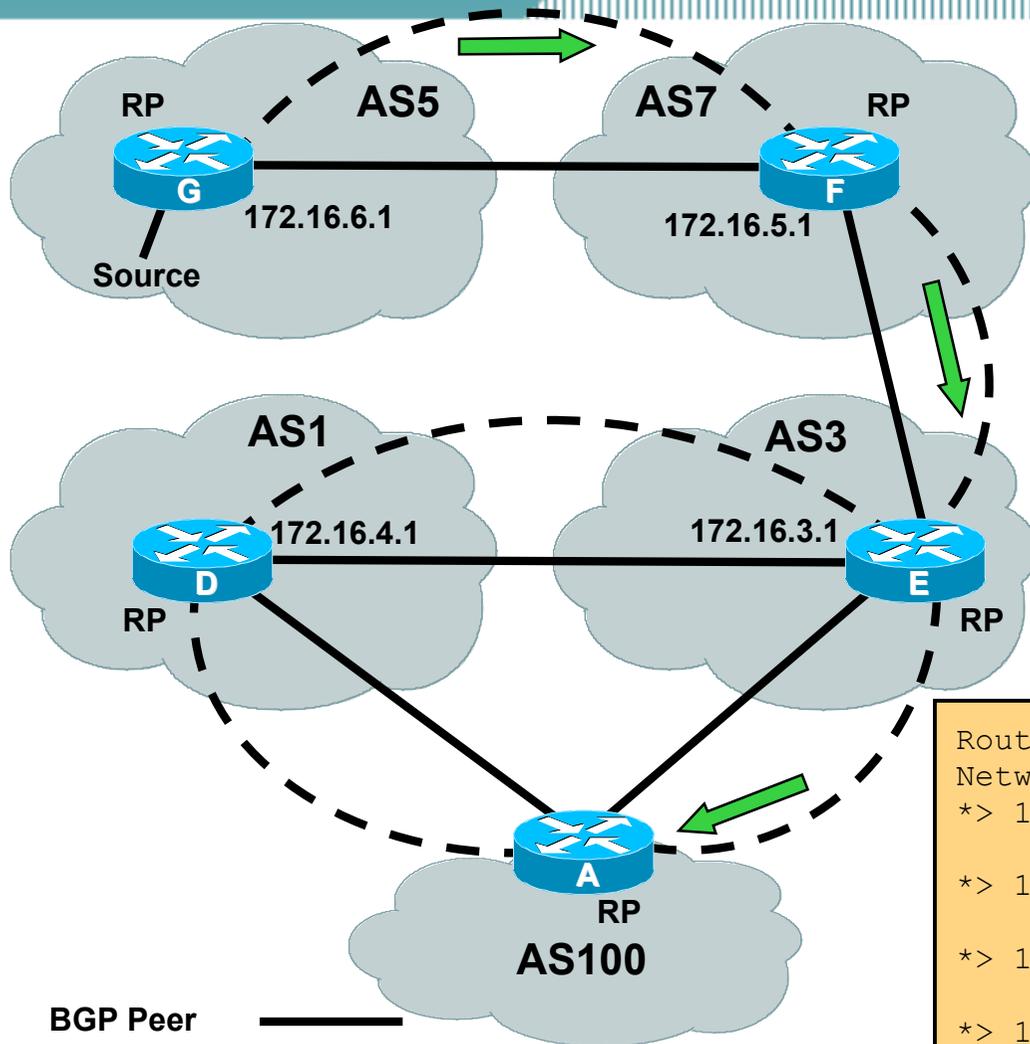
```
show ip bgp address-family ipv4 multicast 172.16.6.1
BGP routing table entry for 172.16.6.0/24, version 8745118
Paths: (1 available, best #1)
 7 5, (received & used)
    172.16.5.1 (metric 68096) from 172.16.1.1 (172.16.1.1)
```

BGP Peer ———  
MSDP Peer - - -  
SA Message →

# RPF Check Rule 2

- **When MSDP peer = eBGP peer**
  - **Find BGP “Best Path” to RP**
    - **Search M-Table first then U-Table.**
      - If no path to Originating RP found, RPF Fails
  - **Rule 2 Test Condition:**
    - **First AS in path to the RP = AS of eBGP peer?**
      - If Yes, RPF Succeeds

# Rule2: MSDP peer = eBGP peer



First-AS in best-path to RP = 3  
AS of MSDP Peer = 3

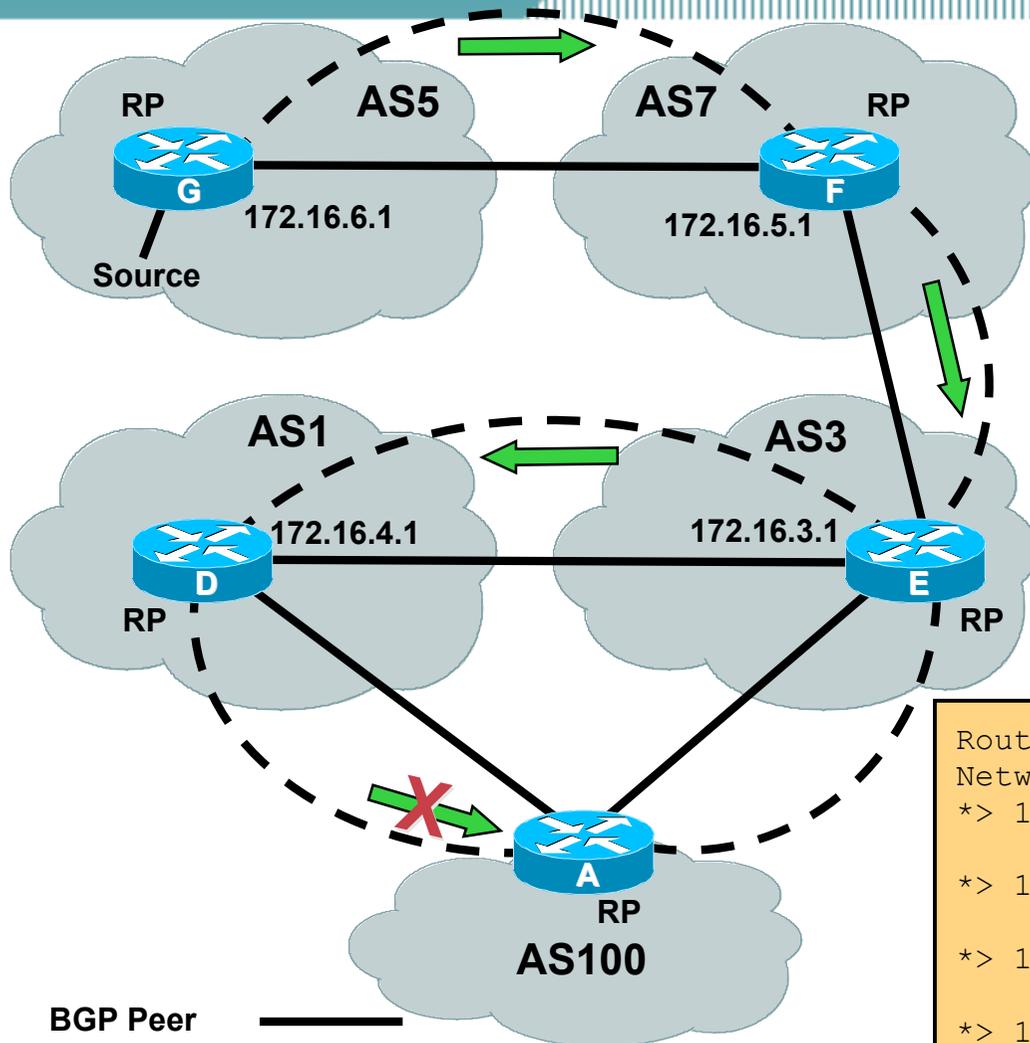
First-AS in best-path to RP = AS of eBGP Peer  
**SA RPF Check Succeeds**

Router A's ipv4 multicast BGP Table

Network	Next Hop	Path
*> 172.16.3.0/24	172.16.3.1	<b>3</b> i
172.16.3.0/24	172.16.4.1	1 3 i
*> 172.16.4.0/24	172.16.4.1	1 i
172.16.4.0/24	172.16.3.1	3 1 i
*> 172.16.5.0/24	172.16.3.1	3 7 i
172.16.5.0/24	172.16.4.1	1 3 7 i
*> 172.16.6.0/24	172.16.3.1	<b>3</b> 7 5 i
172.16.6.0/24	172.16.4.1	1 3 7 5 i

BGP Peer ———  
MSDP Peer - - - -  
SA Message →

# Rule2: MSDP peer = eBGP peer



First-AS in best-path to RP = 3  
AS of eBGP Peer = 1

First-AS in best-path to RP != AS of eBGP Peer  
**SA RPF Check Fails!**

Router A's ipv4 multicast BGP Table

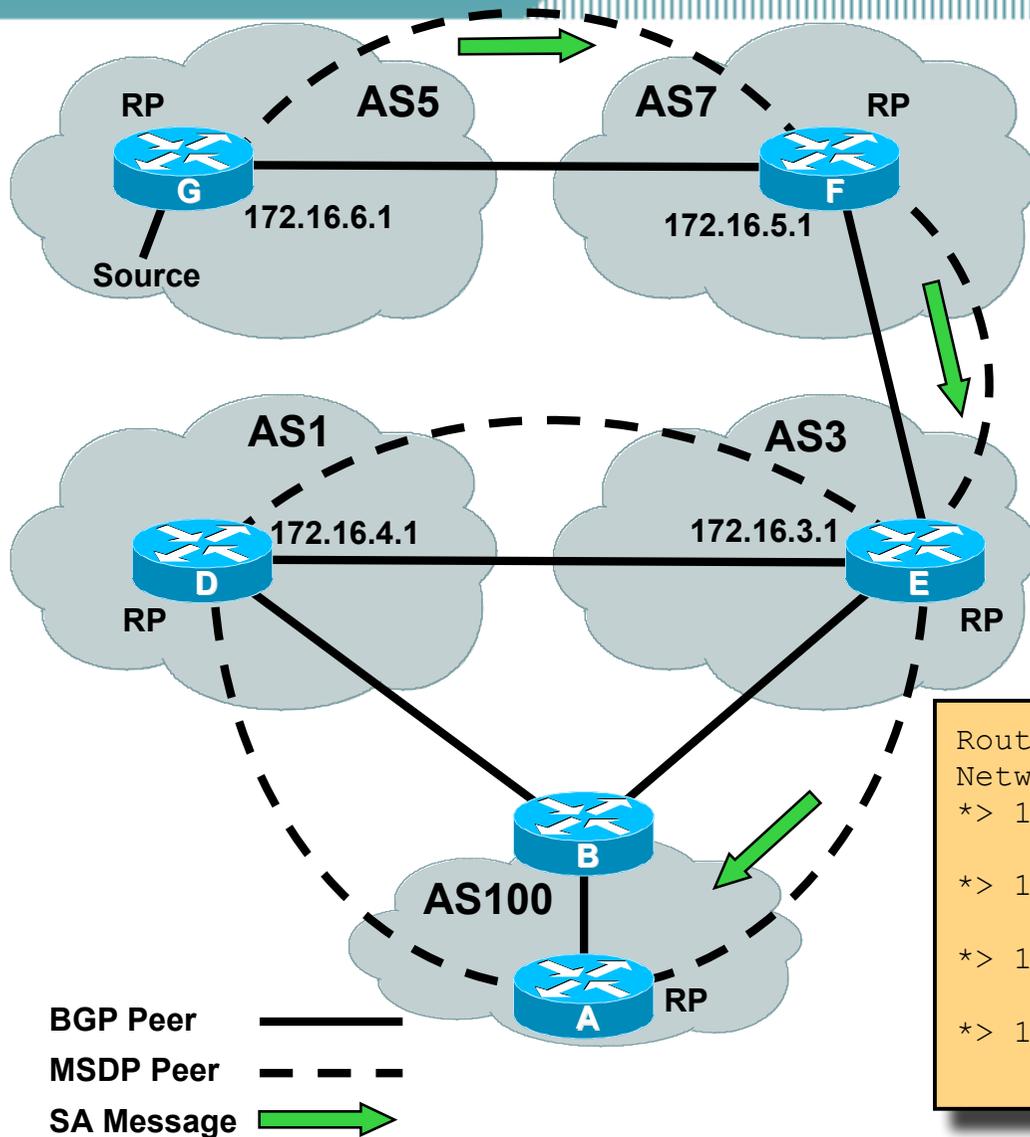
Network	Next Hop	Path
*> 172.16.3.0/24	172.16.3.1	3 i
172.16.3.0/24	172.16.4.1	1 3 i
*> 172.16.4.0/24	172.16.4.1	<b>1</b> i
172.16.4.0/24	172.16.3.1	3 1 i
*> 172.16.5.0/24	172.16.3.1	3 7 i
172.16.5.0/24	172.16.4.1	1 3 7 i
*> 172.16.6.0/24	172.16.3.1	<b>3</b> 7 5 i
172.16.6.0/24	172.16.4.1	1 3 7 5 i

BGP Peer ———  
MSDP Peer - - - -  
SA Message →

# RPF Check Rule 3

- **When MSDP peer != BGP peer**
  - **Find BGP “Best Path” to RP**
    - **Search M-Table first then U-Table.**
      - If no path to Originating RP found, RPF Fails
  - **Find BGP “Best Path” to MSDP peer**
    - **Search M-Table first then U-Table.**
      - If no path to sending MSDP Peer found, RPF Fails
  - **Note AS of sending MSDP Peer**
    - **Origin AS (last AS) in AS-PATH to MSDP Peer**
  - **Rule 3 Test Condition:**
    - **First AS in path to RP = Sending MSDP Peer AS ?**
      - If Yes, RPF Succeeds

# Rule3: MSDP peer != BGP peer



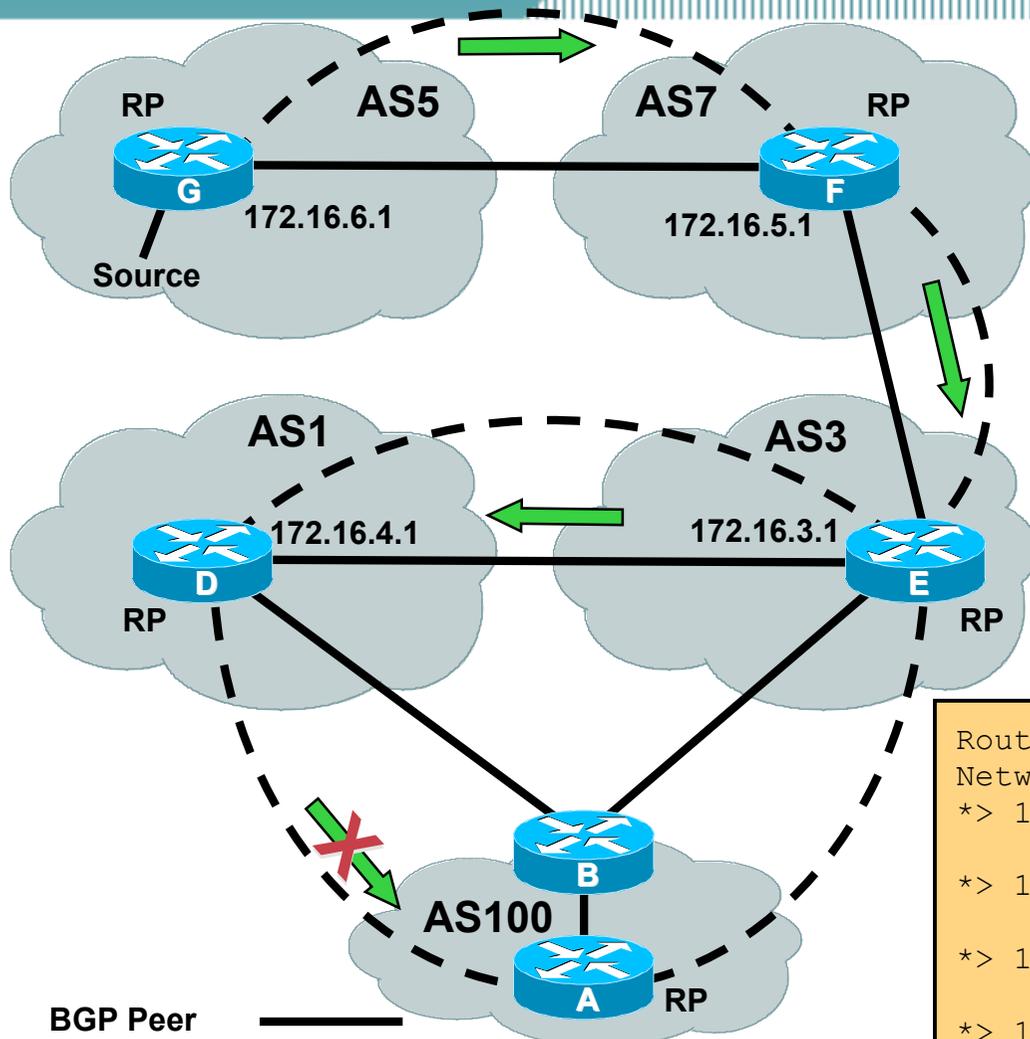
First-AS in best-path to RP = 3  
AS of MSDP Peer = 3

First-AS in best-path to RP = AS of MSDP Peer  
**SA RPF Check Succeeds**

Router A's ipv4 multicast BGP Table

Network	Next Hop	Path
*> 172.16.3.0/24	172.16.3.1	<b>3</b> i
172.16.3.0/24	172.16.4.1	1 3 i
*> 172.16.4.0/24	172.16.4.1	1 i
172.16.4.0/24	172.16.3.1	3 1 i
*> 172.16.5.0/24	172.16.3.1	3 7 i
172.16.5.0/24	172.16.4.1	1 3 7 i
*> 172.16.6.0/24	172.16.3.1	<b>3</b> 7 5 i
172.16.6.0/24	172.16.4.1	1 3 7 5 i

# Rule3: MSDP peer != BGP peer



**BGP Peer** ———  
**MSDP Peer** - - -  
**SA Message** →

**First-AS in best-path to RP = 3**  
**AS of MSDP Peer = 1**

**First-AS in best-path to RP != AS of MSDP Peer**

**SA RPF Check Fails**

Router A's ipv4 multicast BGP Table		
Network	Next Hop	Path
*> 172.16.3.0/24	172.16.3.1	3 i
172.16.3.0/24	172.16.4.1	1 3 i
*> 172.16.4.0/24	172.16.4.1	<b>1</b> i
172.16.4.0/24	172.16.3.1	3 1 i
*> 172.16.5.0/24	172.16.3.1	3 7 i
172.16.5.0/24	172.16.4.1	1 3 7 i
*> 172.16.6.0/24	172.16.3.1	<b>3</b> 7 5 i
172.16.6.0/24	172.16.4.1	1 3 7 5 i

# Agenda

- **MSDP Overview**
- **MSDP Peers**
- **MSDP SA Messages**
- **MSDP Mesh Groups**
- **MSDP State Flags**
- **MSDP Enhancements**

# MSDP Mesh-Groups

- **Optimises SA flooding.**
  - Useful when 2 or more peers are in a group.
  - Requires full mesh of mesh group peers.
- **Reduces amount of SA traffic in the net.**
  - SA's not flooded to other mesh-group peers.
- **Suspends RPF check of SA messages.**
  - When received from a mesh-group peer.
  - SA's always accepted from mesh-group peers.
  - Eliminates need for BGP.

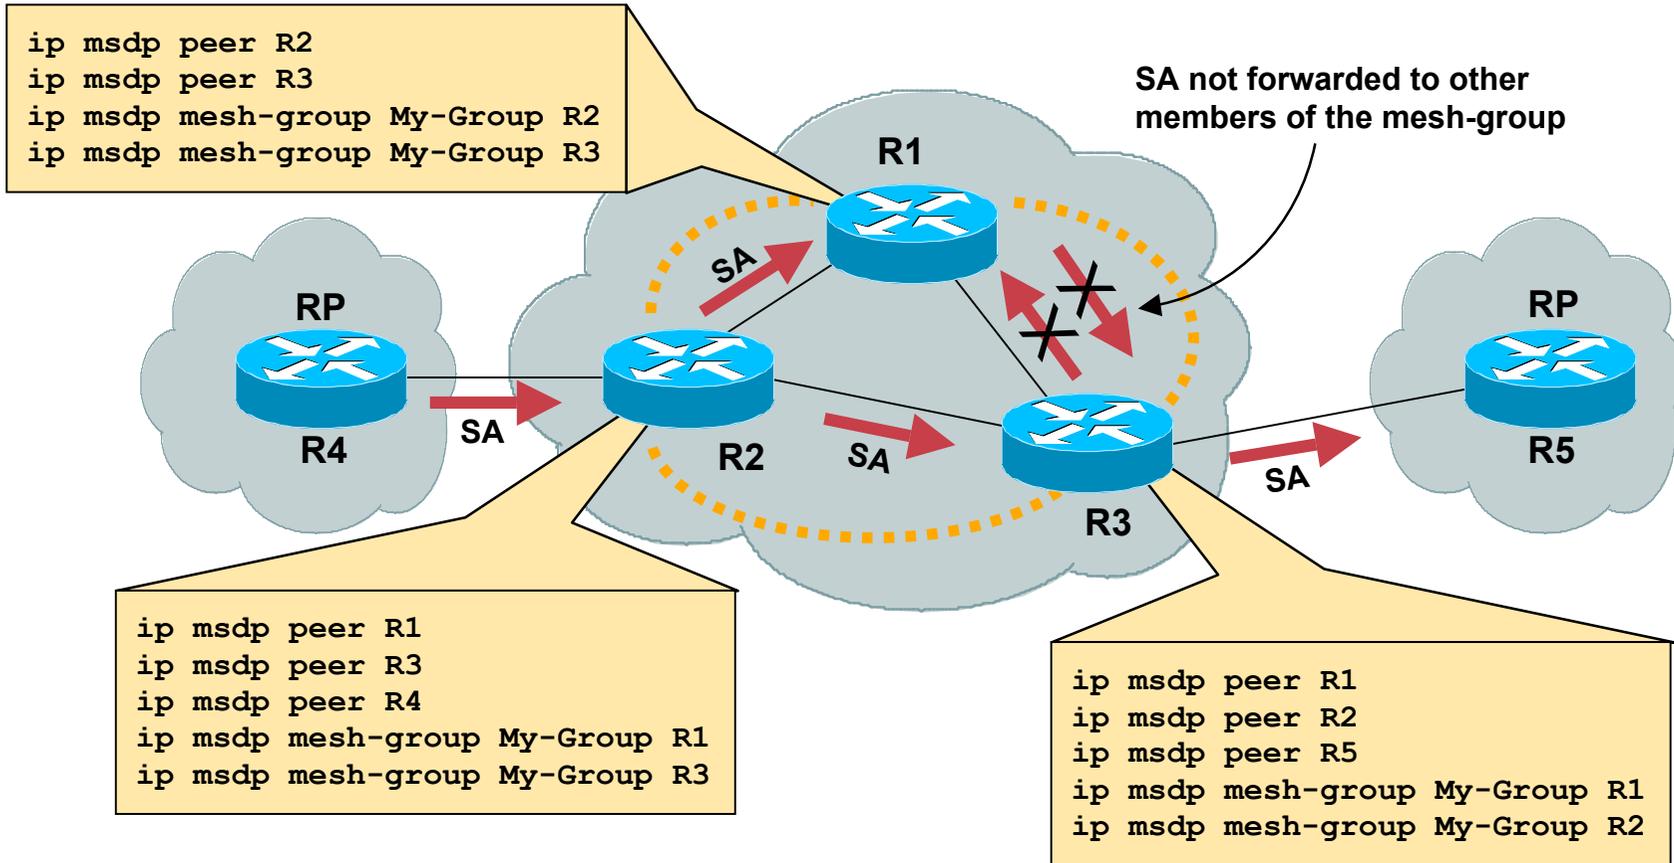
# MSDP Mesh-Groups

- **Configured with:**

```
ip msdp mesh-group <name> <peer-address>
```

- **Peers in the mesh-group must be fully meshed.**
- **Multiple mesh-groups per router are supported.**

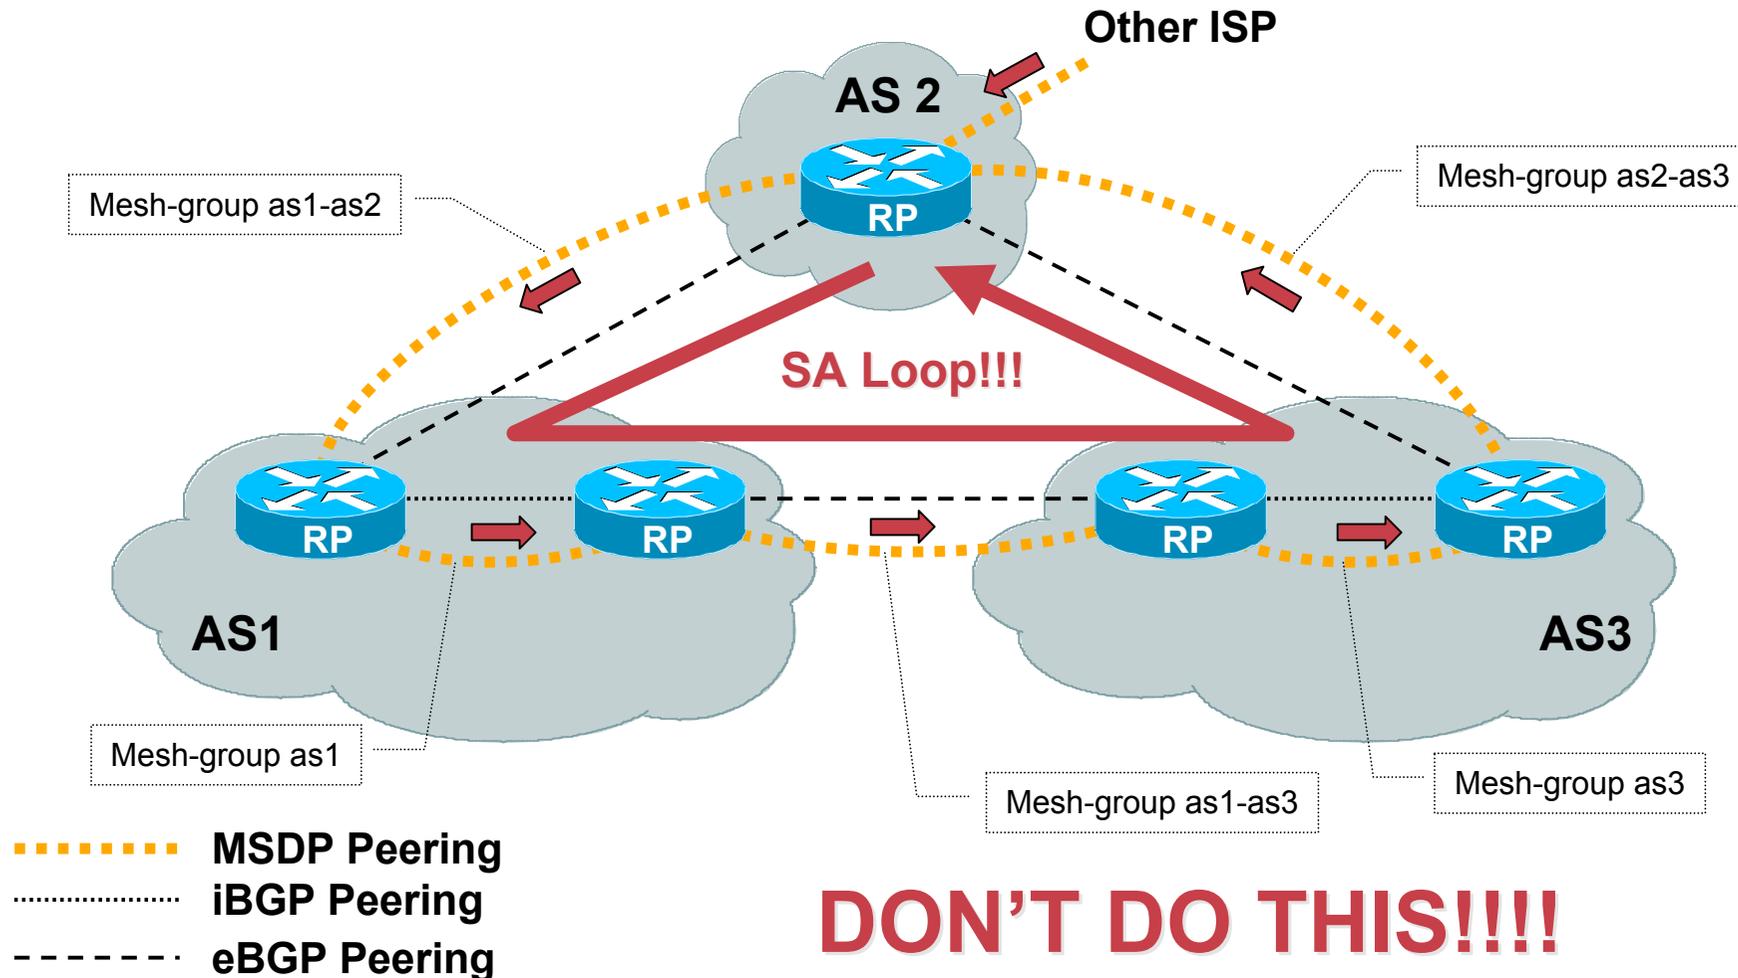
# MSDP Mesh-Group Example



..... MSDP mesh-group peering

# Avoid Mesh-Group Loops!!!

**WARNING: There is no RPF check between Mesh-groups!!!**



# Agenda

- **MSDP Overview**
- **MSDP Peers**
- **MSDP SA Messages**
- **MSDP Mesh Groups**
- **MSDP State Flags**
- **MSDP Enhancements**

# MSDP Mroute Flags

## New 'mroute' Flags for MSDP

```
sj-mbone#show ip mroute summary
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, 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 - Advertised via MSDP
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.2.246.13), 5d17h/00:02:59, RP 171.69.10.13, flags: S
(171.69.185.51, 224.2.246.13), 3d17h/00:03:29, flags: TA
(128.63.58.45, 224.2.246.13), 00:02:16/00:00:43, flags: M
(128.63.58.54, 224.2.246.13), 00:01:16/00:01:43, flags: M
```

“M” flag indicates source was learned via MSDP

“A” flag indicates source is a *candidate* for advertisement by MSDP

# Agenda

- **MSDP Overview**
- **MSDP Peers**
- **MSDP SA Messages**
- **MSDP Mesh Groups**
- **MSDP State Flags**
- **MSDP Enhancements**

# MSDP Enhancements

- **New IOS command**

- `ip msdp new-rpf-rules`

- **MSDP SA RPF check using IGP**

- **Accept SA's from BGP NEXT HOP**

- **Accept SA's from closest peer along the best path to the originating RP**

- **“show ip msdp rpf”**

# MSDP RPF check using IGP

- **When MSDP peer = IGP peer (No BGP)**

**Find best IGP route to RP**

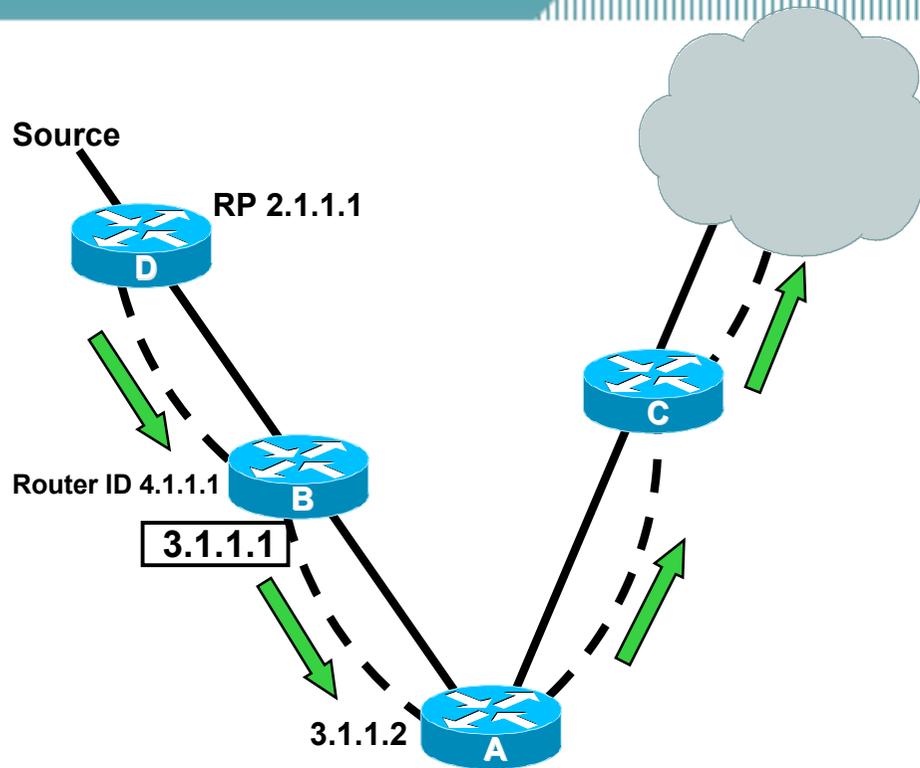
**Search URIB**

**If route to Originating RP found and:**

**If IGP next hop (or advertiser) address for RP is the MSDP peer and in UP state, then that is the RPF peer.**

**If route not found: Fall through to the next rule.**

# IGP Rule: MSDP peer = IGP peer (Next hop)



MSDP Peer = 3.1.1.1

IGP next hop to originating RP = 3.1.1.1

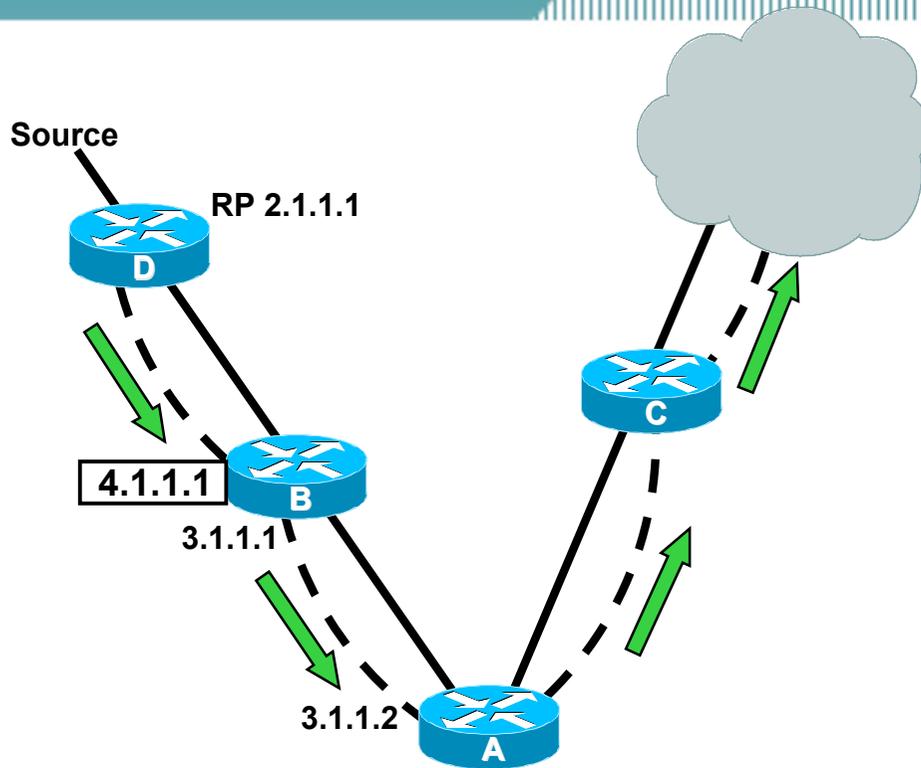
IGP next hop to originating RP = MSDP peer

**SA RPF Check Succeeds**

OSPF neighbor ———  
MSDP Peer - - - -  
SA Message →

```
RouterA#show ip route 2.1.1.1
Routing entry for 2.1.1.0/24
  Known via "ospf 1", distance 110, metric 20, type intra area
  Last update from 3.1.1.1 on Ethernet2, 00:35:10 ago
  Routing Descriptor Blocks:
    * 3.1.1.1, from 4.1.1.1, 00:35:10 ago, via Ethernet2
      Route metric is 20, traffic share count is 1
```

# IGP Rule: MSDP peer = IGP peer (Advertiser)



MSDP Peer = 4.1.1.1

IGP next hop to originating RP = ~~3.1.1.1~~

IGP advertiser to originating RP = 4.1.1.1

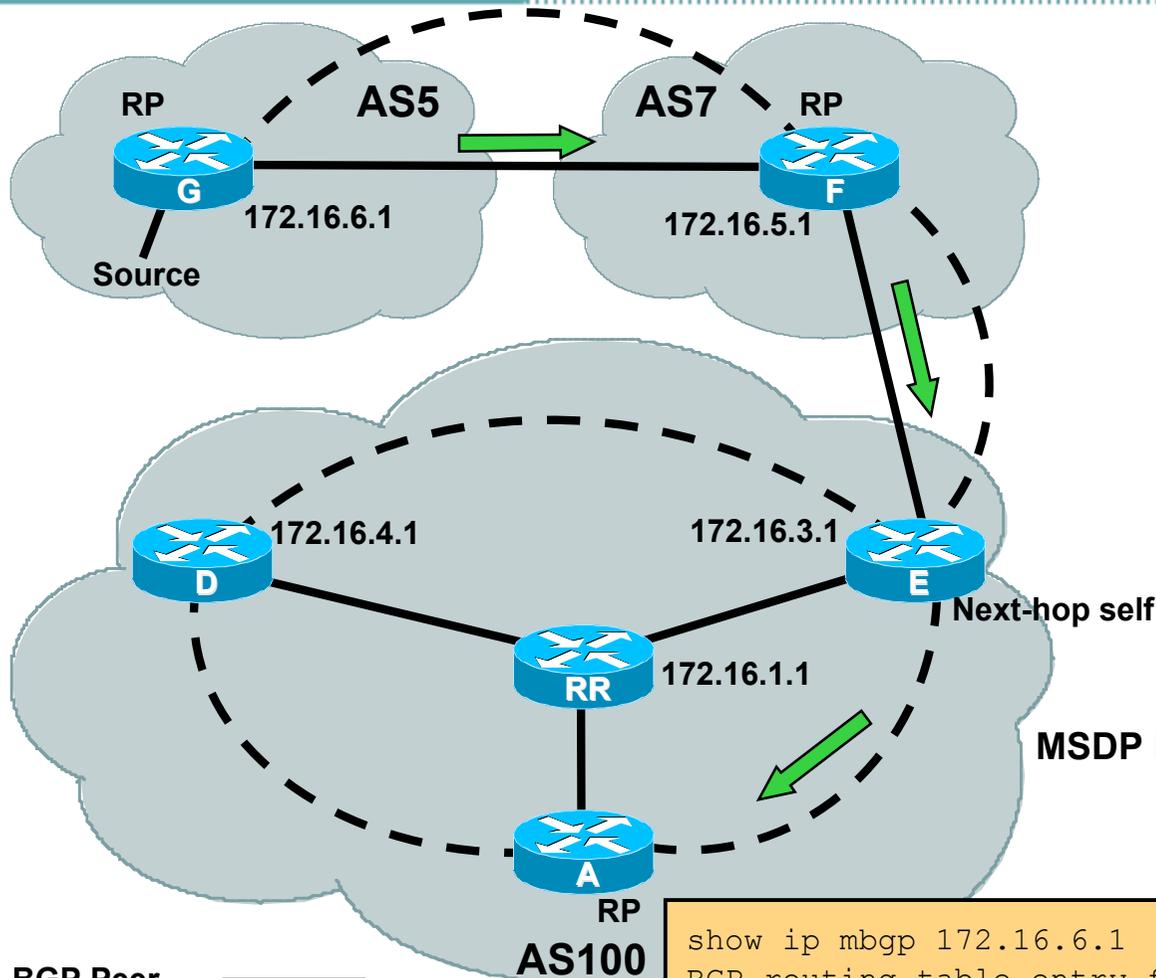
IGP advertiser to originating RP = MSDP peer

**SA RPF Check Succeeds**

- OSPF neighbor
- MSDP Peer
- SA Message

```
RouterA#show ip route 2.1.1.1
Routing entry for 2.1.1.0/24
  Known via "ospf 1", distance 110, metric 20, type intra area
  Last update from 3.1.1.1 on Ethernet2, 00:35:10 ago
  Routing Descriptor Blocks:
    * 3.1.1.1 from 4.1.1.1, 00:35:10 ago, via Ethernet2
      Route metric is 20, traffic share count is 1
```

# SA's accepted from Next Hop



i(m)BGP Peer address = 172.16.1.1  
(Advertiser of next hop)

MSDP Peer address = 172.16.3.1

But, BGP next hop = 172.16.3.1

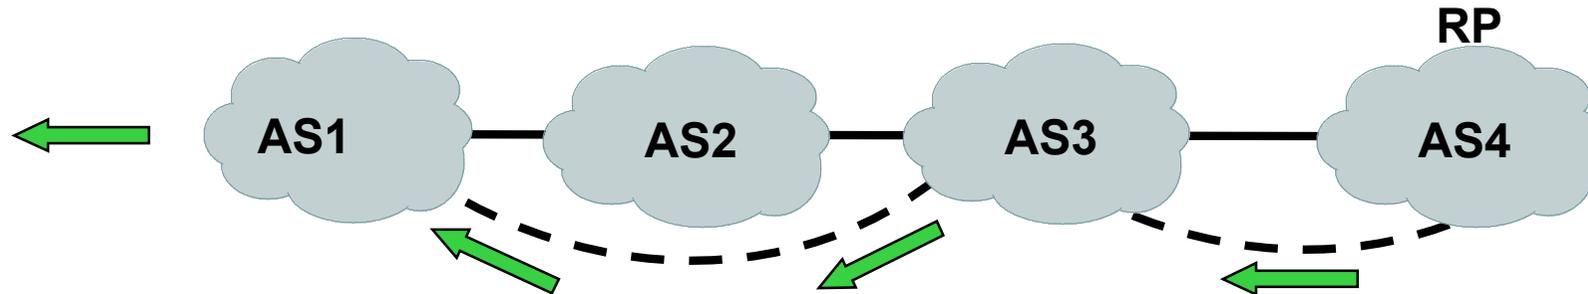
MSDP Peer address = BGP next hop address

**SA RPF Check Succeeds**

BGP Peer ———  
MSDP Peer - - - -  
SA Message →

```
show ip mbgp 172.16.6.1
BGP routing table entry for 172.16.6.0/24, version 8745118
Paths: (1 available, best #1)
7 5, (received & used)
  172.16.3.1 (metric 68096) from 172.16.1.1 (172.16.1.1)
```

# Accept SA along RPF path



**Existing Rule: If first AS in best path to the RP != MSDP peer**

**RPF Fails**

**New code: Choose peer in CLOSEST AS along best AS path to the RP.**

**Loosens rule a bit.**

**RPF Succeeds.**

BGP Peer ———

MSDP Peer - - -

SA Message →

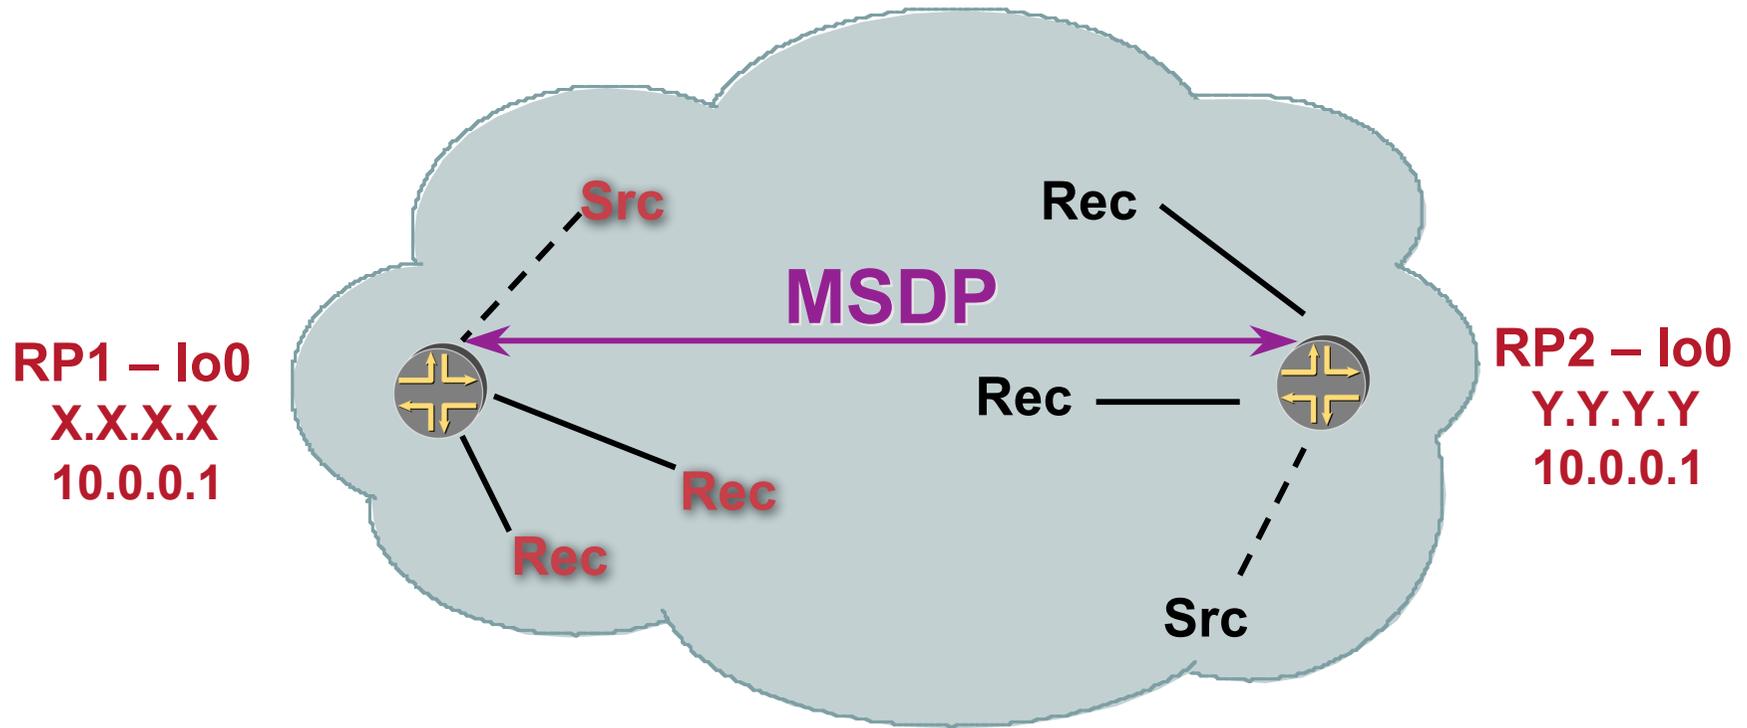
# New MSDP RPF command

```
Router-A# show ip msdp rpf 2.1.1.1
RPF peer information for Router-B (2.1.1.1)
  RPF peer: Router-C (3.1.1.1)
  RPF route/mask: 2.1.1.0/24
  RPF rule: Peer is IGP next hop of best route
  RPF type: unicast (ospf 1)
```

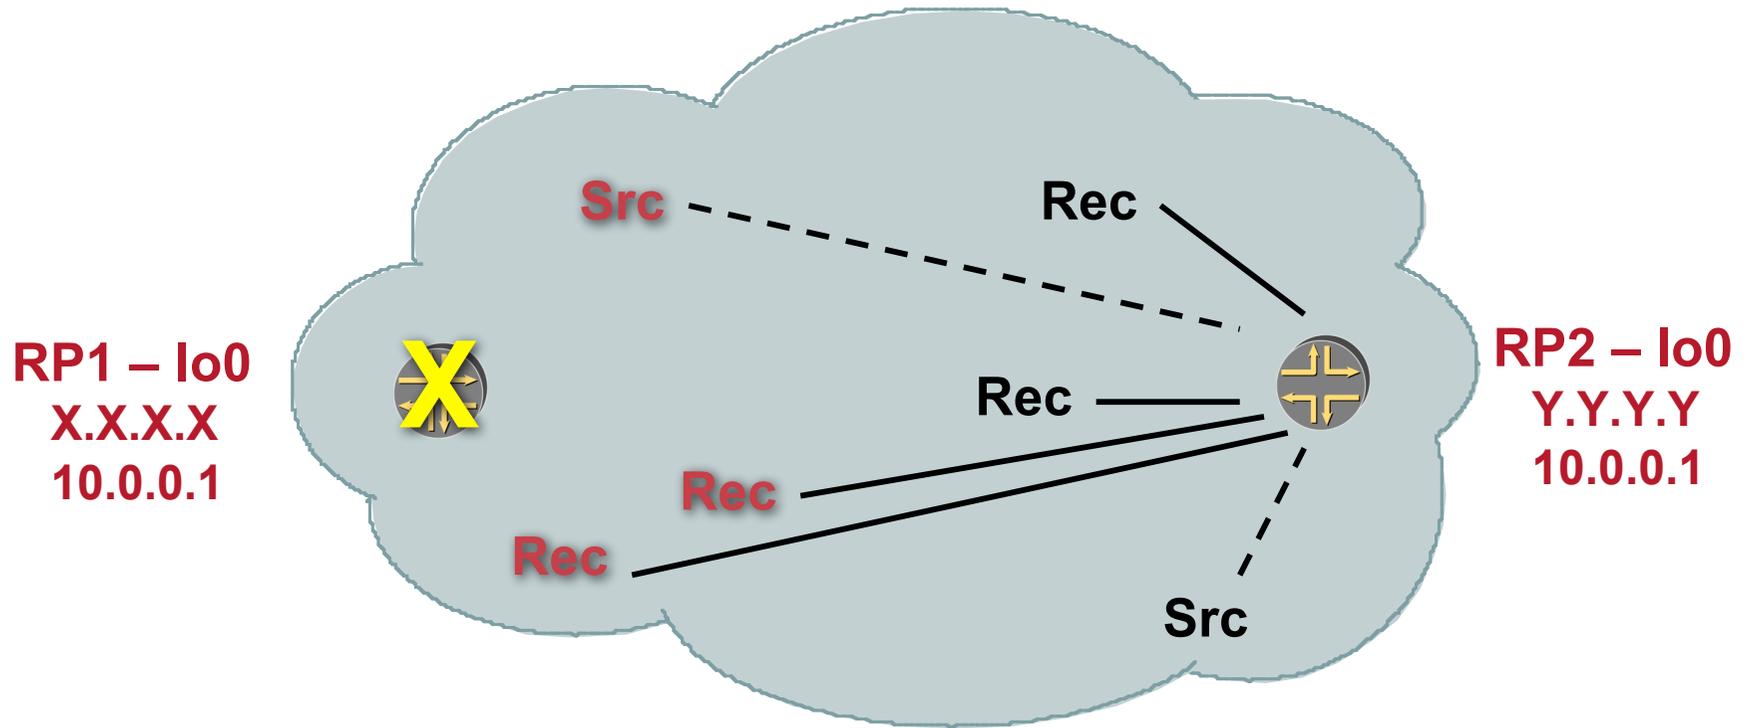
# Anycast-RP

- **draft-ietf-mboned-anycast-rp-08.txt**
- **Within a domain, deploy more than one RP for the same group range**
- **Sources from one RP are known to other RPs using MSDP**
- **Give each RP the same /32 IP address**
- **Sources and receivers use closest RP, as determined by the IGP**
- **Used intra-domain to provide redundancy and RP load sharing, when an RP goes down, sources and receivers are taken to new RP via unicast routing**
  - **Fast convergence!**

# Anycast-RP



# Anycast-RP



# CISCO SYSTEMS

