

Agenda

- Why Multicast VPNs?
- Multicast VPN Solution
- Cisco's Implementation
- Deployment Considerations

Why Multicast VPNs

- Until now only unicast has been supported in MPLS/BGP VPN
- VPN customers need multicast connectivity
 - Applications that require multicast
 - Internet multicast connectivity
- Service Providers want to offer additional services
 - e.g. Video streaming to its VPN customers

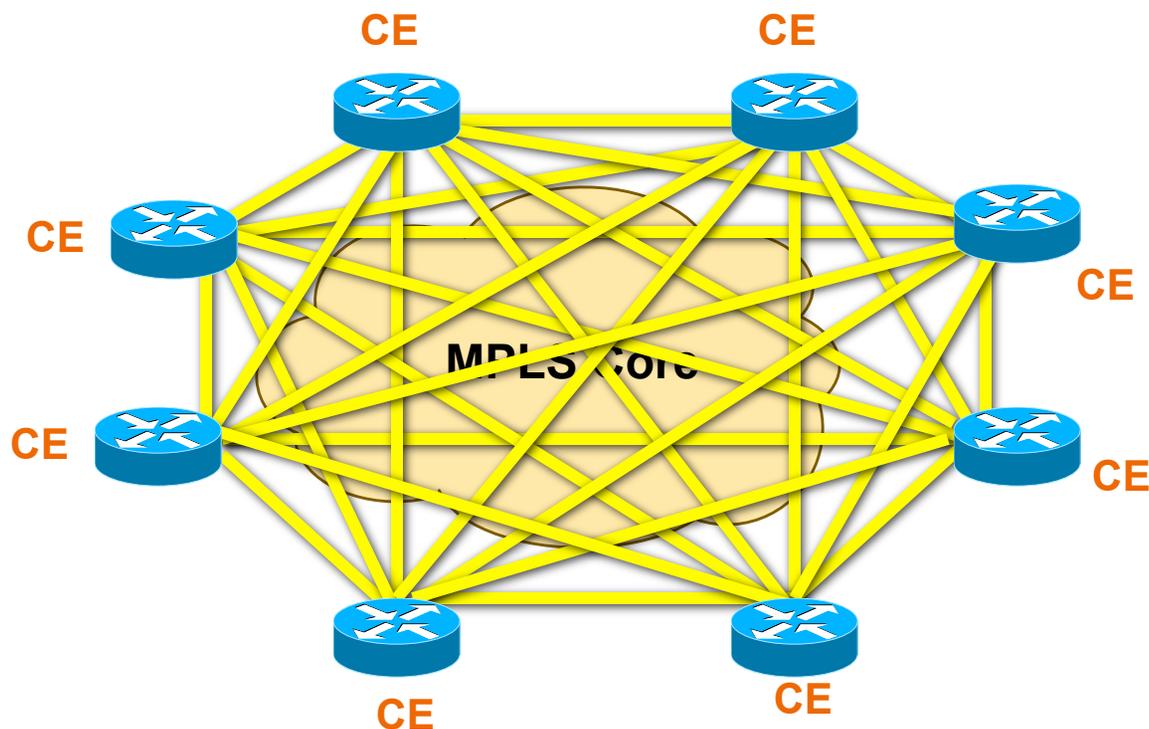
Challenges

- Core (P) routers have no knowledge of VPN source addresses
 - Multicast (PIM) uses the source address to determine the RPF interface
- VPN group address may overlap
- For optimal traffic forwarding multicast routing in the core is required
- Core stability must be assured
 - PIM is a soft-state protocol hence limited # of states can be supported

Why Multicast VPNs

- Workaround has been point-to-point GRE tunnels from CE to CE
- Not scalable with many CE routers
 - Traffic overhead
 - Administrational overhead

Why CE-CE tunneling doesn't scale



Point-to-point tunneling removes the benefits of multicast
Traffic has to be replicated by the CE router for each remote CE router
Traffic in the core will be multiplied by the number of CE routers
Non-congruent unicast/multicast due to RPF issues
A better solution is required

Agenda

- Why Multicast VPNs ?
- Multicast VPN Solution
- Cisco's Implementation
- Deployment Considerations

Multicast VPN solution

- Based on draft-rosen-vpn-mcast.txt
Section 2, “Multicast Domains”
- CEs maintain PIM adjacency with PEs only
- P-network does not hold (S, G)s for individual customers
- Customer multicast groups can overlap

Multicast Domains – Terminology

- **Multicast Domain (MD):** Set of VRFs that can send multicast to each other
- **Multicast VRF (MVRF):** A VRF that supports both unicast and multicast forwarding tables
- **Multicast Tunnel (MT):** Used to carry multicast C-packets among PE routers in a common MD

Multicast Domains – Operation

- An MVRF is assigned to a MD
- Per MD a P-group address is defined
- This P-group address must be unique
- C-packets are encapsulated on the PE routers and send on the MT as P-packets

Source address is address of MP-BGP update-source

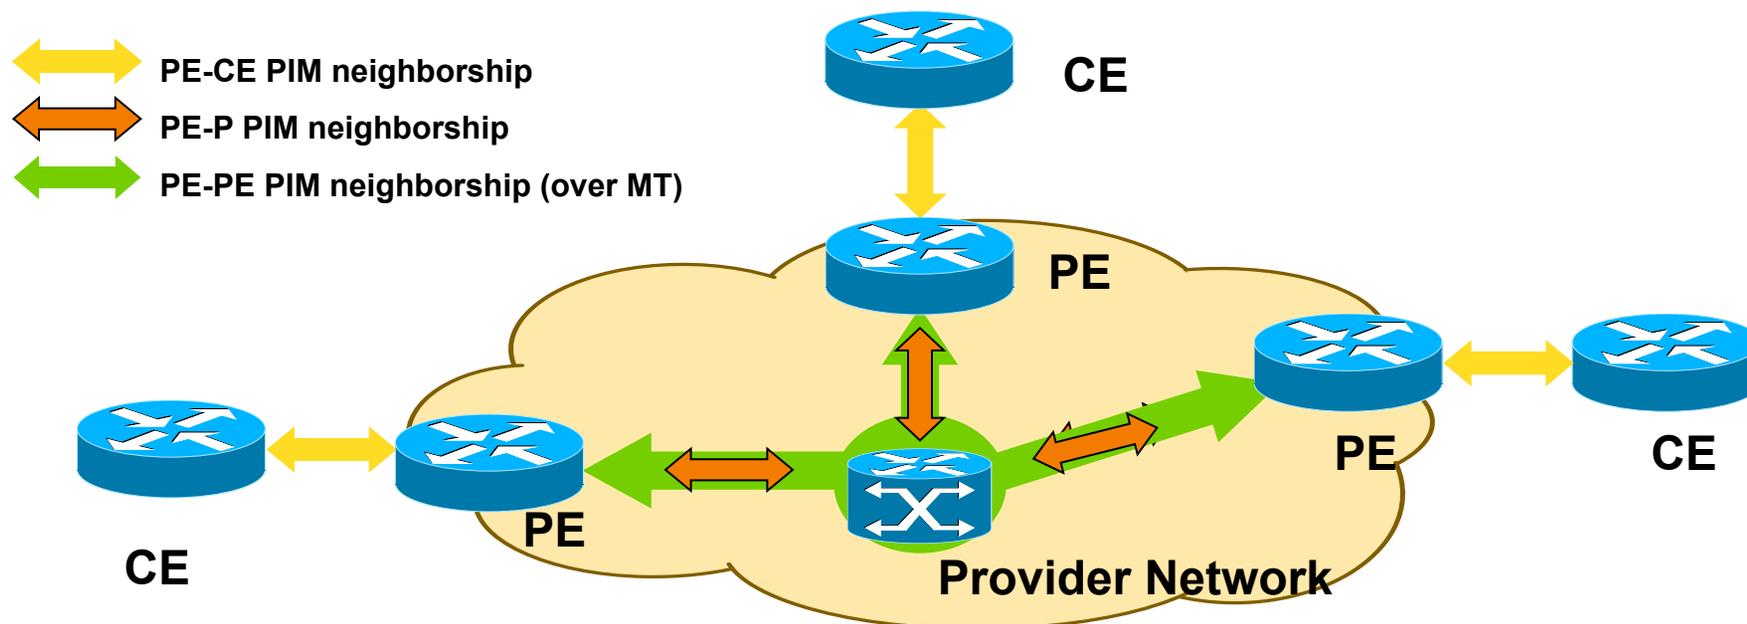
Destination address is P-group address

Encapsulation could be GRE/IPinIP/MPLS

Multicast Domains – Summary

- State in the core is minimized.
 - Optimally one state in the core for all multicast states within a VPN
 - Stability of core is provided
 - Service Provider has control of own destiny
- Traffic replication not optimal as PE routers without interested receivers still receive all multicast traffic per VPN

Multicast Domains – Example



Multicast tunnel is established between PE routers in Provider Network

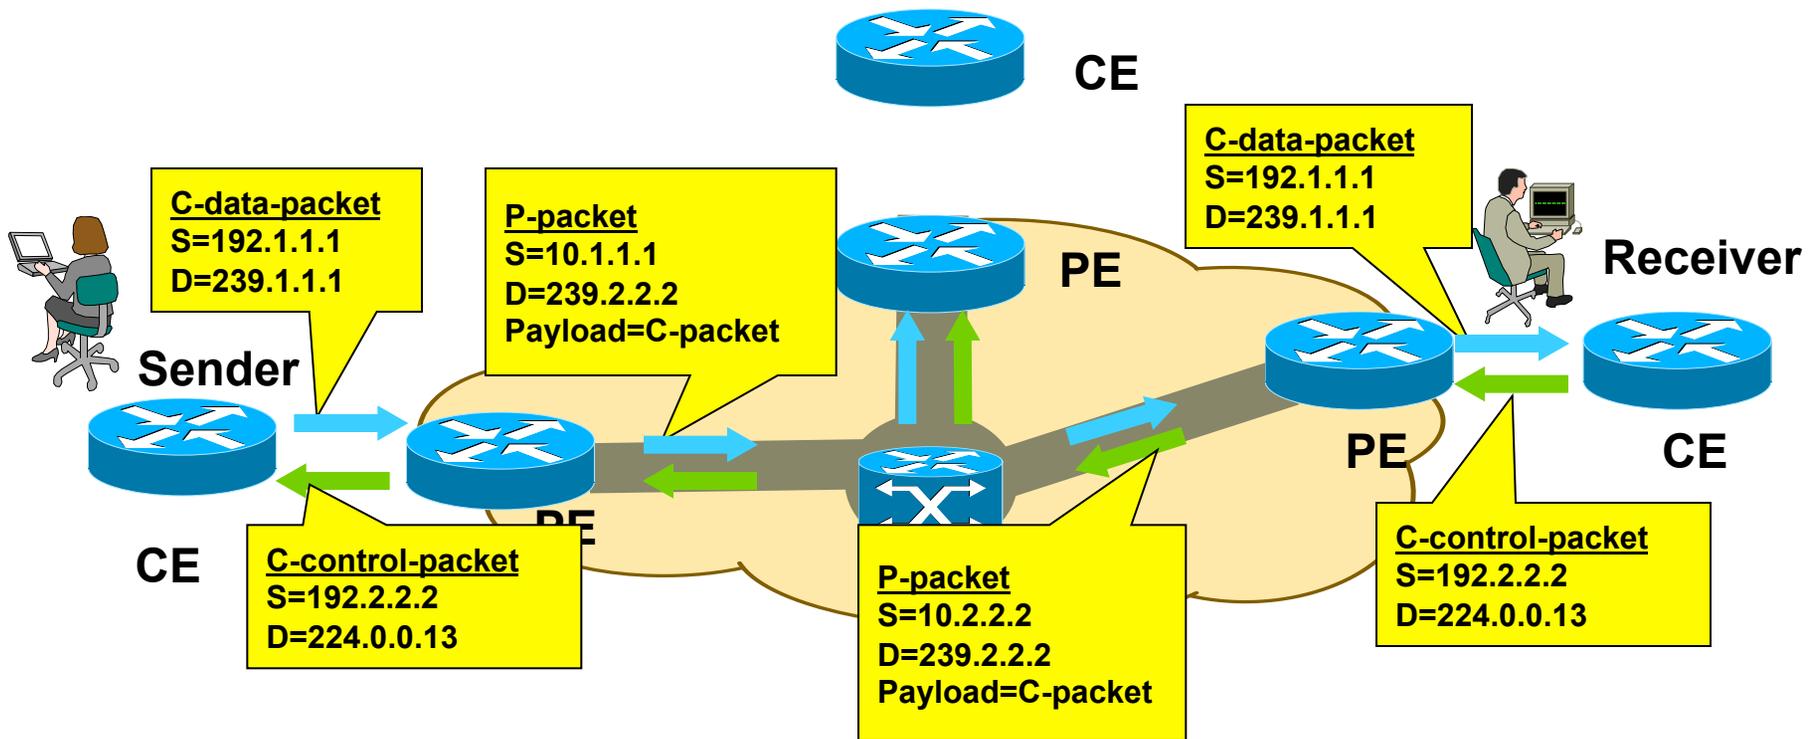
CE router forms PIM neighborship with VRF instance on PE router

PE routers form PIM neighborship with other PE routers over tunnel. This is a VRF specific neighborship

PE routers form PIM neighborship with P routers. This is a global neighborship

Multicast packets from CE routers will be forwarded over multicast tunnel

Multicast Domains – Example



Both customer control and data traffic are sent over the multicast tunnel

P routers only see P packets, so they won't build state for traffic inside the VPN

P packets will go to each PE router that is in the multicast domain

Agenda

- Why Multicast VPNs ?
- Multicast VPN Solution
- Cisco's Implementation
- Deployment Considerations

Cisco's Implementation

- **Overview**
- Configuration of Default-MDT
- MP-iBGP Update
- Multicast Tunnels
- MVRF
- Data flow over the Default-MDT
- Setup of Data-MDT
- Data flow over the Data-MDT

Cisco's Implementation – Overview

- Implementation is based on Multicast Domains.
 - Network stability is key in network operations
 - With enhancement optimal forwarding can be accomplished
 - Provider has control of own destiny

Cisco's Implementation – Overview

- Additional requirements

 - Service provider may have a preference for a PIM operating mode or already has deployed multicast in core

 - VPN customer may have a preference for a PIM operating mode or already has deployed multicast in the network

- Implementation must support all modes

- PIM mode used in the core and VPN should be unrelated

Cisco's Implementation – Overview

- Available PIM modes
 - PIM Bidirectional (PIM-BIDIR)
 - PIM Source Specific Multicast (PIM-SSM)
 - PIM Sparse-Mode (PIM-SM)
 - PIM Dense-Mode (PIM-DM)
- No sane service provider uses PIM-DM in the core, therefore it is not supported as protocol in the core. Neither is any other protocol which is not based on PIM (like dvmrp, mospf etc).

Cisco's Implementation – Overview

- New terminology

MDT: Multicast Distribution Tree – A logical distribution tree between PE routers built using IP Multicast in the core.

Determines mode of operation and addresses of multicast tunnel

Default-MDT: Tree that is always established between *ALL* PE routers belonging to the same VPN and on which control traffic and low rate data traffic is sent.

Data-MDT: Tree that is setup on demand for high rate data traffic.

Cisco's Implementation – Overview

- Additional challenges

Multicast packets in an MVRF are being sent and received on a multicast tunnel that is established between the PE routers.

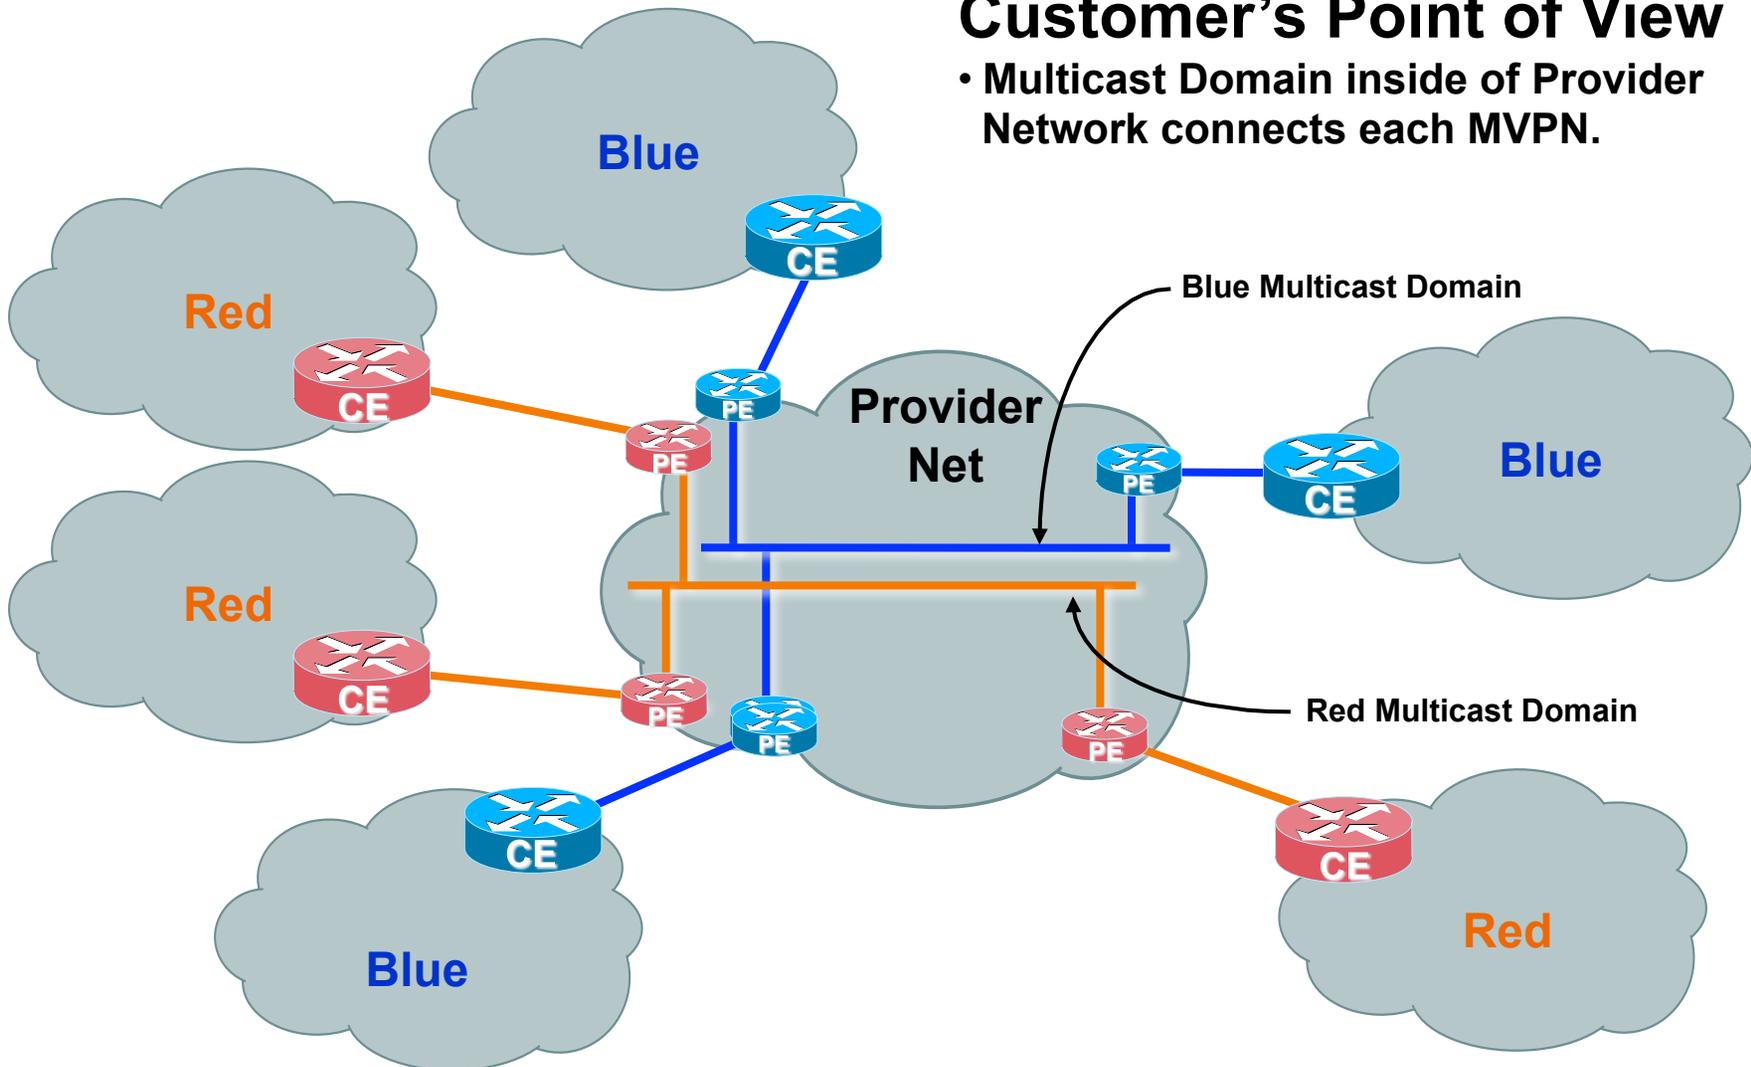
No unicast routing over this tunnel

RPF information in PIM is based on unicast routing information

Cisco's Implementation – Overview

Customer's Point of View

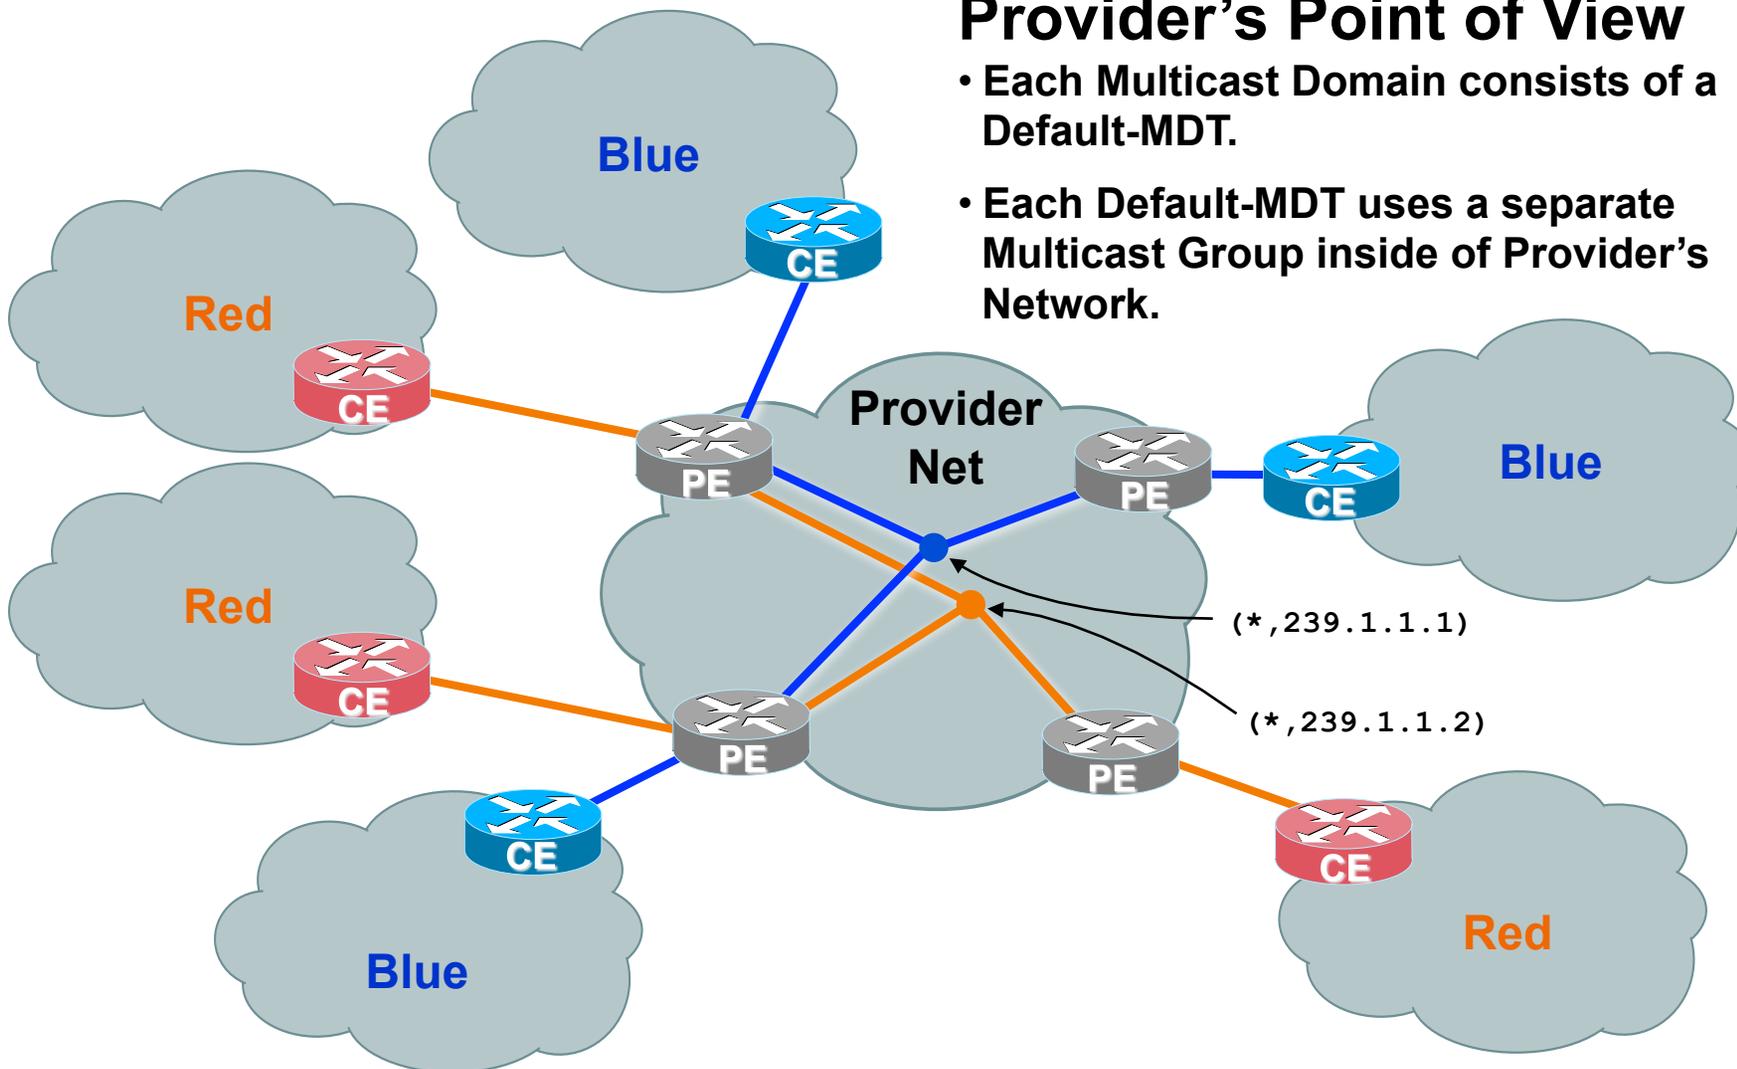
- Multicast Domain inside of Provider Network connects each MVPN.



Cisco's Implementation – Overview

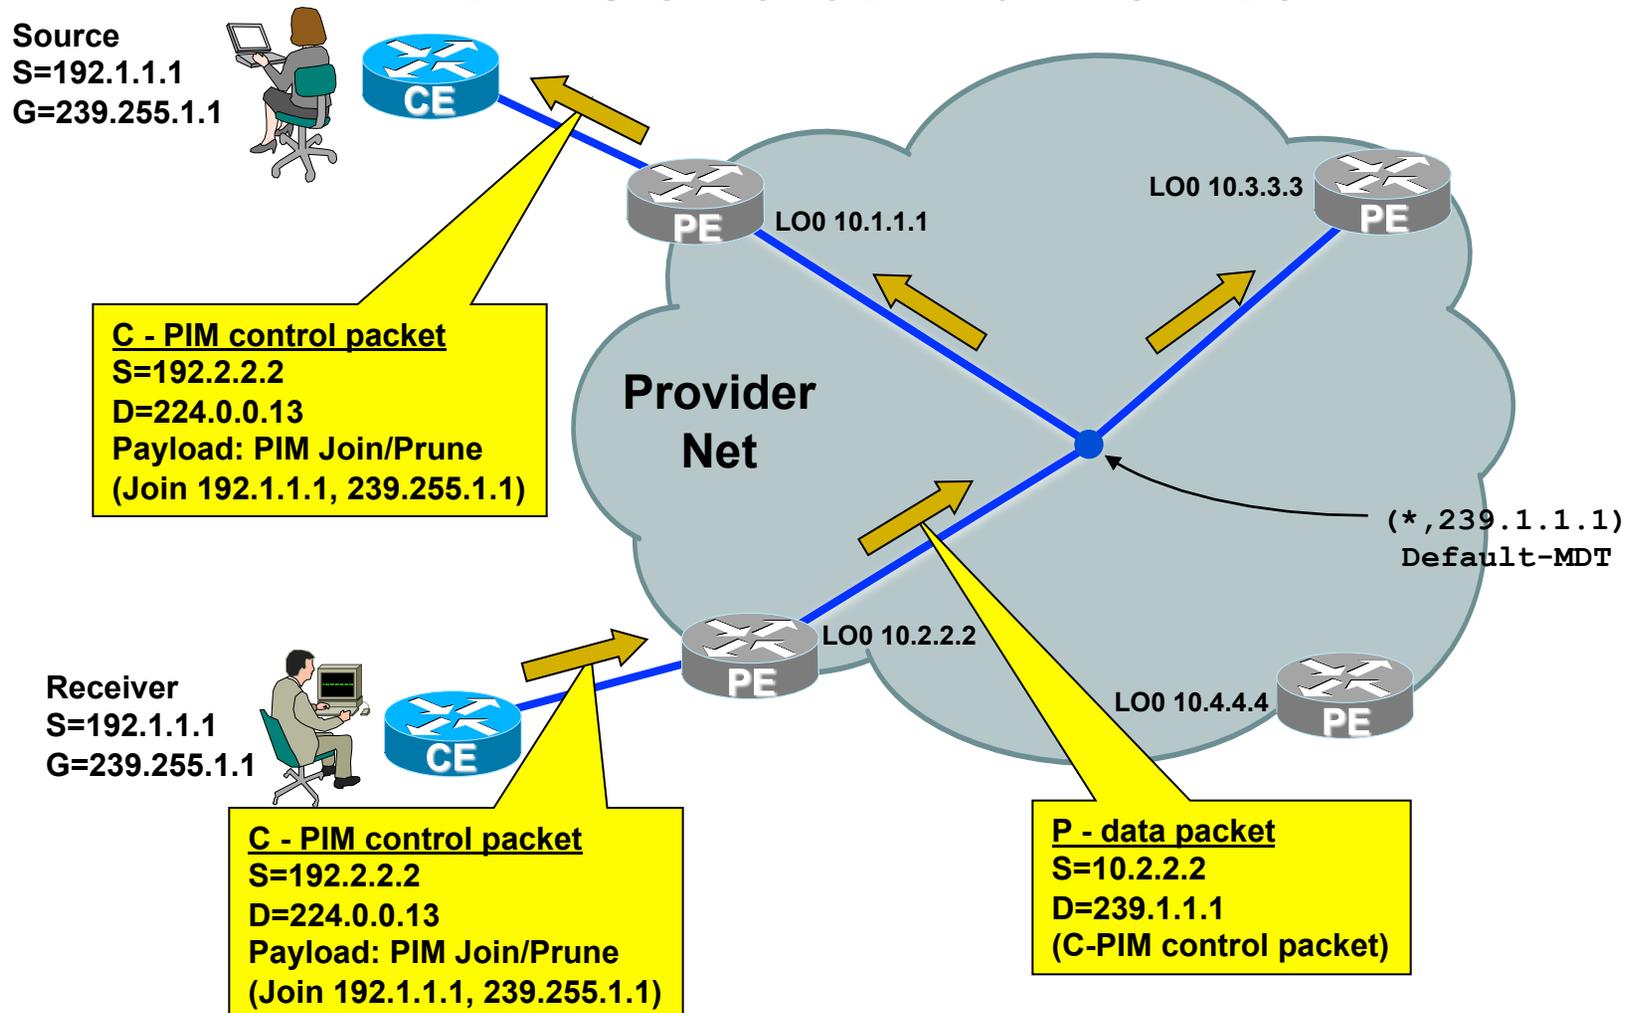
Provider's Point of View

- Each Multicast Domain consists of a Default-MDT.
- Each Default-MDT uses a separate Multicast Group inside of Provider's Network.



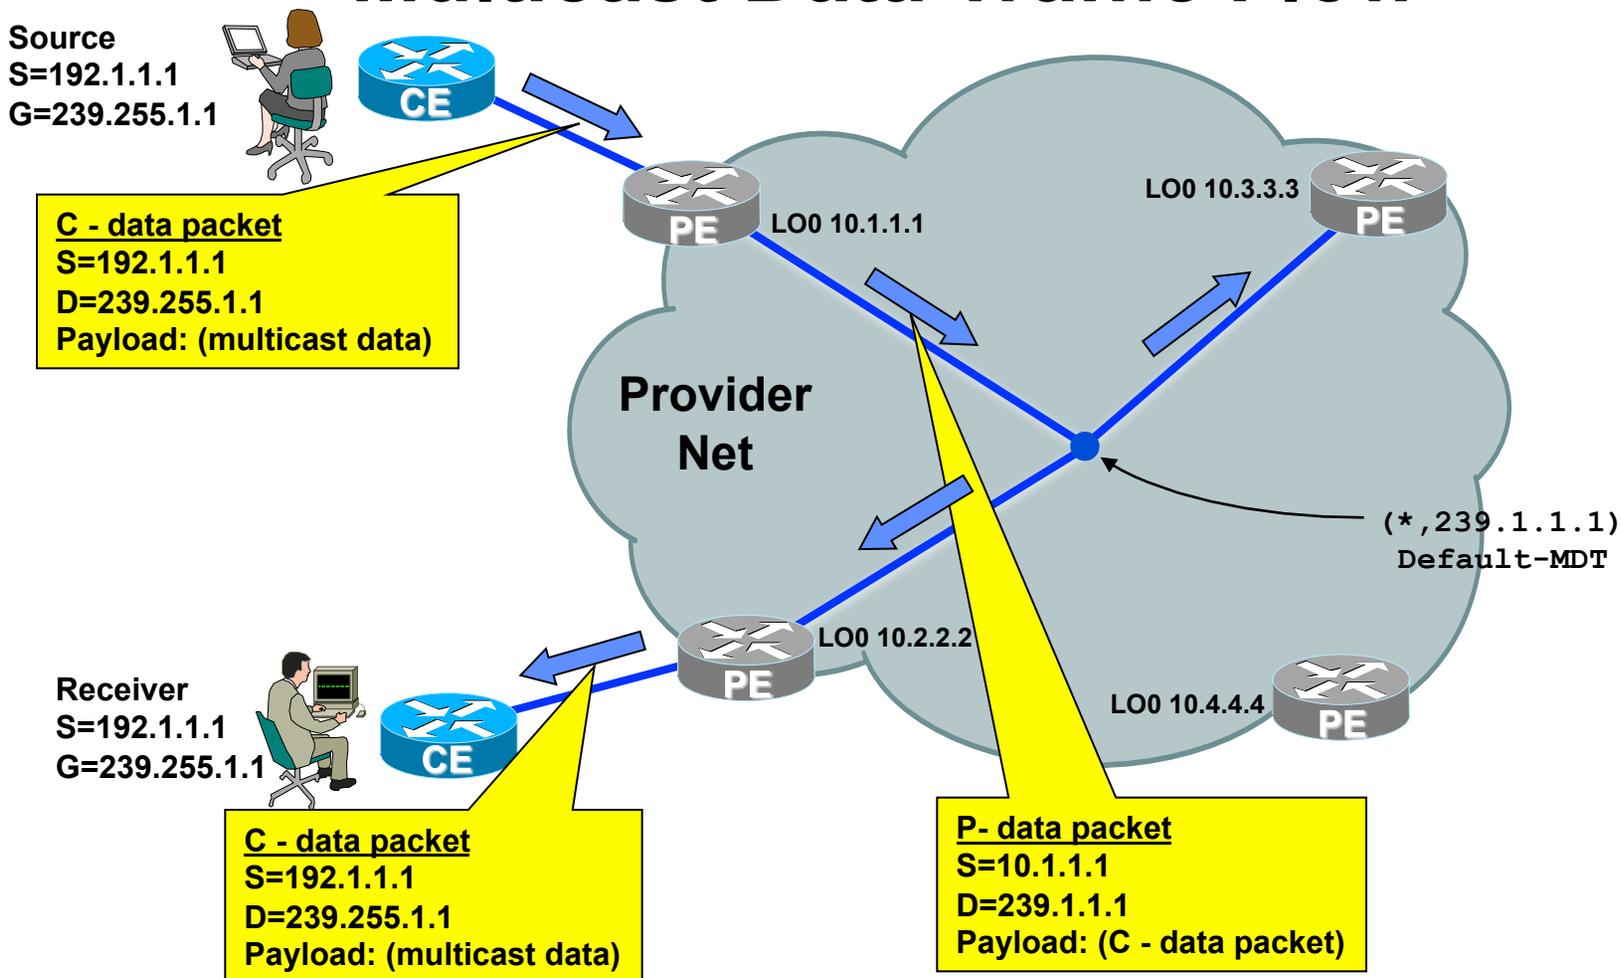
Default MDT – A Closer Look

PIM Control Traffic Flow



Default MDT – A Closer Look

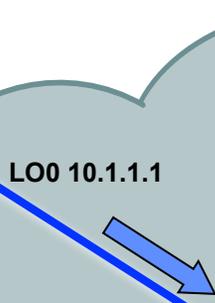
Multicast Data Traffic Flow



Default MDT – A Closer Look

Advantages and Disadvantages

Source
S=192.1.1.1
G=239.255.1.1



Provider
Net

Unwanted Data

(* ,239.1.1.1)
Default-MDT

Receiver
S=192.1.1.1
G=239.255.1.1

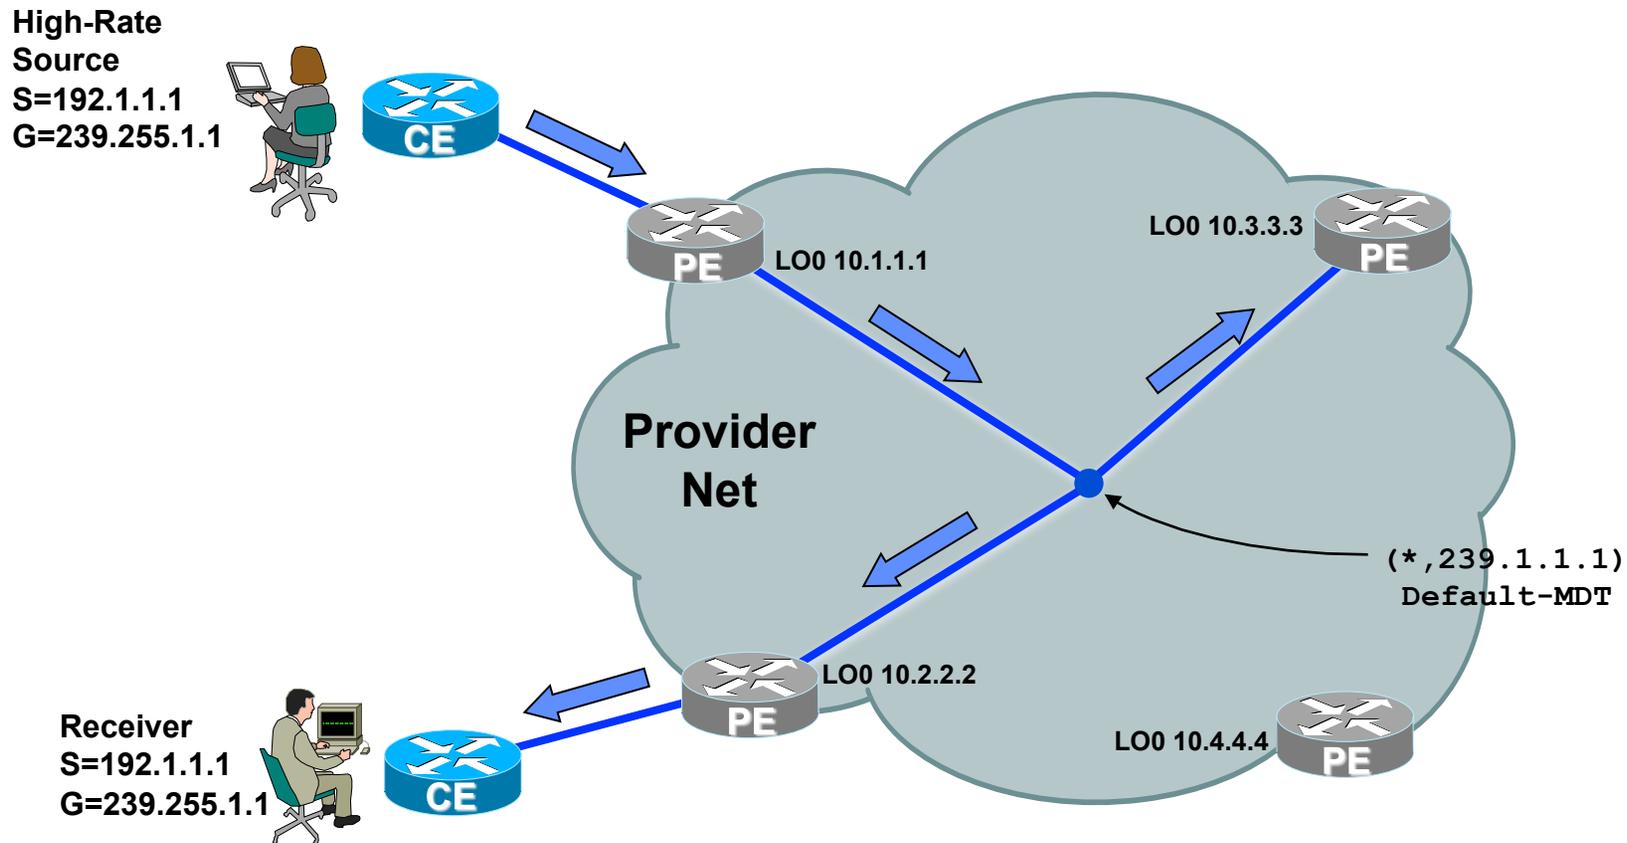


Advantage : Reduces multicast state in the P routers in the core.

Disadvantage : Can result in wasted bandwidth.

Solution : Use separate Data-MDTs for high rate sources.

Data MDTs – Concepts



- Traffic exceeds Data-MDT threshold configured on PE router.

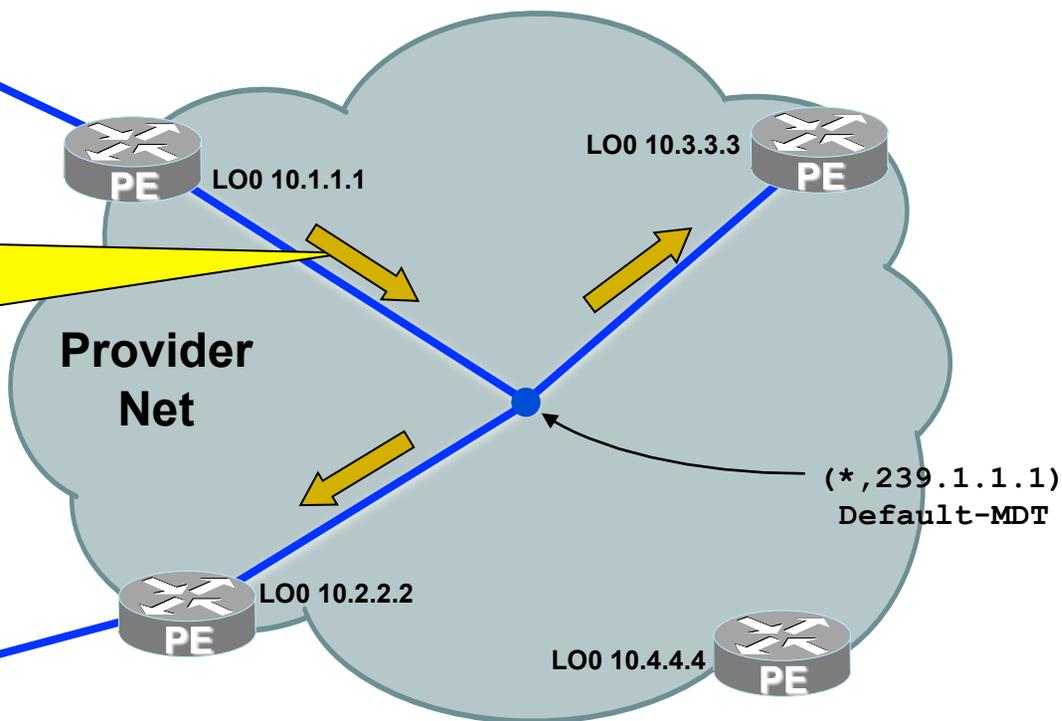
Data MDTs – Concepts

High-Rate
Source
S=192.1.1.1
G=239.1.1.1



P- control packet
S=10.1.1.1
D=224.0.0.13
Payload: (PIM MDT-Data)
S=192.1.1.1, G=239.1.1.1
MDT Group = 239.2.2.1

Receiver
S=192.1.1.1
G=239.1.1.1



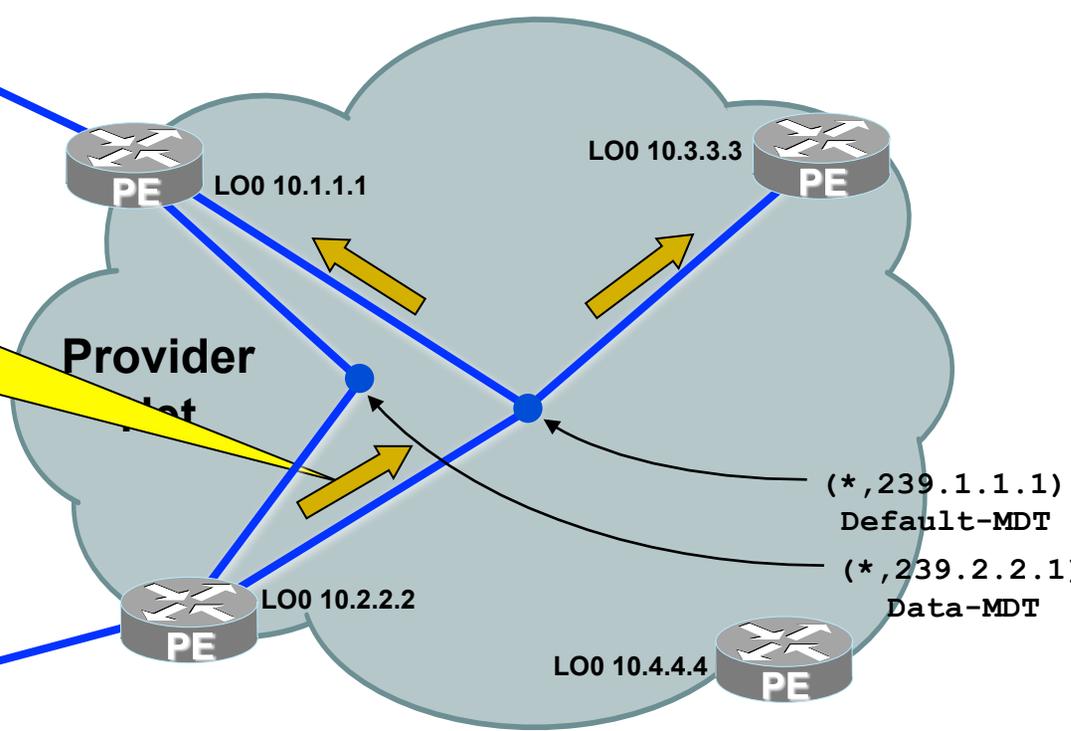
- PE router signals switch to Data-MDT using new group, 239.2.2.1

Data MDTs – Concepts

High-Rate
Source
S=192.1.1.1
G=239.1.1.1



P- control packet
S=10.2.2.2
D=224.0.0.13
Payload: (PIM Join)
S=10.1.1.1, G=239.2.2.1



Receiver
S=192.1.1.1
G=239.1.1.1

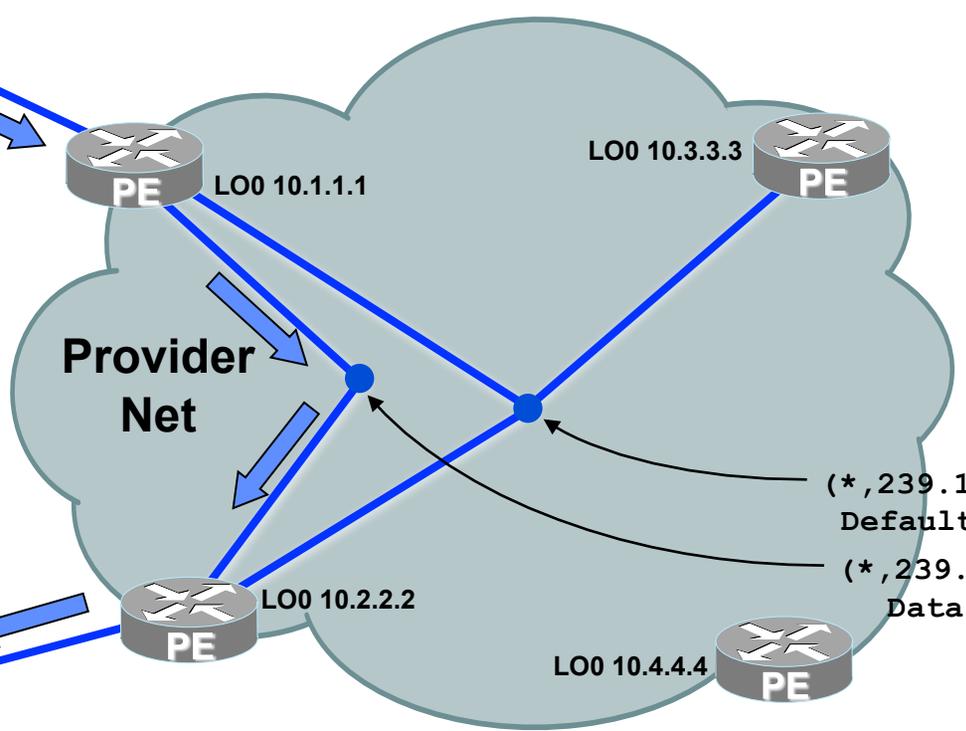


(*, 239.1.1.1)
Default-MDT
(*, 239.2.2.1)
Data-MDT

- PE routers with receivers sends Join to group 239.2.2.1.
- Data-MDT is built using group 239.2.2.1.

Data MDTs – Concepts

High-Rate
Source
S=192.1.1.1
G=239.1.1.1



Receiver
S=192.1.1.1
G=239.1.1.1



- High-rate data begins flowing via Data-MDT.
- Data only goes to PE routers that have receivers.

Data MDTs – Concepts

High-Rate
Source
S=192.1.1.1
G=239.1.1.1



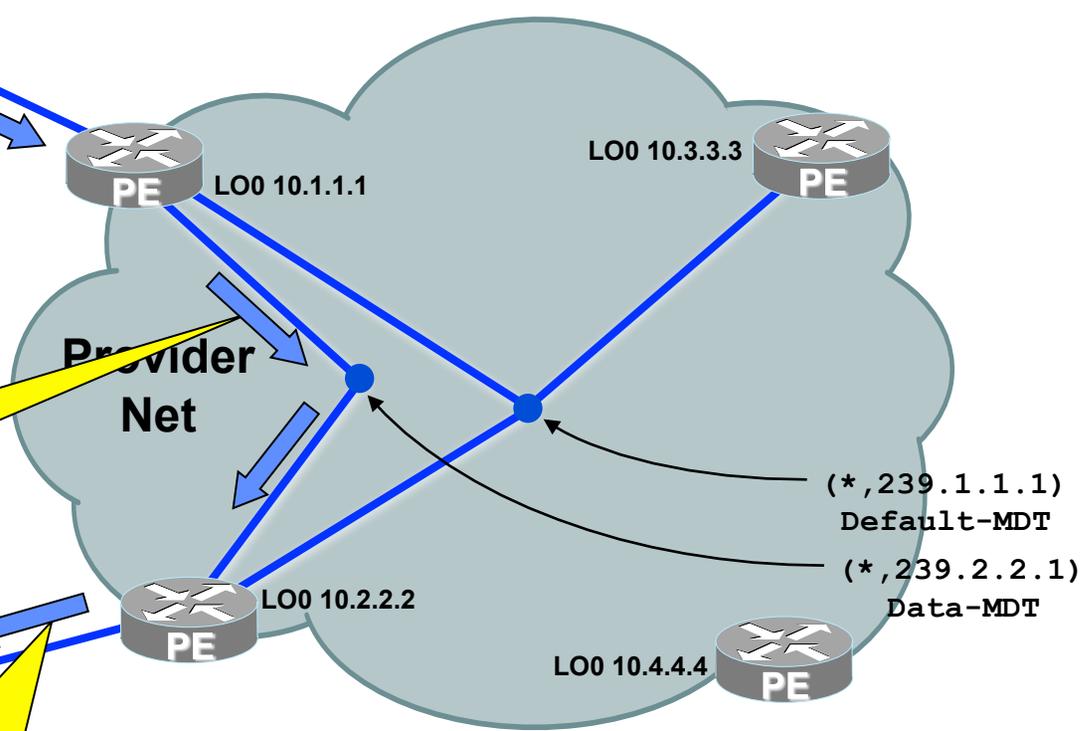
C - data packet
S=192.1.1.1
D=239.1.1.1
Payload: (multicast data)

P- data packet
S=10.1.1.1
D=239.2.2.1
Payload: (C - data packet)

Receiver
S=192.1.1.1
G=239.1.1.1



C - data packet
S=192.1.1.1
D=239.1.1.1
Payload: (multicast data)



(*, 239.1.1.1)
Default-MDT
(*, 239.2.2.1)
Data-MDT

Cisco's Implementation

- Overview
- Configuration of Default-MDT
- MP-iBGP Update
- Multicast Tunnels
- MVRF
- Data flow over the Default-MDT
- Setup of Data-MDT
- Data flow over the Data-MDT

Configuration of Default-MDT

- Default-MDT has two components

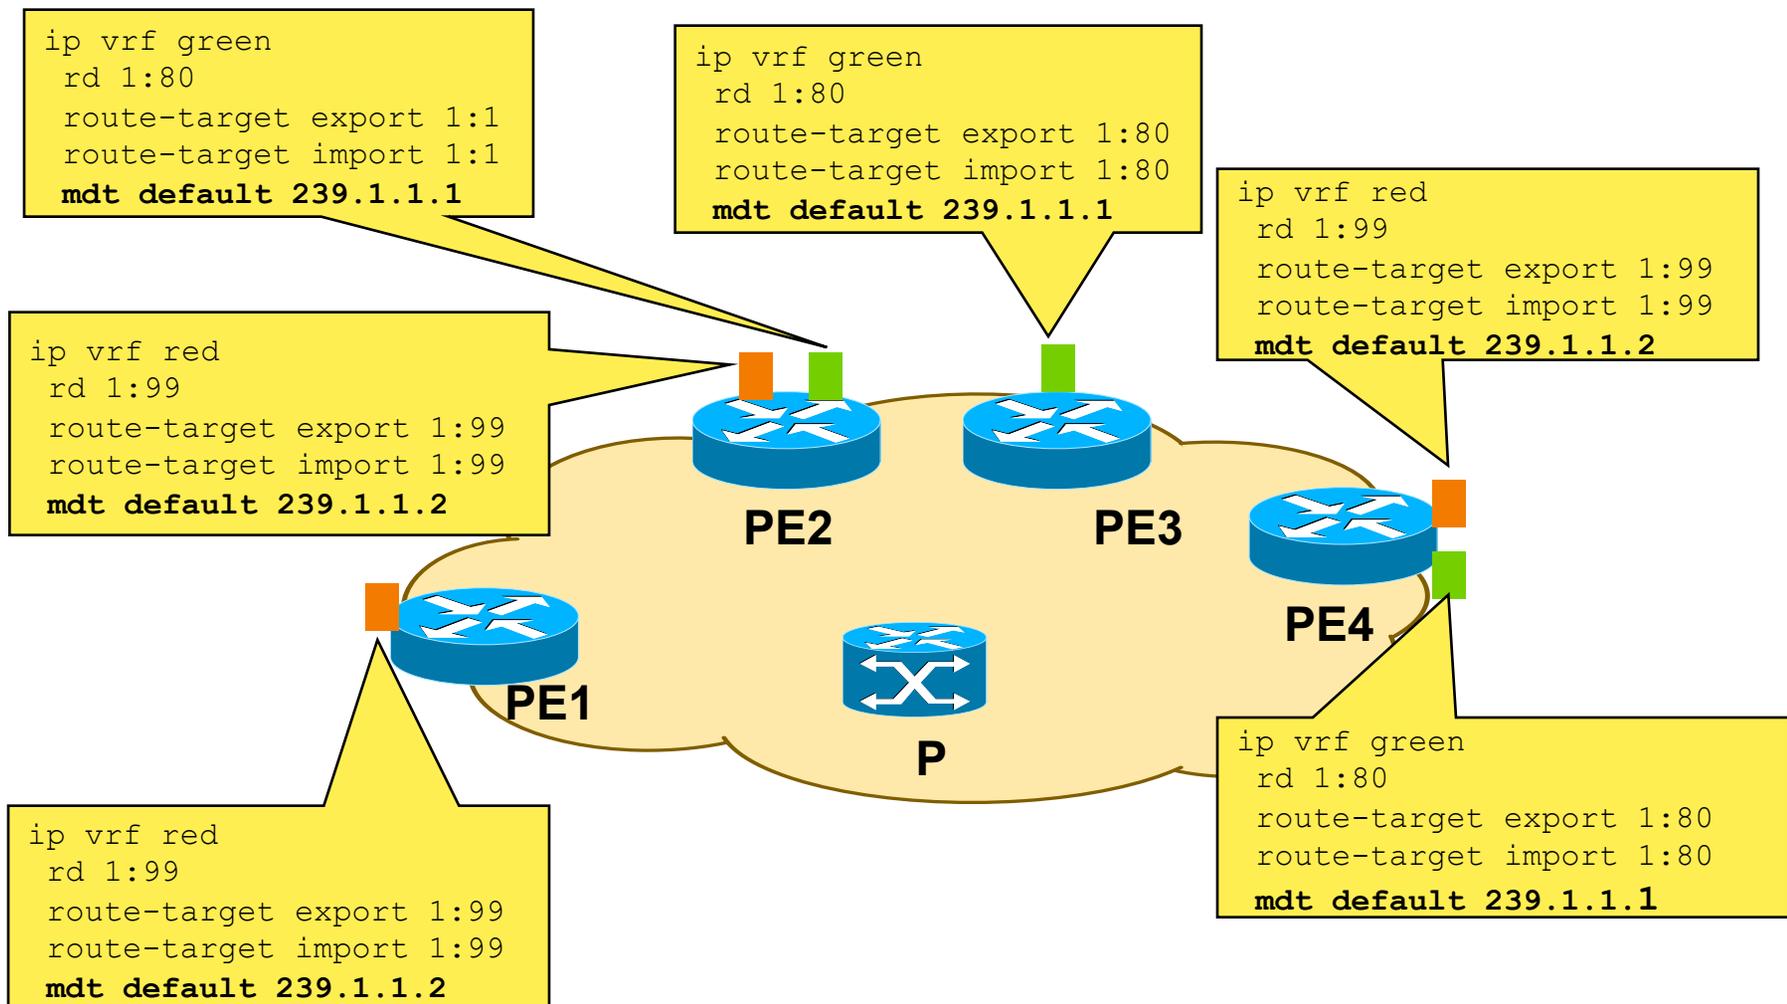
Group address: Configurable

Source address: Update-Source of MP-iBGP session.

Default-MDT Group Address

- Group address must be configured
 - Group address MUST be the same for all MVRFs belonging to the same MVPN
 - Group address MUST be different for all MVRFs belonging to different VPNs that are configured on the same PE router

Default-MDT Group Address – Example



Configuration of Default-MDT

- Default-MDT has two components

Group address: Configurable

Source address: Update-Source of MP-iBGP session

Default-MDT source address

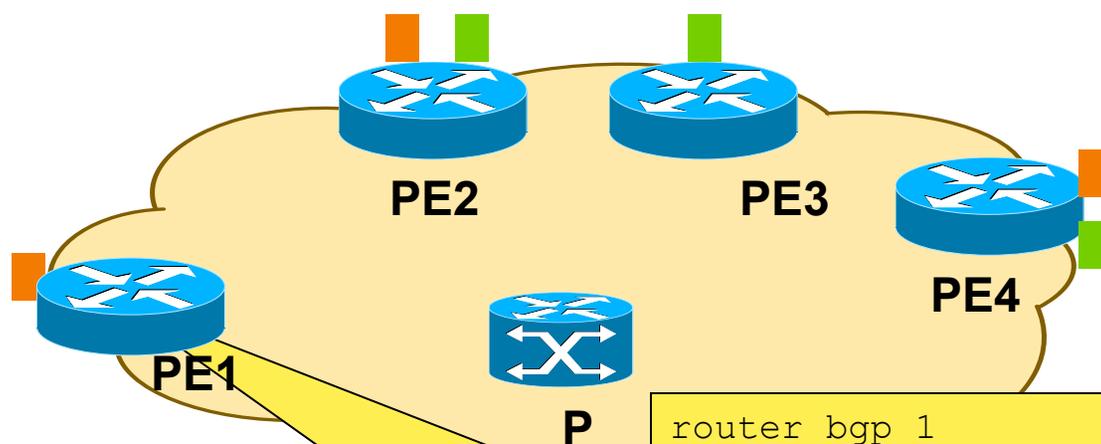
- Source address of the Default-MDT is the address used to source the MP-iBGP sessions with the other PE routers with MVRFs belonging to the same VPN.

Common practice is to use the 'update-source' keyword in the BGP config and choose a loopback interface.

This is a requirement for MVPN

- The update-source interface **MUST** be the same for all MP-iBGP sessions configured on the router for the Default-MDT to be setup properly
- No additional commands are needed

Default-MDT Source Address – Example



```
router bgp 1
  neighbor <PE2> remote-as 1
  neighbor <PE2> update-source Loopback0
  neighbor <PE4> remote-as 1
  neighbor <PE4> update-source Loopback0
  !
  address-family ipv4 vrf RED
  redistribute <IGP>
  exit-address-family
  !
  address-family vpnv4
  neighbor <PE2> activate
  neighbor <PE2> send-community-extended
  neighbor <PE4> activate
  neighbor <PE4> send-community-extended
  exit-address-family
```

Cisco's Implementation

- Overview
- Configuration of Default-MDT
- **MP-iBGP Update**
- Multicast Tunnels
- MVRF
- Data flow over the Default-MDT
- Setup of Data-MDT
- Data flow over the Data-MDT

MP-iBGP Update

- Key is a BGP neighbor relationship with remote PE routers
- PE routers tell other routers that they participate in an MVPN.

This triggers the setup of the Default-MDT

New RD type in VPN-IPv4 address

New extended community

- BGP next hop is used to determine the RPF information

MP-iBGP Update – Additions

- (M)VPN-IPv4 address (12 bytes)

Route Distinguisher - 8 bytes

type-field: 2 bytes

value-field: 6 bytes

New type : 0x02 (Multicast-VPN)

Value field (AS format must be used):

2 bytes ASN

4 bytes assigned number

IPv4 address - 4 bytes

VRF next hop

Used for SSM in the Core)

MP-iBGP Update – Additions

Extended community attribute - 8 bytes

Type Field : 2 bytes

Value Field : 6 bytes

New type: 0x09 (AS format)

Value Field :

2 bytes ASN

4 bytes assigned number (MDT Group address)

OR (currently not supported)

New type: 0x0109 (Address format)

Value Field :

4 bytes IPv4 address (MDT Group address)

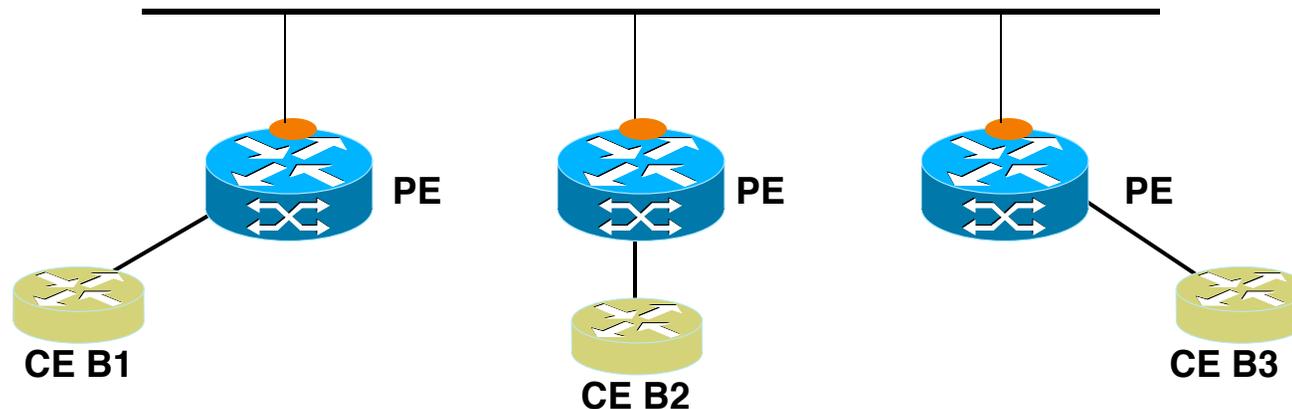
2 bytes assigned number

Cisco's Implementation

- Configuration of Default-MDT
- MP-iBGP Update
- Multicast Tunnels
- MVRF
- Data flow over the Default-MDT
- Setup of Data-MDT
- Data flow over the Data-MDT

Multicast Tunnel Interface (MTI)

MTI 



- Representation of access to the Multicast Domain from an mVRF is by way of the MTI
- MTI is treated as a LAN interface in mVRF
 - Appears as a “TunnelX” interface in the mVRF
 - Multiple PEs seen over same MTI for that mVPN

Multicast Tunnel Interface (MTI)

- MTI is not configurable and takes its properties from interface used for BGP peering (e.g. Loopback 0)
- PIM is always enabled
- No unicast runs over Multicast tunnel interface
 - This affects RPF check
- Traffic forwarded to the MTI is encapsulated
 - At present GRE is the only method available in all switching modes
 - Once a packet is forwarded to the MTI it passes into the global multicast of the SP MTI automatically created when Default-MDT configured

Multicast Tunnel Setup

- Components for Multicast Tunnel Setup

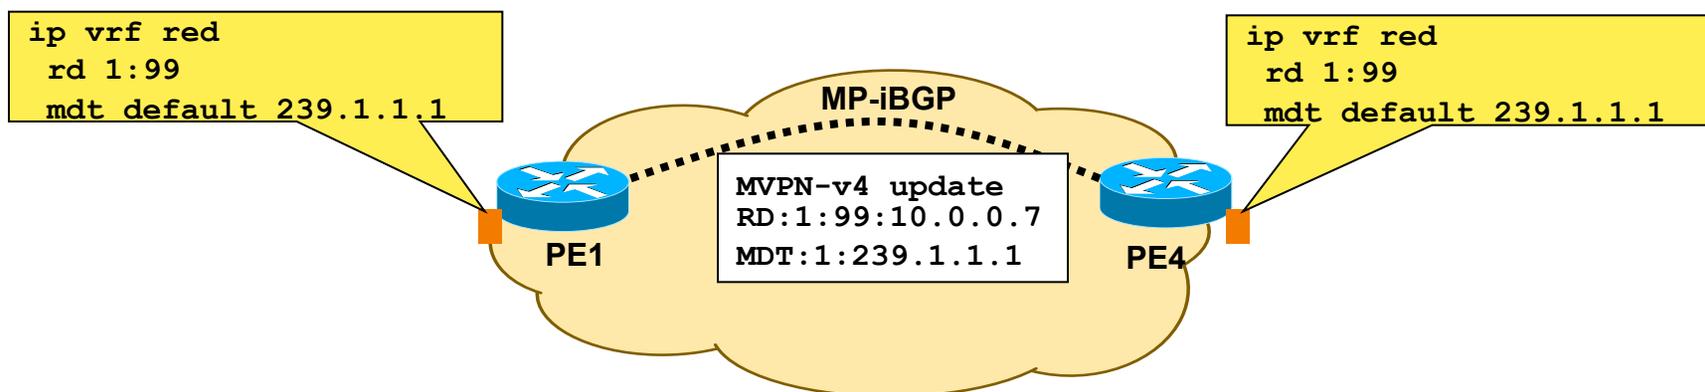
- PIM enabled in the core on all backbone interfaces

- Shared trees or source trees may be used. Each has different scaling characteristics. Later more on this.

- MDT Group address configured

- MP-iBGP session established

Multicast Tunnel Setup – Example



1. PE receives MVPN-IPv4 update
2. PE creates Multicast Tunnel Interface if it has not been created before
3. PE will join the root of the tree with the group address configured in the vrf.
4. The root is the RP for shared trees, and PE's loopback when source trees are used

Multicast Tunnel Setup

BGP received the MVPN-v4 update from the remote peer

```
BGP(2): 10.0.0.7 rcvd UPDATE w/ attr: nexthop 10.0.0.7, origin ?, localpref 100, extended
community RT:1:99 MDT:1:239.233.0.1
BGP(2): 10.0.0.7 rcvd 2:1:99:10.0.0.7/32
```

BGP informs PIM of a new remote MVRF

```
BGP: Inform multicast system about mdt 239.233.0.1 from router-id 10.0.0.7 with next-hop
10.0.0.7 : present
```

PIM receives BGP update, as this is the first peer a multicast tunnel interface is created

```
PIM(1): Received BGP MDT update (10.0.0.7,239.233.0.1) next-hop: 10.0.0.7
PIM(1): Created multicast tunnel interface Tunnel0
```

In the GLOBAL multicast routing table an entry for the MDT-Group address is created

```
MRT(0): Update (*, 239.233.0.1), RPF Null, PC 0x60736800
MRT(0): Add/Update Loopback0/239.233.0.1 to the olist of (*, 239.233.0.1), Forward state
```

PIM generates a join/prune message to setup the MDT

```
PIM(0): Building triggered (*,G) Join / (S,G,RP-bit) Prune message for 239.233.0.1
PIM(0): v2, for RP, Join-list: 172.16.0.21/32, RP-bit, WC-bit, S-bit
PIM(0): Send v2 triggered Join/Prune to RP via 172.16.203.1 (Ethernet0/1)
```

PIM sends hello's on the MT to form a PIM neighborship with the remote PE routers over the tunnel **This neighbor relationship will be visible in the MVRF only**

```
PIM(1): Send v2 Hello on Tunnel0
```

Cisco's Implementation

- Configuration of Default-MDT
- MP-iBGP Update
- Multicast Tunnels
- **MVRF**
- Data flow over the Default-MDT
- Setup of Data-MDT
- Data flow over the Data-MDT

Multicast VRF

- A MVRF is created when multicast routing is enabled for that VRF
- Multicast protocols like IGMP and PIM are configured and operate in the context of an MVRF.
- MVRF only contains the multicast routing information for the VRFs that make up a Multicast Domain.

Multicast VRF – Interface types

- We define three types of interfaces
 - (PNI) Provider Network Interface
 - (CNI) Customer Network Interface
 - (MTI) Multicast Tunnel Interface

Multicast VRF – Interface Types

- Provider Network Interface

 - A PNI is connected to a P router

 - Packets sent and received from a PNI are routed using the global table (Global MVRF)

- Customer Network Interface

 - A CNI is connected to a CE router

 - Packets sent and received from a CNI are routed using the MVRF referenced by the interface

Multicast VRF – Interface Types

- Multicast Tunnel Interface

MTI is a virtual interface dynamically created in (and owned by) the MVRF for each MD the MVRF is assigned to.

- MTI is created when:

MDT Group address configured in VRF

Multicast routing is enabled for VRF

Multicast VRF – Interface Types

- PE routers exchange PIM control messages (via the MTI) and establish PIM neighbor relationships
 - VPN MVRF (*,G) and (S,G) entries have the MTI as the Incoming Interface or in the OIL
- MTI creates (*,G) or (S,G) in Global MVRF
 - A flag (Z) is set indicating local PE is leaf of the MT
 - The OIL points to the MVRFs associated with the MD

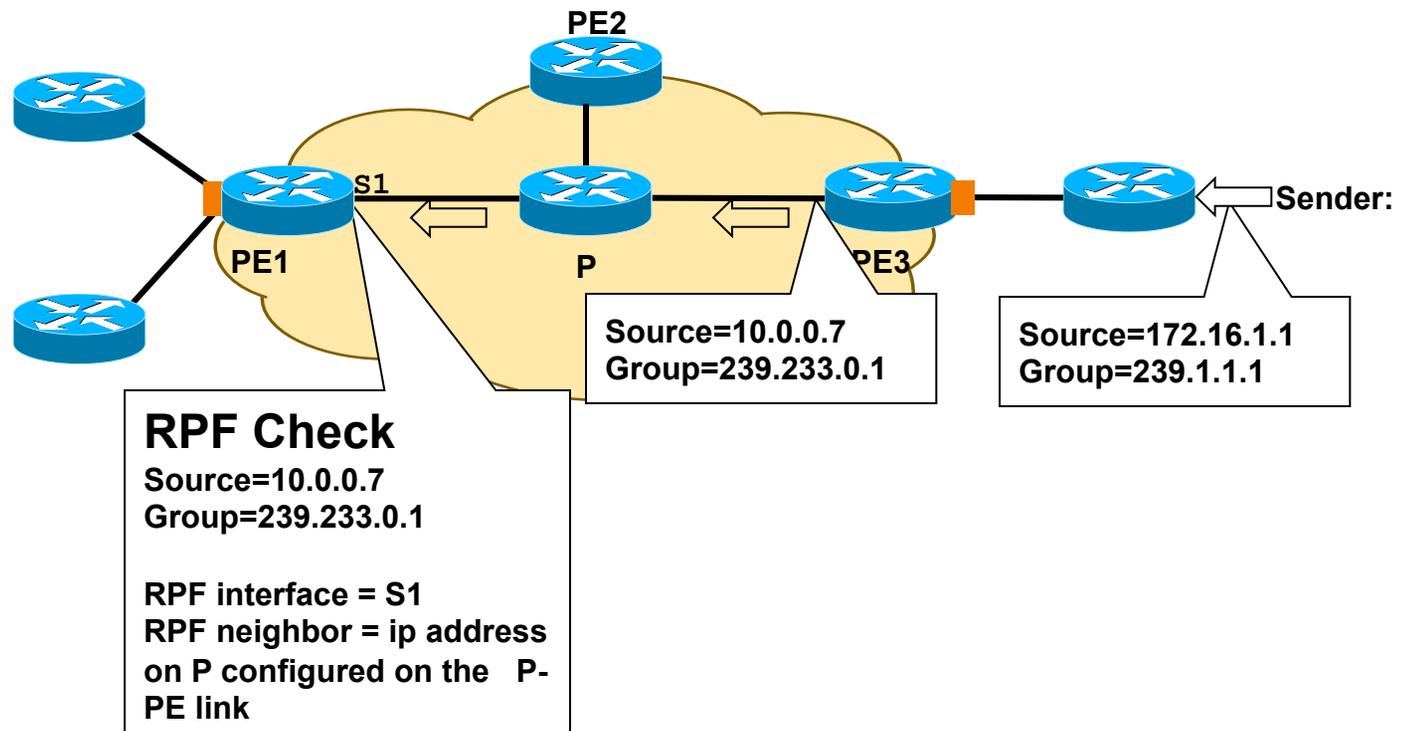
Multicast VRF

- Determining RPF information
- PIM in the MVRF

Determining the RPF information

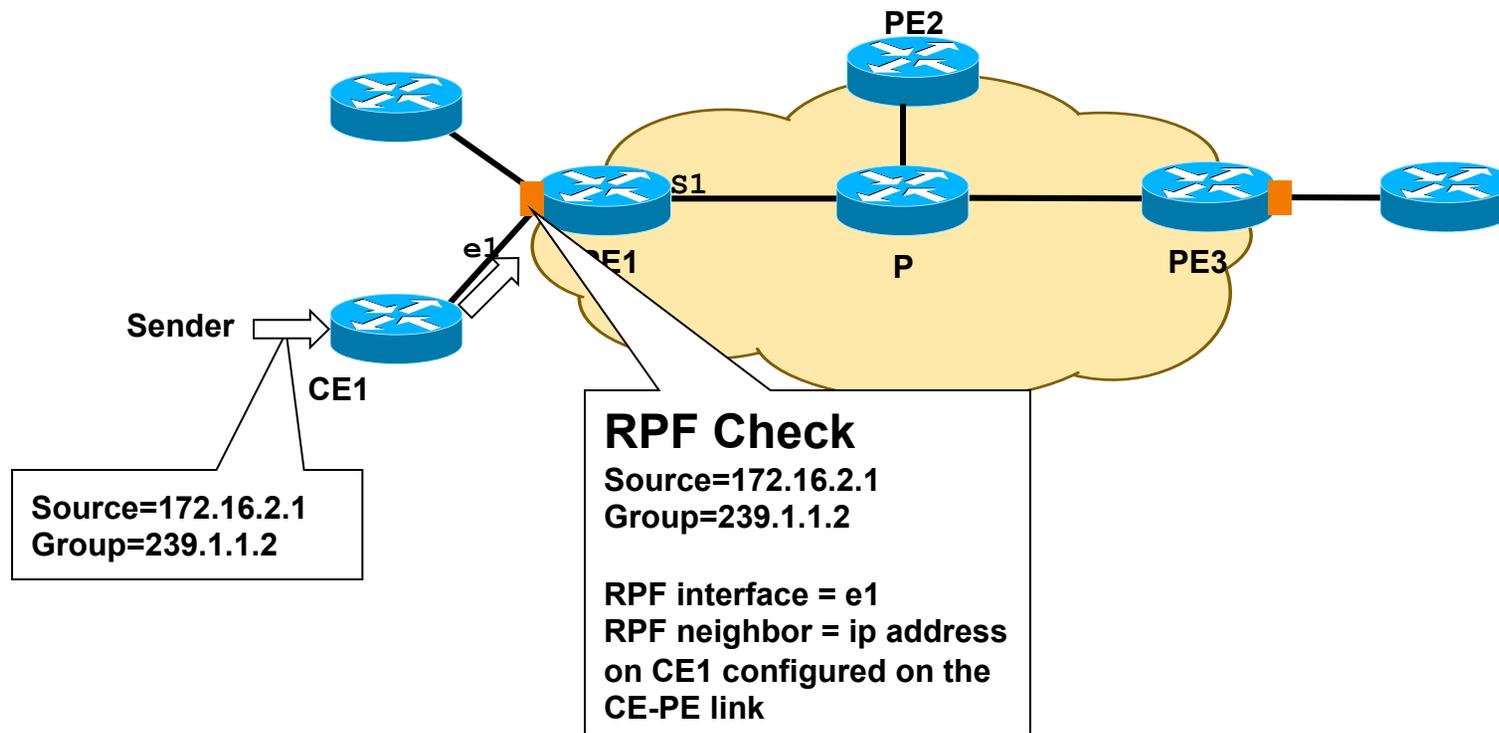
- PIM uses the unicast routing table to determine the RPF information:
 - RPF interface (used to RPF-check packets)
 - RPF neighbor (indicates which upstream PIM neighbor to send PIM join/prune messages)
- Possibilities:
 - RPF lookup in the Global MVRF
 - RPF lookup in an VPN MVRF where:
 - RPF interface is in the same MVRF
 - RPF interface is in the global table

RPF lookup in the Global MVRF



- Procedures for RPF check in global MVRF is the same as on a router that doesn't support MVPN

RPF lookup in an MVRF

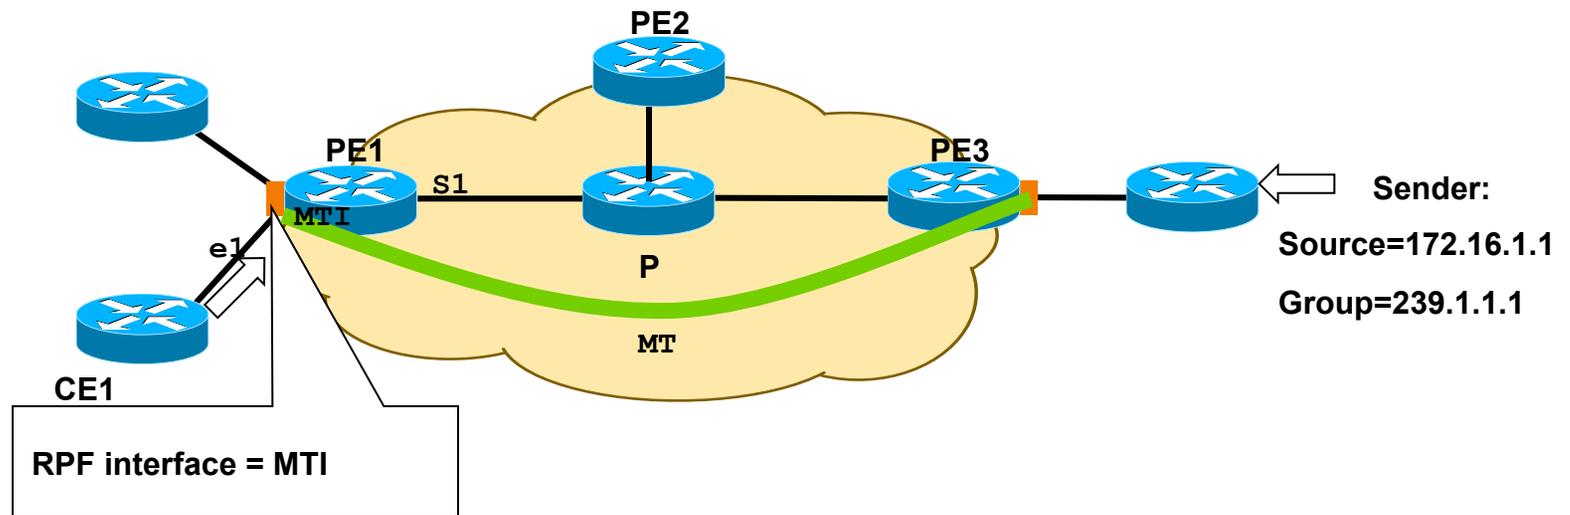


- If lookup returns interface in same MVRF:

RPF interface: The interface returned by the lookup

RPF neighbor: Same rules apply as on a router that doesn't support MVPN to find the RPF neighbor

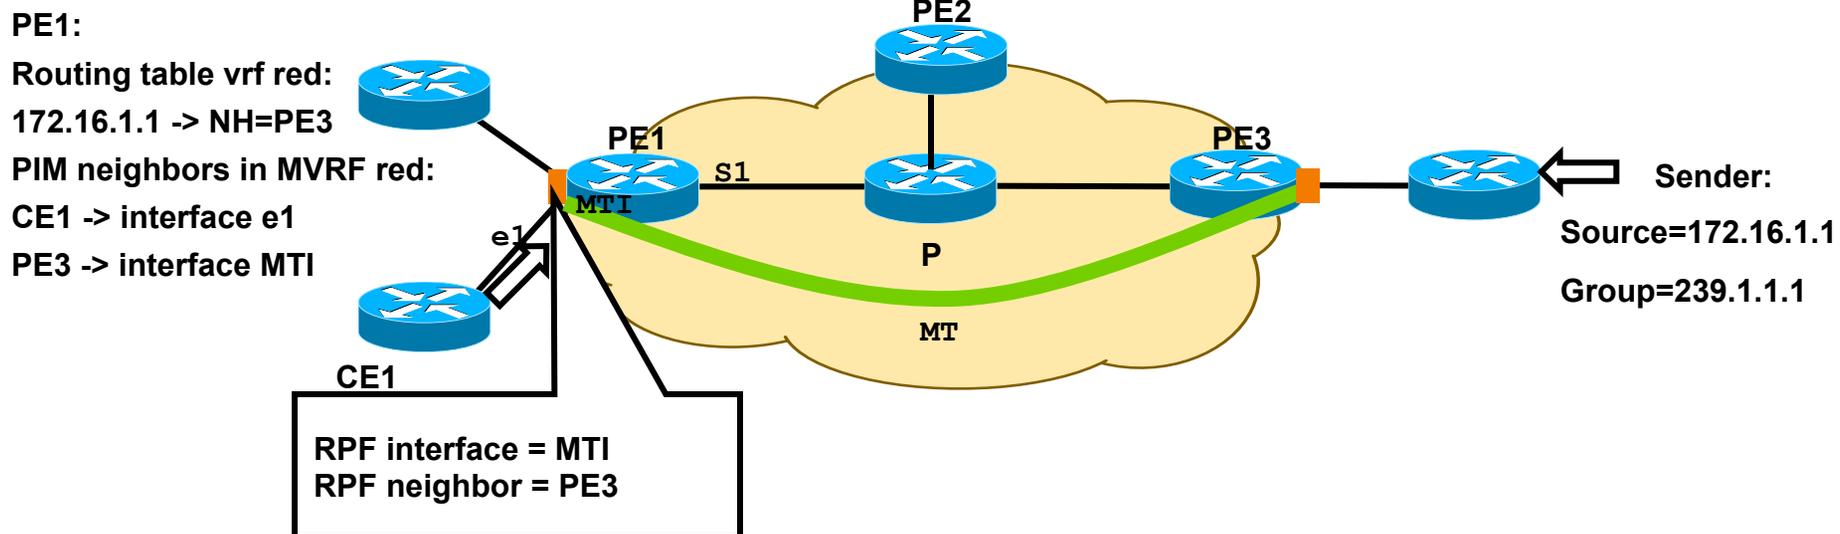
RPF lookup in an MVRF



- If lookup returns interface in Global VRF:
 - RPF-Interface: RPF interface is the MTI created for the MVRF

Note: More than one MTI per MVRF is not yet supported

RPF lookup in an MVRF



- If lookup returns interface in Global VRF:
 - RPF-Neighbor is found when the following two conditions are satisfied:
 - The BGP next-hop to the source exists
 - Next-hop is PIM neighbor for the MD

Multicast VRF

- Determining RPF information
- PIM in the MVRF

PIM in the MVRF

- PIM in the Global MVRF

- Sets up multicast tunnels between PE routers.

- This is transparent for the P routers

- Used for regular multicast services (non-VPN)

PIM in the MVRF

- PIM in a VPN MVRF

Establishes PIM neighbors with remote PE routers over the MTI.

Useful for finding RPF-neighbor and discovery of capabilities of remote router (like BIDIR etc).

Establishes PIM neighbors with CE routers

Creates MVRF specific multicast forwarding entries

Discovers VPN specific RP information

Auto-RP

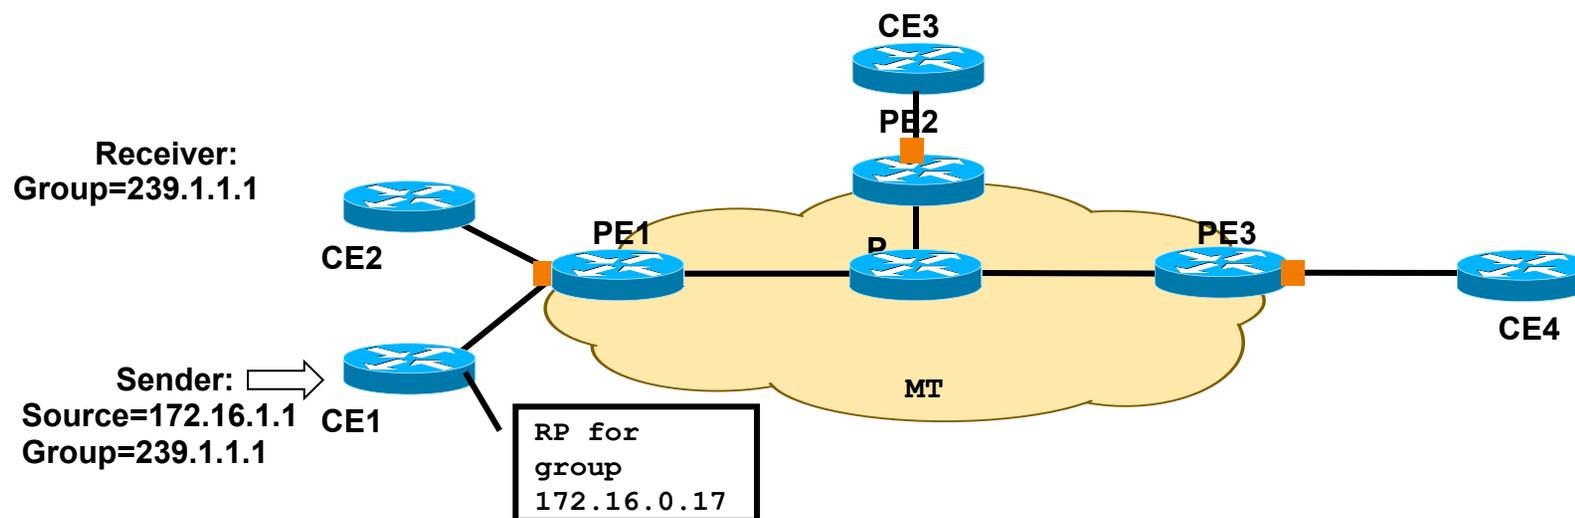
BSR

Static

Cisco's Implementation

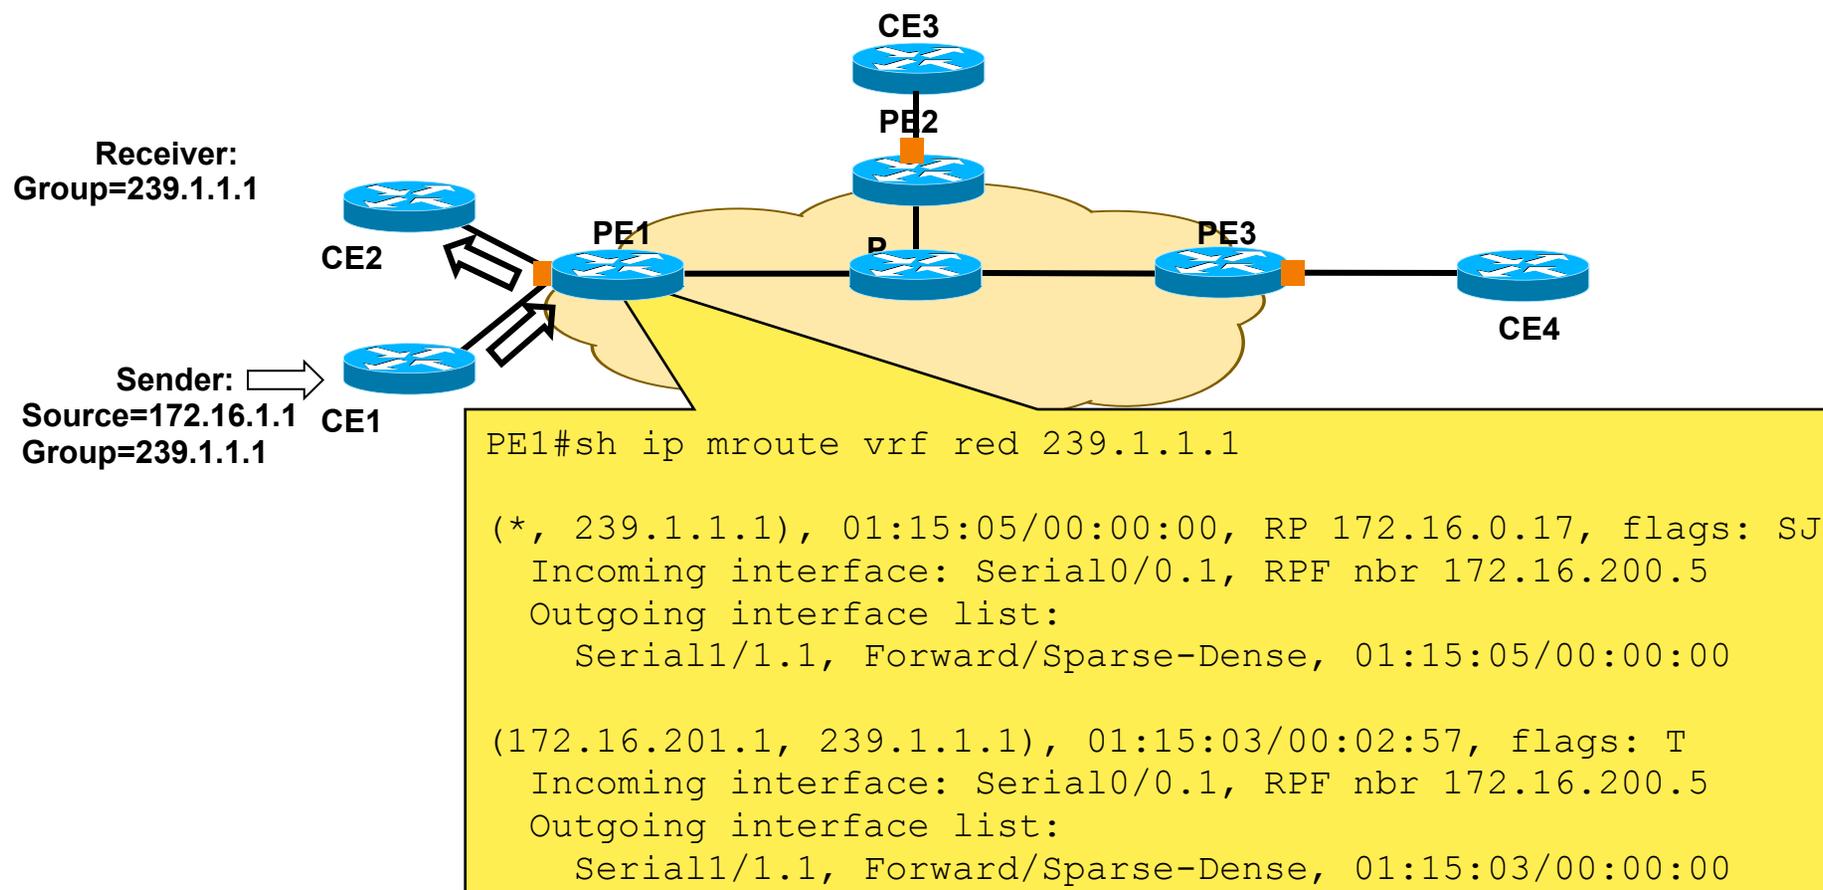
- Configuration of Default-MDT
- MP-iBGP Update
- Multicast Tunnels
- MVRF
- Data flow over the Default-MDT
- Setup of Data-MDT
- Data flow over the Data-MDT

Data flow over the Default-MDT



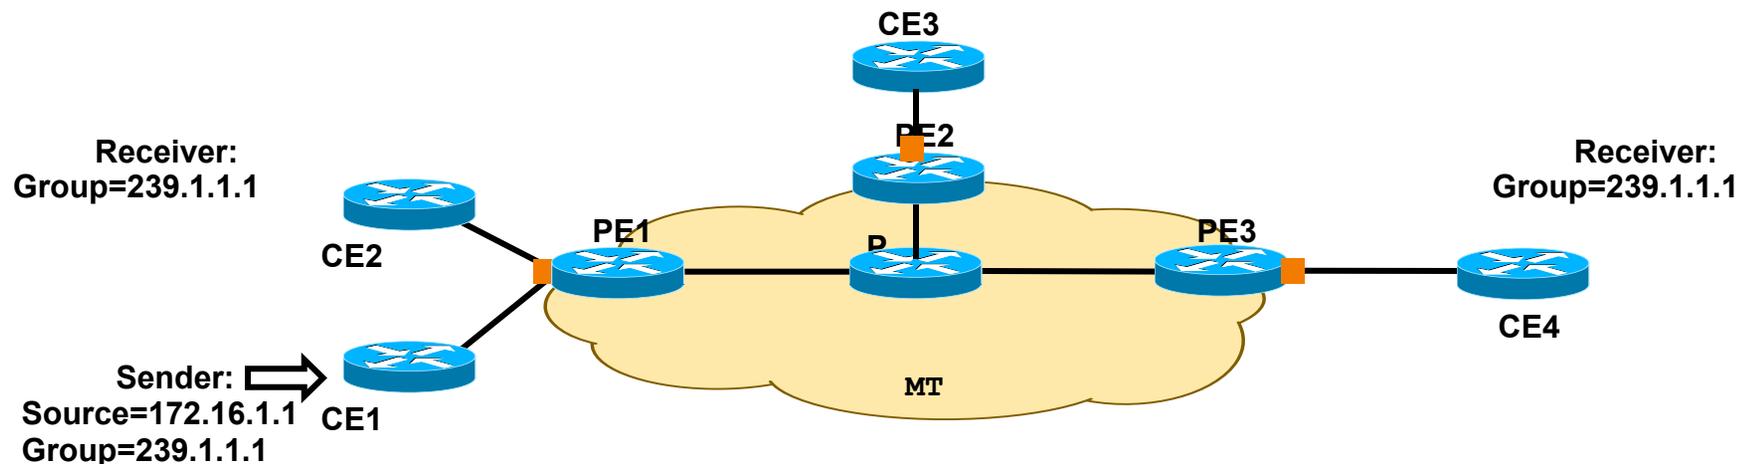
- **Topology above will be used to discuss the following scenarios:**
 - One receiver in the same MVRF
 - One receiver over the Multicast tunnel
 - Second receiver over the Multicast tunnel

Data flow over the Default-MDT



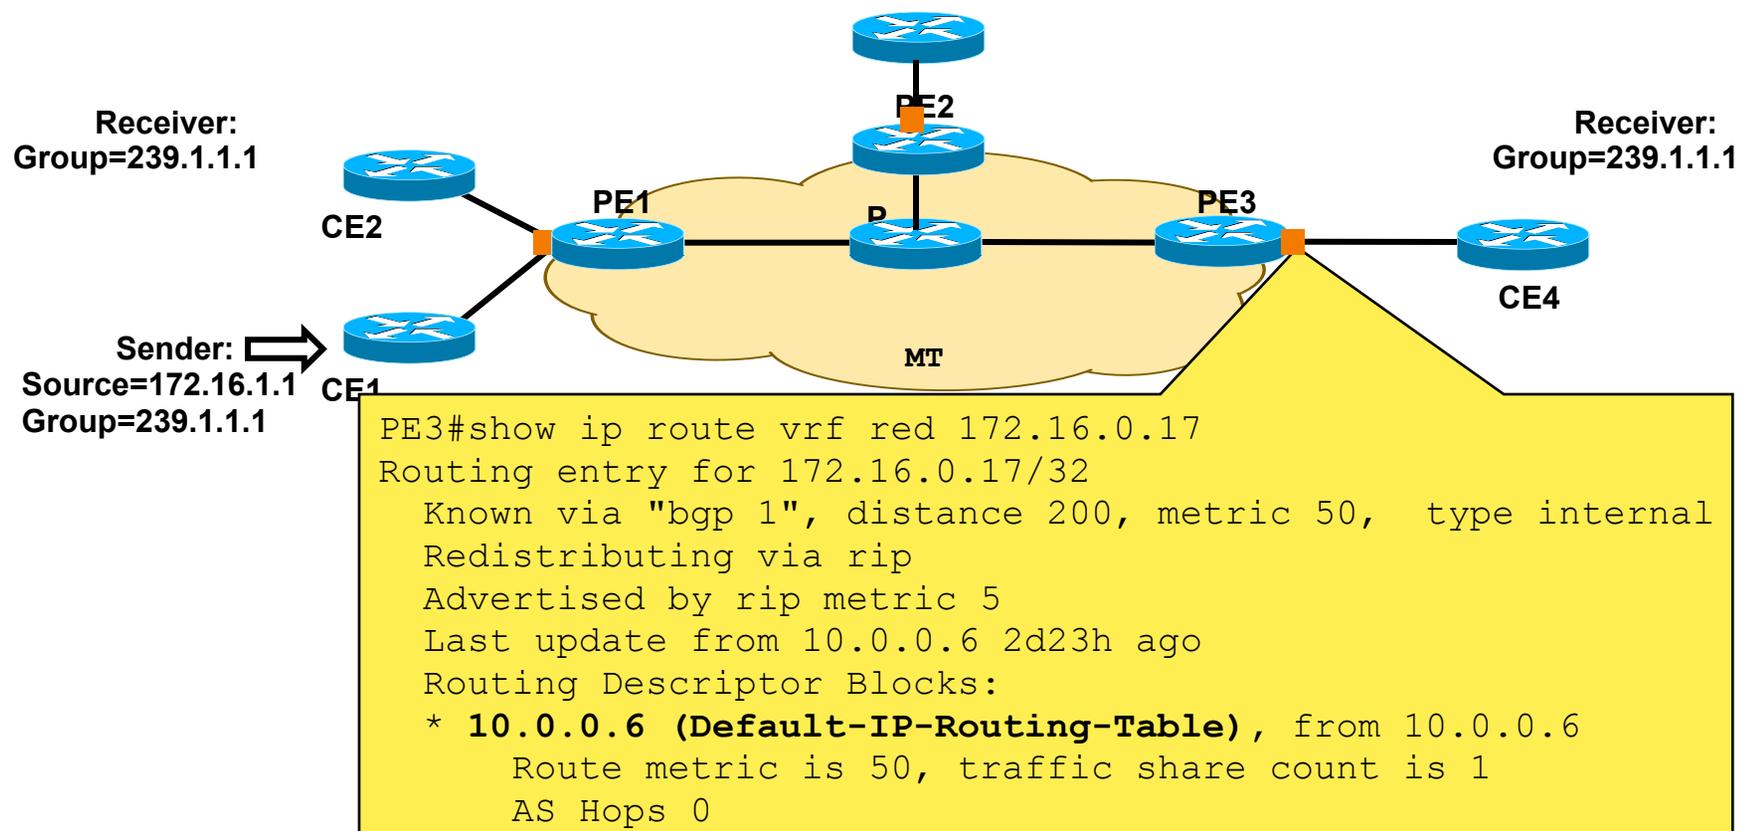
- **One receiver in the same MVRF:**
 - On PE1 the interface to CE2 will be added in the OIL for the VRF specific multicast routing table.

Data flow over the Default-MDT



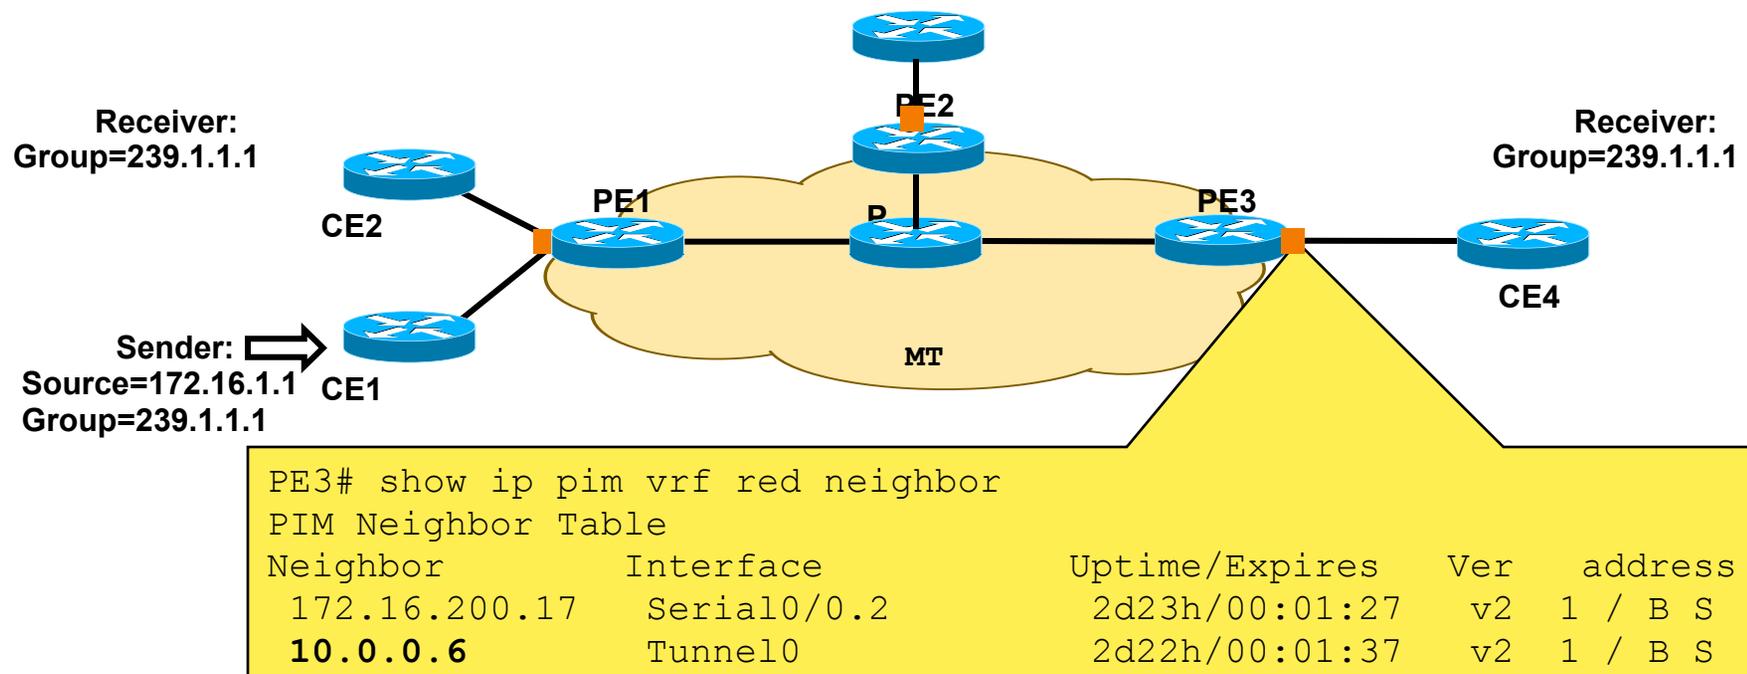
- **One receiver over the Multicast tunnel :**
 - PIM join from CE4 arrives on PE3
 - PE3 determines RPF interface for source (RP)
 - Join is encapsulated and sent over multicast tunnel
 - P router will only have state for the default-mdt groups
 - Encapsulated join arrives on PE1, which is leaf of the multicast tunnel and has the 'Z' flag set for this entry
 - MTI is added to the OIL in the MVRF on PE1 and traffic is forwarded

Finding the RPF-interface



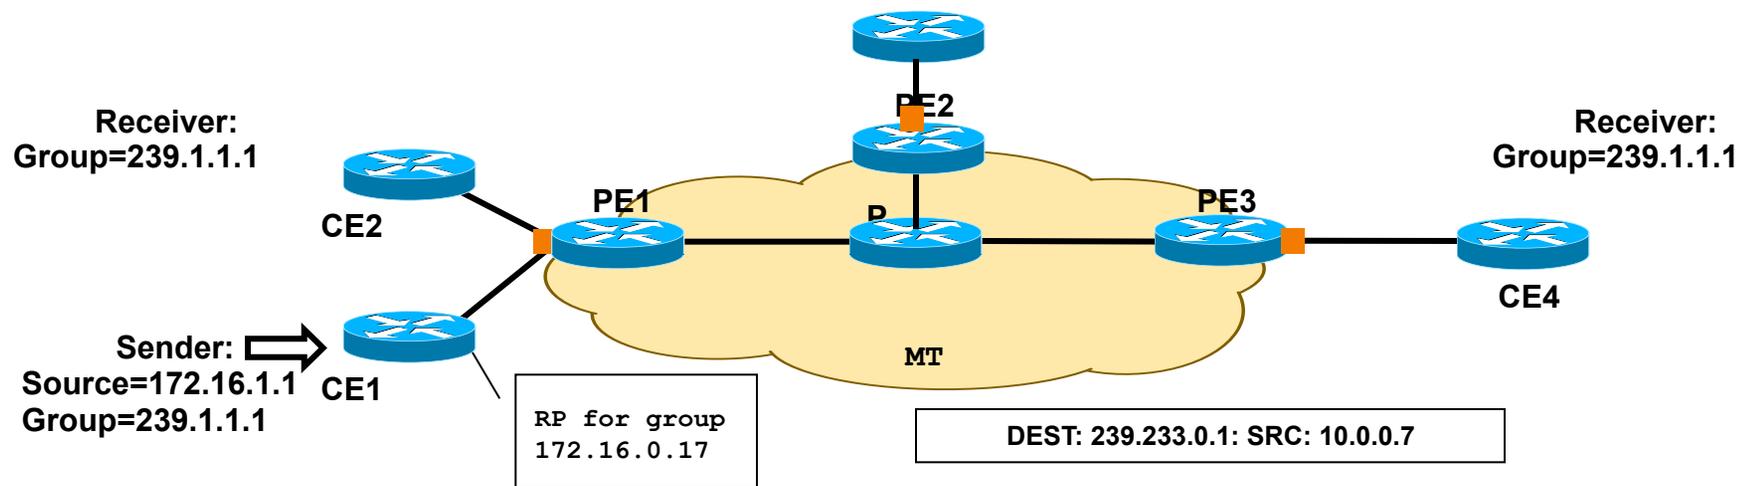
- One receiver over the Multicast tunnel :
RPF lookup returns BGP next-hop from Global MVRF
Therefore, the MTI is our RPF-interface

Finding the RPF-neighbor



- One receiver over the Multicast tunnel :
BGP next-hop is also our PIM neighbor on this MVRF
Therefore, 10.0.0.6 (PE1) is our RPF-neighbor

Data flow over the Default-MDT

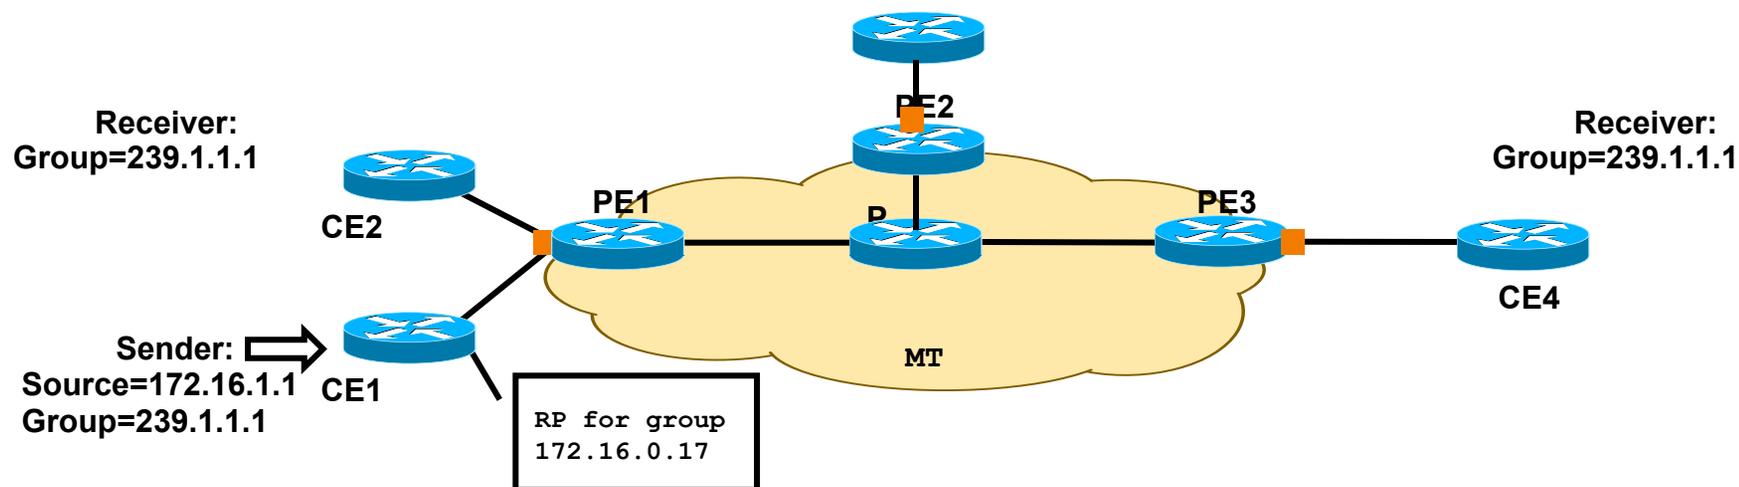


With the RPF-interface and RPF-neighbor determined, PE3 can successfully:

- Encapsulate the PIM-join/prune message
- Encapsulate the packets and send it out on the multicast tunnel

```
*Mar 5 00:48:55.567: PIM(1): Received v2 Join/Prune on Serial0/0.2 from 172.16.200.17, to us
*Mar 5 00:48:55.567: PIM(1): Join-list: (*, 239.1.1.1) RP 172.16.0.17
*Mar 5 00:48:55.567: MRT(1): Create (*, 239.1.1.1), RPF Null, PC 0x607573F4
*Mar 5 00:48:55.567: PIM(1): Check RP 172.16.0.17 into the (*, 239.1.1.1) entry, RPT-bit set, WC-bit set, S-bit set
*Mar 5 00:48:55.571: MRT(1): Add/Update Serial0/0.2/224.0.0.2 to the olist of (*, 239.1.1.1), Forward state
*Mar 5 00:48:55.571: PIM(1): Add Serial0/0.2/172.16.200.17 to (*, 239.1.1.1), Forward state
*Mar 5 00:48:55.571: PIM(1): Send v2 Join on Tunnel0 to 10.0.0.6 for (172.16.0.17/32, 239.1.1.1), WC-bit, RPT-bit, S-bit
```

Data flow over the Default-MDT



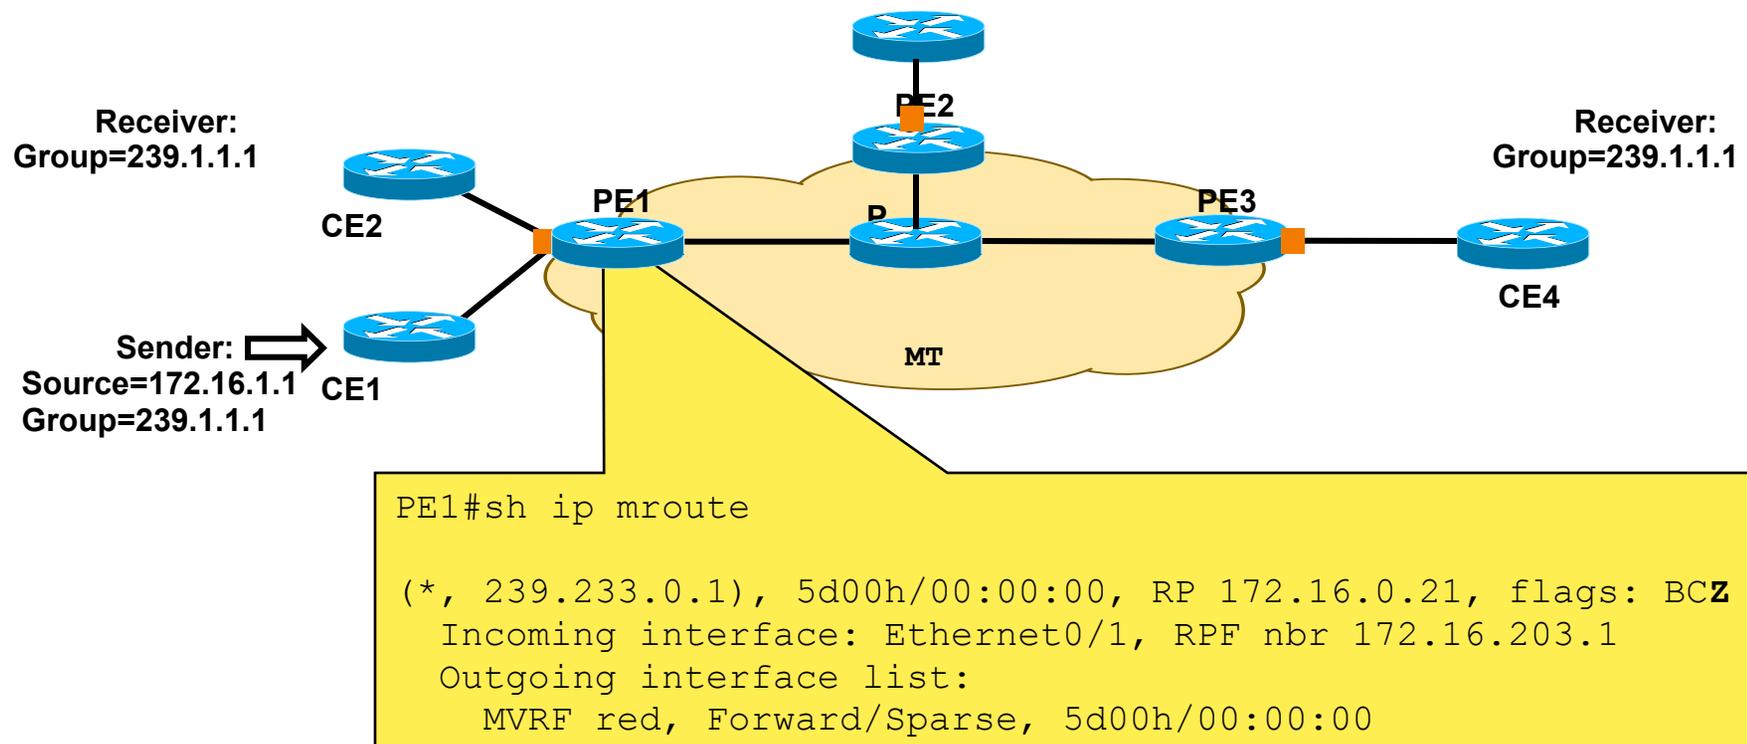
The P router will replicate the P packet to all of its outgoing interfaces in the OIL for the default-group.

PE2 will discard this packet

PE1 will accept, deencapsulate the packet and forward it.

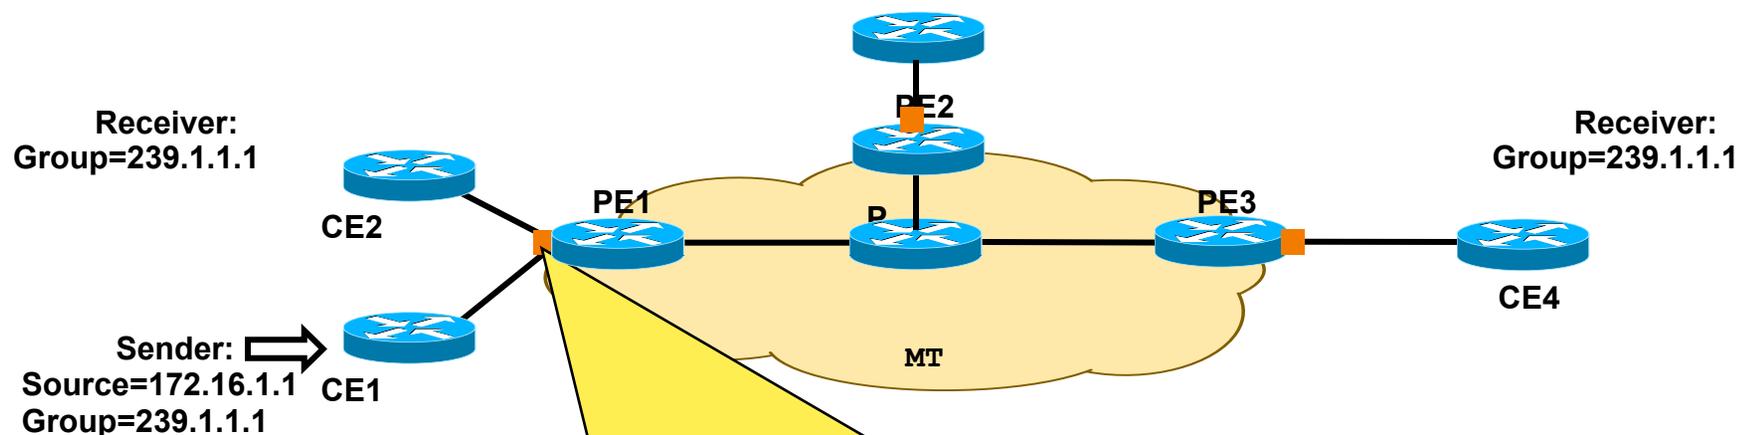
PE1 will add the MTI to the OIL of the group 239.1.1.1

Data flow over the Default-MDT



PE1 has in its global multicast routing table the Z flag set, meaning the router is leaf of the multicast tunnel

Data flow over the Default-MDT



```
PE1# show ip mroute vrf red
(*, 239.1.1.1), 1d03h/00:00:00, RP 172.16.0.17, flags: SJC
Incoming interface: Serial0/0.1, RPF nbr 172.16.200.5
Outgoing interface list:
Serial1/1.1, Forward/Sparse-Dense, 1d03h/00:00:00
Tunnel0, Forward/Sparse-Dense, 00:01:56/00:02:32
```

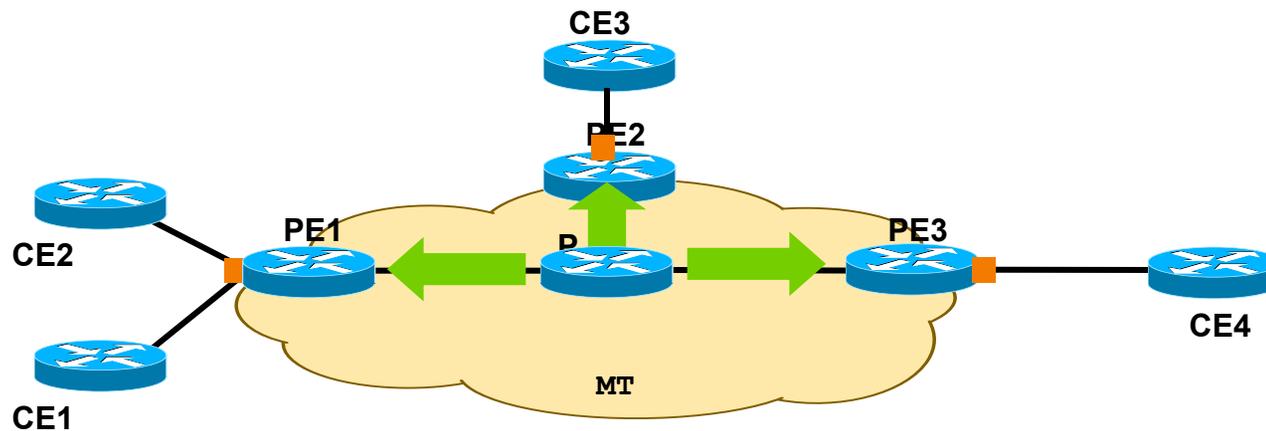
PE1 adds the Multicast Tunnel Interface to the OIL of the VRF specific multicast routing table.

```
PIM(1): Received v2 Join/Prune on Tunnel0 from 10.0.0.7, to us
PIM(1): Join-list: (*, 239.1.1.1) RP 172.16.0.17, RPT-bit set, WC-bit set, S-bit set
PIM(1): Add Tunnel0/10.0.0.7 to (*, 239.1.1.1), Forward state
```

Cisco's Implementation

- Configuration of Default-MDT
- MP-iBGP Update
- Multicast Tunnels
- MVRF
- Data flow over the Default-MDT
- Setup of Data-MDT
- Data flow over the Data-MDT

Setup of Data MDT



With only a default MDT the traffic will be replicated to all PE routers.

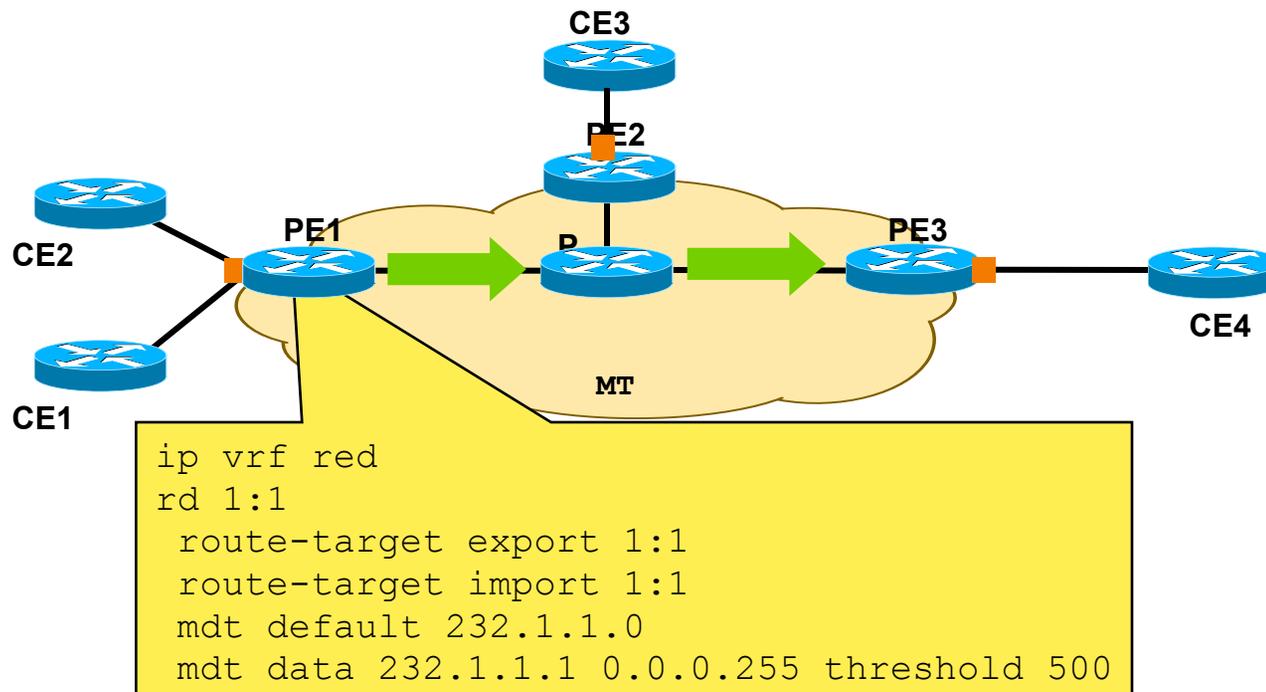
Advantage: Preserves state

Disadvantage: Bandwidth inefficient

A separate DATA MDT can be created for sources that exceed a preset threshold

Useful for high bandwidth sources.

Setup of Data MDT



PE1 is configured with a threshold. E.g. 500 kbps

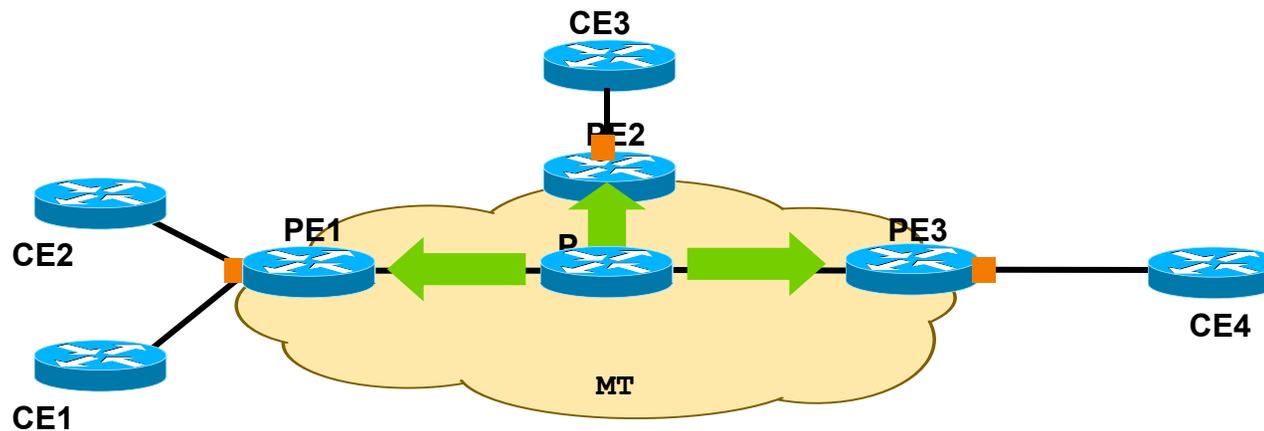
Source behind CE1 exceeds this threshold

PE1 signals to all other routers that it will use a new distribution tree for this source

Only PE routers with interest in this group will join this new tree

Traffic will only go to PE routers that need it.

Setup of Data MDT



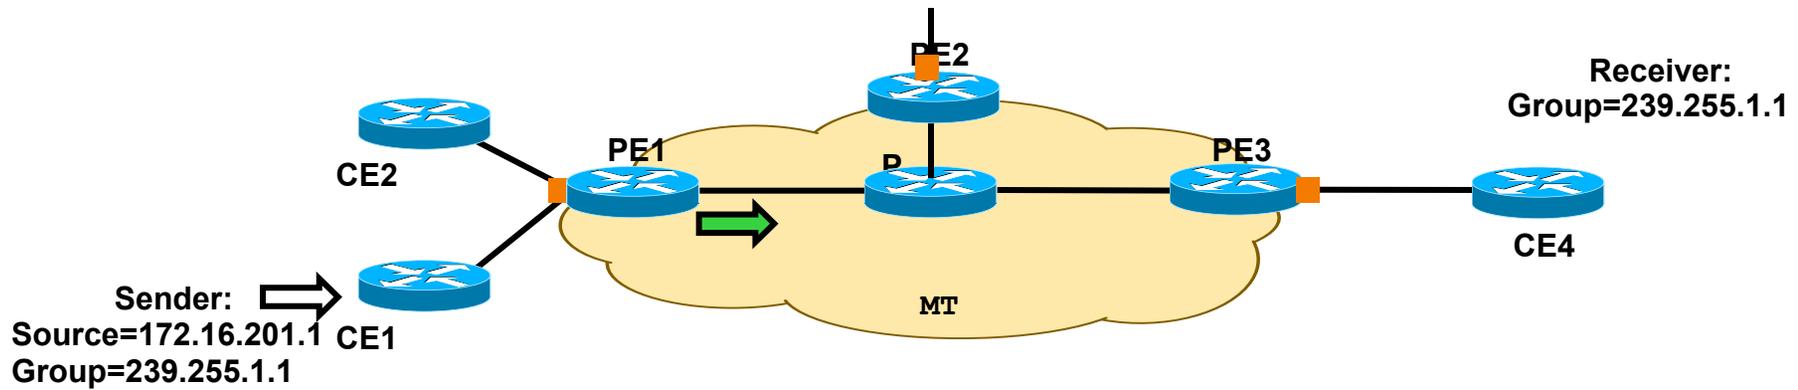
Signaling is done over default-MDT

New UDP message introduced, sent to ALL-PIM-ROUTERS (224.0.0.13)

UDP port number = 3232

0		16	31
Type	Length		
Source			
Group			
MDT-Data			

Setup of Data MDT – Example



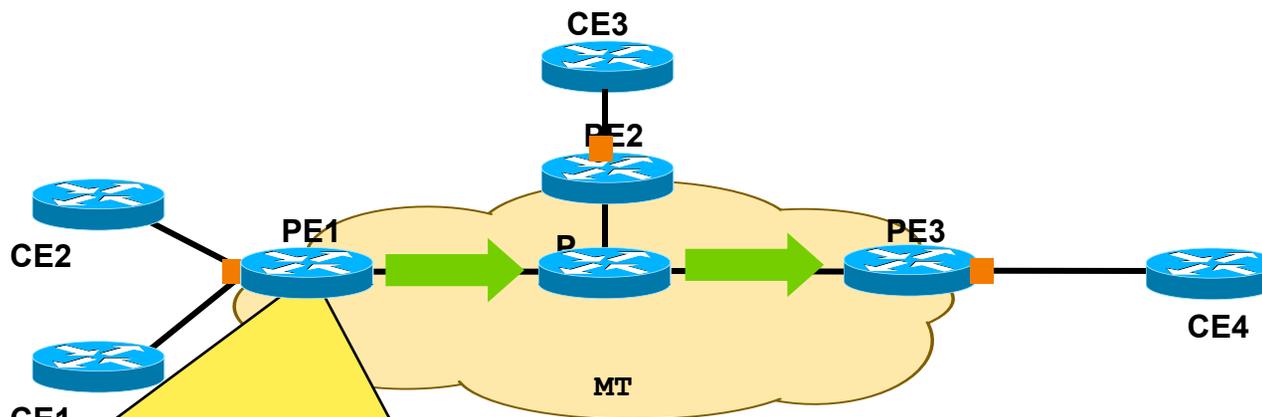
PE1 signals use of new DATA MDT for high bandwidth source.

```
PIM(1): MDT threshold exceeded for (172.16.201.1,239.255.1.1)
PIM(1): MDT join sent for (172.16.201.1,239.255.1.1) MDT:232.0.1.0 Tunnel0
PIM(0): Received v2 Join/Prune on Ethernet0/1 from 172.16.203.1, to us
PIM(0): Join-list: (10.0.0.6/32, 232.0.1.0), S-bit set
PIM(0): Add Ethernet0/1/172.16.203.1 to (10.0.0.6/32, 232.0.1.0), Forward state
```

Cisco's Implementation

- Configuration of Default-MDT
- MP-iBGP Update
- Multicast Tunnels
- MVRF
- MVPN Data flow
- Setup of Data-MDT
- VPN traffic over the Data-MDT

VPN traffic over the Data-MDT



```
PE1# show ip mroute active
Active IP Multicast Sources - sending >= 4 kbps

Group: 232.0.1.0, (?)
Source: 10.0.0.6 (?)
Rate: 50 pps/588 kbps(1sec), 588 kbps(last 50 secs), 586 kbps(life avg)
```

Only interested PE routers receive traffic !

Tradeoff: More state in core network, efficient bandwidth usage

Agenda

- Why Multicast VPNs?
- Multicast VPN Solution
- Cisco's Implementation
- Deployment Considerations

Deployment Considerations

- Provider only considerations

 - What do I need to do to setup my MVPN Core

 - What do I need to do to secure my MVPN core

- Provider-Customer interaction

 - What do I need to do to connect my Multicast customer

 - What do I need to do to secure my customer's multicast traffic

Deployment Considerations

- Provider only considerations
 - PIM-Mode inside the core
 - Platform specific information
 - Many others, not discussed today

Deployment Considerations

- Provider-Customer interaction
 - CE-PE Protocol choice
 - Group to PIM-Mode mapping
 - Again, many others

Deployment Considerations

- Provider only considerations
 - PIM-Mode inside the core
 - Platform specific information

Provider only Considerations

- PIM Mode inside the core

- Bi-directional trees for Default-MDT

- Each PE is a receiver and sender to the group

- Typical many to many application

- One state for each VPN

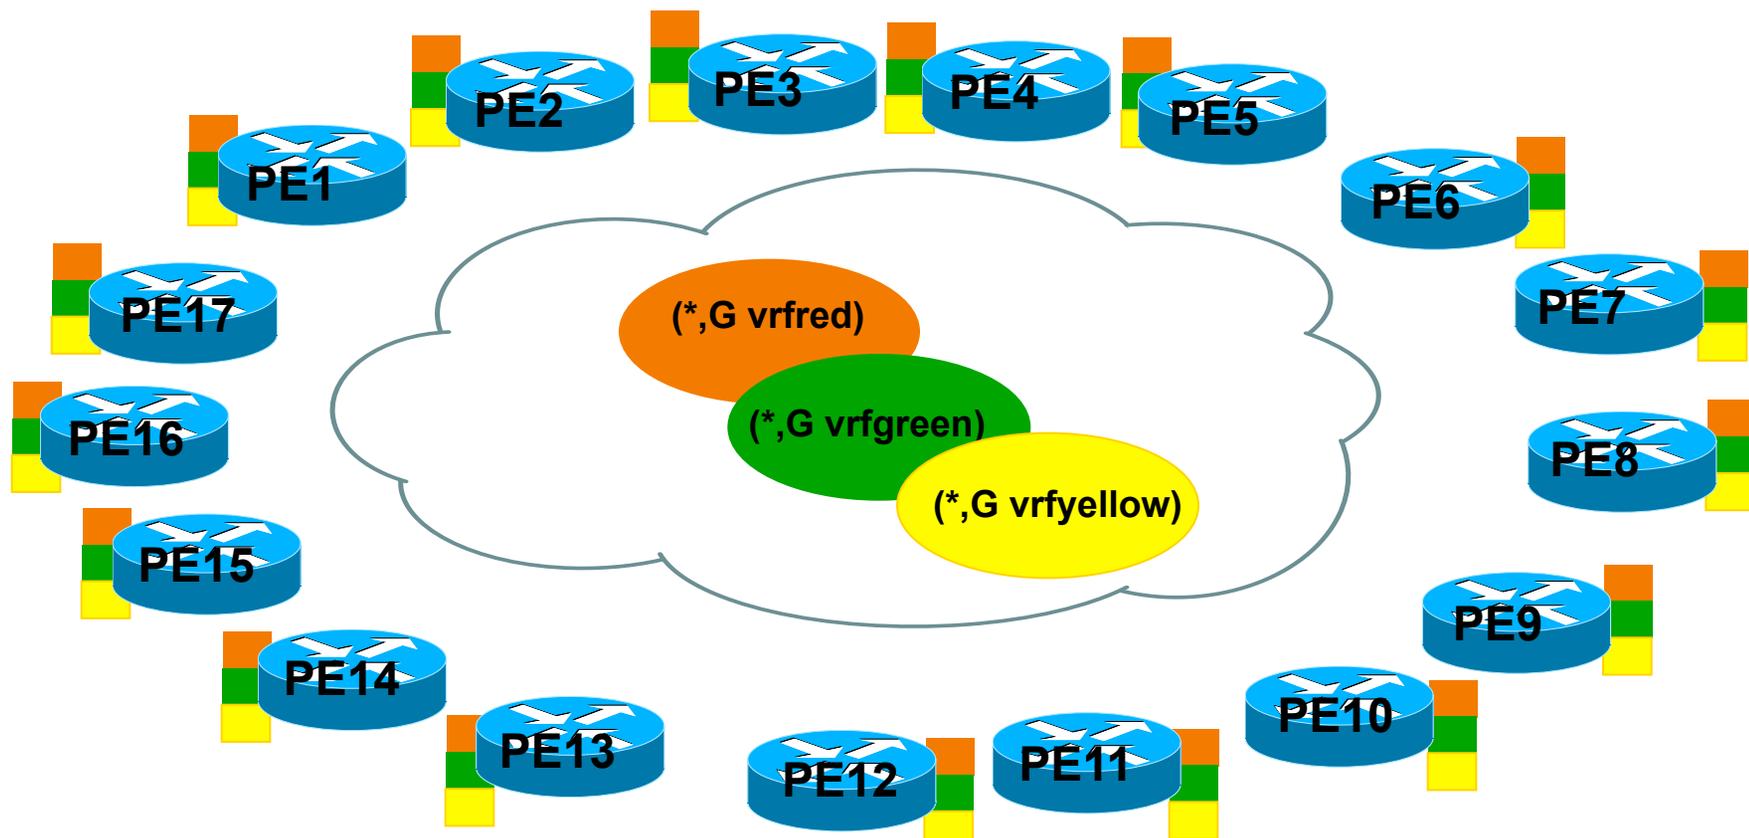
- SSM for Data-MDT

- Only one PE is sender to the group

- Typical one to many application

PIM Mode inside the core

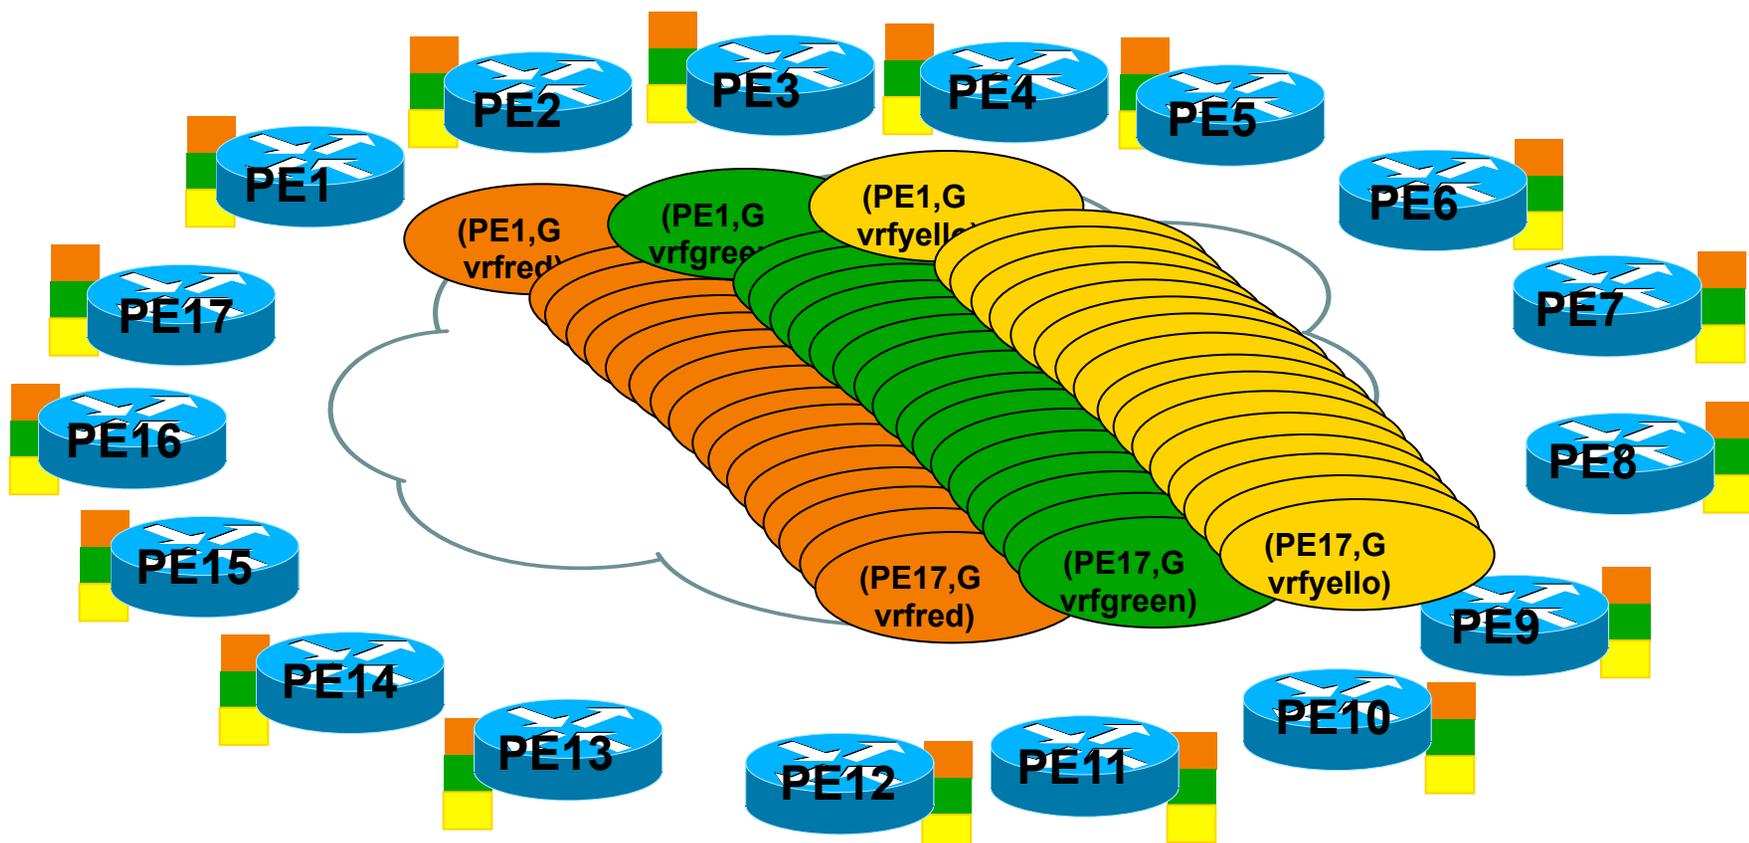
Bidirectional PIM for the Default-MDT



- Only one state per VPN in global table
- 3 in total in this example

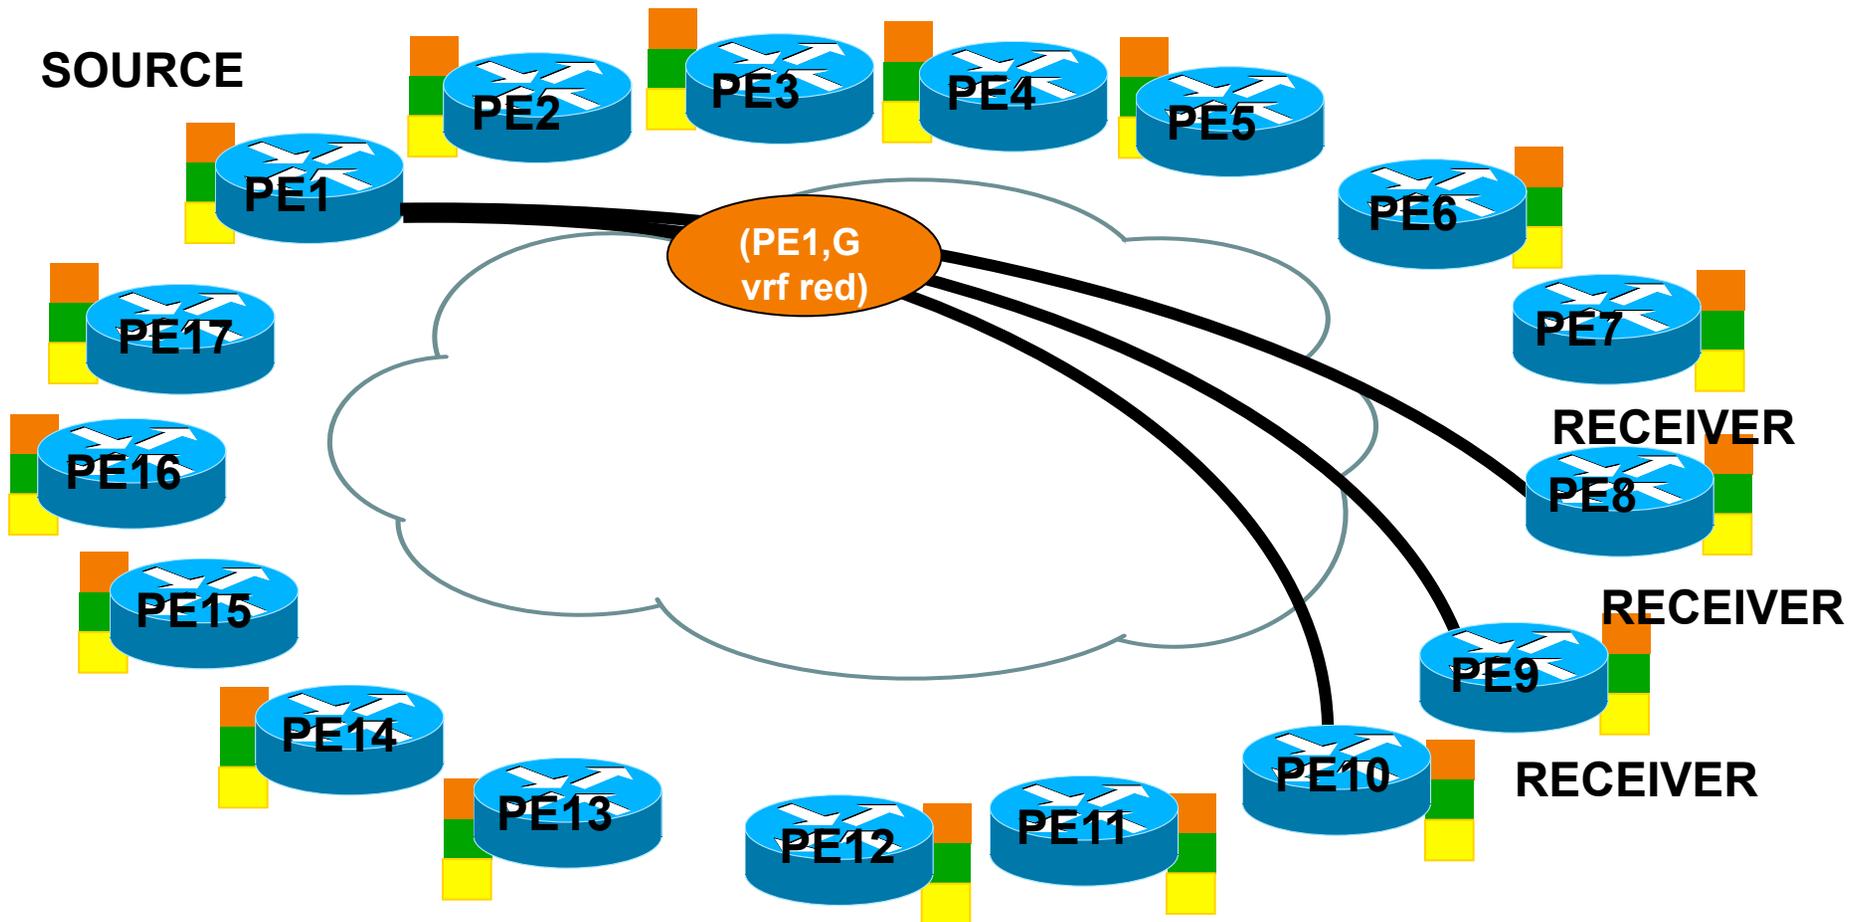
PIM Mode inside the core

PIM-SM or PIM-SSM for the default-MDT



- One state per VRF per PE in global table.
- 51 in total in this example

Use only SSM for Data-MDTs !!!



- Always one to many.
- Group ranges can be identical per VRF per PE

Deployment Considerations

- Provider only considerations
 - PIM-Mode inside the core
 - Platform specific information

Deployment Considerations

- Provider-Customer interaction
 - CE-PE Protocol choice
 - Group to PIM-Mode mapping

CE-PE Protocol Choice

- Does the customer use Dense-mode?
- Does the customer use Auto-RP?
 - ‘ip pim auto-rp listener’ could be useful

Deployment Considerations

- Provider-Customer interaction
 - CE-PE Protocol choice
 - Group to PIM-Mode mapping

Group to PIM Mode Mapping

- Does the customer use ssm ?

Configure the ssm range the customer uses

```
ip pim vrf <x> ssm range <acl>
```

- Does the customer use bidir ?

Configure the bidir rp the customer uses

```
ip pim vrf <x> rp-address <acl> bidir
```

Summary

- Cisco's implementation allows optimal forwarding while preserving state inside core network
- No upgrade of core routers necessary
- Tradeoff state vs bandwidth is under provider's control