

Protocol Independent Multicast - Dense Mode (PIM-DM)

FEATURE OVERVIEW AND CONFIGURATION GUIDE

Introduction

Protocol Independent Multicast - Dense Mode (PIM-DM) is a data-driven multicast routing protocol, which builds source-based multicast distribution trees that operate on the Flood-and-Prune principle. It requires unicast-reachability information, but does not depend on a specific unicast routing protocol.

Products and software version that apply to this guide

This guide applies to AlliedWare Plus™ products that support PIM-DM, running version **5.4.4** or later.

To see whether your product supports PIM-DM, see the following documents:

- The [product's Datasheet](#)
- The [AlliedWare Plus Datasheet](#)
- The product's [Command Reference](#)

These documents are available from the above links on our website at alliedtelesis.com.

Feature support may change in later software versions. For the latest information, see the above documents.

Content

Introduction.....	1
Products and software version that apply to this guide	1
Characteristics of PIM-DM	3
PIM-DM Terminology	3
PIM-DM Configuration.....	5
Configuration example	5
Verifying configuration.....	8
PIM Dense Mode Show Commands.....	9

Characteristics of PIM-DM

PIM Dense Mode (PIM-DM) is a significantly less complex protocol than PIM Sparse Mode (PIM-SM). PIM-DM works on the principle that it is probable that any given multicast stream will have at least one downstream listener. PIM-DM is ideal where many hosts subscribe to receive multicast packets, so most of the PIM Routers receive and forward all multicast packets.

Where PIM-SM only forwards a multicast stream when requested, PIM-DM always floods any new multicast stream that arrives at the PIM Router and only stops flooding the multicast stream on a given link if it is explicitly told to, by receiving a Prune message from the downstream PIM Router.

PIM-DM does not include the concepts of Rendezvous Points, which are used in PIM-SM. PIM-SM explicitly builds unidirectional shared trees rooted at a Rendezvous Point (RP) per group. PIM-DM implicitly builds shortest-path trees by flooding multicast traffic domain wide, then Prunes back branches of the tree where no receivers are available. As with PIM-SM, so does PIM-DM also use Reverse Path Forwarding (RPF) to stop loops for packet forwarding for PIM Routers receiving multicast packets.

PIM-DM Terminology

See the below descriptions of the terms and concepts used to describe the PIM-DM protocol:

- PIM router** Any Layer 3 routing device that is running PIM, such as an Allied Telesis managed Layer 3 switch or Allied Telesis router.
- Reverse path forwarding** Reverse Path Forwarding (RPF) is the mechanism that PIM uses to make sure it does not forward multicast streams around in loops. If a set of PIM Routers are connected in a loop, and each PIM Router is forwarding a given multicast stream, then eventually the multicast stream would be forwarded right around the loop.

To prevent this from happening, PIM makes use of the fact that the unicast routing tables in a set of PIM Routers should converge into a loop-free tree of paths to any given destination.

When a PIM Router receives a multicast stream from source address **SourceA** through an interface **IF1**, it checks whether **IF1** is the interface the PIM Router would use to reach **SourceA**. The PIM Router will only forward the multicast stream if **IF1** is the interface the PIM Router would use to reach **SourceA**.

RPF determines whether the interface is correct by consulting unicast routing tables. This ensure that the multicast stream is forwarded in a loop-free manner back up the tree of unicast paths that lead to the source.

Forwarding multicast packets PIM Routers forward a given multicast stream onto all PIM enabled IP interfaces that have not received a Prune for the given multicast stream. As with unicast routing, the PIM Router decrements the TTL (Time To Live) in each packet that the PIM Router forwards. The packet is discarded if the TTL is decremented to 0.

However, unlike unicast routing, the destination MAC addresses of the packets are not altered as they are forwarded by the PIM Router. The destination MAC addresses remain set to the multicast MAC addresses that correspond to the destination group address of the multicast stream.

Upstream Towards the Source.

Downstream Anything other than the upstream interface for that group.

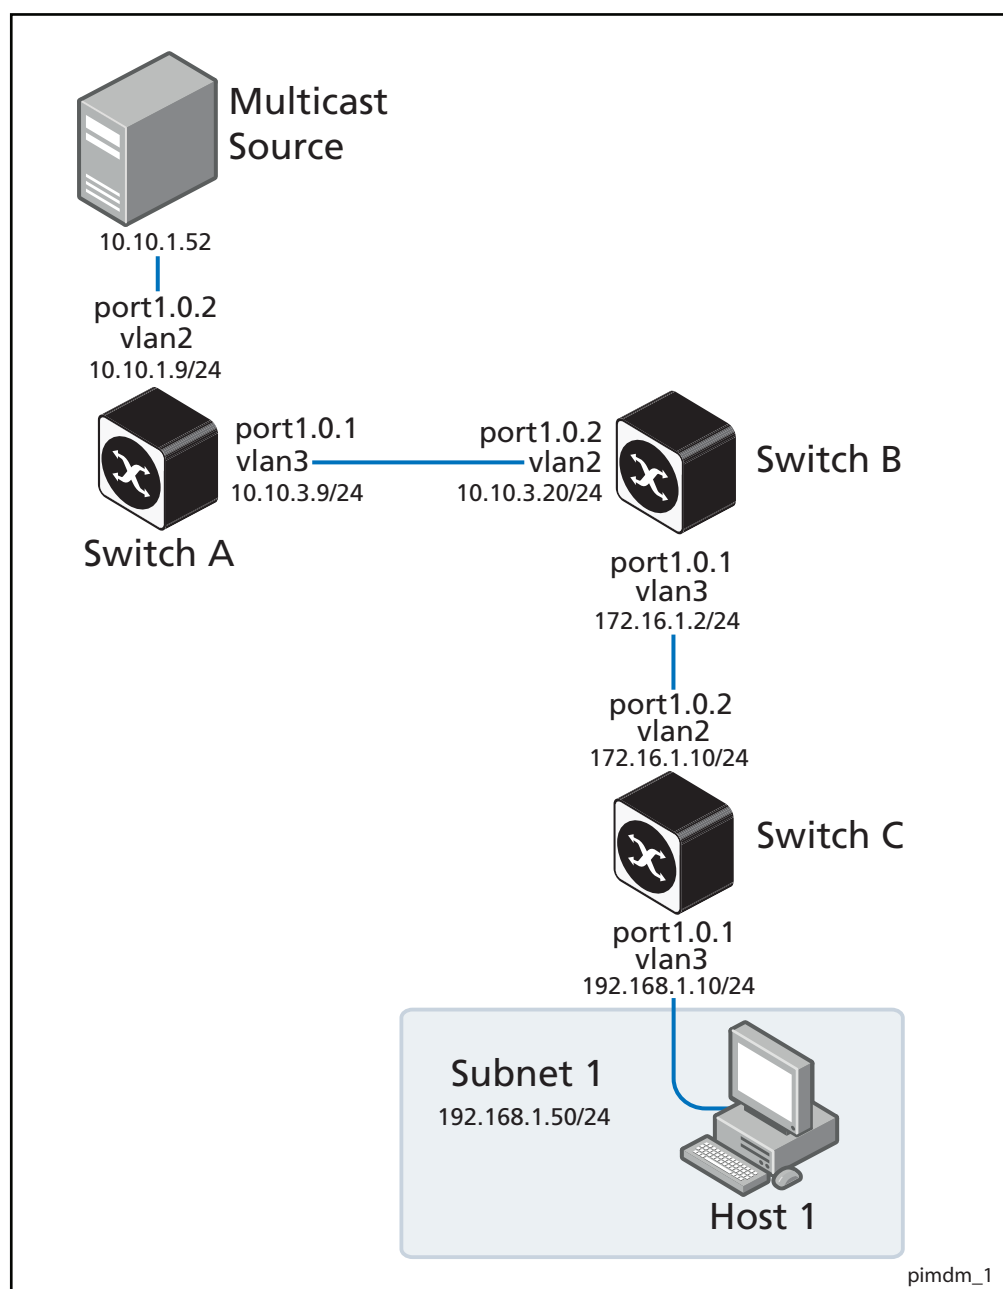
PIM-DM Configuration

The main requirement is to enable PIM-DM on the desired interfaces. This section provides a PIM-DM configuration example for a relevant scenario. The configuration uses Allied Telesis managed Layer 3 switches as the PIM routers. Three PIM routers are connected in a chain, and a multicast client is attached to the third PIM router:

Configuration example

In this example, the address of the multicast source is 10.10.1.52. The following figure displays the network topology used in this example:

Figure 1: Network topology example



The steps involved in the forwarding of the multicast streams for this sample configuration are:

- Switch A**
1. When the PIM Routers start, they use the exchange of PIM Hello packets for PIM neighbor relationships with each other. Then each PIM Router becomes aware of the location of its PIM neighbors.
 2. As a multicast stream arrives from the source to Switch A, it performs an RPF check on the source IP address of the multicast stream. Switch A determines the best route to the source IP address (10.10.1.52) is the receiving interface, so it forwards the multicast stream to its only PIM neighbor.
 3. Switch A creates an (S, G) (Source, Group) entry in its PIM-DM forwarding table. Any further packets from the same source, which are destined to be forwarded to the same group, will be automatically forwarded without an RFP (Reverse Path Forwarding) check.
- Switch B**
4. When the multicast stream arrives at Switch B, it performs the same steps (2 and 3) as Switch A. This results in Switch B also having an (S, G) entry for the multicast stream in its PIM forwarding table, and the multicast stream is forwarded to Switch C.
- Switch C**
5. When the multicast stream arrives at Switch C, it will perform an RPF check on the multicast stream as it arrives, and accept it.

This PIM Router does not have any downstream PIM Routers, but if Switch C has received an IGMP report from the client to request this multicast stream, Switch C will forward the multicast stream out port1.0.1, but no other ports.

If the client leaves the group, and Switch C has no other attached clients requesting the group, then Switch C will send a Prune message upstream, resulting in Switch A and Switch B stopping forwarding the multicast stream to Switch C.

Switch A configuration See the following configuration output for Switch A:

```
hostname Switch A
vlan database
vlan 2 state enable
vlan 3 state enable
interface vlan2
ip address 10.10.1.9/24
ip igmp
ip pim dense-mode
!
interface vlan3
ip address 10.10.3.9/24
ip igmp
ip pim dense-mode
!
interface port1.0.1
switchport access vlan 3
!
interface port1.0.2
switchport access vlan 2
!
ip multicast-routing
!
```

Switch B
configuration

See the following configuration output for Switch B:

```
hostname Switch B
vlan database
vlan 2 state enable
vlan 3 state enable
interface vlan2
ip address 10.10.3.20/24
ip igmp
ip pim dense-mode
!
interface vlan3
ip address 172.16.1.2/24
ip igmp
ip pim dense-mode
!
interface port1.0.1
switchport access vlan 3
!
interface port1.0.2
switchport access vlan 2
!
ip multicast-routing
!
```

Switch C
configuration

See the following configuration output for Switch C:

```
hostname Switch C
vlan database
vlan 2 state enable
vlan 3 state enable
interface vlan2
ip address 172.16.1.10/24
ip igmp
ip pim dense-mode
!
interface vlan3
ip address 192.168.1.10/24
ip igmp
ip pim dense-mode
!
interface port1.0.1
switchport access vlan 3
!
interface port1.0.2
switchport access vlan 2
!
ip multicast-routing
!
```

Verifying configuration

Use the following commands to verify the interface details and multicast routing table.

Interface details The **show ip pim dense-mode interface** command displays the interface details for Switch C.

```

Address          Interface VIFindex Ver/   Nbr
                port1.0.1 0       Mode  Count
192.168.1.10    port1.0.1 0       v2/D  0
172.16.1.10     port1.0.2 2       v2/D  1

```

IP multicast routing table The **show ip mroute** command displays the IP multicast routing table (for Switch C).

```

IP Multicast Routing Table
Flags: I - Immediate Stat, T - Timed Stat, F - Forwarder installed
Timers: Uptime/Stat Expiry
Interface State: Interface (TTL)

(10.10.1.52, 224.0.1.3), uptime 00:00:15
Owner PIM-DM, Flags: F
  Incoming interface: port1.0.2
  Outgoing interface list:
    port1.0.1 (1)

```

IP PIM-DM multicast routing table The **show ip pim dense-mode mroute** command displays the IP PIM-DM multicast routing table (for Switch C).

```

PIM-DM Multicast Routing Table
(10.10.1.52, 224.0.1.3)
  RPF Neighbor: 172.16.1.2, Nexthop: 172.16.1.2, port1.0.2
  Upstream IF: port1.0.2
    Upstream State: Forwarding
    Assert State: NoInfo
  Downstream IF List:
    port1.0.1, in 'olist':
      Downstream State: NoInfo
      Assert State: NoInfo

```


PIM Dense Mode Show Commands

First, check that PIM is active on the interfaces you believe you have configured as PIM interfaces.

show ip pim dense-mode interface

```
awplus-1#show ip pim dense-mode interface
Total configured interfaces: 3 Maximum allowed: 100
Total active interfaces: 3
Address Interface VIFIndex Ver/ Nbr
Mode Count
192.168.1.1 vlan1 2 v2/D 0
192.168.2.1 vlan2 3 v2/D 1
192.168.3.1 vlan3 4 v2/D 1
```

The output of this command shows us the IP address and VLAN of the interfaces on the switch which have PIM-DM mode configured, as well as the PIM Version number (v2) and the PIM Mode—D (Dense). We can see from the Neighbor count (Nbr Count) that there are no PIM Dense Mode neighbors on vlan1 (which has the Multicast Server on it only) and one PIM-DM neighbor on each of vlan2 and vlan3.

The VIF Index is an internal index for the interface used by the switch to identify which interface it is dealing with.

show ip pim dense-mode interface detail

```
awplus-1#show ip pim dense-mode interface detail
vlan1 (vif-id: 2):
Address 192.168.1.1
Hello period 30 seconds, Next Hello in 10 seconds
Max Graft Retries: infinite
GenId: 762299093
Ext-Srcs as Directly-Connected: Disabled
Over-ride interval 2500 milli-seconds
Propagation-delay 1000 milli-seconds
Neighbors: none
```

Let's now look at the meaning of the information in this output.

■ vlan1 (vif-id: 2):

Again, we see the VIF index value referred to. Another way to see the mapping of the Multicast Virtual Interface (MVIF) values is by using the **show ip mvif** command:

```
awplus-1#show ip mvif
Interface Vif Owner TTL Local Remote Uptime
Idx Module Address Address
vlan1 2 PIM-DM 1 192.168.1.1 0.0.0.0 04:17:43
vlan2 3 PIM-DM 1 192.168.2.1 0.0.0.0 04:17:43
vlan3 4 PIM-DM 1 192.168.3.1 0.0.0.0 04:17:43
```

- **Address 192.168.1.1**

The IP address of vlan1 is 192.168.1.1.

- **Hello period 30 seconds, Next Hello in 5 seconds**

When we enable PIM-DM on an interface, the router starts to send multicast PIM hello messages to 224.0.0.13 in order to form adjacency with PIM neighbors. We can see that the Hello interval is set to the default interval of 30 seconds. The interface is due to send another Hello message in 5 seconds.

- **Max Graft Retries: infinite**

The PIM-DM protocol requires that a Graft message receive an ACK message (or an implicit ACK in the form of a state refresh message). A router should resend a Graft message if it does not receive an ACK message. The default is for the switch to keep resending Graft messages until it gets an acknowledgment (infinite retries). This can be modified by using the `ip pim max-graft-retries` command.

- **GenId: 1106721597**

The Generation Identifier. This is a random 32-bit value that is sent whenever the neighbor activates PIM on the interface. The value associated with a given interface is sent in the Hello packets transmitted from this interface. It can be used to determine when the neighbor has been reactivated after a failure.

- **Ext-Srcs as Directly-Connected: Disabled**

By default, if a PIM router receives data on an interface, it requires that the source IP address of the data is in the same subnet as the receiving interface, or that the data has come via a PIM neighbor. Sometimes you may wish to over-rule this requirement.

The command `ip pim ext-srcs-directly-connected` configures an interface to treat all multicast traffic arriving on the interface as though it was sent from a host directly connected to the interface. This can be useful in cases where the traffic is arriving into a subnet via a mechanism other than PIM. The default is disabled.

- **Over-ride interval 2500 milli-seconds**

If all routers on a LAN support the LAN Prune Delay option, then the PIM routers on that LAN will use the values received to adjust their Join/Prune Over-ride Interval on that interface.

To avoid synchronization of Prune Over-ride (Join) messages when multiple downstream routers share a multi-access link, sending of these messages is delayed by a small random amount of time. The value 2.5 seconds means that each router will use a random timer value between 0 and 2.5 seconds.

- **Propagation-delay 1000 milli-seconds**

The expected delay in the transfer of PIM messages across the VLAN interface that it is attached to. The default is 1 second and is configurable by using the `ip pim propagation-delay` command.

- **Neighbors: none**

Here we see the PIM Dense Mode neighbors that this VLAN has, if any. As switch awplus-1 vlan1 is only connected to the Multicast Source, it does not have any PIM neighbors.

Having checked that the interfaces look right, let's see if PIM is forwarding streams.

show ip pim dense-mode mroute

```
awplus-1# show ip pim dense-mode mroute
(192.168.1.50, 225.1.1.1)
Source directly connected on vlan1
State-Refresh Originator State: Originator
Upstream IF: vlan1
Upstream State: Forwarding
Assert State: NoInfo
Downstream IF List:
vlan2, in 'olist':
Downstream State: NoInfo
Assert State: NoInfo
vlan3:
Downstream State: Pruned
Assert State: NoInfo
```

Here we can see the PIM-DM multicast routing table entries. Let's understand the meaning of the information in the PIM mroute entry (from the PIM-DM Multicast Routing Table) shown above.

- (192.168.1.50, 225.1.1.1)

We see that group 225.1.1.1 is being sent from source 192.168.1.50 (our multicast server).

- **Source directly connected on vlan1**

There is no router between the host that is sending this group and a port in vlan1 on the switch.

- **State-Refresh Originator State: Originator**

This tells us that this switch will be sending State Refresh Control Messages for this group. This is a periodic reminder, multicast to all other PIM-DM routers in the network, that this group is still available.

- **Upstream IF: vlan1**

This tells us the interface via which the source of this group is to be reached.

- **Upstream State: Forwarding**

If the switch has any ports that want to receive this group, then the Upstream State will be 'Forwarding'. If there are no ports that want to receive this group, then the Upstream State will be 'Pruned'.

This sounds a little paradoxical, as the forwarding happens out of the downstream interface(s), not out the upstream interface. However, the state machine defined in the RFC uses the name 'Forwarding' for the state of an upstream interface that is channeling the stream to active downstream interfaces.

- **Assert State: NoInfo**

By default the Assert state of an interface will be in the 'NoInfo' state.

If the switch receives multicast traffic on a non-RPF interface, Assert messages are sent to decide which switch will be the one to forward the multicast group onto the shared LAN.

The switch first sets its Assert state to 'Winner' and starts the Assert negotiation. The switch that receives Assert messages with better credentials than itself will set its Assert state to 'Loser'.

But if no traffic for this group has been received on a non-RPF interface, the Assert process is not triggered, and the Assert state stays at 'NoInfo'.

- **Downstream IF List:**

The downstream interface list will contain all the PIM interfaces on the router except the upstream interface for this stream. Irrespective of whether or not the stream is being forwarded to one of these interfaces.

- **vlan2, in 'olist':**

From this 'olist', we can see that this group is being forwarded to vlan2.

- **Downstream State: NoInfo**

'NoInfo' means that this group is being forwarded out this interface.

If there are no listeners on this interface and it is therefore receiving Prune messages, it will show as 'Pruned'.

- **vlan3:**

Downstream State: Pruned

The switch has received Prune messages on this interface (there are no listeners) and has set the Downstream state to 'Pruned'. The stream is not being forwarded out this interface.

Let us also check the state of the switch's neighbor relationships.

show ip pim dense-mode neighbor detail

The neighbor information can be seen in another format with this command.

```
awplus#show ip pim dense-mode neighbor detail
Neighbor 192.168.2.2 (vlan2)
Up since 02:02:19, Expires in 00:01:27
Neighbor 192.168.3.2 (vlan3)
Up since 02:02:18, Expires in 00:01:29
```

show ip pim dense-mode nexthop

This command shows the next-hops of the unicast routes to the Sources and Receivers in the network.

```
awplus-1#show ip pim dense-mode nexthop
Destination Nexthop Nexthop Nexthop Metric Pref
Num Addr Interface
192.168.1.50 1 0.0.0.0 vlan1 3 1
192.168.6.100 1 192.168.2.2 vlan2 3 110
```

The **Destination** field shows the IP address of the Multicast Source (192.168.1.50) and the Multicast Client (192.168.6.100).

The **Nexthop Num** shows us that the nexthop is 1 hop away.

The **Nexthop Addr** shows as '0.0.0.0' for the Multicast Source 192.168.1.50, as this is directly connected to vlan1 on this switch.

The **Nexthop Addr** of the client 192.168.5.100 is that of awplus-2's vlan2 interface address.

The **Metric** and **Preference** are derived from the Unicast IP route table. The Preference is equivalent to the Administrative Distance (AD) and is a measure of how trusted the source of the routing information is—in the unicast routing table a directly connected interface has a Preference of 0, and the Preference for a route learned via OSPF is 110.

The Metric shows the cost of this path to the network.

We might have expected that the Preference and Metric of a directly connected route would be 0, as is shown in the unicast routing table, however in AlliedWare Plus, 3 is the default PIM-DM Metric if the metric from the unicast route table is 0, and a Preference of 1 is the default if the preference from the unicast route table is 0.

Just to illustrate this, here are the IP route table entries from which the information in the **show ip pim dense-mode next-hop** output is derived:

```
awplus-1#show ip route 192.168.1.50
Routing entry for 192.168.1.0/24
Known via "connected", distance 0, metric 0, best
* is directly connected, vlan1
```

```
awplus-1#show ip route 192.168.6.100
Routing entry for 192.168.6.0/24
Known via "ospf", distance 110, metric 3, best
Last update 00:58:09 ago
* 192.168.2.2, via vlan2
```