

# Ingress Multicast Path Management

---

## In This Chapter

This section provides information about Ingress Multicast Path Management (IMPM).

Topics in this section include:

- [Applicability on page 2406](#)
- [Summary on page 2407](#)
- [Overview on page 2408](#)
- [Configuration on page 2413](#)
- [Conclusion on page 2448](#)

## Applicability

The information in this section is applicable to all Alcatel-Lucent 7x50 platforms configured with IOM1/2, IOM3-XP and IMMs line cards, but not to the SR-1, ESS-1, 7710 SR c-4/12 or the 7750 SR c-4/12. The configuration was tested on release 9.0R6. There are no pre-requisites for this configuration.

## Summary

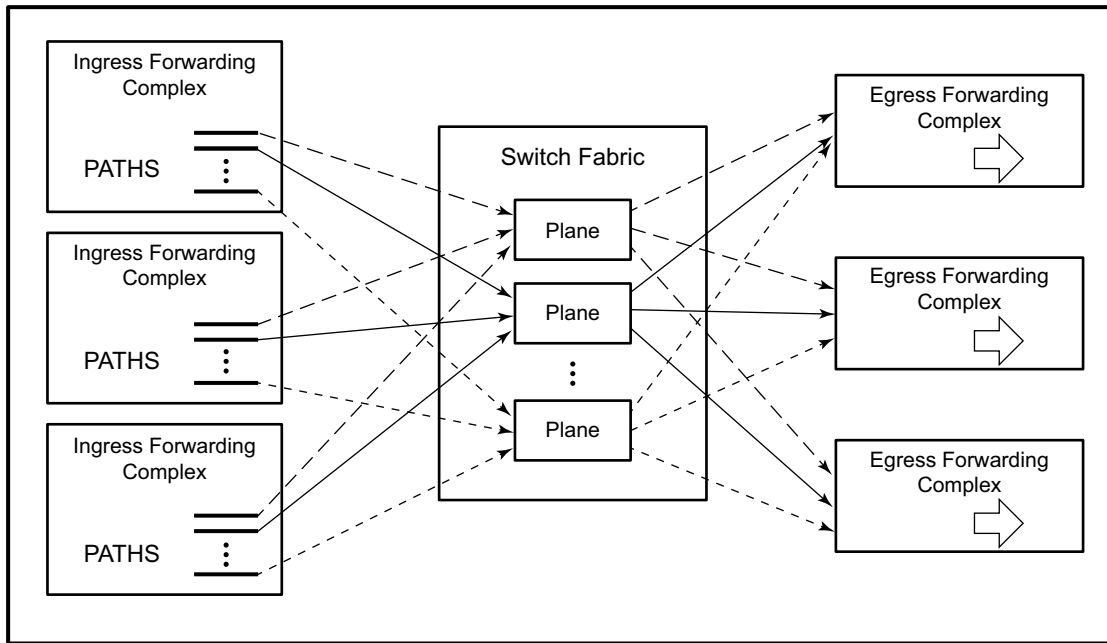
Ingress Multicast Path Management (IMPM) optimizes the IPv4 and IPv6 multicast capacity on the applicable systems with the goal of achieving the maximum system-wide IP multicast throughput. It controls the delivery of IPv4/IPv6 routed multicast groups and of VPLS (IGMP and PIM) snooped IPv4 multicast groups, which usually relate to the distribution of IP TV channels.

A description is also included of the use of IMPM resources by point-to-multipoint LSP IP multicast traffic, and policed ingress routed IP multicast or VPLS broadcast, unknown or multicast traffic. The system capacity for these traffic types can be increased even with IMPM disabled.

## Overview

IMPM introduces the concept of paths on a line card (IOM/IMM) which connect to planes on the chassis switch fabric (Figure 371).

IMPM monitors the ingress rate of IP multicast channels (S,G<sup>1</sup> multicast streams) on line card paths and optimizes the use of the capacity of each switch fabric plane. Its goal is to forward as many channels as possible through the system in order to make maximum use of the switch fabric planes without incurring multicast packet loss. IMPM achieves this by moving entire multicast channels between the line card paths, and therefore between switch fabric planes, to achieve an optimal packing of channels onto path/planes. These actions take into consideration the total ingress multicast traffic being received by all line cards with IMPM enabled and a configured preference of each channel.



OSSG725

**Figure 371: IOM/IMM Paths Connecting to Switch Fabric Planes**

There are three types of path: primary, secondary and ancillary paths (the ancillary path is specific to the IOM1/2 and is discussed in [Ancillary Path on page 2428](#)).

1. S,G refers to an individual multicast stream by referencing the source (S) and multicast group (G) used by the stream.

When a new channel is received on a line card for which there is an egress join, its traffic is initially placed on to a secondary path by default. IMPM monitors the channel's traffic rate and, after an initial monitoring period, can move the channel to another path (usually a primary path) on which sufficient capacity exists for the channel. IMPM constantly monitors all of the ingress channels and therefore keeps a picture of the current usage of all the line card paths and switch fabric planes. As new channels arrive, IMPM assigns them onto available capacity, which may involve moving existing channels between paths (and planes). If a channel stops, IMPM will be aware that more capacity is now available. If the traffic rate of any channel(s) changes, IMPM will continue to optimize the use of the path/planes.

In the case where there is insufficient plane capacity available for all ingress channels, entire channel(s) are blackholed (dropped) rather than allowing a random degradation across all channels. This action is based on a configurable channel preference with the lowest preference channel being dropped first. If path/plane capacity becomes available, then the blackholed channel(s) can be re-instated.

---

## Paths and Planes

Each path connects to one plane on the switch fabric which is then used to replicate multicast traffic to the egress line cards. Further replication can occur on the egress line card but this is not related to IMPM so is not discussed.

Each plane has a physical connection to every line card and operates in full duplex mode allowing the possibility for traffic received on a plane from one line card path to be sent to every other line card, and back to the ingress line card (note that traffic is actually only sent to the line cards where it may exit). Therefore a given plane interconnects all line cards which allow ingress multipoint traffic from a line card with a path connected to this plane to be sent to multiple egress line cards.

Traffic could be sent by only one line card, or by multiple line cards simultaneously, on to a given plane. The total amount of traffic on a path or plane cannot exceed the capacity of that path or plane, respectively.

There could be more planes available on the switch fabrics than paths on the line cards. Conversely, there could be more total line card paths than planes available on the switch fabrics. In the latter case, the system distributes the paths as equally as possible over the available planes and multiple paths would be assigned to a given plane. Note that multiple paths of either type (primary or secondary) can terminate on a given plane.

The number of paths available per line card depends on the type of line card used whereas the number of planes on a system depends on the chassis type, the chassis mode (**a**, **b**, **c**, **d**) and the number of installed switch fabrics.

To clarify these concepts, consider a system with the following hardware installed.

## Paths and Planes

```
A:PE-1# show card
=====
Card Summary
=====
Slot      Provisioned      Equipped      Admin   Operational      Comments
Card-type  Card-type       State         State
-----
6         iom3-xp         iom3-xp       up      up
7         imm8-10gb-xfp   imm8-10gb-xfp up      up
8         iom3-xp         iom3-xp       up      up
A         sfm4-12         sfm4-12       up      up/active
B         sfm4-12         sfm4-12       up      up/standby
=====
A:PE-1#
```

Output 1 shows the mapping of paths to switch fabric planes.

```
A:PE-1# show system switch-fabric high-bandwidth-multicast
=====
Switch Fabric
=====
                Cap:                Planes:
Slot/Mda Min  Max  Hbm Grp  Hi | Lo
-----
6/1      100% 100% No  0      1 0 3 4 5 6 7 8 9 10 11 12 13 14 15 | 16
7/1      100% 100% No  0     19 17 20 21 22 23 24 25 26 27 28 29 30 31 32 | 33
8/1      100% 100% No  0     35 34 36 37 38 39 40 41 42 43 44 45 46 47  0 |  1
A         100% 100% No  0      2 |  2
B         100% 100% No  0      2 |  2
=====
A:PE-1#
```

### Output 1: Paths and Planes in Chassis Mode d

This system has two SF/CPM4s and is using chassis mode **d**, this creates 24 planes per SF/CPM4 to give a total of 48 planes which are numbered 0-47. The IOM3-XP/IMMs have 16 paths each which are connected to different planes. The SF/CPM4s together use a single plane and an additional plane (18, which is not in the output above) is used by the system itself. As there are more paths (3x16=48) in this configuration than available planes (48-2[system planes 2,18]=46), some planes are shared by multiple paths, namely planes 0 and 1. Note that the path to plane mapping can change after a reboot or after changing hardware.

The following output shows the equivalent information if an IOM2 is added to this configuration in slot 5. In order for the IOM2 to be recognized, the system must be changed to use chassis mode **a**, **b** or **c**.

```
A:PE-1# show card
=====
Card Summary
=====
Slot      Provisioned      Equipped      Admin   Operational      Comments
Card-type  Card-type       State         State
-----
```

```

5          iom2-20g          iom2-20g          up          up
6          iom3-xp           iom3-xp           up          up
7          imm8-10gb-xfp     imm8-10gb-xfp     up          up
8          iom3-xp           iom3-xp           up          up
A          sfm4-12          sfm4-12          up          up/active
B          sfm4-12          sfm4-12          up          up/standby
=====
A:PE-1#

```

The following output shows the mapping of the line card paths to the switch fabric planes with the IOM2 installed.

```

A:PE-1# show system switch-fabric high-bandwidth-multicast
=====
Switch Fabric
=====
          Cap:                Planes:
Slot/Mda  Min  Max  Hbm Grp  Hi | Lo
-----
5/1       100% 100% No  0    1 | 0
5/2       100% 100% No  0    4 | 3
6/1       100% 100% No  0    6 5 7 8 9 10 11 12 13 14 15 0 1 3 4 | 5
7/1       100% 100% No  0    7 6 8 9 10 11 12 13 14 15 0 1 3 4 5 | 6
8/1       100% 100% No  0    8 7 9 10 11 12 13 14 15 0 1 3 4 5 6 | 7
A         100% 100% No  0    2 | 2
B         100% 100% No  0    2 | 2
=====
A:PE-1#

```

**Output 2: Paths and Planes in Chassis Mode a/b/c**

Now that the system is not in chassis mode **d**, in fact it is in mode **a** (but the output would be the same in modes **b** or **c**) the SF/CPM4s each create 8 planes giving a total of 16, numbered 0-15. One plane (2) is used by the SF/CPM4s, leaving 15 (0-1,3-15) planes for connectivity to the line card paths. Each IOM2 forwarding complex has 2 paths, so the paths of the IOM2 in slot 5 are using planes 0 and 1, and 3 and 4. Note that there are now fewer planes available and more paths, so there is more sharing of planes between paths than when chassis mode **d** was used.

## IMPM Managed Traffic

IMPM manages IPv4/IPv6 routed multicast traffic and VPLS (IGMP and PIM) snooped IPv4 multicast traffic, traffic that matches a  $\langle *,G \rangle$  or a  $\langle S,G \rangle$  multicast record in the ingress forwarding table. It manages IP multicast traffic on a bud LSR when using point-to-multipoint (P2MP) LSPs but it does not manage IP protocol control traffic or traffic using multipoint-shared queuing. Traffic being managed by IMPM involves IMPM monitoring and potentially moving the related channels between paths/planes. The unmanaged traffic rates are also monitored and taken into account in the IMPM algorithm.

## IMPM Managed Traffic

Care should be taken when using the mrouter-port configuration in a VPLS service. This creates a (\*,\*) multicast record and consequently all multicast channels that are not delivered locally to a non-mrouter port will be treated by IMPM as a single channel.



## Configuration

This section covers:

- [IMPM on an IOM3-XP/IMM on page 2416](#)
  - [IMPM on an IOM1/2 on page 2426](#)
  - [IMPM Not Enabled on page 2429](#)
- 

## Prerequisites

As IMPM operates on IPv4/IPv6 routed or VPLS IGMP/PIM snooped IPv4 multicast traffic, some basic multicast configuration must be enabled. This section uses routed IP multicast in the global routing table which requires IP interfaces to be configured with PIM and IGMP. The configuration uses a PIM rendezvous point and static IGMP joins. The following is an example of the complete configuration of one interface.

```
configure
router
  interface "int-IOM3-1"
    address 172.16.6.254/24
    port 6/2/1
  exit
  igmp
    interface "int-IOM3-1"
      static
        group 239.255.0.1
        starg
      exit
    exit
  exit
  no shutdown
exit
pim
  interface "int-IOM3-1"
  exit
  rp
    static
      address 192.0.2.1
      group-prefix 239.255.0.0/16
    exit
  exit
  exit
  no shutdown
exit
exit
```

## Prerequisites

One interface is configured on each line card configured in the system, as shown in the following output, but omitting their IGMP and PIM configuration.

```
configure
  router
    interface "int-IMM8"
      address 172.16.3.254/24
      port 7/2/1
    exit
    interface "int-IOM2"
      address 172.16.1.254/24
      port 5/2/1
    exit
    interface "int-IOM3-1"
      address 172.16.2.254/24
      port 6/2/1
    exit
    interface "int-IOM3-2"
      address 172.16.4.254/24
      port 8/2/1
    exit
  exit
```

## Configuring IMPM

The majority of the IMPM configuration is performed under the **mcast-management** CLI nodes and consists of:

- a. The bandwidth-policy for characteristics relating to the IOM/IMM paths. This is applied on an IOM3-XP/IMM fp<sup>2</sup>, or an IOM1/2 MDA, under ingress mcast-path-management, with a bandwidth-policy named default being applied by default.

→ IOM1/2

```
config# card slot-number mda mda-slot
      ingress
          mcast-path-management
              bandwidth-policy policy-name
```

→ IOM3-XP/IMM

```
config# card slot-number fp [1]
      ingress
          mcast-path-management
              bandwidth-policy policy-name
```

- b. The multicast-info policy for information related to the channels and how they are handled by the system. To facilitate provisioning, parameters can be configured under a three level hierarchy with each level overriding the configuration of its predecessor:

→ Bundle: a group of channels

→ Channel: a single channel or a non-overlapping range of channels

→ Source-override: channels from a specific sender

```
config# mcast-management multicast-info-policy policy-name [create]
      bundle bundle-name [create]
          channel ip-address [ip-address] [create]
              source-override ip-address [create]
```

This policy is applied where the channel enters the system, so under router or service (vpls or vprn); the latter allows the handling of channels to be specific to a service, even if multiple services use overlapping channel addresses.

```
config# router multicast-info-policy policy-name
```

```
config# service vpls service-id multicast-info-policy policy-name
```

```
config# service vprn service-id multicast-info-policy policy-name
```

A default multicast-info-policy is applied to the above when IMPM is enabled.

- c. The chassis-level node configures the information relating to the switch fabric planes.

```
config# mcast-management chassis-level
```

---

2. fp is the system term for a forwarding complex on an IOM3-XP/IMM.

In addition, the command `hi-bw-mcast-src` (under an IOM3-XP/IMM fp or an IOM1/2 MDA) can be used to control the path to plane mapping among forwarding complexes.

## IMPM on an IOM3-XP/IMM

IMPM is enabled on IOM3-XP/IMMs on under the card/fp CLI node as follows

```
config# card slot-number fp 1 ingress mcast-path-management no shutdown
```

## IOM3-XP/IMM Paths

16 paths are available on an IOM3-XP/IMM when IMPM is enabled which can be either primary paths or secondary paths. By default the 16 paths are divided into 15 primary paths and 1 secondary path, as can be seen using the following command with IMPM enabled only on slot 6 (this corresponds to the plane assignment in Output 1):

```
*A:PE-1# tools dump mcast-path-mgr cpm
McPathMgr[6][0]: 0xf33b0a00
PATH:
                                         PLANE:
Type SGs      InUseBW  AvailBW  TotalBw  ID  SGs  InUseBW  AvailBW  TotalBw
P      1          0          -          -   1    1          0  2000000  2000000
P      1          0          -          -   0    1          0  2000000  2000000
P      1          0          -          -   3    1          0  2000000  2000000
P      1          0          -          -   4    1          0  2000000  2000000
P      1          0          -          -   5    1          0  2000000  2000000
P      1          0          -          -   6    1          0  2000000  2000000
P      1          0          -          -   7    1          0  2000000  2000000
P      1          0          -          -   8    1          0  2000000  2000000
P      1          0          -          -   9    1          0  2000000  2000000
P      1          0          -          -  10    1          0  2000000  2000000
P      1          0          -          -  11    1          0  2000000  2000000
P      1          0          -          -  12    1          0  2000000  2000000
P      1          0          -          -  13    1          0  2000000  2000000
P      1          0          -          -  14    1          0  2000000  2000000
P      1          0          -          -  15    1          0  2000000  2000000
S      1          0          -          -  16    1          0  1800000  1800000
B      0          0          -          -   -    -          -    -          -
*A:PE-1#
```

**Output 3: Paths/Planes on IOM3-XP/IMM**

The left side of the output displays information about the paths (type {P=primary, s=secondary or B=blackholed}, number of “S,G”s and bandwidth in use (the path bandwidth cannot be set on an IOM3-XP/IMM, so the path available and total bandwidth always shows “-“)) and the right side displays similar information about the associated planes (this will be a combination of the information for all paths connected to this plane). Note that one SG is always present on each path; this is used by the system and relates to the unmanaged traffic.

The primary/secondary paths are also highlighted in the planes section of Output 1, the primary paths being connected to the planes on the left of the “|” and the secondary paths to its right. There is a default primary path and a default secondary path; these correspond to the left-most plane and right-most plane for each line card, respectively.

Primary paths are used by:

- Expedited IES, VPLS and VPRN service ingress non-managed multipoint traffic (using the SAP based queues). This uses the default primary path.
- Expedited network ingress non-managed multipoint traffic (using the network interface queues). This uses the default primary path.
- Managed multicast explicit path primary channels (using the primary paths managed multipoint queue)
- All managed multicast dynamic path channels when the primary paths or multicast planes are not at their limit (using the primary paths managed multipoint queue)
- Highest preference managed multicast dynamic path channels when the primary paths or multicast planes are at their limit (using the primary paths managed multipoint queue)
- Non-managed P2MP LSP IP multicast traffic. This does not require IMPM to be enabled, so is discussed later in [IMPM Not Enabled on page 2429](#).
- Non-managed expedited ingress policed multipoint traffic. This does not require IMPM to be enabled, so is discussed in [IMPM Not Enabled on page 2429](#).

Secondary paths are used by:

- Best-Effort IES, VPLS and VPRN service ingress non-managed multipoint traffic (using the SAP based queues). This uses the default secondary path.
- Best-Effort network ingress non-managed traffic (using the network interface multipoint queues). This uses the default secondary path.
- Managed multicast explicit path secondary channels (using the secondary paths managed multipoint queue)
- Lower preference managed multicast dynamic path channels when the primary paths or multicast planes are at their limit (using the secondary paths managed multipoint queue)
- Non-managed best-effort ingress policed multipoint traffic. This does not require IMPM to be enabled, so is discussed in [IMPM Not Enabled on page 2429](#).

When IMPM is enabled, the managed traffic does not use the standard multipoint queues but instead is placed onto a separate set of shared queues which are associated with the 16 paths. These queues are instantiated in an access ingress pool (called MC Path Mgmt, see “show output” section) which exists by default (this pool can be used even when IMPM is not enabled – see section “IMPM not enabled”). Statistics relating to traffic on these queues are reflected back to the standard ingress multipoint queues for accounting and troubleshooting purposes. Note that non-

managed traffic continues to use the standard ingress multipoint queues, with the exception of P2MP LSP IP multicast traffic and policed multipoint traffic.

The size of the pool by default is 10% of the total ingress pool size, the reserved CBS is 50% of the pool and the default slope policy is applied. Care should be taken when changing the size of this pool as this would affect the size of other ingress pools on the line card.

```
config# mcast-management bandwidth-policy policy-name [create]
      mcast-pool percent-of-total percent-of-buffers
              resv-cbs percent-of-pool
              slope-policy policy-name
```

It is possible to configure the parameters for the queues associated with both the primary and secondary paths, and also the number of secondary paths available, within the bandwidth-policy.

```
config# mcast-management bandwidth-policy policy-name create
      t2-paths
        primary-path
          queue-parameters
            cbs percentage
            hi-priority-only percent-of-mbs
            mbs percentage
        secondary-path
          number-paths number-of-paths [dual-sfm number-of-paths]
          queue-parameters
            cbs percentage
            hi-priority-only percent-of-mbs
            mbs percentage
```

The number of primary paths is 16 minus the number of secondary paths (at least one of each type must exist). The number-paths parameter specifies the number of secondary paths when only one switch fabric is active, while the dual-sfm parameter specifies the same value when two switch fabrics are active.

Packets are scheduled out of the path/multicast queue as follows:

- Traffic sent on primary paths is scheduled at multicast high priority while that on secondary paths is scheduled at multicast low priority.
- For managed traffic, the standard ingress forwarding class/prioritization is not used, instead IMPM managed traffic prioritization is based on a channel's preference (described in [Channel Prioritization and Blackholing Control on page 2425](#)). Egress scheduling is unchanged.

Congestion handling (packet acceptance into the path/multicast queue):

- For non-managed traffic, this is based on the standard mechanism, namely the packet's enqueueing priority is used to determine whether the packet is accepted into the path multipoint queue depending on the queue mbs/cbs and the pool shared-buffers/reserved-buffers/WRED.

For managed traffic, the congestion handling is based upon the channel's preference (described later) and the channel's cong-priority-threshold which is configured in the multicast-info-policy (here under a bundle).

```
config# mcast-management multicast-info-policy policy-name [create]
      bundle bundle-name [create]
          cong-priority-threshold preference-level
```

When the preference of a channel is lower than the cong-priority-threshold setting, the traffic is treated as low enqueueing priority, when it is equal to or higher than the cong-priority-threshold it is treated as high enqueueing priority. The default cong-priority-threshold is 4.

---

### IOM3-XP/IMM Planes

The capacity per plane for managed traffic is by default 2Gbps for a primary path and 1.8Gbps for a secondary path. The logic behind a reduced default on the secondary is to leave capacity for new streams in case the default secondary is fully used by managed streams.

The plane capacities can be configured as follows, note that this command configures the plane bandwidth associated with primary/secondary paths as seen by each line card, the TotalBw on the right side of Output 3:

```
config# mcast-management chassis-level
      per-mcast-plane-limit megabits-per-second [secondary megabits-per-second]
          [dual-sfm megabits-per-second [secondary-dual-sfm megabits-per-second]]
```

The first parameter defines the capacity for a primary path, the second a secondary path and the dual-sfm configures these capacities when two switch fabrics are active. The maximum plane capacity is 4Gbps but for the release used here it should only be configured on 7750 SR-12 or 7450 ESS-12 systems populated with SF/CPM4(s) and 100G FP IMMs; for all other hardware combinations the maximum should be 2Gbps. Note that secondary plane capacity cannot be higher than that of the primary plane.

These values can be tuned to constrain the amount of managed multicast traffic in favour of non-managed multicast and unicast traffic.

On the IOM3-XP/IMM line cards there is no separate control of the line card path capacity, the capacity is only constrained by the plane.

---

### IOM3-XP/IMM Path to Plane Mapping

By default all fps (line cards for IOM3-XP/IMM) are configured into the default (zero) group as seen in Output 1 and the system distributes the paths as equally as possible over the available

planes. This default works well if there is a low volume of multicast traffic (compared to the plane capacity), or if there is a higher volume multicast entering only one line card where the ingress capacity does not exceed that provided by the planes the line card is connected to.

If there are more paths than planes and, for example, there is a high bandwidth multicast channel entering two different line cards it could happen that both line cards select the same plane for two paths that are used. This would result in one of the channels being blackholed if the plane capacity is exceeded, effectively reducing the available multicast capacity from that line card. In order to avoid this situation, it is possible to configure the paths in to different groups and the system will attempt to use different planes for each group.

Output 1 and Output 2 show examples of how the paths are mapped to planes.

In both cases there are more paths than planes so some planes are shared by multiple paths. The following command provides control of this mapping.

```
config# card slot-number fp [1]
      hi-bw-mcast-src [alarm] [group group-id] [default-paths-only]
```

If an fp is configured into a non-zero group (range: 1 to 32), the system will attempt to assign dedicate planes to its paths compared to other line cards in different non-zero groups. This action is dependent on there being sufficient planes available. If two line cards are assigned to the same group, they will be assigned the same planes. The **default-paths-only** parameter performs the assignment optimization only for the default primary and secondary paths and is only applicable to IOM3-XP/IMMs. The alarm keyword causes an alarm to be generated if some planes are still shared with fps in a different group.

An example of the use of this command is shown later.

Note: When VPLS IGMP and PIM snooped traffic is forwarded to a spoke or mesh SDP, by default it is sent by the switch fabric to all line card forwarding complexes on which there is a network IP interface. This is due to the dynamic nature of the way that a spoke or mesh SDP is associated with one or more egress network IP interfaces. If there is an active spoke/mesh SDP for the VPLS service on the egress forwarding complex, the traffic will be flooded on that spoke/mesh SDP, otherwise it will be dropped on the egress forwarding complex. This can be optimized by configuring an inclusion list under the spoke or mesh SDP defining which MDAs this traffic should be flooded to.

```
config>service>vpls# [spoke-sdp|mesh-sdp] sdp-id:vc-id egress
      mfib-allowed-mda-destinations
      [no] mda mda-id
```

The switch fabric flooding domain for this spoke or mesh SDP is made up only of the MDAs that have been added to the list. An empty list implies the default behavior.

It is important to ensure that the spoke or mesh SDP can only be established across the MDAs listed, for example by using RSVP with an explicit path.



## IMPM Operation on IOM3-XP/IMM

This section covers:

- [Principle of Operation on page 2421](#)
- [Monitoring Traffic Rates on page 2422](#)
- [Channel Prioritization and Blackholing Control on page 2425](#)

### Principle of Operation

Where IMPM is enabled, it constantly monitors the line cards for ingress managed traffic.

When a new channel arrives it will be placed by default on to the default secondary path. IMPM determines the ingress point for the channel and then monitors the traffic of the channel within its monitoring period in order to calculate the rate of the channel. The system then searches the multicast paths/planes attached to the line card for available bandwidth. If there is sufficient capacity on such a path/plane, the channel is moved to that plane. Planes corresponding to primary paths are used first, when there is no capacity available on any primary path/plane a secondary path/plane is used (unless the channel is explicitly configured onto a specific path type – see the following description).

If the required bandwidth is unavailable, the system will then look for any channels that ingress this or other line cards that could be moved to a different multicast plane in order to free up capacity for the new channel. Any channel that is currently mapped to a multicast plane available to the ingress line card is eligible to be moved to a different multicast plane.

If an eligible existing channel is found, whether on this or another line card, that existing channel is moved without packet loss to a new multicast plane. If necessary, this process can be repeated resulting in multiple channels being moved. The new multicast channel is then mapped to the multicast plane previously occupied by the moved channels, again this normally is using a primary path.

If no movable channel is found, then lower preference channel(s) on any ingress line card that share multicast planes with the ingress line card of the new channel can be blackholed to free up capacity for the new channel. It is also possible to both blackhole some channels and move other channels in order to free up the required capacity. If no lower preference channel is found and no suitable channel moves are possible, the new channel will be blackholed.

If required, channels can be explicitly configured to be on either a primary or secondary path. This can be done for a bundle of channels, for example

```
config# mcast-management multicast-info-policy policy-name [create]
      bundle bundle-name [create]
          explicit-sf-path {primary|secondary|ancillary}
```

Note that the ancillary path is not applicable to the IOM3-XP/IMM line cards, however, it is discussed in the section relating to the IOM1/2. If a channel on an IOM3-XP/IMM is configured onto the ancillary path it will use a primary path instead.

One secondary path on an IOM3-XP/IMM is used as a default startup path for new incoming channels. If a large amount of new channel traffic could be received within the monitoring period, it is possible that the plane associated with the default secondary path is over loaded before IMPM has time to monitor the channels' traffic rate and move the channels to a primary path (and a plane with available capacity). This can be alleviated by configuring the following command:

```
config# mcast-management chassis-level
      round-robin-inactive-records
```

When round-robin-inactive-records is enabled, the system redistributes new channels (which are referenced by inactive S,G records) among all available line card multicast (primary, secondary) paths and their switch fabric planes.

---

## Monitoring Traffic Rates

The monitored traffic rate is the averaged traffic rate measured over a monitoring period. The monitoring period used depends on the total number of channels seen by IMPM, the minimum is a 1 second interval and the maximum a 10 seconds interval.

The way in which the system reacts to the measured rate can be tuned using the following command:

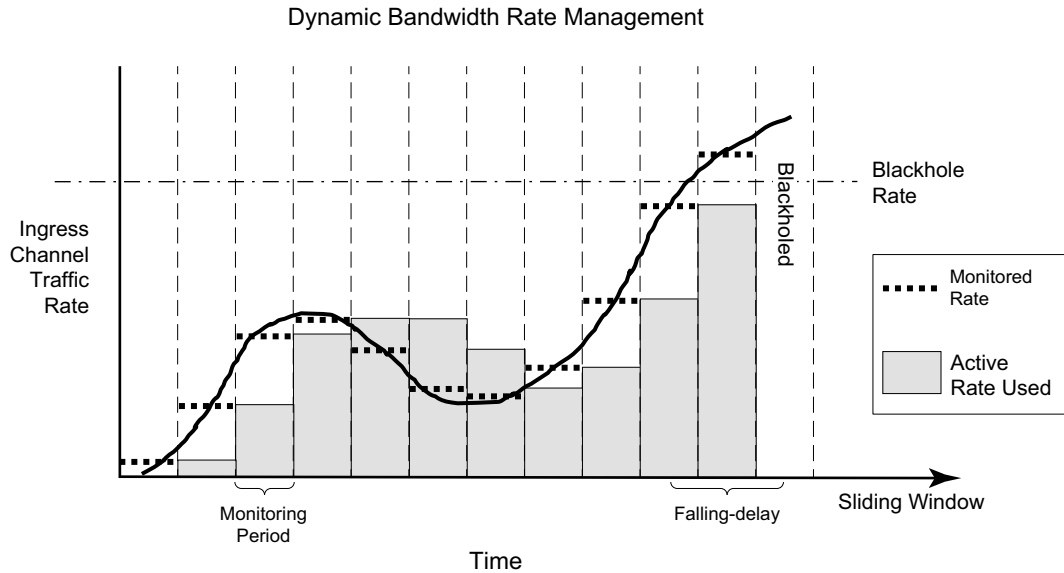
```
config# mcast-management multicast-info-policy policy-name [create]
      bundle bundle-name [create]
      bw-activity {use-admin-bw|dynamic [falling-delay seconds]}
                  [black-hole-rate kbps]
```

The default is to use the dynamic bandwidth rate, in which case a channel's active traffic rate is determined based on the measured monitored rates. IMPM then makes a decision of how to process the channel as follows.

If the channel was un-managed, IMPM will attempt to place the channel on a path/plane with sufficient available bandwidth.

If the channel was already managed, IMPM determines the highest monitored traffic rate (within a given monitoring period) in a sliding window defined by the falling-delay. This highest chosen monitored rate is then used to re-assess the placement of the channel on the path/planes; this may cause IMPM to move the channel. This mechanism prevents the active rate for a channel being reduced due to a momentarily drop in traffic rate. The default value for falling-delay is 30 seconds, with a range of 10-3600 seconds.

The above logic is shown in [Figure 372](#) (for simplicity, the falling-delay is exactly twice the monitoring period). It can be seen that the active rate used when the traffic rate decreases follows the highest monitored rate in any falling-delay period.



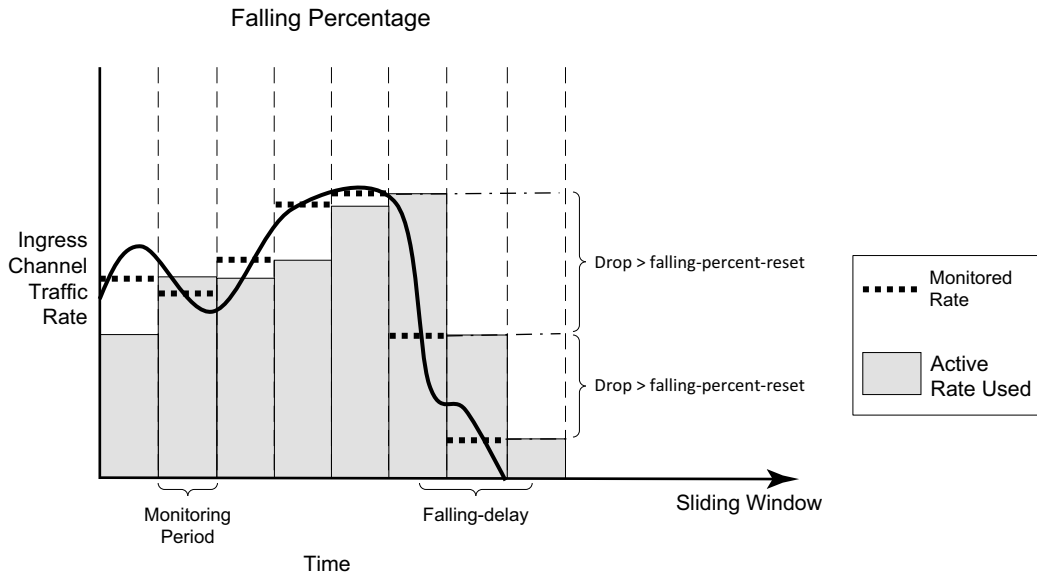
OSSG726

**Figure 372: Dynamic Bandwidth Rate Management**

By using the sliding window of monitored rate measurements in the dynamic bandwidth measurement mode, IMPM delays releasing capacity for a channel in its calculations when the channel's rate has been reduced. This allows IMPM to ignore temporary fluctuations in a channel's rate. It is possible to tune this for cases where the reduction in a channel's rate is large by using the falling-percent-reset parameter. The default for the falling-percent-reset is 50%. Setting this to 100% effectively disables it.

```
config# mcast-management bandwidth-policy policy-name create
      falling-percent-reset percent-of-highest
```

When the monitored rate falls by a percentage which is greater or equal to falling-percent-reset, the rate used by IMPM is immediately set to this new monitored rate. This allows IMPM to react faster to significant reductions in a channel's rate while at the same time avoiding too frequent reallocations due to normal rate fluctuations. An example of the falling-percent-reset is shown in [Figure 373](#). In the last two monitoring periods, it can be seen that the active rate used is equal to the monitored rate in the previous periods, and not the higher rate in the previous falling-delay window.



OSSG727

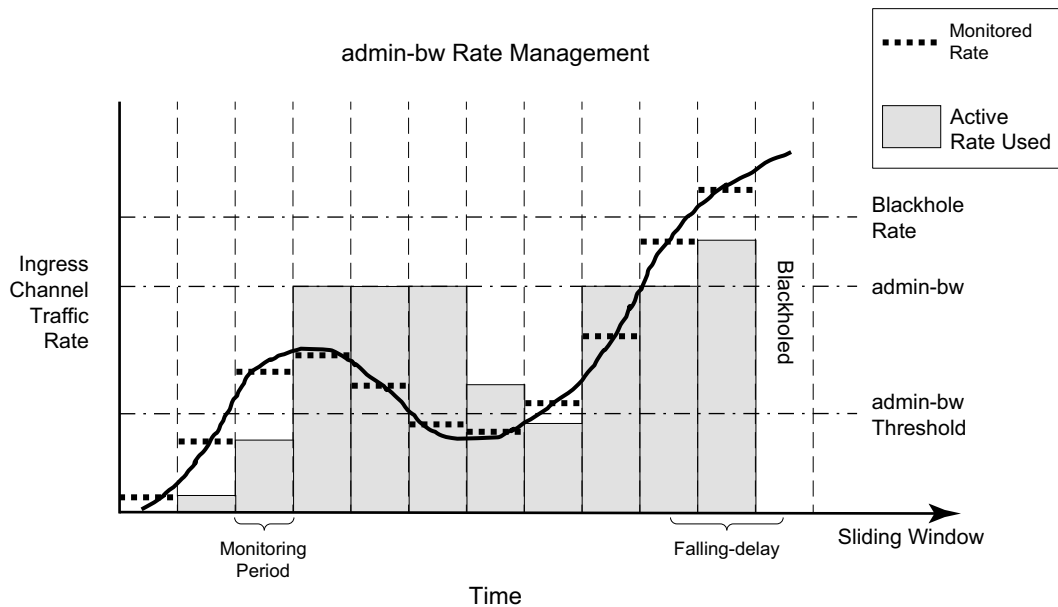
**Figure 373: Falling-Percent-Reset**

The rate management can be further tuned based on the expectation that the channel bandwidth will fluctuate around a given rate. When the bw-activity is set to use-admin-bw within the multicast-info-policy, the following parameters come into play.

```
config# mcast-management multicast-info-policy policy-name [create]
    bundle bundle-name [create]
    admin-bw kbps

config# mcast-management bandwidth-policy policy-name create
    admin-bw-threshold kilo-bits-per-second
```

IMPM will use the rate configured for the admin-bw if the monitored rate is above the admin-bw-threshold but below or equal to the admin-bw in the sliding window of the falling-delay. Whenever the monitored rate is below the admin-bw-threshold or above the admin-bw, IMPM uses the dynamic rate management mechanism. The admin-bw-threshold needs to be smaller than the admin-bw, with the latter being non-zero. This is shown in [Figure 374](#) (for simplicity, the falling-delay is exactly twice the monitoring period). It can be seen that while the monitored rate stays between the admin-bw-threshold and the admin-bw, the active rate used is set to the admin-bw.



OSSG728

**Figure 374: Admin-Bw Rate Management**

Finally, IMPM also takes into consideration the unmanaged traffic rate on the primary and secondary paths associated with SAP/network interface queues when determining the available capacity on these paths/planes. This is achieved by constantly monitoring the rate of this traffic on these queues and including this rate in the path/plane capacity usage. IMPM must be enabled on the ingress card of the unmanaged traffic otherwise it will not be monitored.

## Channel Prioritization and Blackholing Control

IMPM decides which channels will be forwarded and which will be dropped based on a configured channel preference. The preference value can be in the range 0-7, with 7 being the most preferred and the default value being 0.

When there is insufficient capacity on the paths/planes to support all ingress multipoint traffic, IMPM uses the channel preferences to determine which channels should be forwarded and which should be blackholed (dropped).

This is an important distinction compared to the standard forwarding class packet prioritization; by using a channel preference, an entire channel is either forwarded or blackholed, this allows IMPM to avoid congestion having a small impact on multiple channels at the cost of entire channels being lost.

## Configuring IMPM

The channel preference is set within the multicast-info-policy, for example at the bundle level, with the settable values being 1-7:

```
config# mcast-management multicast-info-policy policy-name [create]
      bundle bundle-name [create]
      preference preference-level
```

The channel preference is also used for congestion control in the line card path queues – see “congestion handling” in the section on “IOM3-XP/IMM Paths” above.

Blackhole protection can also be enabled using the bw-activity command, shown above in the “Monitoring traffic rates” section. Regardless of which rate monitoring mechanism is used, a channel can be blackholed if the monitored rate exceeds the black-hole-rate, in which case the channel will be put immediately on the blackhole list and its packets dropped at ingress. This channel will no longer consume line card path or switch fabric capacity. The intention of this parameter is to provide a protection mechanism in case channels consume more bandwidth than expected which could adversely affect other channels.

The black-hole-rate can range from 0 to 40000000kbps, with no black-hole-rate by default. This protection is shown in the last monitoring period of both Figure 2 and Figure 4. Note that it will take a falling-delay period in which the channel’s rate is always below the black-hole-rate in order for the channel to be re-instated unless the reduction in the rate is above the falling-percent-reset.

---

### IMPM on an IOM1/2

As most of the principles when using IMPM on an IOM1/2 compared to on an IOM3-XP/IMM are the same and are described above, this section focuses only on the difference between the two.

Note that an IOM1 and IOM2 have two independent 10G forwarding complexes; in both cases there is a single MDA per forwarding complex, consequently some aspects of IMPM are configured under the mda CLI node.

IMPM is enabled on an IOM1/2 under the MDA CLI node as follows:

```
config# card slot-number mda mda-slot ingress mcast-path-management no shutdown
```

---

### IOM1/2 Paths

Each forwarding complex has three paths: one primary and one secondary path, and another type called the ancillary path which is IOM1/2 specific. The paths can be seen using the following command with IMPM enabled only on slot 5 MDA 1 and MDA 2 (referenced as [0] and [1] respectively):

```
A:PE-1# tools dump mcast-path-mgr cpm
```

```

McPathMgr[5][0]: 0xf33b0a00
PATH:
Type SGs      InUseBW  AvailBW  TotalBw  ID  SGs  InUseBW  AvailBW  TotalBw
P      1          0  2000000  2000000  1    1          0  2000000  2000000
S      1          0  1500000  1500000  0    1          0  1800000  1800000
A      0          0  5000000  5000000  -    -          -    -          -
B      0          0          -          -    -    -          -    -          -
McPathMgr[5][1]: 0xf33b3198
PATH:
Type SGs      InUseBW  AvailBW  TotalBw  ID  SGs  InUseBW  AvailBW  TotalBw
P      1          0  2000000  2000000  4    1          0  2000000  2000000
S      1          0  1500000  1500000  3    1          0  1800000  1800000
A      0          0  5000000  5000000  -    -          -    -          -
B      0          0          -          -    -    -          -    -          -
A:PE-1#

```

### Output 4: Paths/Planes on IOM1/2

The primary and secondary paths function as on the IOM3-XP/IMM, specifically for:

- Traffic usage.
- Associated queues instantiated in the ingress “MC Path Mgmt” ingress pool.
- Packet scheduling.
- Congestion handling.

The queue parameters can be configured within the bandwidth-policy in a similar way to the IOM3-XP/IMM (note that the IOM3-XP/IMM equivalent for this is under the t2-paths CLI node). The bandwidth-policy is then applied under the respective MDA.

```

config# mcast-management bandwidth-policy policy-name create
  primary-path
    queue-parameters
      cbs percentage
      hi-priority-only percent-of-mbs
      mbs percentage
  secondary-path
    queue-parameters
      cbs percentage
      hi-priority-only percent-of-mbs
      mbs percentage

```

The IOM1/2 allows capacity control on the paths themselves, which is not possible on the IOM3-XP/IMM. This is achieved using the following commands.

```

config# mcast-management bandwidth-policy policy-name create
  primary-path
    path-limit megabits-per-second
  secondary-path
    path-limit megabits-per-second

```

## Configuring IMPM

The maximum path-limit for both the primary and secondary path is 2Gbps with a default of 2Gbps for the primary path and 1.5Gbps for the secondary path. The capability to set a path limit for the IOM1/2 can be seen when comparing Output 3 with Output 4; in the latter the “AvailBW” and “TotalBw” for the “PATH” shows the path limit.

In addition to setting the path limits in the bandwidth-policy, they can also be overridden on a given MDA.

```
config# card slot-number mda mda-slot
      ingress
          mcast-path-management
              primary-override
                  path-limit megabits-per-second
              secondary-override
                  path-limit megabits-per-second
```

The achievable capacity will be the minimum of the path’s path-limit and the plane’s per-mcast-plane-limit.

### The ancillary path

The ancillary path allows managed multicast traffic to be forwarded through the switch fabric as unicast and so is not constrained to the path or plane capacities. This is achieved using ingress replication, in order to send a channel to multiple destination forwarding complexes (DFCs), the ingress forwarding complex creates and forwards one copy of each packet to each DFC connected to the switch fabric.

However, the total replication capacity available for the ancillary path is constrained to 5G to prevent it impacting the unicast or primary/secondary path capacities. This means that the total amount of ancillary capacity usable can be calculated from (note that the first copy sent is not included in this capacity, hence the “-1”):

$$5\text{Gbps}/(\text{number\_of\_switch\_fabric\_DFCs} - 1)$$

Taking an example shown later, if some channels ingress an IOM2 and egress 2 IOM3-XP (1 DFC each) and 1 IMM8 (1 DFC) to give a total of 3 egress DFCs, then total ancillary capacity available is

$$5\text{Gbps}/(3-1) = 2.5\text{Gbps}.$$

This would allow, for example, approximately 250 channels at 10Mbps each to use the ancillary path.

Due to the relationship between ancillary capacity and number of DFCs, the system will prefer the ancillary path as default whenever a channel enters an IOM1/2 and egresses on up to 3 DFCs. If there are 4 or more egress DFCs for the channel, then the primary path is preferred. The determination is performed on a per channel basis.



The configuration parameters relating to the primary and secondary paths are also available for the ancillary path.

```
config# mcast-management bandwidth-policy policy-name create
  ancillary-path
    queue-parameters
      cbs percentage
      hi-priority-only percent-of-mbs
      mbs percentage

config# mcast-management bandwidth-policy policy-name create
  ancillary-path
    path-limit megabits-per-second

config# card slot-number mda mda-slot
  ingress
    mcast-path-management
      ancillary-override
        path-limit megabits-per-second
```

---

### IOM1/2 Planes

The capacity per plane for managed traffic is by default 2Gbps for a primary path and 1.8Gbps for a secondary path. Note that the default IOM1/2 secondary path limit is 1.5Gbps. A maximum of 2Gbps should be configured for either path type using the per-mcast-plane-limit (as shown for the IOM3-XP/IMM) when an IOM1/2 is being used with IMPM.

As the ancillary path does not use the switch fabric planes, there is no associated plane limit.

---

### IOM1/2 Path to Plane Mapping

The hi-bw-mcast-src command function is the same for IOM1/2 line cards as for IOM3-XP/IMM line cards, as described above.

---

### IMPM Operation on IOM1/2

This is exactly the same as the operation as for the IOM3-XP/IMM, see above.

---

### IMPM Not Enabled

When IMPM is not enabled most multipoint traffic on an IOM1/2 and IOM3-XP/IMM can use only one primary path and one secondary path per forwarding complex. When ingress multipoint arrives it is placed on a multipoint queue and these queues are connected either to a primary path (if the queue is expedited) or a secondary path (if the queue is best-effort) depending on the

ingress QOS classification applied. Standard ingress forwarding class/scheduling prioritization is used.

The capacity of the primary and secondary paths is 2Gbps, unless the system is a 7750 SR-12 or 7450 ESS-12 populated with SF/CPM4(s) and 100G FP IMMs in which case the capacity is 4Gbps.

In Output 1 and Output 2, the primary path is associated with the left-most plane and the secondary path is associated with the right-most plane for each line card.

There are exceptions to the above on the IOM3-XP/IMM line cards for

- Point-to-multipoint LSP IP multicast traffic
- Policed ingress routed IP multicast or VPLS broadcast, unknown or multicast traffic

---

### Point-to-Multipoint (P2MP) LSP IP Multicast Traffic

IMPM will manage traffic on a P2MP LSP for any IP multicast channel that is delivered locally, for example, the system is a bud LSR. However, non-managed P2MP LSP IP multicast traffic will also make use of the primary paths, regardless of whether IMPM is enabled or not.

For each primary queue created in the MC IMPM Path pool, an additional queue is created to carry non-managed P2MP LSP IP multicast traffic. The non-managed P2MP LSP IP multicast traffic is automatically distributed across all primary paths based on a modulo N function of the 10 least significant bits of channel destination group address, where N is the number of primary paths. Note that the number of primary paths can be changed with IMPM enabled or disabled by applying a bandwidth-policy which sets the number of secondary paths.

---

### Policed Ingress Routed IP Multicast or VPLS Broadcast, Unknown Or Multicast Traffic

Routed IP multicast or VPLS broadcast, unknown or multicast traffic passing through ingress hardware policers on the IOM3-XP/IMM can also use the IMPM managed queues, with IMPM enabled or disabled. If this traffic is best-effort (forwarding classes BE, L2, AF, L1) it will use the secondary paths, if it is expedited (forwarding classes H2, EF, H1, NC) it will use the primary paths. Note that this traffic uses the shared ingress policer-output-queues which have a fixed forwarding class to queue mapping).

When IMPM is not enabled, this traffic is non-managed and 1 secondary path plus 15 primary paths are available (the default). Consequently, extra capacity is only available for the expedited traffic, which could use up to 15 planes worth of switch fabric capacity. If extra capacity is required for best-effort traffic, a bandwidth-policy allocating more secondary paths can applied to the line card even without IMPM being enabled.

The policed ingress routed IP or VPLS broadcast, unknown or multicast traffic is distributed across the paths using a standard LAG hash algorithm (as described in the LAG and ECMP Hashing section in the Interface Configuration guide).

For both of these exceptions, it is recommended to reduce the managed traffic primary/secondary plane limits (using per-mcast-plane-limit) in order to allow for the non-managed traffic.

## Show Output

This section includes the show output related to IMPM. The first part covers generic output and uses IOM3-XP/IMMs and chassis mode **d**. The second part includes an IOM2 and so uses chassis mode **a**.

### IOM3-XP/IMM and Generic Output

The system has an IOM3-XP in slots 6 and 8, with an IMM8 is slot 7.

```
A:PE-1# show card
=====
Card Summary
=====
```

Slot	Provisioned Card-type	Equipped Card-type	Admin State	Operational State	Comments
6	iom3-xp	iom3-xp	up	up	
7	imm8-10gb-xfp	imm8-10gb-xfp	up	up	
8	iom3-xp	iom3-xp	up	up	
A	sfm4-12	sfm4-12	up	up/active	
B	sfm4-12	sfm4-12	up	up/standby	

```
=====
A:PE-1#
```

The status of IMPM on a given card can be shown as follows:

```
*A:PE-1# show card 6 detail
=====
Card 6
=====
```

Slot	Provisioned Card-type	Equipped Card-type	Admin State	Operational State	Comments
6	iom3-xp	iom3-xp	up	up	

```
FP 1 Specific Data
  hi-bw-mc-srcEgress Alarm      : 2
  hi-bw-mc-srcEgress Group     : 0
  mc-path-mgmt Admin State     : In Service
  Ingress Bandwidth Policy     : default
```

IMPM is enabled on the fp, it is using the default bandwidth-policy and is using the default hi-bw-mcast-src group (0).

The MC Path Mgmt pool is created by default with the default settings.

```
*A:PE-1# show pools 6/1
=====
Type      Id      App.      Pool Name      Actual ResvCBS  PoolSize
Admin ResvCBS
-----
MDA       6/1     Acc-Ing  MC Path Mgmt  18816          37632
50%
=====
*A:PE-1#
```

The default bandwidth-policy can be shown, giving the default parameters for the MC Path Pool and the associated queues.

```
*A:PE-1# show mcast-management bandwidth-policy "default" detail
=====
Bandwidth Policy Details
=====
Policy          : default
-----
Admin BW Thd    : 10 kbps          Falling Percent RST: 50
Mcast Pool Total : 10              Mcast Pool Resv Cbs: 50
Slope Policy    : default
Primary
Limit          : 2000 mbps      Cbs              : 5.00
Mbs            : 7.00             High Priority     : 10
Secondary
Limit          : 1500 mbps      Cbs              : 30.00
Mbs            : 40.00          High Priority     : 10
Ancillary
Limit          : 5000 mbps     Cbs              : 65.00
Mbs            : 80.00          High Priority     : 10
T2-Primary
Cbs            : 5.00             Mbs              : 7.00
High Priority   : 10
T2-Secondary
Cbs            : 30.00          Mbs              : 40.00
High Priority   : 10          Paths(Single/Dual) : 1/1
=====
Bandwidth Policies : 1
=====
*A:PE-1#
```

The defaults for the multicast-info-policy can be seen in configuration mode.

```
*A:PE-1# configure mcast-management
*A:PE-1>config>mcast-mgmt# info detail
-----
multicast-info-policy "default" create
no description
bundle "default" create
```

## IOM3-XP/IMM and Generic Output

```

no cong-priority-threshold
no description
no ecmp-opt-threshold
no admin-bw
no preference
no keepalive-override
no explicit-sf-path
bw-activity dynamic falling-delay 30
no primary-tunnel-interface
exit
exit

```

The paths/planes on an IOM3-XP/IMM can be shown here for card 6.

```

*A:PE-1# tools dump mcast-path-mgr cpm
McPathMgr[6][0]: 0xf33b0a00
PATH:

```

PATH:				PLANE:					
Type	SGs	InUseBW	AvailBW	TotalBw	ID	SGs	InUseBW	AvailBW	TotalBw
P	1	0	-	-	4	1	0	2000000	2000000
P	1	0	-	-	3	1	0	2000000	2000000
P	1	0	-	-	5	1	0	2000000	2000000
P	1	0	-	-	6	1	0	2000000	2000000
P	1	0	-	-	7	1	0	2000000	2000000
P	1	0	-	-	8	1	0	2000000	2000000
P	1	0	-	-	9	1	0	2000000	2000000
P	1	0	-	-	10	1	0	2000000	2000000
P	1	0	-	-	11	1	0	2000000	2000000
P	1	0	-	-	12	1	0	2000000	2000000
P	1	0	-	-	13	1	0	2000000	2000000
P	1	0	-	-	14	1	0	2000000	2000000
P	1	0	-	-	15	1	0	2000000	2000000
P	1	0	-	-	16	1	0	2000000	2000000
P	1	0	-	-	0	1	0	2000000	2000000
S	1	0	-	-	1	1	0	1800000	1800000
B	0	0	-	-	-	-	-	-	-

```

*A:PE-1#

```

Notice the plane total bandwidth is by default 2000Mbps for the primary paths and 1800Mbps for the secondary path, as can also be seen using this output.

```

*A:PE-1# show mcast-management chassis
=====
Chassis Information
=====

```

BW per MC plane	Single SFM	Dual SFM
Primary Path	2000	2000
Secondary Path	1800	1800

```

-----
MMRP Admin Mode          Disabled
MMRP Oper Mode           Disabled
Round Robin Inactive Records Disabled
=====
*A:PE-1#

```

The Round Robin Inactive Records is disabled. The MMRP (Multiple MAC Registration Protocol) modes relate to the use of the MC Path Mgmt queues for MMRP traffic. When this is enabled, normal IMPM behavior is suspended so it is not in the scope of this configuration note.

A single channel (239.255.0.2) is now sent into interface int-IOM3-1 on port 6/2/1 with static IGMP joins on interfaces int-IMM8, int-IOM3-1 and int-IOM3-2. The current forwarding rate can be seen.

```
*A:PE-1# show router pim group detail
=====
PIM Source Group ipv4
=====
Group Address      : 239.255.0.2
Source Address     : 172.16.2.1
RP Address         : 192.0.2.1
Flags              : spt, rpt-prn-des   Type                : (S,G)
MRIB Next Hop     : 172.16.2.1
MRIB Src Flags    : direct              Keepalive Timer Exp: 0d 00:03:14
Up Time           : 0d 00:00:16         Resolved By         : rtable-u

Up JP State       : Joined              Up JP Expiry        : 0d 00:00:00
Up JP Rpt        : Pruned              Up JP Rpt Override  : 0d 00:00:00

Register State    : Pruned              Register Stop Exp   : 0d 00:00:59
Reg From Anycast RP: No

Rpf Neighbor      : 172.16.2.1
Incoming Intf     : int-IOM3-1
Outgoing Intf List : int-IMM8, int-IOM3-1, int-IOM3-2

Curr Fwding Rate  : 9873.0 kbps
Forwarded Packets : 18017                Discarded Packets   : 0
Forwarded Octets  : 24899494            RPF Mismatches      : 0
Spt threshold    : 0 kbps                ECMP opt threshold  : 7
Admin bandwidth  : 1 kbps

=====
```

From the two sets of output below it can be seen that this is using the default primary path and switch fabric plane 4.

```
*A:PE-1# tools dump mcast-path-mgr cpm
McPathMgr[6][0]: 0xf33b0a00
PATH:
PLANE:
Type SGs      InUseBW  AvailBW  TotalBw  ID  SGs  InUseBW  AvailBW  TotalBw
P      2      9895    -        -    4    2      9895    1990105 2000000
P      1        0    -        -    3    1        0    2000000 2000000
P      1        0    -        -    5    1        0    2000000 2000000
P      1        0    -        -    6    1        0    2000000 2000000
P      1        0    -        -    7    1        0    2000000 2000000
P      1        0    -        -    8    1        0    2000000 2000000
P      1        0    -        -    9    1        0    2000000 2000000
P      1        0    -        -   10    1        0    2000000 2000000
P      1        0    -        -   11    1        0    2000000 2000000
P      1        0    -        -   12    1        0    2000000 2000000
```

## IOM3-XP/IMM and Generic Output

```

P      1      0      -      - 13      1      0 2000000 2000000
P      1      0      -      - 14      1      0 2000000 2000000
P      1      0      -      - 15      1      0 2000000 2000000
P      1      0      -      - 16      1      0 2000000 2000000
P      1      0      -      -  0      1      0 2000000 2000000
S      1      0      -      -  1      1      0 1800000 1800000
B      0      0      -      -  -      -      -      -      -

```

\*A:PE-1#

```
*A:PE-1# show system switch-fabric high-bandwidth-multicast
```

```
=====
Switch Fabric
=====
```

```

          Cap:                Planes:
Slot/Mda Min  Max  Hbm Grp  Hi | Lo
-----
6/1      100% 100% No  0   4 3 5 6 7 8 9 10 11 12 13 14 15 16 0 | 1
7/1      100% 100% No  0   19 17 20 21 22 23 24 25 26 27 28 29 30 31 32 | 33
8/1      100% 100% No  0   35 34 36 37 38 39 40 41 42 43 44 45 46 47 0 | 1
A         100% 100% No  0    2 | 2
B         100% 100% No  0    2 | 2
=====

```

\*A:PE-1#

The information about the channel can be seen using this command.

```

*A:PE-1# show mcast-management
channel [router router-instance|vpls service-id|service-name service-name]
        [mda slot[/mda]]
        [group ip-address [source ip-address]]
        [path path-type]
        [detail]

```

The output for the channel being sent is as follows.

```

*A:PE-1# show mcast-management channel
=====
Multicast Channels
=====
Legend :  D - Dynamic  E - Explicit
=====
Source Address          Slot/Cpx  Current-Bw  Path      D/E
Group Address          Highest-Bw Plane
-----
172.16.2.1             6/1      9873        Primary   D
239.255.0.2            9873        4
=====
Multicast Channels : 1
=====
*A:PE-1#

```



```

*A:PE-1# show mcast-management channel detail
=====
Multicast Channels
=====
-----
Source Address      : 172.16.2.1
Group Address       : 239.255.0.2
-----
Slot/Complex        : 6/1                Current Bw       : 9873 kbps
Dynamic/Explicit    : Dynamic            Current Path     : Primary
Oper Admin Bw       : 0 kbps             Current Plane    : 4
Ing last highest    : 9873                Preference      : 0
Black-hole rate     : None                Ing sec highest  : 9873
Time remaining      : 30 seconds           Blackhole        : No
-----
Multicast Channels : 1
=====
*A:PE-1#

```

The channel is using the dynamic bandwidth activity measurement and the current bandwidth, last highest and second last highest rates are shown (which are the same as this traffic is from a traffic generator).

The Time remaining is the time remaining in the current falling-delay period. This is reset to the falling-delay every time the last highest bandwidth gets updated, when it reaches zero the value of last highest bandwidth will be replaced with second highest bandwidth and the second highest bandwidth will be set to the value of current bandwidth.

The Oper Admin Bw displays the value used for the admin-bw for this channel.

A subset of this information can be seen using this tools command.

```

*A:PE-1# tools dump mcast-path-mgr channels slot 6
=====
Slot: 6 Complex: 0
=====
-----
Source address      CurrBw  Plane PathType  Path Pref
Group address       PathBw  Repl  Exp           BlkHoleBw
-----
172.16.2.1          9873    4     primary    0    0
239.255.0.2         9873    2     none       0
Unmanaged traffic   0        4     primary    0    8
slot: 6 cmplx: 0 path: 0  0        0     none       0
Unmanaged traffic   0        3     primary    1    8
slot: 6 cmplx: 0 path: 1  0        0     none       0
Unmanaged traffic   0        5     primary    2    8
slot: 6 cmplx: 0 path: 2  0        0     none       0
Unmanaged traffic   0        6     primary    3    8
slot: 6 cmplx: 0 path: 3  0        0     none       0
Unmanaged traffic   0        7     primary    4    8
slot: 6 cmplx: 0 path: 4  0        0     none       0
Unmanaged traffic   0        8     primary    5    8
slot: 6 cmplx: 0 path: 5  0        0     none       0
Unmanaged traffic   0        9     primary    6    8
slot: 6 cmplx: 0 path: 6  0        0     none       0
-----

```

## IOM3-XP/IMM and Generic Output

```

Unmanaged traffic          0      10  primary  7  8
  slot: 6 cplx: 0 path: 7  0      0  none    0  8
Unmanaged traffic          0      11  primary  8  8
  slot: 6 cplx: 0 path: 8  0      0  none    0  8
Unmanaged traffic          0      12  primary  9  8
  slot: 6 cplx: 0 path: 9  0      0  none    0  8
Unmanaged traffic          0      13  primary 10  8
  slot: 6 cplx: 0 path:10  0      0  none    0  8
Unmanaged traffic          0      14  primary 11  8
  slot: 6 cplx: 0 path:11  0      0  none    0  8
Unmanaged traffic          0      15  primary 12  8
  slot: 6 cplx: 0 path:12  0      0  none    0  8
Unmanaged traffic          0      16  primary 13  8
  slot: 6 cplx: 0 path:13  0      0  none    0  8
Unmanaged traffic          0      0   primary 14  8
  slot: 6 cplx: 0 path:14  0      0  none    0  8
Unmanaged traffic          0      1   secondary 15  8
  slot: 6 cplx: 0 path:15  0      0  none    0  8
*A:PE-1#

```

The bandwidth activity monitoring is now changed to use an admin-bw of 12Mbps with a blackhole rate of 15Mbps.

```

*A:PE-1# configure mcast-management
*A:PE-1>config>mcast-mgmt# info
-----
      bandwidth-policy "bandwidth-policy-1" create
        admin-bw-threshold 8000
      exit
      multicast-info-policy "multicast-info-policy-1" create
        bundle "default" create
        exit
        bundle "bundle-1" create
          channel "239.255.0.1" "239.255.0.16" create
            admin-bw 12000
            bw-activity use-admin-bw black-hole-rate 15000
          exit
        exit
      exit
-----
*A:PE-1>config>mcast-mgmt# exit all

*A:PE-1# show mcast-management channel group 239.255.0.2 detail
=====
Multicast Channels
=====
-----
Source Address      : 172.16.2.1
Group Address       : 239.255.0.2
-----
Slot/Complex        : 6/1          Current Bw         : 9873 kbps
Dynamic/Explicit    : Dynamic      Current Path       : Primary
Oper Admin Bw       : 12000 kbps   Current Plane     : 4
Ing last highest    : 12000        Preference        : 0
Black-hole rate     : 15000 kbps   Ing sec highest   : 12000
Time remaining      : 30 seconds    Blackhole         : No

```

## Ingress Multicast Path Management

```

=====
Multicast Channels : 1
=====
*A:PE-1#

*A:PE-1# tools dump mcast-path-mgr cpm
McPathMgr[6][0]: 0xf33b0a00
PATH:
                                PLANE:
Type SGs      InUseBW   AvailBW   TotalBw  ID   SGs   InUseBW   AvailBW   TotalBw
P       2      12000    -         -     4     2      12000    1988000  2000000
P       1         0        -         -     3     1         0      2000000  2000000
P       1         0        -         -     5     1         0      2000000  2000000
P       1         0        -         -     6     1         0      2000000  2000000
P       1         0        -         -     7     1         0      2000000  2000000
P       1         0        -         -     8     1         0      2000000  2000000
P       1         0        -         -     9     1         0      2000000  2000000
P       1         0        -         -    10     1         0      2000000  2000000
P       1         0        -         -    11     1         0      2000000  2000000
P       1         0        -         -    12     1         0      2000000  2000000
P       1         0        -         -    13     1         0      2000000  2000000
P       1         0        -         -    14     1         0      2000000  2000000
P       1         0        -         -    15     1         0      2000000  2000000
P       1         0        -         -    16     1         0      2000000  2000000
P       1         0        -         -     0     1         0      2000000  2000000
S       1         0        -         -     1     1         0      1800000  1800000
B       0         0        -         -     -     -         -         -         -
*A:PE-1#

```

Now the system treats the channel as though it is using 12Mbps capacity even though its current rate has not changed.

If the rate is increased above the blackhole rate, the channel is blackholed and an alarm is generated.

```

*A:PE-1#
11 2011/10/21 01:40:13.21 UTC MINOR: MCPATH #2001 Base Black-hole-rate is reached
"Channel (172.16.2.1,239.255.0.2) for vRtr instance 1 slot/cplx 6/1 has been blackholed."
*A:PE-1# show mcast-management channel group 239.255.0.2 detail
=====
Multicast Channels
=====
-----
Source Address      : 172.16.2.1
Group Address       : 239.255.0.2
-----
Slot/Complex        : 6/1
Dynamic/Explicit    : Dynamic
Oper Admin Bw       : 12000 kbps
Ing last highest    : 19480
Black-hole rate     : 15000 kbps
Time remaining      : 23 seconds
Current Bw          : 19458 kbps
Current Path        : Blackhole
Current Plane       : N/A
Preference          : 0
Ing sec highest     : 19469
Blackhole           : Yes
=====
Multicast Channels : 1
=====
*A:PE-1#
*A:PE-1# tools dump mcast-path-mgr cpm

```

## IOM3-XP/IMM and Generic Output

```

McPathMgr[6][0]: 0xf33b0a00
PATH:
PLANE:
Type SGs      InUseBW   AvailBW   TotalBw   ID   SGs      InUseBW   AvailBW   TotalBw
P      1          0         -         -     4     1          0   2000000   2000000
P      1          0         -         -     3     1          0   2000000   2000000
P      1          0         -         -     5     1          0   2000000   2000000
P      1          0         -         -     6     1          0   2000000   2000000
P      1          0         -         -     7     1          0   2000000   2000000
P      1          0         -         -     8     1          0   2000000   2000000
P      1          0         -         -     9     1          0   2000000   2000000
P      1          0         -         -    10     1          0   2000000   2000000
P      1          0         -         -    11     1          0   2000000   2000000
P      1          0         -         -    12     1          0   2000000   2000000
P      1          0         -         -    13     1          0   2000000   2000000
P      1          0         -         -    14     1          0   2000000   2000000
P      1          0         -         -    15     1          0   2000000   2000000
P      1          0         -         -    16     1          0   2000000   2000000
P      1          0         -         -     0     1          0   2000000   2000000
S      1          0         -         -     1     1          0   1800000   1800000
B      1      19480         -         -     -     -          -     -         -
*A:PE-1#

```

The output displayed above shows an alarm generated for a channel being blackholed due to the channel rate reaching the configured black-hole-rate. The example output below is an alarm for a channel being blackholed due to insufficient bandwidth being available to it.

```

7 2011/10/22 21:53:43.54 UTC MINOR: MCPATH #2001 Base No bandwidth available
"Channel (172.16.2.1,239.255.0.2) for vRtr instance 1 slot/cplx 6/1 has been blackholed."

```

Note that the following alarm relates to a dummy channel used to account for the unmanaged traffic. However, this traffic is never actually blackholed.

```

6 2011/10/21 00:27:58.00 UTC MINOR: MCPATH #2002 Base
"Channel (0.0.0.0,0.6.0.0) for unknown value (2) instance 0 slot/cplx 6/1 is no longer
being blackholed."

```

Alarms are also generated when all paths of a given type reach certain thresholds.

For primary and secondary paths, the two path range limit thresholds are

- Full: less than 5% capacity is available

```

9 2011/10/24 22:53:27.02 UTC MINOR: MCPATH #2003 Base
"The available bandwidth on secondary path on slot/cplx 6/1 has reached its maximum
limit."

```

- Not full: more than 10% of the path capacity is available.

```

10 2011/10/24 22:53:48.02 UTC MINOR: MCPATH #2004 Base
"The available bandwidth on secondary path on slot/cplx 6/1 is within range lim-
its."

```

A maximum of one alarm is generated for each event (blackhole start/stop, path full/not full) within a 3 second period. So, for example, if multiple channels are blackholed within the same 3 second period only one alarm will be generated (for the first event).

The effect of using the hi-bw-mcast-src command is illustrated below. Firstly, line card 6 is configured into group 1.

```
*A:PE-1# configure card 6 fp hi-bw-mcast-src group 1 alarm
*A:PE-1# show system switch-fabric high-bandwidth-multicast
=====
Switch Fabric
=====
          Cap:                Planes:
Slot/Mda Min  Max  Hbm Grp  Hi | Lo
-----
6/1      100% 100% Yes  1    4 3 5 6 7 8 9 10 11 12 13 14 15 16 0 | 1
7/1      100% 100% No   0   19 17 20 21 22 23 24 25 26 27 28 29 30 31 32 | 33
8/1      100% 100% No   0   35 34 36 37 38 39 40 41 42 43 44 45 46 47 0 | 1
A         100% 100% No   0    2 | 2
B         100% 100% No   0    2 | 2
=====
*A:PE-1#
```

The plane assignment has not changed (though it is possible that the system could re-arrange the planes used by card 6) and there are still planes (0,1) shared between card 6 and card 8.

Now card 8 is configured into group 2 (note that IMPM is only enabled on card 6 here).

```
*A:PE-1# configure card 8 fp hi-bw-mcast-src group 2 alarm
*A:PE-1# show system switch-fabric high-bandwidth-multicast
=====
Switch Fabric
=====
          Cap:                Planes:
Slot/Mda Min  Max  Hbm Grp  Hi | Lo
-----
6/1      100% 100% Yes  1    4 3 5 6 7 8 9 10 11 12 13 14 15 16 0 | 1
7/1      100% 100% No   0   19 17 20 21 22 23 24 25 26 27 28 29 30 31 32 | 33
8/1      100% 100% Yes  2   35 34 36 37 38 39 40 41 42 43 44 45 46 47 17 | 19
A         100% 100% No   0    2 | 2
B         100% 100% No   0    2 | 2
=====
*A:PE-1#
```

There are no longer planes shared between cards 6 and 8.

If card 7 is configured into group 3, the following is seen.

```
*A:PE-1# configure card 7 fp hi-bw-mcast-src group 3 alarm
*A:PE-1#
7 2011/10/21 00:35:50.95 UTC MINOR: CHASSIS #2052 Base Mda 6/1
"Class MDA Module : Plane shared by multiple multicast high bandwidth taps"

8 2011/10/21 00:35:50.95 UTC MINOR: CHASSIS #2052 Base Mda 6/2
```

## IOM3-XP/IMM and Generic Output

```
"Class MDA Module : Plane shared by multiple multicast high bandwidth taps"
```

```
9 2011/10/21 00:35:50.97 UTC MINOR: CHASSIS #2052 Base Mda 7/1
```

```
"Class MDA Module : Plane shared by multiple multicast high bandwidth taps"
```

```
10 2011/10/21 00:35:50.97 UTC MINOR: CHASSIS #2052 Base Mda 7/2
```

```
"Class MDA Module : Plane shared by multiple multicast high bandwidth taps"
```

```
*A:PE-1# show system switch-fabric high-bandwidth-multicast
```

```
Switch Fabric
```

```
=====  
Cap:          Planes:  
Slot/Mda Min  Max  Hbm Grp  Hi | Lo  
-----  
6/1      100% 100% Yes 1    4 3 5 6 7 8 9 10 11 12 13 14 15 16 0 | 1  
7/1      100% 100% Yes 3    21 20 22 23 24 25 26 27 28 29 30 31 32 33 0 | 1  
8/1      100% 100% Yes 2    35 34 36 37 38 39 40 41 42 43 44 45 46 47 17 | 19  
A         100% 100% No  0     2 | 2  
B         100% 100% No  0     2 | 2  
=====
```

```
*A:PE-1#
```

There are insufficient planes to allow each card/group to have dedicated planes. Planes 0 and 1 are still shared between cards 6 and 7, generating the associated alarms.

A common example of the use of the **hi-bw-mcast-src** command would be when cards 6 and 8 have uplink ports on which high bandwidth multicast channels could be received. It would be desired to have these cards use different planes. To achieve this, card 7 could be configured into group 1, as follows.

```
*A:PE-1# configure card 7 fp hi-bw-mcast-src group 1 alarm
```

```
*A:PE-1# show system switch-fabric high-bandwidth-multicast
```

```
Switch Fabric
```

```
=====  
Cap:          Planes:  
Slot/Mda Min  Max  Hbm Grp  Hi | Lo  
-----  
6/1      100% 100% Yes 1    4 3 5 6 7 8 9 10 11 12 13 14 15 16 0 | 1  
7/1      100% 100% Yes 1    4 3 5 6 7 8 9 10 11 12 13 14 15 16 0 | 1  
8/1      100% 100% Yes 2    35 34 36 37 38 39 40 41 42 43 44 45 46 47 17 | 19  
A         100% 100% No  0     2 | 2  
B         100% 100% No  0     2 | 2  
=====
```

```
*A:PE-1#
```

Now it can be seen that card 7 shares the same planes as card 6, but more importantly card 6 has no planes in common with card 8.

Note that when traffic is received on card 6, it will also be seen on the same plane (not path) on card 7. In the example below, traffic can be seen on plane 4 which is used by both cards 6 and 7, but only card 6 has non-zero InUseBW path capacity.

```
*A:PE-1# tools dump mcast-path-mgr cpm
McPathMgr[6][0]: 0xf33b0a00
PATH:
Type SGs      InUseBW  AvailBW  TotalBw  PLANE:
ID  SGs      InUseBW  AvailBW  TotalBw
P   2        9707    -        -        4   3        9707    1990293  2000000
P   1         0      -        -        3   2         0      2000000  2000000
...
McPathMgr[7][0]: 0xf33b3198
PATH:
Type SGs      InUseBW  AvailBW  TotalBw  PLANE:
ID  SGs      InUseBW  AvailBW  TotalBw
P   1         0      -        -        4   3        9707    1990293  2000000
P   1         0      -        -        3   2         0      2000000  2000000
```

When IMPM managed traffic is received on SAPs (in an IES, VPLS or VPRN service) it can be seen against a specific queue counter (Off. Managed) in the SAP stats. The following output shows where sap 7/2/1:3 belongs to a VPLS service using igmp-snooping. A similar counter is not available for policer statistics.

```
*A:PE-1# show service id 2 sap 7/2/1:3 stats
=====
Service Access Points(SAP)
=====
-----
Sap per Queue stats
-----
                Packets                Octets
Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio      : 0                    0
Off. LoPrio      : 0                    0
Dro. HiPrio      : 0                    0
Dro. LoPrio      : 0                    0
For. InProf      : 0                    0
For. OutProf     : 0                    0

Ingress Queue 11 (Multipoint) (Priority)
Off. HiPrio      : 0                    0
Off. LoPrio      : 0                    0
Off. Managed     : 149410                209771640
Dro. HiPrio      : 0                    0
Dro. LoPrio      : 0                    0
For. InProf      : 149410                209771640
For. OutProf     : 0                    0

Egress Queue 1
For. InProf      : 0                    0
For. OutProf     : 0                    0
Dro. InProf      : 0                    0
Dro. OutProf     : 0                    0
=====
*A:PE-1#
```

## IOM1/2 Specific Output

The system is configured with the following cards and is in chassis mode **a**. As can be seen, an IOM2 is in slot 5.

```
A:PE-1# show card
=====
Card Summary
=====
```

Slot	Provisioned Card-type	Equipped Card-type	Admin State	Operational State	Comments
5	iom2-20g	iom2-20g	up	up	
6	iom3-xp	iom3-xp	up	up	
7	imm8-10gb-xfp	imm8-10gb-xfp	up	up	
8	iom3-xp	iom3-xp	up	up	
A	sfm4-12	sfm4-12	up	up/active	
B	sfm4-12	sfm4-12	up	up/standby	

```
=====
A:PE-1#
```

IMPM is enabled on MDA 1 and 2 of the IOM2 in slot 5, with a primary, secondary and ancillary path.

```
A:PE-1# show mcast-management mda
=====
MDA Summary
=====
```

S/C	Policy	Type	In-use-Bw	Admin
5/1	default	Primary	0 Kbps	up
	default	Secondary	0 Kbps	up
	default	Ancillary	0 Kbps	up
5/2	default	Primary	0 Kbps	up
	default	Secondary	0 Kbps	up
	default	Ancillary	0 Kbps	up
6/2	default	Primary	0 Kbps	down
	default	Secondary	0 Kbps	down
	default	Ancillary	0 Kbps	down
7/1	default	Primary	0 Kbps	down
	default	Secondary	0 Kbps	down
	default	Ancillary	0 Kbps	down
7/2	default	Primary	0 Kbps	down
	default	Secondary	0 Kbps	down
	default	Ancillary	0 Kbps	down
8/2	default	Primary	0 Kbps	down
	default	Secondary	0 Kbps	down
	default	Ancillary	0 Kbps	down

```
=====
A:PE-1#
```



The path/plane usage can be shown.

```
*A:PE-1# show system switch-fabric high-bandwidth-multicast
=====
Switch Fabric
=====
          Cap:                Planes:
Slot/Mda Min  Max  Hbm Grp  Hi | Lo
-----
5/1      100% 100% No  0   1 | 0
5/2      100% 100% No  0   4 | 3
6/1      100% 100% No  0   6 5 7 8 9 10 11 12 13 14 15 0 1 3 4 | 5
7/1      100% 100% No  0   7 6 8 9 10 11 12 13 14 15 0 1 3 4 5 | 6
8/1      100% 100% No  0   8 7 9 10 11 12 13 14 15 0 1 3 4 5 6 | 7
A        100% 100% No  0   2 | 2
B        100% 100% No  0   2 | 2
=====
```

```
*A:PE-1#
*A:PE-1# tools dump mcast-path-mgr cpm
McPathMgr[5][0]: 0xf33b0a00
PATH:
          PLANE:
Type SGs      InUseBW  AvailBW  TotalBw  ID  SGs      InUseBW  AvailBW  TotalBw
P      1          0    2000000  2000000  1   1          0    2000000  2000000
S      1          0    1500000  1500000  0   1          0    1800000  1800000
A      0          0    5000000  5000000  -   -          -     -         -
B      0          0     -         -         -   -          -     -         -
McPathMgr[5][1]: 0xf33b3198
PATH:
          PLANE:
Type SGs      InUseBW  AvailBW  TotalBw  ID  SGs      InUseBW  AvailBW  TotalBw
P      1          0    2000000  2000000  4   1          0    2000000  2000000
S      1          0    1500000  1500000  3   1          0    1800000  1800000
A      0          0    5000000  5000000  -   -          -     -         -
B      0          0     -         -         -   -          -     -         -
*A:PE-1#
```

The path range limit alarm thresholds for the ancillary path are

- Full: less than 2% capacity is available
- Not full: more than 4% of the path capacity is available.

A single channel (239.255.0.1) is now sent into interface int-IOM2 on port 5/2/1 with static IGMP joins on interfaces int-IMM8, int-IOM3-1 and int-IOM3-2. The current forwarding rate can be seen.

```
*A:PE-1# show router pim group 239.255.0.1 detail
=====
PIM Source Group ipv4
=====
Group Address      : 239.255.0.1
Source Address     : 172.16.1.1
RP Address         : 192.0.2.1
Flags              : spt, rpt-prn-des   Type                : (S,G)
MRIB Next Hop     : 172.16.1.1
```

## IOM1/2 Specific Output

```
MRIB Src Flags      : direct           Keepalive Timer Exp: 0d 00:02:44
Up Time             : 0d 00:07:45      Resolved By        : rtable-u

Up JP State         : Joined           Up JP Expiry       : 0d 00:00:00
Up JP Rpt          : Pruned           Up JP Rpt Override : 0d 00:00:00

Register State     : Pruned           Register Stop Exp  : 0d 00:00:32
Reg From Anycast RP: No

Rpf Neighbor       : 172.16.1.1
Incoming Intf      : int-IOM2
Outgoing Intf List : int-IMM8, int-IOM3-1, int-IOM3-2

Curr Fwding Rate   : 9734.8 kbps
Forwarded Packets  : 591874            Discarded Packets  : 0
Forwarded Octets   : 817969868        RPF Mismatches     : 0
Spt threshold      : 0 kbps           ECMP opt threshold : 7
Admin bandwidth    : 1 kbps
```

-----

As there are only 3 (<4) DFCs, the ancillary path is used.

```
*A:PE-1# show mcast-management channel
```

```
=====
Multicast Channels
=====
```

```
Legend : D - Dynamic E - Explicit
```

```
=====
Source Address          Slot/Cpx  Current-Bw  Path      D/E
Group Address          Highest-Bw Plane
-----
172.16.1.1             5/2      9729       Ancillary D
239.255.0.1           9740     -          -
=====
```

```
Multicast Channels : 1
=====
```

```
*A:PE-1#
```

```
*A:PE-1# show mcast-management channel detail
```

```
=====
Multicast Channels
=====
```

```
Source Address      : 172.16.1.1
Group Address       : 239.255.0.1
```

```
-----
Slot/Complex        : 5/2              Current Bw          : 9729 kbps
Dynamic/Explicit    : Dynamic          Current Path        : Ancillary
Oper Admin Bw       : 0 kbps           Current Plane       : N/A
Ing last highest    : 9740             Preference          : 0
Black-hole rate     : None             Ing sec highest     : 9740
Time remaining      : 27 seconds        Blackhole           : No
=====
```

```
Multicast Channels : 1
=====
```

```
*A:PE-1#
```

If another join caused this traffic to be switched via an additional DFC, the system would place the channel on the primary path.

The ancillary path being used can also be seen as follows.

```
*A:PE-1# tools dump mcast-path-mgr cpm
McPathMgr[5][0]: 0xf33b0a00
PATH:
Type SGs      InUseBW  AvailBW  TotalBW  ID  SGs  InUseBW  AvailBW  TotalBW
P      1          0  2000000  2000000  1   1      0  2000000  2000000
S      1          0  1500000  1500000  0   1      0  1800000  1800000
A      0          0  5000000  5000000  -   -      -    -        -
B      0          0    -        -        -   -      -    -        -
McPathMgr[5][1]: 0xf33b3198
PATH:
Type SGs      InUseBW  AvailBW  TotalBW  ID  SGs  InUseBW  AvailBW  TotalBW
P      1          0  2000000  2000000  4   1      0  2000000  2000000
S      1          0  1500000  1500000  3   1      0  1800000  1800000
A      1      19480  4980520  5000000  -   -      -    -        -
B      0          0    -        -        -   -      -    -        -
*A:PE-1#
```

Note that the bandwidth shown on the ancillary path on the second MDA is approximately two times that of the ingress traffic, this matches the algorithm described earlier for the ancillary path. This can also be seen in the next output, there is the original channel traffic plus two replications (Repl).

```
*A:PE-1# tools dump mcast-path-mgr channels
=====
Slot: 5 Complex: 0
=====
Source address          CurrBw  Plane PathType  Path Pref
Group address          PathBw  Repl  Exp      BlkHoleBw
-----
Unmanaged traffic      0       1    primary  0      8
  slot: 5 cmplx: 0 path: 0  0       0    none     0
Unmanaged traffic      0       0    secondary 1  8
  slot: 5 cmplx: 0 path: 1  0       0    none     0
=====
Slot: 5 Complex: 1
=====
Source address          CurrBw  Plane PathType  Path Pref
Group address          PathBw  Repl  Exp      BlkHoleBw
-----
172.16.1.1             9740   48   ancillary 16     0
239.255.0.1            19480  2    none     0
Unmanaged traffic      0       4    primary  0      8
  slot: 5 cmplx: 1 path: 0  0       0    none     0
Unmanaged traffic      0       3    secondary 1  8
  slot: 5 cmplx: 1 path: 1  0       0    none     0
*A:PE-1#
```

## Conclusion

This section has described the configuration of Ingress Multicast Path Management which optimizes IPv4 and IPv6 multicast capacity to achieve the maximum system-wide IP multicast throughput. It can be used for both routed IPv4/IPv6 and VPLS (IGMP and PIM) snooped IPv4 multicast groups, which usually relate to the distribution of IP TV channels.