

Application Note

Using Multicast Call Admission Control for IPTV Bandwidth Management

Managing Multicast Bandwidth in IPTV Networks Using Multicast Call Admission Control in the Edge Router

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408.745.2000
1.888 JUNIPER
www.juniper.net

Table of Contents

Introduction	3
Scope	3
Design Considerations	3
IGMP Signaling and MCAC	4
MCAC for IGMP Passthrough	5
Example Scenario Explaining MCAC for IGMP Passthrough	5
Lab Analysis for C-VLAN and Port Multicast Admission Control for IGMP Passthrough	6
MCAC for IGMP Proxy or Snooping	8
Example Scenario Explaining Multicast Admission Control for IGMP Proxy	10
Lab Analysis for C-VLAN and Port Multicast Admission Control for IGMP Proxy/Snooping	11
Group Priorities for Multicast Admission Control	12
Summary	13
About Juniper Networks	14

Introduction

A key requirement in IPTV networks is the ability to deliver high-quality video streams to each household. For linear/broadcast TV in packet networks, IP multicast is used to allow the network to copy and forward copies of the same source stream to a large number of viewers. The set-top box (STB) sends IGMP join messages that terminate at an Access Node (AN) or Broadband Services Router (BSR). In turn, the AN or BSR responds by forwarding the requested multicast group (television channel) to the subscriber who made the request. Frequently, the access network is overbuilt to ensure that every channel can be delivered to every access node.

Instead of overbuilding the network, IP multicast can preserve bandwidth by sending multicast groups across the network where required. This is most critical in xDSL networks that can only accommodate a limited number of channels due to bandwidth constraints to the subscriber. In addition, the rise in unicast video—such as video on demand (VOD), replay TV or streaming video downloaded from the Internet—can result in the link to the AN becoming a congestion point where the bandwidth consumed by multicast IPTV must be restricted.

Taking advantage of this feature of IPTV requires that the network be able to prevent an interface between the AN and BSR to become congested, which could happen if the amount of bandwidth consumed by multicast IPTV is not bounded. In addition to protecting the xDSL bandwidth by limiting the IPTV traffic to each subscriber, it may also be necessary to limit the bandwidth consumed by this traffic on other links, such as going to the DSLAM.

Scope

This document details MCAC, which is implemented in JUNOS software used in Juniper's E-series BSRs. It prevents network oversubscription due to excessive IPTV traffic by providing an upper limit on the amount of bandwidth available for delivering multicast IPTV to each DSLAM and each subscriber. Its use is described for the cases of both a non-replicating and replicating AN. Each model is a valid deployment option but will affect the BSR's capabilities for MCAC.

Design Considerations

This application note describes a router-based method for Multicast Call Admission Control (MCAC). Instead of sending all multicast groups to the AN or planning for worst case per-user or per-AN bandwidth consumption scenarios, the BSR can help manage resources when there is the possibility of network oversubscription due to multicast joins.

The BSR processes each IGMP request, whether on a per-port, per-M-VLAN or C-VLAN basis, and determines if network resources remain to accept the join request. If the join request is accepted, the group is replicated and forwarded out the requesting interface. If the join request is denied due to the exhaustion of network resources, the multicast group is not forwarded out the interface.

The Juniper Networks E-series Broadband Service Routing Platform (E-series), based on Juniper Networks JUNOS™ software, uses its MCAC feature to define multicast group bandwidth and match each join request against this table and the interface limit to determine if a join request should be processed.

IGMP Signaling and MCAC

One simple design for an IPTV deployment is to make the edge or BSR part of every multicast tree. The edge router then processes IGMP joins to replicate and send copies of video into the access network and ultimately to the home¹.

The IGMP protocols are well defined, but how they are implemented in the access network is not. In particular, there are two fundamental models² for IGMP that are covered in this document:

- IGMP passthrough, in which the AN is not part of the multicast replication process and the BSR is the last replication point
- IGMP transparent snooping/proxy, in which the AN provides distributed replication and receives multicast groups over a shared multicast VLAN or router interface as part of a C-VLAN or S-VLAN architecture

The choice of model plays a major role in how bandwidth resources are managed. For example, with IGMP snooping and IGMP proxy, the AN manages join processing on a per-subscriber basis. In these cases, the BSR does not manage per-subscriber replication and therefore cannot provide multicast resource control on a per-user basis.

On the other hand, IGMP passthrough multicasting locates per-subscriber IGMP processing—and replication decisions—in the BSR, which can then manage network resources on both a per-user and per-port basis.

Table 1: Edge Architectures and Type of MCAC Supported

	C-VLAN	C-VLAN with M-VLAN	S-VLAN
Per subscriber	Yes	No	No
Per AN/DSLAM	Yes	Yes	Yes
Per BSR port	Yes	Yes	Yes

Table 1 summarizes the type of MCAC granularity or hierarchy that can be supported in JUNOS software. When the BSR is not the last replication point, in either the C-VLAN with M-VLAN or S-VLAN models, then the BSR can only provide aggregate resource management on a per-AN or per-port basis. Therefore, the C-VLAN with M-VLAN and S-VLAN architectures are identical for multicast while varying in how they handle unicast services.

For the case of per-user or per-AN multicast bandwidth control, JUNOS software can limit the amount of bandwidth consumed by multicast traffic using its MCAC capability. Multicast admission control maps each IGMP join to a maximum bandwidth limit to determine if replication is allowed. This augments quality of service (QoS) in situations where it is possible to congest the interface by oversubscribing subscriber bandwidth or endpoints issuing joins outside of a predefined service profile.

An admission bandwidth is assigned to each multicast group. This multicast admission control engine uses this value to determine whether the next join request would result in the total bandwidth requested exceeding the admission bandwidth limit. Each multicast group can be statically configured with an admission bandwidth or the system can dynamically measure each group bandwidth when received.

While it is generally not desired to lock out any TV channel, there could be channels less important than others for which a sporadic block for a limited amount of time might be acceptable. For example, pay TV channels cannot be blocked in any circumstance, while some free channels might be subject to MCAC. To accomplish this, JUNOS software Release 8-2-0 enriched the basic MCAC feature by adding the concept of prioritization of groups. The range of multicast prefixes or multicast addresses that belong to a specific subnet can now be given a priority value different than the default one, which is 0.

¹For more information on use of IGMP in IPTV networks, see Introduction to IGMP for IPTV Networks, www.juniper.net/solutions/literature/white_paper

²For more information on these models, see Wireline Broadband Access Design: IGMP and VLANs, www.juniper.net/solutions/literature/solution_brief

At the port level, it is now possible to specify two thresholds for the maximum outgoing multicast bandwidth. The first threshold can be reached by honoring and forwarding any multicast group of any priority; the second threshold can be reached by adding only multicast groups with a priority higher than 0. This is to make sure that TV channels that have to be transmitted anyway have a way to go through. That is, the difference in bandwidth between the two thresholds should be greater than or equal to the sum of the equivalent requested bandwidths of the TV channels that cannot be blocked.

In the unlikely case that a threshold is hit, note that no more IGMP joins are honored until the sum of the multicast bandwidth falls below a predefined threshold. This is to avoid the circumstance in which subscribers who cannot watch the requested TV channel keep insisting on sending joins for the same channel, resulting in the multicast bandwidth limit being continuously reached. It is expected that subscribers will stop zapping once they find a TV channel that is not blocked. That threshold should help provide relief from such blocking situations.

MCAC for IGMP Passthrough

The first model detailed for MCAC is IGMP passthrough. In this model, the BSR will receive IGMP joins from each household and do replication on a per-household basis. Since the BSR is processing per-user IGMP, it can determine the amount of multicast sent over each user interface. The BSR can also limit the amount of multicast sent out to all users via the port connected to the AN.

Example Scenario Explaining MCAC for IGMP Passthrough

Figure 1 highlights this model across two households. The BSR keeps a local mapping of multicast groups and their bandwidth. This table can be statically created or dynamically measured.

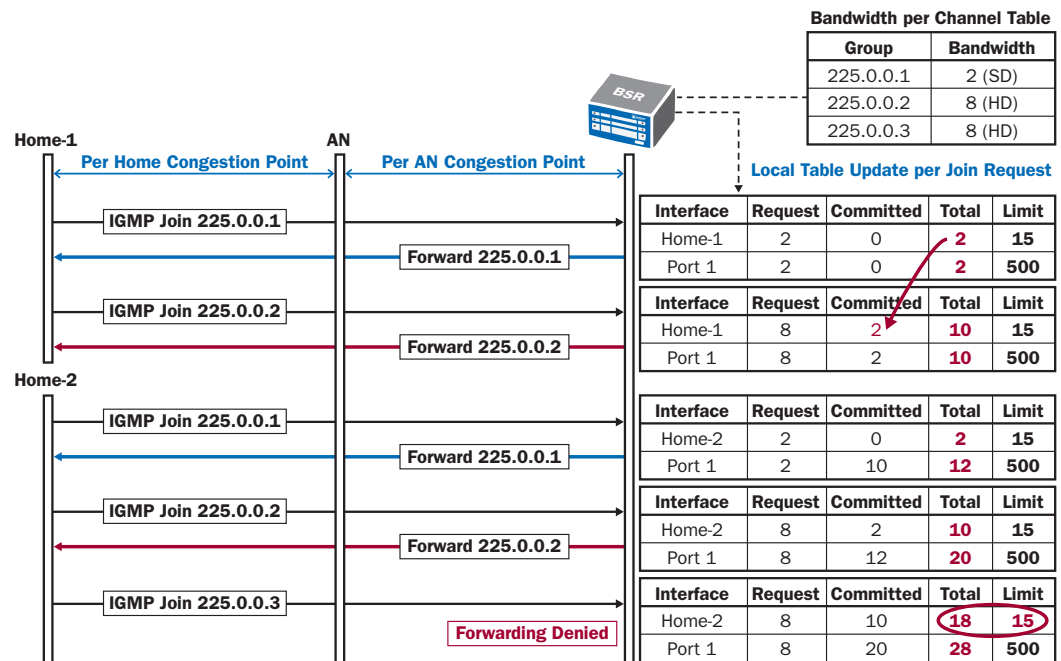


Figure 1: Example of MCAC Applied to an IGMP Passthrough Scenario

The BSR keeps an admission control resource table per user and per port. Each user is set with a limit of 15 Mbps of multicast while the aggregate across all users is set at the port level to 500 Mbps. The user limits can be set based on planned service models, DSL access loop speed or other parameters. The port limit is set based on the planned amount of multicast consumed by all users on a given interface.

As each IGMP join request comes in, the BSR sums the amount of currently committed bandwidth and the request amount. As long as this amount is below the configured limit then the interface is sent a copy of the multicast group. The resource state table is then updated with a new committed rate that includes the new multicast group. This is performed for both the home and port entries in the table. The per-home values are unique to each C-VLAN interface while port admission control is shared across all users on the same port.

If the bandwidth limit is exceeded, then the BSR will deny the join request and not send a copy of the multicast out of that interface. Note that this capability is at the network level. If a join request is denied, then the STB will not receive any type of feedback notification directly from the router. The router will generate a trap that can be used for future capacity planning or integrated to the video application for direct STB feedback messages. Providing feedback to the user is outside the scope of this paper.

Lab Analysis for C-VLAN and Port Multicast Admission Control for IGMP Passthrough

For subscriber-level admission control, each household maps to a C-VLAN interface that is configured with a maximum amount of multicast bandwidth allowed to be replicated across that interface using the following JUNOS software command configured under the C-VLAN interface:

```
interface gigabitEthernet 15/0/0.600
...
ip multicast admission-bandwidth-limit 9000000
```

The configuration part in common with all the C-VLANs and/or ports where multicast bandwidth has to be checked before accepting the IGMP join is essentially made of three parts: a route-map that maps a multicast address or a range to a bandwidth value (either static or dynamically measured, that is, “adaptive”), an access-list that filters the multicast address or a range of them recalled in the route-map and a command to enable the MAC function. Here follows an example of a basic configuration:

```
access-list "SDTV" permit ip any 224.0.107.0 0.0.0.255
!
route-map "per-group-mcast-bandwidth" permit 10
  match ip address "SDTV"
  set admission-bandwidth adaptive
!
ip multicast-routing bandwidth-map "per-group-mcast-bandwidth"
```

In this test, a maximum limit of 9 Mbps is used so that with the 3.75 Mbps SDTV video stream, only two groups can be joined at the same time. The `show ip mroute`³ output shows that the first two groups (224.0.17.104, 224.0.17.107) were joined successfully.

E320_MRA_ER_0#show ip mroute

```
...
(11.1.3.26, 224.0.17.104) uptime 0 00:14:56
  Admission bandwidth: 3907000 bps (adaptive)
  QoS bandwidth: 3907000 bps (adaptive)
  RPF route: 11.1.3.0/24, incoming interface GigabitEthernet2/1/3
    neighbor 11.1.6.6, owner IsIs (ECMP route)
  Incoming interface list:
    GigabitEthernet2/1/3 (11.1.6.1/30), Accept/Pim (RPF IIF)
  Outgoing interface list:
```

³This command is used because the MCAC feature has been implemented in the Multicast Group Table Manager (MGTM) component rather than in the IGMP one. MGTM is the component responsible for protocols such as PIM.

```

GigabitEthernet15/0/0.600 (192.168.15.1/24), Forward/Pim, 0
00:15:00/never
(11.1.3.26, 224.0.17.107) uptime 0 03:04:59
Admission bandwidth: 3266000 bps (adaptive)
QoS bandwidth: 3266000 bps (adaptive)
RPF route: 11.1.3.0/24, incoming interface GigabitEthernet3/1/2
neighbor 11.1.6.6, owner IsIs (ECMP route)
Incoming interface list:
GigabitEthernet3/1/2 (11.1.6.5/30), Accept/Pim (RPF IIF)
Outgoing interface list:
GigabitEthernet15/0/0.600 (192.168.15.1/24), Forward/Pim, 0
03:05:00/never

```

However, when the third group (224.0.17.103) join request is received by the BSR, it is not replicated because the configured interface admission limit (*intf-adm-limit*) has been exceeded:

```

(11.1.3.20, 224.0.17.103) uptime 0 00:00:09
Admission bandwidth: 3460000 bps (adaptive)
QoS bandwidth: 3460000 bps (adaptive)
RPF route: 11.1.3.0/24, incoming interface GigabitEthernet3/1/2
neighbor 11.1.6.6, owner IsIs (ECMP route)
Incoming interface list:
GigabitEthernet3/1/2 (11.1.6.5/30), Accept/Pim (RPF IIF)
Outgoing interface list:
GigabitEthernet15/0/0.600 (192.168.15.1/24), Blocked
(intf-adm-limit)/Pim, 0 00:00:09/never

```

This group is shown as *blocked*, denoting that group replication is not occurring due to an interface administrative limit (*intf-adm-limit*). This type of blocking only impacts a single subscriber who attempts to join beyond his 9-Mbps C-VLAN bandwidth limit.

In addition to the per C-VLAN bandwidth limits, JUNOS software also provides admission bandwidth hierarchy by allowing port-level bandwidth limits using the following global command:

```

mroute port <# slot #> // <# bay #> / <# port #> admission-bandwidth-limit <#
bandwidth #>

```

Instead of an *intf-adm-limit*, the replication is blocked with a port admission limit (*port-adm-limit*):

```

(11.1.5.66, 224.0.20.11) uptime 0 00:32:20
Admission bandwidth: 3864000 bps (adaptive)
QoS bandwidth: 3864000 bps (adaptive)
RPF route: 11.1.5.0/24, incoming interface GigabitEthernet3/1/2
neighbor 11.1.6.6, owner IsIs (ECMP route)
Incoming interface list:
GigabitEthernet3/1/2 (11.1.6.5/30), Accept/Pim (RPF IIF)
Outgoing interface list:
GigabitEthernet15/0/7.10 (192.168.12.10/24), Blocked
(port-adm-limit)/Icmp, 0 00:32:20/never
GigabitEthernet15/0/5.1 (192.168.40.1/24), Forward/Icmp, 0 00:32:20/
never

```

This will affect subscribers sending a join request that exceeds the port limit, regardless of their C-VLAN multicast admission bandwidth.

MCAC for IGMP Proxy or Snooping

When the AN is replicating multicast, the simplest method for traffic planning is to send a copy of all multicast groups from the BSR to the AN, ensuring there is no congestion at the network edge. For example, if there is 500 Mbps of offered multicast traffic, the entire 500 Mbps is sent to the AN. This works for small multicast group counts or if there is a surplus of bandwidth to the AN. As the network evolves, however, sending all groups to the AN becomes inefficient and can even result in the reduction of bandwidth for revenue-generating unicast applications.

This inefficiency is highlighted in the following graphs. Figure 2 shows the offered load of multicast received by the BSR from the core network. The average multicast load is 343 Mbps. This is a steady average across all days of the week and is based solely on the number of multicast groups (aka TV channels) and their corresponding bandwidths. This value only changes based on the number of groups offered, and the encoding rate changes either lower to optimized encoding schemes or higher based on industry adoption of HDTV. This value is independent of user viewing.




Figure 2: Offered Multicast Load for IPTV

Actual user multicast consumption was measured at the BSR MVLAN for a user base of 2000 households as shown in Figure 3. The average bandwidth is only 141 Mbps, about 40 percent of the offered load. The peak is measured at 170 Mbps, still about 50 percent of the offered multicast load. This highlights the bandwidth savings that can be achieved by not planning for all multicast groups to be pushed to the AN at all times.




Figure 3: Aggregate Multicast Bandwidth for 2000 Households*

*The data set used in this graph was selected to illustrate the maximum growth in subscribers over a short period of time (about 3 weeks). Doing so eliminates any seasonal influence on subscriptions.

Even as user count grows, it can be seen that the network load is mildly impacted by the increase in user count. In this study, the graph in Figure 4 is based on 2350 users. The additional 350 users had the impact of increasing the peak value by only 4 Mbps—equivalent to 1 to 2 standard definition channels. So while the user base increased 17.5 percent, the multicast demand was only increased by 2 percent.




Figure 4: Aggregate Multicast Bandwidth for 2350 Households*

This study highlights the benefit of only sending watched channels down to the AN and reusing bandwidth for other applications. Since the bandwidth is reused for traffic with equal or higher priority, then there is a need to make sure the overall multicast traffic does not exceed a specific threshold. Failing to do so could result in congestion with the catastrophic consequence that all multicast channels get affected. Here is where MCAC helps by monitoring multicast bandwidth utilization and denying an IGMP join for a group that would make the overall multicast bandwidth cross a pre-established threshold.

Blocking of IGMP joins with MCAC in this aggregate case should almost never occur and should only be deployed after thorough and extended monitoring of the behavior of IPTV subscribers. In fact, MCAC helps if and only if the variance of the multicast bandwidth utilization at a specific hour is not far from the average across different days, weeks or months for the same hour.

*The data set used in this graph was selected to illustrate the maximum growth in subscribers over a short period of time (about 3 weeks). Doing so eliminates any seasonal influence on subscriptions.

Example Scenario Explaining Multicast Admission Control for IGMP Proxy

The second model detailed for MCAC is IGMP proxy or snooping. In this model, the AN will receive and process IGMP joins from each household. Now the BSR is only processing per-AN IGMP and can only make resource determinations regarding the amount of multicast sent over each M-VLAN interface to the AN.

Figure 5 highlights this model across two households. The BSR keeps a local mapping of multicast groups and their bandwidth. This table can be statically created or dynamically measured.

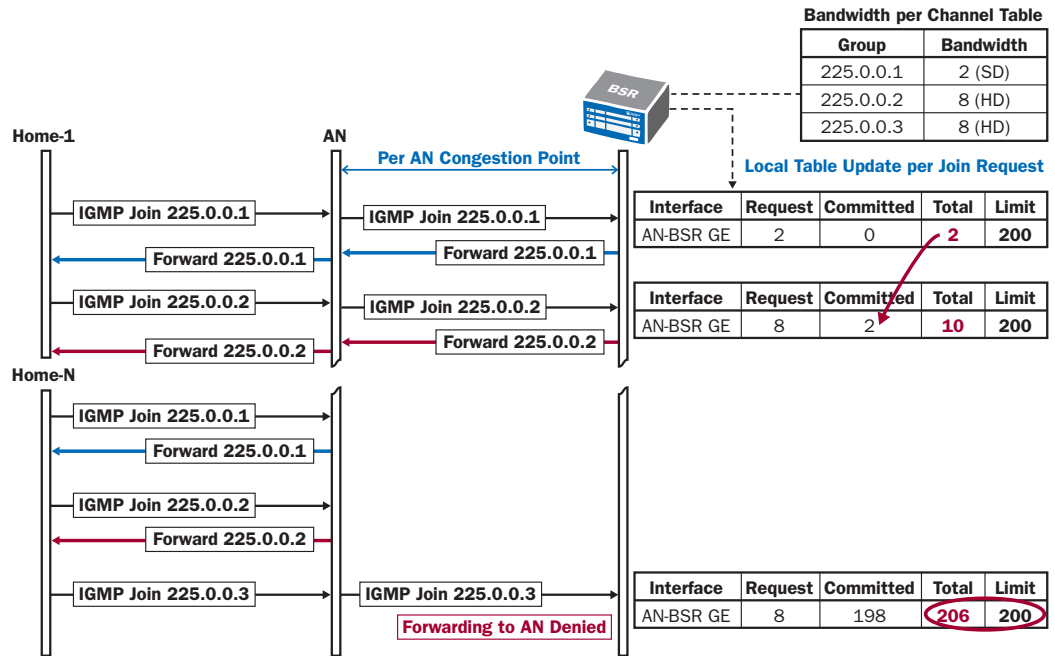


Figure 5: MCAC for Port or M-VLAN Across N Households

The BSR keeps an admission control resource table per M-VLAN mapping to an AN. In this example each M-VLAN is set with a limit of 200 Mbps of multicast. The M-VLAN limit can be set based on planned traffic engineering values.

As each IGMP join request comes in, the BSR sums the amount of currently committed bandwidth and the request amount. As long as this amount is below the configured limit then the interface is sent a copy of the multicast group. The resource state table is then updated with a new committed rate that includes the new multicast group.

If the limit is exceeded, then the BSR will deny the join request and not send a copy of the multicast out that interface.

Since the BSR only sees unique join requests per group, the resource admission control only applies to unique groups per AN. If both homes 1 and 2 join the same group, then only a single resource check happens in the BSR.

If the bandwidth limit is exceeded, then the BSR will deny the join request and not send a copy of the multicast out of that interface. Note that this capability is at the network level. If a join request is denied, then the STB will not receive any type of feedback notification directly from the router. The router will generate a trap that can be used for future capacity planning or integrated to the video application for direct STB feedback messages. Providing feedback to the user is outside the scope of this paper.

Lab Analysis for C-VLAN and Port Multicast Admission Control for IGMP Proxy/Snooping

For the IGMP proxy model, only M-VLAN multicast admission control is configured with a maximum amount of multicast bandwidth allowed to be replicated across that interface using the following JUNOS software command under the M-VLAN subinterface:

```
ip multicast admission-bandwidth-limit 200000000
```

Hereafter it is assumed that the AN forwards proxied or snooped IGMP messages onto only the M-VLAN. However, there might be cases where IGMP joins can be received by the BNG on a different interface than that of the M-VLAN. In these cases a port-based command like the following one should be used in order to track all of the IGMP requests:

```
mroute port <# slot #>/<# bay #>/<# port #> admission-bandwidth-limit <#
bandwidth #>
```

In this test a maximum limit of 200 Mbps is used so that with the 3.5 to 4.0 Mbps SDTV video streams only three unique groups can be joined, although multiple subscribers can all watch the same three streams. The *show ip mroute* output shows that the first three groups (224.0.20.103, 224.0.20.106, 224.0.20.107) were joined successfully.

```
E320_MRA_ER_0#show ip mroute
...
(11.1.5.66, 224.0.20.103) uptime 0 00:00:38
  Admission bandwidth: 4005000 bps (adaptive)
  QoS bandwidth: 4005000 bps (adaptive)
  RPF route: 11.1.5.0/24, incoming interface GigabitEthernet3/1/2
    neighbor 11.1.6.2, owner IsIs (ECMP route)
  Incoming interface list:
    GigabitEthernet2/1/3 (11.1.6.1/30), Discard/Pim
    GigabitEthernet3/1/2 (11.1.6.5/30), Accept/Pim (RPF IIF)
  Outgoing interface list:
    GigabitEthernet15/0/0.222 (192.168.13.1/24), Forward/Pim, 0 00:00:38/
never

(11.1.5.66, 224.0.20.106) uptime 0 00:00:41
  Admission bandwidth: 3864000 bps (adaptive)
  QoS bandwidth: 3864000 bps (adaptive)
  RPF route: 11.1.5.0/24, incoming interface GigabitEthernet2/1/3
    neighbor 11.1.6.2, owner IsIs (ECMP route)
  Incoming interface list:
    GigabitEthernet2/1/3 (11.1.6.1/30), Accept/Pim (RPF IIF)
  Outgoing interface list:
    GigabitEthernet15/0/0.222 (192.168.13.1/24), Forward/Pim, 0 00:00:41/
never

(11.1.5.66, 224.0.20.107) uptime 0 00:00:21
  Admission bandwidth: 4252000 bps (adaptive)
  QoS bandwidth: 4252000 bps (adaptive)
  RPF route: 11.1.5.0/24, incoming interface GigabitEthernet3/1/2
    neighbor 11.1.6.2, owner IsIs (ECMP route)
  Incoming interface list:
    GigabitEthernet2/1/3 (11.1.6.1/30), Discard/Pim
```

```

GigabitEthernet3/1/2 (11.1.6.5/30), Accept/Pim (RPF IIF)
Outgoing interface list:
  GigabitEthernet15/0/0.222 (192.168.13.1/24), Forward/Pim, 0 00:00:21/
  never

```

However, when the fourth group (224.0.20.108) join request is received by the BSR it will exceed the admission bandwidth limit, resulting in a non-replicated group:

```

(11.1.5.66, 224.0.20.108) uptime 0 00:00:04
Admission bandwidth: 3752000 (adaptive)
QoS bandwidth: 3752000 bps (adaptive)
RPF route: 11.1.5.0/24, incoming interface GigabitEthernet2/1/3
  neighbor 11.1.6.2, owner IsIs (ECMP route)
Incoming interface list:
  GigabitEthernet2/1/3 (11.1.6.1/30), Accept/Pim (RPF IIF)
Outgoing interface list:
  GigabitEthernet15/0/0.222 (192.168.13.1/24), Blocked (intf-adm-limit)/
  Pim, 0 00:00:05/never

```

This group is shown as *blocked*, denoting that group replication is not occurring due to an interface administrative limit (*intf-adm-limit*). This type of blocking impacts any subscriber issuing a join request for a new multicast group that must be delivered over the M-VLAN to the AN. If the AN proxy join request is beyond the 15 Mbps M-VLAN bandwidth limit, then AN replication cannot occur.

If per-subscriber admission control limits are required, this must be performed in the AN. The AN used in the test scenario does not support this feature and therefore was not tested. The BSR cannot provide per-subscriber multicast control since it is not responsible for per-subscriber replication.

Group Priorities for Multicast Admission Control

Instead of just using a simple group-bandwidth table for multicast admission control, JUNOS software provides the extra capability of prioritizing groups that are part of the decision process.

By default, any entry in the route-map that is used to map a bandwidth to a multicast address or range of addresses also maps a priority of 0 to those addresses. If some of the groups cannot be blocked in any circumstance, then a priority higher than 0 should be associated with those groups.

```

!
route-map "per-group-mcast-bandwidth" permit 10
  match ip address "PAY-TV-SDTV"
  set admission-bandwidth adaptive
  set priority 1
route-map "per-group-mcast-bandwidth" permit 20
  match ip address "SDTV"
  set admission-bandwidth adaptive
!

```

Note when the port-level command is used:

```

mroute port <# slot #>/<# bay #>/<# port #> admission-bandwidth-limit
<# max-bandwidth #> priority-bandwidth-limit <# priority-bandwidth #>
hysteresis <# value #>

```

All groups are accepted until the *priority-bandwidth* level has been reached. Once that threshold has been hit, only those groups for which the route-map *per-group-mcast-bandwidth* associates a priority 1 are honored (unless already served below the *priority-bandwidth level*) until the second bandwidth threshold *max-bandwidth* is reached. After that all groups are blocked.

The hysteresis value is a percentage number and it relates to the *priority-bandwidth*. Once the *priority-bandwidth-limit* has been reached, in order to accept again non-prioritized channels multicast bandwidth has to go below that percentage.

Please note that there is an implicit, non-configurable hysteresis value also for the *admission-bandwidth-limit*, which is 85 percent.

Show commands are similar to those already seen in this paper. The following is an example of an accepted prioritized group:

```
E320_MRA_ER_0#show ip mroute
...
(11.1.3.26, 224.0.17.104) uptime 0 00:14:56
  Admission bandwidth: 3907000 bps (adaptive)
  QoS bandwidth: 3907000 bps (adaptive)
  Priority: 1
  RPF route: 11.1.3.0/24, incoming interface GigabitEthernet2/1/3
    neighbor 11.1.6.6, owner IsIs (ECMP route)
  Incoming interface list:
    GigabitEthernet2/1/3 (11.1.6.1/30), Accept/Pim (RPF IIF)
  Outgoing interface list:
    GigabitEthernet15/0/0.600 (192.168.15.1/24), Forward/Pim, 0
00:15:00/never
```

Summary

There are a range of trade-offs involved in designing an IPTV network that delivers content using both unicast and multicast technologies. By choosing the right architecture at key points in the network, service providers can oversubscribe the edge network and still deliver the quality of experience (QoE) needed to satisfy customer demand and thus to compete effectively. The network designer must also be aware of the capabilities that are needed in key network components—in both the data and control planes—to support the desired architecture.

Juniper Networks has the expertise and product portfolio that service providers need to design the next-generation networks that can deliver a high QoE and maximize investment in infrastructure and administrative staff. To learn more, go to www.juniper.net.

About Juniper Networks

Juniper Networks, Inc. is the leader in high-performance networking. Juniper offers a high-performance network infrastructure that creates a responsive and trusted environment for accelerating the deployment of services and applications over a single network. This fuels high-performance businesses. Additional information can be found at www.juniper.net.

**CORPORATE HEADQUARTERS
AND SALES HEADQUARTERS FOR
NORTH AND SOUTH AMERICA**
Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or 408.745.2000
Fax: 408.745.2100
www.juniper.net

**EUROPE, MIDDLE EAST, AFRICA
REGIONAL SALES HEADQUARTERS**
Juniper Networks (UK) Limited
Building 1
Aviator Park
Station Road
Aldlestone
Surrey, KT15 2PG, U.K.
Phone: 44.(0).1372.385500
Fax: 44.(0).1372.385501

EAST COAST OFFICE
Juniper Networks, Inc.
10 Technology Park Drive
Westford, MA 01886-3146 USA
Phone: 978.589.5800
Fax: 978.589.0800

ASIA PACIFIC REGIONAL SALES HEADQUARTERS
Juniper Networks (Hong Kong) Ltd.
26/F, Cityplaza One
1111 King's Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803

Copyright 2008 Juniper Networks, Inc. All rights reserved. Juniper Networks, the Juniper Networks logo, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. JUNOS and JUNOSe are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

**To purchase Juniper Networks solutions, please
contact your Juniper Networks sales representative
at 1-866-298-6428 or authorized reseller.**