

# Multicast Deployment Guide

## *Foundry Enterprise Network Solutions*

April 2006

### Table of Contents

Summary .....	2
Overview of Foundry's Multicast Solution .....	2
Addressing .....	3
Protocols .....	3
RP Configuration .....	4
Controlling Network Traffic .....	4
Layer 2 Core, Layer 3 Edge .....	4
Layer 3 Core, Layer 2 Edge .....	6
Scaling the Network with MSDP .....	7
Conclusion .....	7

## Summary

Multicast Services are increasingly finding their way into the day to day life of small-to-medium and enterprise businesses that require its services for essential communications. Not only is multicast needed for its more traditional uses like video-conferencing and IPTV, but it is an essential building block for the next generation of applications like software distribution, multimedia, and leading-edge “push” applications like news and sports updates and real-time stock quotes. Network multicasting is also crucial for universities for distance learning, digital video libraries, online collaboration tools, and other types of advanced applications important to research and education.

Today's IP Multicast applications fall under one of three categories:

- One-to-many. This is when one source is sending out to multiple receivers, and includes applications like database updates, news feeds, lectures, and monitoring.
- Many-to-one: This is when any number of receivers sends data back to a source, and either side can generate the request. This is used in applications like data collection, auctions, resource discovery, and polling.
- Many-to-many: This is when any number of hosts is sending and receiving from the same multicast addresses. This is used for multimedia conferencing, distance learning, chat groups, and interactive music sessions.

With all these varying applications and uses, it would seem that deploying multicast would be a complex undertaking. However, since the traffic is running over a regular IP network, Small-to-Medium and Enterprise customers can continue to use their standard Ethernet network. Because the voice, data, and video services are all using a common core, the infrastructure costs do not need to take into account separate networks for each service or application. The IP network infrastructure is simplified so that companies can scale bandwidth and provide new services relatively easily.

Foundry's product portfolio provides advanced features to help build flexible IP Multicast networks that are scalable, optimize bandwidth, and enable truly distributed applications. This document will cover the multicast features that are supported in Foundry's FastIron SuperX Family of products as well as the FastIron Edge Switch X and FastIron Workgroup Switch X products. The purpose of this document is to discuss the key points that need to be addressed when deploying a multicast network, and how Foundry's IronWare features can be used to optimize multicast traffic flow throughout.

## Overview of Foundry's Multicast Solution

Foundry's proprietary Metro Ring Protocol (MRP) provides a rapidly-converging Layer 2 or Layer 3 backbone made up of FastIron SuperX switches. In a Layer 2 core, MRP coupled with PIM Snooping optimizes network bandwidth because multicast traffic will only be forwarded out ports with interested receivers. In a Layer 3 core, PIM Routing and Passive Multicast Route Insertion (PMRI) can be used together to ensure that multicast streams are only forwarded out ports with interested receivers and unwanted traffic is dropped in hardware.

In the aggregation layer where 10GbE ports can be used to bring in traffic to and from the core, Protected Link can be used for added availability. For larger networks that are expanding, Super Aggregated VLANs (SAV) or Q-in-Q can be used to scale up to sixteen million VLANs. The edge devices can provide either Layer 2 or Layer 3 routing functionalities that can take advantage of PIM Routing and Passive Multicast Route Insertion to ensure that multicast streams are forwarded only where they are wanted and unwanted traffic is dropped in hardware. The X-Series switches also support Source Specific Multicast (SSM), which helps enable Layer 3 customer equipment to interoperate with Layer 3 signaling and prevent distributed denial-of-service (DDoS) attacks.

Finally, the X-Series provides security measures like inbound, ACL-based rate-limiting which can be enabled per-port or per-VLAN. In addition, each individual switch can be configured to protect against control plane attacks, malicious man-in-the-middle attacks, and unauthorized user access.

Following are some of the issues that need to be addressed when deploying a multicast network, from assigning addresses to deciding where the RP should reside. When evaluating which multicast protocol to deploy, it is necessary to take into account what kind of topology already exists, where the receivers are positioned throughout the network, and the amount of multicast traffic that needs to be forwarded.

## Addressing

One of the first decisions that must be made when deploying Multicast is determining the address of the multicast stream or streams. The Internet Assigned Numbers Authority (IANA) has assigned some multicast group addresses as well-known addresses, and these are reserved. Well-known addresses include some of the following:

224.0.0.1 All hosts on a subnet  
224.0.0.2 All routers on a subnet  
224.0.0.9 RIPv2  
224.0.1.1 Network Time Protocol (NTP)  
224.0.0.13 Protocol Independent Multicast (PIM)

The IP addresses are then mapped to MAC addresses such that only the lowest 23 bits of the multicast group ID are copied to the ethernet address. This means that all multicast addresses will begin with 01:00:5e and the rest will be filled by the lowest 23 bits of the IP address. For example, a multicast IP of 225.8.20.153 will map to a multicast MAC of 01:00:5e:08:14:99 and 224.128.0.1 will map to a multicast MAC of 01:00:5e:00:00:01.

Notice that this last MAC address is the same MAC that the “all hosts” IP address uses. If 224.128.0.1 was used as a video stream in a network, problems would arise because the switch would recognize the MAC address as a “known” address and it would send all that traffic to the CPU. For that reason, it is crucial to be aware of what MAC is associated with the chosen IP so that a known MAC does not get used for multicast data traffic.

## Protocols

Foundry offers both protocol-dependent and protocol-independent options for routing multicast traffic through a Layer 3 network:

**Distance Vector Multicast Routing Protocol (DVMRP)** was the first multicast routing protocol developed, and it is not a highly scalable option. DVMRP must calculate and exchange its own RIP-like routing metrics, and since it operates in dense-mode, it will flood data through the network and then prune off branches so that every router in the network must maintain the state for every source.

**Protocol Independent Multicast (PIM)** is supported in three different modes on the FastIron SuperX and FastIron Edge X switches: Dense, Sparse, and Source Specific modes:

***Dense Mode:*** Dense mode floods all branches of the network periodically and then prunes branches that have no multicast group receivers. This mode is most effective in environments that have receivers densely distributed throughout the network and where it is likely that all subnets will have interested receivers. It is important to keep in mind that dense mode requires far more available bandwidth in the network because all multicast streams will be flooded to every subnet.

The primary benefit of dense mode is that it is easy to configure and easy to maintain. In addition, there is little overhead in the control messaging required to maintain PIM-DM and only one path is established for each stream.

***Sparse Mode:*** Sparse mode makes use of a Rendez-Vous point, which registers all of the sources and receivers in the network. Through their Designated Routers, a receiver will then send a message to the RP letting it know it is interested in receiving the traffic. Originally, the traffic will flow through the RP to get from source to receiver, but then a shortest-path tree will be created so that the traffic flows along the most efficient path through the network.

Sparse mode is beneficial because it does not waste bandwidth by flooding multicast traffic to uninterested receivers. Instead, the RP waits for an explicit join from the receiver before forwarding any traffic. In addition, PIM Sparse Mode is highly scalable. Individual PIM-SM domains can be connected using MBGP or MSDP to provide native multicast services over the Internet.

***Source-Specific Mode:*** Source Specific mode is similar to Sparse mode, but instead of simply joining a group, the receiver uses IGMPv3 messages to join a group sent from a specific source. This way, instead of the receiver getting multiple copies of the same stream from multiple sources, it will only receive traffic from the one source to which it joins.

The benefits of SSM are clear: Not only is the amount of unwanted traffic in the network reduced exponentially, but because each multicast group is associated with a particular host, different hosts can be assigned the same multicast address for different streams. This greatly increases the number of multicast groups that can be used in the network. Another added benefit of SSM is that it increases security by reducing the possibility of a rogue source disrupting the traffic from a legitimate source.

These days, almost all multicast deployments use some variation of PIM. While PIM-DM is good for smaller networks that have a lot of receivers on various subnets and bandwidth to spare, PIM-SM is better for larger, more complex multicast deployments. PIM-SSM adds a new angle to PIM by making bandwidth usage more efficient and security more available.

## RP Configuration

There are a couple considerations when configuring the Rendez-Vous point in a multicast network: Static or dynamic RP? And where should the RP be located?

Configuring dynamic RPs is almost always the better option. Assigning a dynamic Rendez-Vous point ensures that if the primary RP goes down, the multicast traffic does not need to stop altogether. Instead, a backup can come up and take over the work of the primary RP. Static RP is most useful for troubleshooting purposes in the event that dynamic RP configuration is not working properly.

The physical location of the RP is not generally a factor. Since Sparse Mode builds a shortest-path tree once the first packet is received by the last-hop router, the RP will not be a part of the data path for very long.

## Controlling Network Traffic

Foundry's FastIron SuperX and FESX Family of switches offer several alternatives for controlling multicast traffic throughout the network for both Layer 2 and Layer 3 cores.

### *Layer 2 Core, Layer 3 Edge*

In Layer 2 core networks, multicast flooding is an issue because the default behavior of a pure Layer 2 switch is to flood multicast traffic on the broadcast domain. Take the case of an MRP ring made up of several SuperX switches that each connect to one branch of a medium- to large enterprise office. If multicast traffic is being sent between any two of those branches, every branch in the business could potentially receive all of those streams, regardless of whether or not they have interested receivers.

A simplified illustration of how this would work is in Figure 1 below.

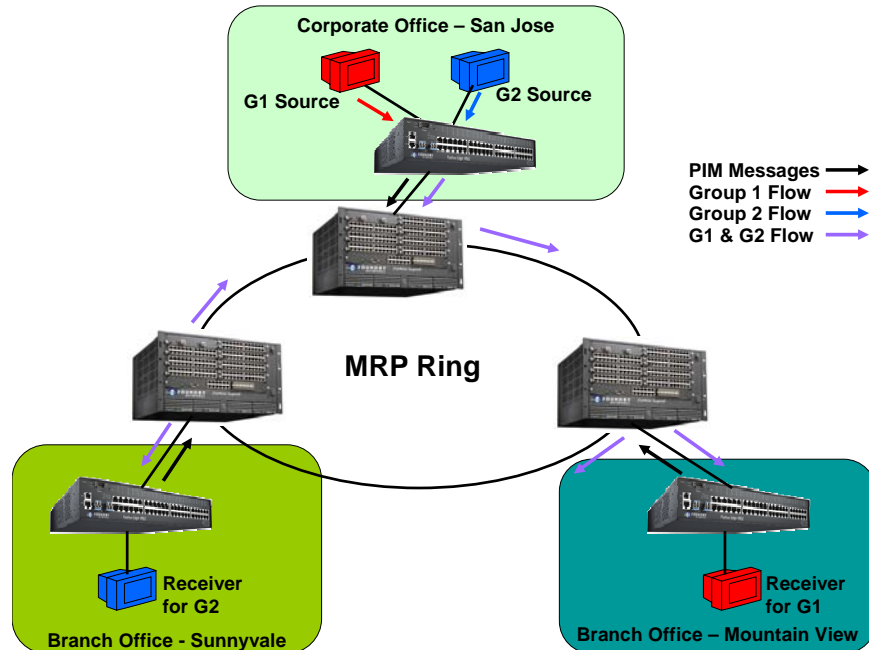


Figure 1 – Multicast Flooding in an L2 Core

While the Branch office switches are sending PIM joins to the Corporate Office switch, the problem is that the Layer 2 MRP switches do not pay attention to PIM and therefore they will simply forward all the multicast traffic throughout the VLAN.

Enabling PIM Snooping on the MRP switches, however, prevents this from happening. With PIM Snooping, each SuperX will inspect each PIM message that it sees to determine which ports have interested receivers downstream of them. This way, all the switches connected to branches that do not want the multicast traffic will not forward them downstream, and only those devices that want the traffic will receive it. This is illustrated in Figure 2 below.

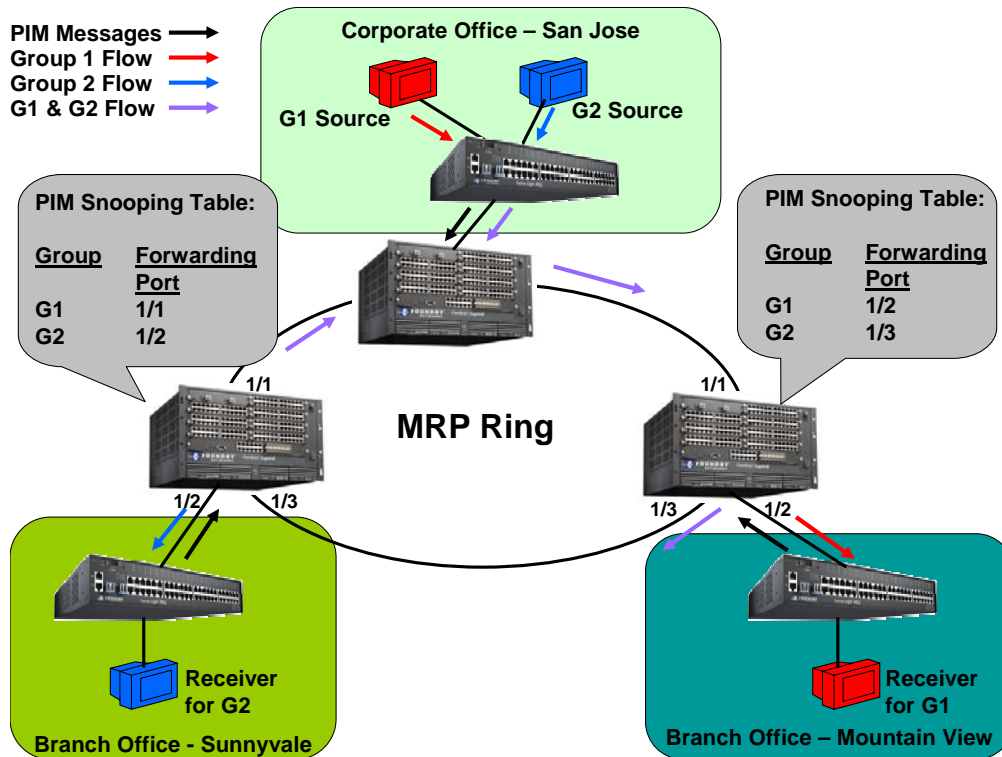


Figure 2 – PIM Snooping in a L2 Core

In that same scenario, however, if the branch offices have several routers, some with interested receivers and others without, it would be possible for some switches to still receive undesired traffic. Passive Multicast Route Insertion is available on the FES-X and SuperX to ensure that that traffic does not hit the CPU and is instead dropped in hardware.

### ***Layer 3 Core, Layer 2 Edge***

If the core of the network is a Layer 3 MRP ring, PIM-Sparse or PIM-SSM can be configured on each switch to ensure that the traffic is only forwarded to the appropriate downstream routers. In this case, the Layer 3 switches will use PIM messaging to notify their PIM neighbors of which, if any, streams they are interested in.

At the Layer 2 edge, IGMP Snooping is required to minimize multicast flooding on those devices. IGMP Snooping forces the switch to inspect all IGMP messages, and the switch then builds a forwarding table based on which ports sent IGMP joins for which groups. This way, when multicast traffic comes in from the router port, it will only be forwarded to those ports that are in the forwarding table for that particular group.

Figure 3 below shows an example using these features:

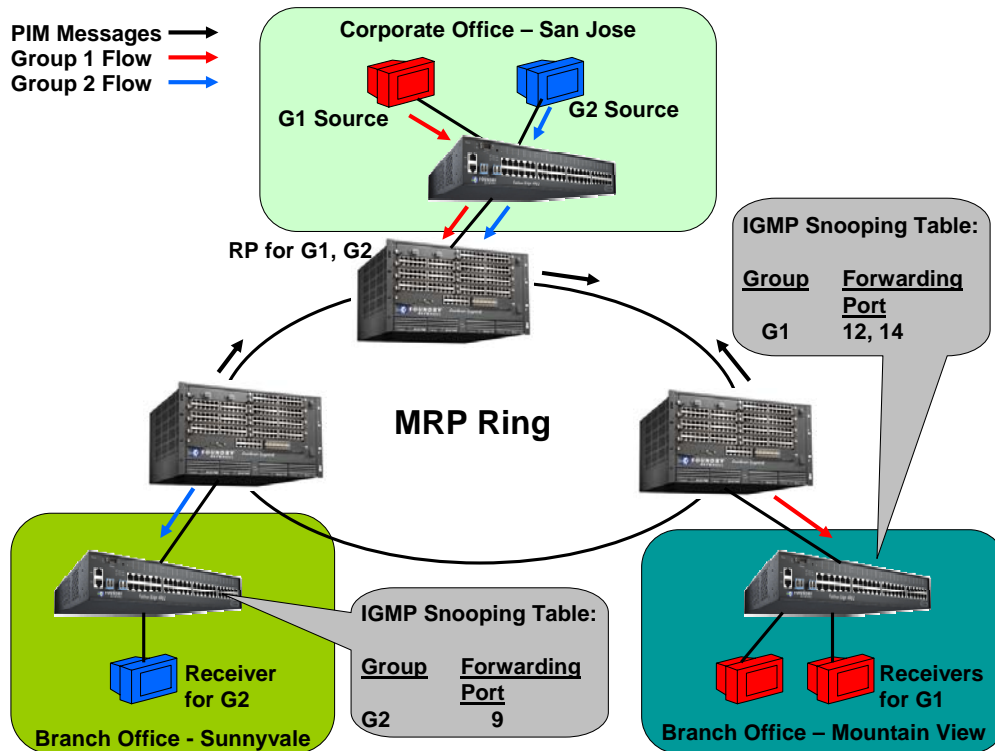


Figure 3 – PIM-SM and IGMP Snooping in a L3 Core

Regardless of which topology is used, Foundry's advanced multicast features ensure that flooding in the network is limited so that bandwidth is used as efficiently as possible. PIM, PIM Snooping, PMRI, and IGMP Snooping are all critical to optimizing network performance in an IP Multicasting environment.

## Scaling the Network with MSDP

As businesses expand and networks grow in size, it could become necessary to connect PIM domains and allow multicast applications to reach offices across the core or even across the network. Foundry's BigIron RX and BigIron Jetcore chassis both support Multicast Source Discovery Protocol (MSDP) which allows for just that.

Using only PIM for expansion causes some difficulty because it requires that everyone on both sides of the network be able to reach the RP. This is one case where placement of the RP would be important. However, MSDP simplifies the expansion by allowing each PIM domain to remain separate and have its own RP.

The RP in each domain will establish an MSDP peering either with RPs in other domains or with border routers leading to other domains. When an RP learns about a new multicast source in its own domain (using the normal PIM registration process), it will then send a Session Advertisement to all of its MSDP peers letting them know about this new stream. This way multicast traffic can be received at all corners of the network without having to reconfigure each already-existing PIM domain.

## Conclusion

As IP Multicasting applications are increasing in size and scale, enterprise networks can continue to provide their employees with the ability to use this technology fairly easily. Since the multicast traffic can use the already-existing infrastructure, there is no added cost from having to build or maintain an entire separate network. In addition, multicast traffic can be controlled and bandwidth can be optimized by taking advantage of Foundry's advanced features, including PIM Snooping and PIM-Source Specific Multicast. Multicast domains can also be

expanded across networks- or the internet- through the use of MSDP. As multicast applications become more and more a part of our daily life at home and at work, it is important for the network administrators to ensure that the network stays up and that it is designed to forward multicast as optimally as possible.

Mirah Sederlof  
Product Marketing Engineer

Foundry Networks, Inc.  
Headquarters  
4980 Great America Parkway  
Santa Clara, CA 95054

U.S. and Canada Toll-free: (888) TURBOLAN  
Direct telephone: +1 408.586.1700  
Fax: +1 408.586.1900  
Email: [info@foundrynet.com](mailto:info@foundrynet.com)  
Web: <http://www.foundrynet.com>

Foundry Networks, AccessIron, BigIron, EdgeIron, FastIron, IronPoint, IronView IronWare, JetCore, NetIron, ServerIron, Terathon, TurboIron, and the "Iron" family of marks are trademarks or registered trademarks of Foundry Networks, Inc. in United States and other countries. All other trademarks are the properties of their respective owners.

Although Foundry has attempted to provide accurate information in these materials, Foundry assumes no legal responsibility for the accuracy or completeness of the information. More specific information is available on request from Foundry. Please note that Foundry's product information does not constitute or contain any guarantee, warranty or legally binding representation, unless expressly identified as such in a duly signed writing.

©2006 Foundry Networks, Inc. All Rights Reserved.