WHITE PAPER

EMERGING MULTICAST VPN APPLICATIONS

Drivers for Scalable Next-Generation Multicast VPN

Table of Contents

Table of Figures

Executive Summary

Several emerging multicast applications require virtualization to support service delivery, traffic separation, and new business models. A more scalable and flexible approach to implementing multicast VPN (MVPN) is being deployed in service provider, financial, and enterprise networks to effectively support these emerging applications, while achieving greater economies of scale.

Introduction

The volume of multicast traffic has been growing primarily based on the emergence of video-based applications. There are a growing number of Layer 3 VPN customers who have IP multicast traffic. (These Layer 3 VPNs are commonly known as 2547 VPNs and are based on the original RFC 2547.) Service providers who have an installed base of Layer 3 VPN for unicast service are looking at upselling with new media-rich solutions to increase ARPU while achieving operational efficiency. The significant interest in IPTV services and wholesale business models are driving the need to consider more scalable ways to deliver multicast services. Similarly, the delivery of multicast-dependent financial services requires scalable and reliable MVPN infrastructures. In addition to Layer 3 MPLS VPN services, service providers and cable operators are beginning to offer virtual private LAN service (VPLS) to small and medium businesses, as well as to large enterprises. There is also a growing momentum towards IPv6-enabled services based on new mobile applications. This diversity and breadth of services pose a challenge for operators to create an infrastructure that supports Layer 2 VPNs, Layer 3 VPNs, IPv4, IPv6, unicast, and multicast traffic. The difficulty is particularly true for virtual services that require complex control and data plane operations. Another challenge is to support emerging multicast applications incrementally on top of existing Layer 3 VPN and VPLS infrastructures without adding operational complexity.

Background on Multicast VPN

MVPN is a technology to deploy multicast service in an existing VPN or as part of a transport infrastructure. Essentially, multicast data is transmitted between private networks over a VPN infrastructure by encapsulating the original multicast packets.

Layer 3 BGP-MPLS VPNs are widely deployed in today's networks. RFC 4364, which supersedes RFC 2547, describes protocols and procedures for building BGP-MPLS VPNs for forwarding VPN unicast traffic only. An incremental approach for deploying multicast services can use the same technology as used for deploying Layer 3 VPN for unicast services. This approach can reduce the operational and deployment risk, as well as qualification cost and OPEX, resulting in a higher ROI. As multicast applications, such as IPTV and multimedia collaboration, gain popularity, and as the number of networks with different service needs merge over a shared MPLS infrastructure, the demand for a scalable, reliable MVPN service is rapidly increasing.

Basic Implementation Model

For MVPNs, a mechanism is needed to carry MVPN multicast routing information (control plane). This information is carried from the provider edge (PE) routers connected to the sites that contain the receivers to the PEs connected to the sites containing the sources. This path allows the receivers to inform the sources that the receiver sites want to receive traffic from the sender sites.

As well, a mechanism is needed to carry multicast traffic (data plane). This information is carried from the PE routers connected to the sites that contain the sources to the PEs connected to the sites that contain the receivers. This path enables the flow of multicast traffic from the sources to the receivers.

Typically, an MVPN provider network carries multicast routing information from the receivers to the sources and carries multicast data traffic from the sources to the receivers. Different MVPNs may use the same address space (RFC 1918), including an IP multicast addressing space. It would be ideal to use the same route distinguisher mechanism as used by 2547 VPN for unicast to support the overlapping address space.

Example MVPN Operational Model

In Figure 1, the control plane path for MVPN X is from the PEs connected to the receiver sites (Site A and Site C) to the PEs connected to the sender sites (Site B and Site C). The data plane path for MVPN X is from the PEs connected to the sender sites (Site B and Site C) to the PEs connected to the receiver sites (Site A and Site C). Here, Site C is both a receiver and sender site.

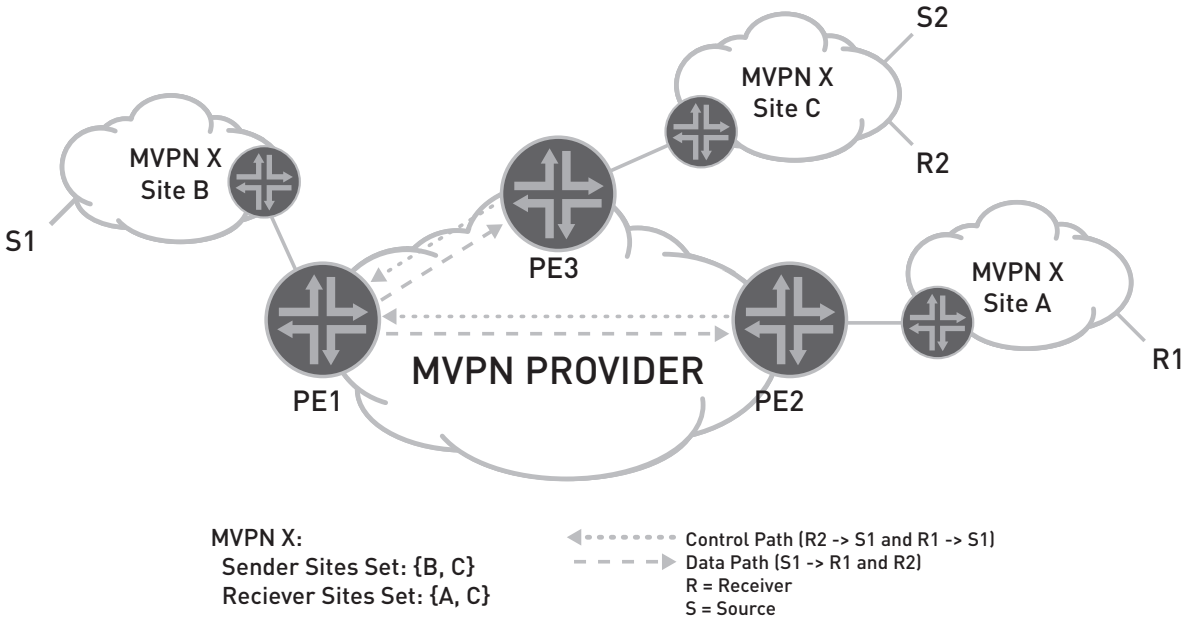

Figure 1: MVPN operational model

Issues with Prior Schemes

Traditionally, VPN multicast traffic is overlaid on top of a BGP-MPLS network based on Draft-Rosen, which specifies a virtual router architecture. This model uses PIM-SM for exchanging control plane information, as well as setting up multicast forwarding state on the Layer 3 VPN infrastructure. The customer domain multicast (C-multicast) protocol information, typically customer PIM (C-PIM) join/prune messages, received from local customer edge (CE) routers is propagated to other PEs using these PE-PE PIM sessions across the VPN-specific virtual network. The most important issue here is that the PIM sessions are between the VRFs. Thus, for a given MPVN, a PE maintains a PIM session with every other PE that has membership in that MVPN. This complexity poses a significant scaling challenge. For example, with 1000 MVPNs per PE and 100 sites per MVPN, there would be 100,000 PIM neighbors per PE, which results in 3300 PIM hellos/second.

On the MVPN provider network, VPN multicast control and data traffic are forwarded via multipoint GRE tunnels. These tunnels are signaled by a different instance of PIM protocol running on the provider side (P-PIM). The GRE header used for tunneling VPN multicast data and control traffic is a multicast group address assigned to that VPN by the provider. This multicast group address gives the GRE the multipoint property as these tunnels are, in fact, VPN multicast distribution trees (MDTs).

With this virtual router model, control and data plane scaling issues arise and MVPN providers must maintain two different routing and forwarding mechanisms for VPN unicast and multicast services (Table 1). As a result, there is an effort by the IETF Layer 3 VPN working group to specify the next-generation MVPNs (also referred to as NG MVPNs). The working group published an IETF draft (draft-ietf-l3vpn-2547bis-mcast) that is a superset of the previous specifications for MVPNs. There are two main drafts: draft-ietf-l3vpn-2547bis-mcast-bgp, which outlines the procedures of the NG MVPN, and draft-ietf-l3vpn-2547bis-mcast, which describes the building blocks for the different options.

NG MVPN Approach

The NG MVPN drafts introduced a BGP-based control plane that is modeled after its highly successful counterpart of the VPN unicast control plane. NG MVPNs adopted two important properties of unicast BGP-MPLS VPNs:

The BGP protocol is used for distributing all necessary routing information to enable VPN multicast service. This protocol allows service providers to leverage their knowledge and investment in BGP-MPLS VPN unicast services to offer VPN multicast services.

The use of BGP for distributing C-multicast routes results in the control traffic exchange being out-of-band from the data plane. This implementation allows for the separation of the control and data plane protocols and makes it easier to leverage newer transport technologies, such as point-to-multipoint (P2MP) MPLS, in delivering MVPN services.

The BGP-based NG MVPN control plane lends itself naturally to supporting flexible topologies, such as extranet and hub and spoke, as well as IPv6 support. IPv6 NG MVPN provides the ability to naturally use MPLS encapsulation for IPv6 multicast. It also uses the same model as IPv6 VPN (as defined in RFC 4659) for unicast. Thus, service providers are ensured of a smooth integration of IPv6 multicast services with an existing IPv4 NG MVPN or IPv6 unicast VPN model. BGP MVPNs also provide multihoming support for connecting a multicast source to two PEs, thus enabling sub-second failover from a PE node failure. The autodiscovery of MVPN members available with the BGP approach provides a high degree of automation for establishing provider tunnels that are used for carrying MVPN data among PEs.

Table 1: Comparison of Draft-Rosen and NG MVPN

	DRAFT-ROSEN	NG MVPN
Transport	PIM-SM GRE	Different MVPNs can use different tunneling technologies (P2MP MPLS or PIM-SM GRE).
Signaling	PIM	BGP, same model as unicast Layer 3 VPN. Supports autodiscovery of routes.
PE-PE signaling sessions	Each PE needs a separate PIM adjacency with each remote PE per VRF.	Each PE uses existing IBGP sessions, which may only require sessions with the route reflectors.
VPN traffic aggregation	No ability to aggregate multiple MVPNs into a single inter-PE tunnel.	It is possible to aggregate multiple (S,G) of a given MVPN into a single selective tunnel and aggregate multiple P2MP LSPs using P2MP LSP hierarchy.
Inter-AS operations	Inter-AS/inter-provider operations options B and C (as defined in RFC 4364) require PEs in different ASes/providers to have direct PIM routing peering.	NG MVPN seamlessly works with all three options (A, B, and C as defined in RFC 4364) available for inter-AS unicast. It also has the concept of segmented inter-AS trees that allows each AS to independently run a different tunneling technology.
Provider Tunnel (P-tunnel) mesh requirement	Required between the PEs, which forces providers to sell an MVPN service where every customer site can be a source and a receiver.	P-tunnel mesh requirement is removed. Allows providers to support MVPN customers where multicast sources can be limited to a subset of its sites. Provides the flexibility to build pricing models for an MVPN service based on sites connected to either sources or receivers or both.

Migration Options

Network operators want to ensure their existing MVPN deployments can easily migrate to the next-generation solution. The flexibility of NG MVPN allows the BGP control plane numerous scaling advantages. For example, as part of one migration approach, an MVPN VRF on a PE could run both Draft-Rosen and BGP-MVPN control plane while the data plane continues to be a PIM-SM GRE. This approach allows all the PEs to be upgraded to the BGP-MVPN control plane while they continue to run the Draft-Rosen control plane. Then the Draft-Rosen control plane can be deconfigured one <PE, VRF> at a time to seamlessly migrate all the PEs to the BGP-MVPN control plane. Thus, providers could migrate to NG MVPN one MVPN at a time and/or one PE at a time. The data plane can then be migrated seamlessly to P2MP MPLS over time.

NG MVPN also inherently supports the Carrier-of-Carrier model very similar to the 2547 Layer 3 Unicast VPN model. This effective technique can be used for migrating from Draft-Rosen islands to NG MVPN. The Draft-Rosen VPN islands can seamlessly be carried over NG MVPN, which gives providers the option to gain experience with the NG MVPN without disruption of existing deployments. This approach allows providers to grow the NG MVPN carrier AS while limiting the growth of the Draft-Rosen VPN islands.

There are several MVPN deployment considerations discussed in draft-rekhter-mboned-mvpn-deploy-00.

Next-Generation Multicast in VPLS

VPLS, a key enabler for delivering multipoint Ethernet services and Layer 2 backhaul of DSL or mobile traffic, requires an efficient mechanism to transport multicast traffic. By default, VPLS implementations use ingress replication, which does not offer bandwidth efficiency for multicast traffic. The use of P2MP LSPs with BGP VPLS allows replication on the network only where it is required.

In this approach, each PE needs to tell the other PEs the identity of the RSVP-signaled P2MP LSP on which it will send the traffic (multicast, broadcast, or unknown) for a particular BGP-VPLS instance. It sets up the control plane by putting the RSVP session object of the P2MP LSP in a BGP update.

The key advantage is that VPLS leverages the BGP control plane used for NG MVPN to enable P2MP LSPs. Further, BGP autodiscovery allows for the setup of dynamic P2MP LSPs such that the leaves of the P2MP RSVP-TE LSP do not need to be statically configured. In other words, when PE autodiscovers new PE members of a VPLS instance through BGP, it automatically adds a new leaf to the corresponding P2MP LSP.

Providers who want to offer multicast virtualization services for Layer 3 VPN and VPLS have an option to use a common control plane (BGP) and data plane (MPLS) framework that leads to a consistent service model and simplified operations.

Emerging MVPN Applications

There are several multicast applications driving the deployment of new MVPNs and requiring the level of scale that is particularly addressed by the NG MVPN solution. Some of the key emerging applications being discussed include the following:

- Layer 3 VPN multicast service offered by service providers to enterprise customers
- Video transport applications for separation/virtualization between different customers
  - IPTV wholesale
  - Multiple content providers attached to the same network
- Distribution of media-rich financial services or enterprise multicast services
- Multicast backhaul over a metro network

Layer 3 VPN Multicast Service

The BGP-MPLS Layer 3 VPN (2547 VPN) is widely deployed for unicast service. Service providers want to extend 2547 VPN service offerings to include support for IP multicast.

In this model (Figure 2), the MVPN provider PE has the VPN service instance configured to learn information about multicast sources and receivers. The actual sources and receivers reside on the customer sites that are connected to the MVPN provider, and the information about them is propagated to all PEs participating in the MVPN.

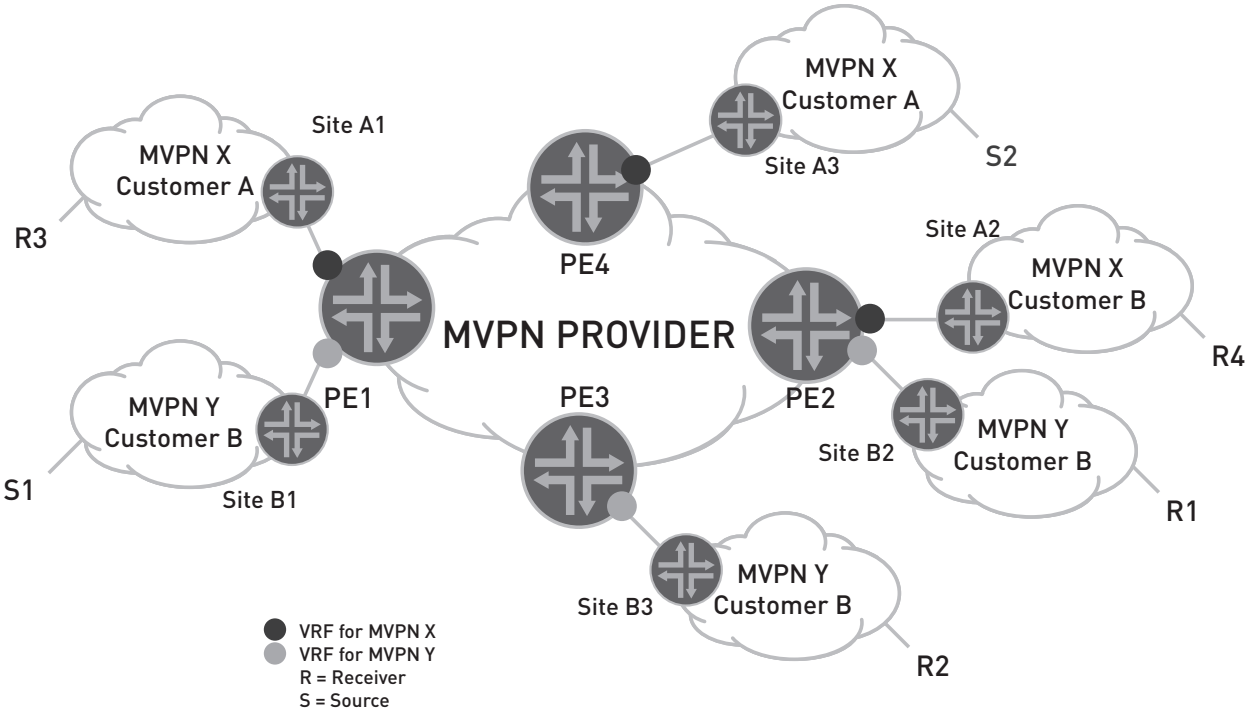

Figure 2: Layer 3 VPN multicast service

Some of the implementation goals for this type of service that are effectively achieved by NG MVPN include:

- Follow the same architecture/model as 2547 VPN unicast. There is no need to have the virtual router model (Draft-Rosen) for multicast and the 2547 VPN model for unicast.
- Reuse 2547 VPN unicast mechanisms, with extensions, as necessary. Thus, there is no need to have PIM/GRE for multicast and BGP MPLS for unicast.
- Retain as much flexibility and scalability as possible of the 2547 VPN unicast architecture.

From a functional view within the MVPN provider, the following operations need to happen:

- Communication between PEs of which multicast tree type will be used as tunnels for carrying MVPN traffic and (if necessary) which particular tree. This operation is easily done by autodiscovery mechanisms in the case of BGP-based NG MVPNs.
- Exchange of VPN-IPv4 unicast routes between PEs (multicast sources in the customer domain would be among these routes). This operation is done by standard route target parameters that are also used for Layer 3 unicast VPNs.
- Discovery of the location of multicast receivers for each multicast group in each MVPN. For NG MVPN, this operation is again easily facilitated by BGP using MCAST-VPN NLRI.

Broadcast Video and IPTV Wholesale

MVPN is an effective IPTV distribution mechanism. In particular, the NG MVPN scheme is a superior solution if an MVPN provider wants P2MP MPLS in the core and PIM islands at the edge of the network, with traffic sent into P2MP LSPs determined by PIM join messages from the edge. This implementation allows for an efficient integration between the core and edge multicast domains.

The other primary driver for using MVPN for video transport applications is to separate the multicast traffic across customer or administrative boundaries, as in the case of wholesale.

- If the receiver sites exist in separate customer sites or separate service provider networks, then it is required to separate the video traffic being distributed to the receiver sites.
- If the sources reside in separate administrative domains (for example, in separate content provider networks), then it is required to separate the traffic path going from the individual sources to the subscribing receiver sites.

From an implementation perspective, the following two cases are applicable.

Case 1: Video Multicast Distribution from a Receiving Site to the End Users—PIM Islands in the Access Network

In this model (Figure 3), PIM join messages from edge routers trigger multicast flow from the ingress PE into the core. This flow is facilitated by BGP-based NG MVPN mechanisms. The discovery of multicast receivers (R) by the remote PEs is done via integration between PIM at the edge and BGP between PE routers. In Figure 3, PE1 receives a PIM join message that triggers PE1 to send BGP C-multicast route information (location of receiver) using an MCAST-VPN NLRI. On receiving the BGP C-multicast routes from PE1, PE3 sends a PIM join message towards the local CE router (CE3) to which the source is connected.

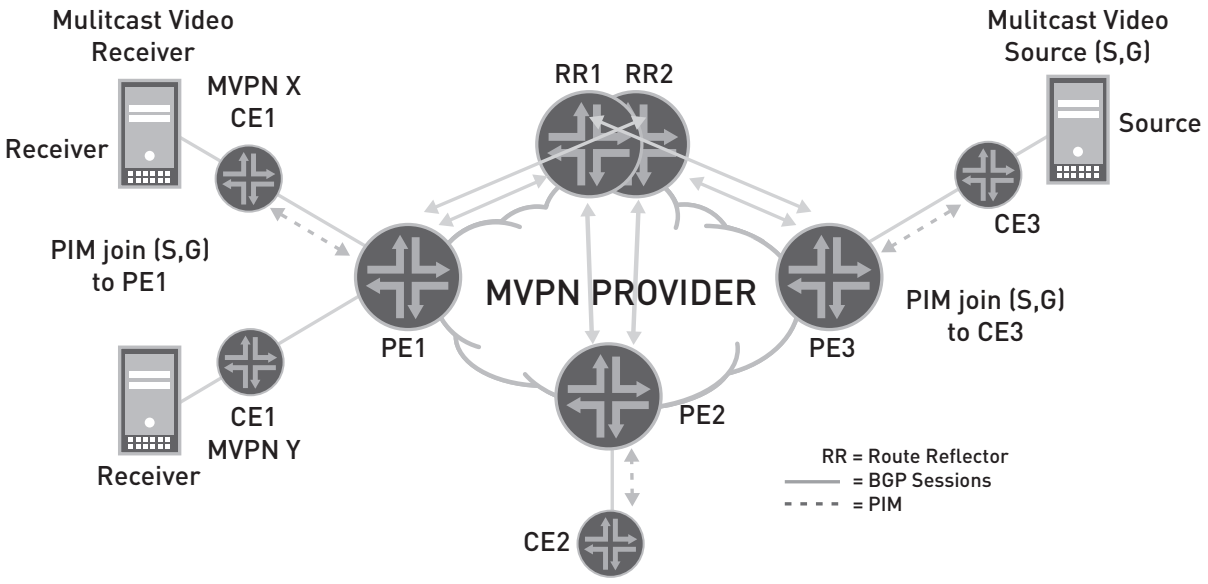

Figure 3: Video multicast distribution—PIM Islands in the access network

Essentially, the MVPN protocols designed to carry multicast routing information from the receiver sites to the source sites enable the dynamic discovery of multicast receivers for video distribution in the network. Note that receiver sites may be in a separate administrative domain, which would require virtualization on the PE router as shown for MVPN X and MVPN Y in Figure 3.

Case 2: Video Multicast Distribution from a Video Head-End to Receiving Sites—Multiple Content Providers

Sites within a given multicast VPN might be within the same organization or in different organizations, which means that a multicast VPN can be either an intranet or an extranet. In this model (Figure 4), the video multicast sources may be located in separate content provider networks. The receivers could subscribe to content from a specific content provider, making this scenario an extranet. The MVPN provider is connected to the content providers and offers video multicast services to end users. The PEs on the MVPN provider network learn about the sources and receivers using MVPN mechanisms. These PEs can use selective trees as the multicast distribution mechanism in the backbone, which carries traffic belonging only to a specified set of one or more multicast groups, from one or more multicast VPNs. As a result, this model facilitates the distribution of content from multiple providers on a selective basis if desired.

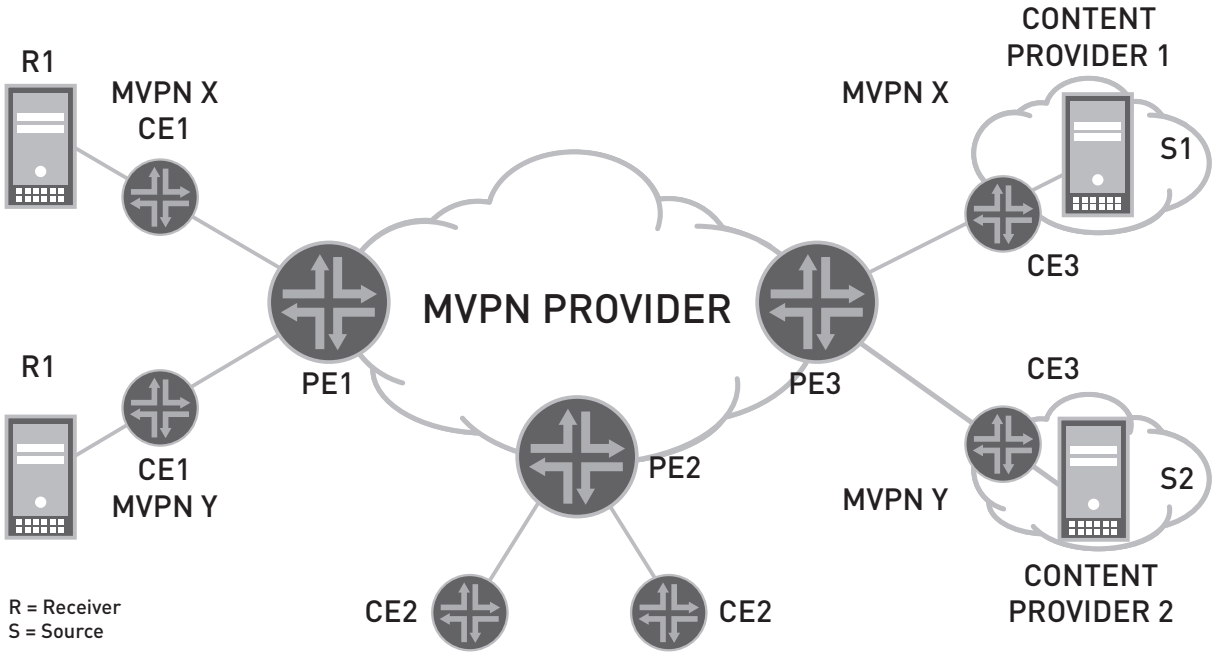

Figure 4: Video multicast distribution—multiple content providers

Enterprise and Financial Services Infrastructures

Large enterprise networks are deploying MPLS for their internal infrastructure for service convergence, as well as virtualization for administrative separation. The growing demand to transport services, such as media-rich collaboration services, video conferencing, and financial trading applications, requires the use of high-performance multicast over an MPLS infrastructure.

The delivery of financial data (Figure 5), such as financial market data, stock ticker values, and financial TV channels, is a perfect example of an application that must deliver the same data stream to hundreds and potentially thousands of end users. Some large financial institutions have a business model of a full-scale service provider that offers a multicast distribution network for delivery of financial services. In some cases, the content is generated and hosted within the financial service provider, while in other instances it is sourced from dedicated content providers. The content distribution mechanisms largely rely on multicast within the financial provider network. In this case, there could also be an extensive multicast topology within the customer networks (such as brokerage firms and banks) to enable further distribution of content and for trading applications. Financial service providers require MVPN for traffic separation between its customers accessing the content. Based on the mission-critical nature of the content delivered, it is critical for the MVPN provider to have a highly reliable, high-performance network. BGP-based NG MVPN allows for PE redundancy using multihoming capabilities, as well as MPLS restoration mechanisms using P2MP MPLS transport.

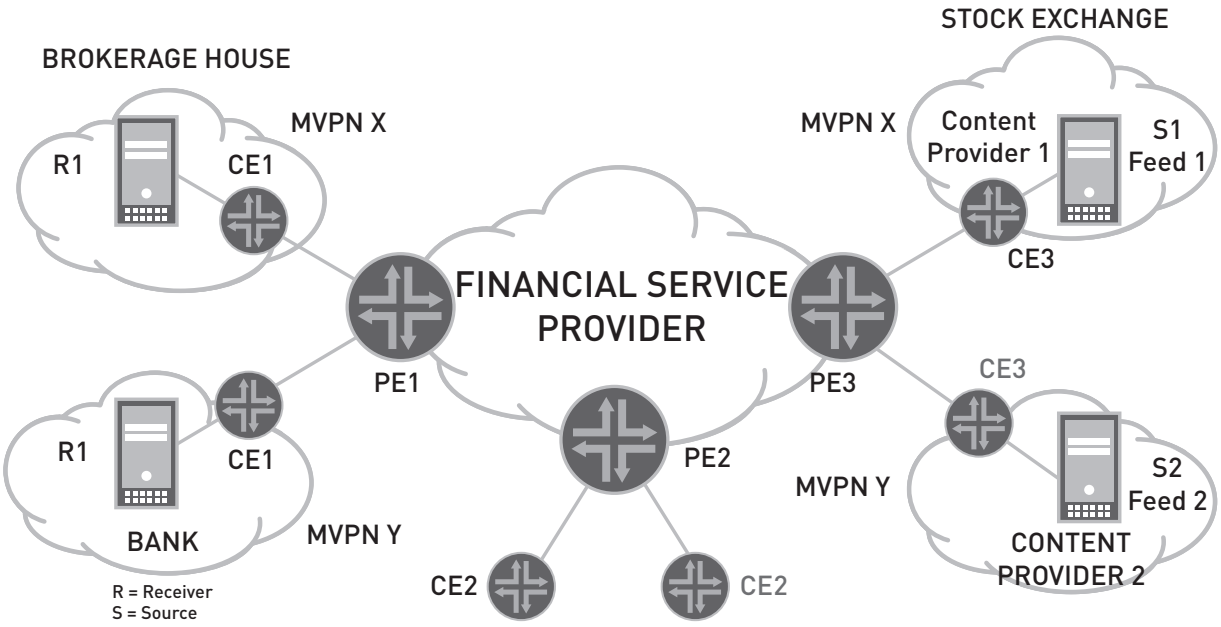

Figure 5: Financial service provider MVPN network

Multicast Backhaul Over a Metro Network

There may be two possible options to use backhaul in a provider network: 1) for backhauling customer traffic to a centralized service delivery edge, and 2) to extend the reach of services when the metro is not owned by the MVPN provider. There are at least two choices that a provider has for metro backhaul. One choice is VPLS and the other is 2547 VPNs. If 2547 VPNs are used for the backhaul, then multicast can be transported using NG MVPN technology over P2MP LSPs, which also supports selective trees. If VPLS is used especially for Layer 2 traffic, then multicast can be optimized using P2MP LSPs.

If VPLS is used in the metro for backhaul, it is also possible to terminate it into 2547 VPNs both for unicast and multicast at the service edge. Each of the two backhaul options allows an MVPN provider to extend its service reach deeper into the metro network. When the metro network is not owned by the MVPN provider, the backhaul option could be offered as a transport service by the metro provider to the MVPN provider. Thus, having these implementation options allows service providers to employ different business models based on traffic and topology profiles.

Conclusion

The emergence of multicast applications is driving service providers and enterprise networks to implement more scalable solutions, such as NG MVPN. Juniper Networks® is a leading innovator of these next-generation solutions. Juniper Networks is not only driving standards bodies, but has also created a software toolkit that includes options for enabling customers to scale, as well migrate from, existing deployments.

Acronyms

ARPU	average revenue per user
AS	autonomous system
BGP	Border Gateway Protocol
CE	customer edge
C-multicast	customer multicast
C-PIM	customer PIM
GRE	generic routing encapsulation
IBGP	internal BGP
IETF	Internet Engineering Task Force
LSP	label switched path
MCAST	multicast
MDT	multicast distribution tree
MPLS	Multiprotocol Label Switching
MVPN	multicast VPN
NG	next generation
NLRI	network layer reachability information
OPEX	operating expenditures
P2MP	point-to-multipoint
PE	provider edge
PIM	Protocol Independent Multicast
PIM SM	Protocol Independent Multicast sparse mode
P-PIM	provider PIM
ROI	return on investment
RSVP-TE	Resource Reservation Protocol traffic engineering
S,G	source and group
VLAN	virtual LAN
VPLS	virtual private LAN service
VPN	virtual private network
VRF	VPN routing and forwarding

For More Information

Advanced Topics in NGEN MVPNs. Rahul Aggarwal. MPLS World Congress, 2007.

BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs, draft-ietf-l3vpn-2547bis-mcast-bgp-05.txt.
www.ietf.org/internet-drafts/draft-ietf-l3vpn-2547bis-mcast-bgp-05.txt

BGP/MPLS IP Virtual Private Networks (VPNs) – RFC 4364.
www.ietf.org/rfc/rfc4364.txt?number=4364

Carrying Mission-Critical Traffic over MPLS Networks. Sean Clarke, Julian Lucek, and Andre Stiphout. MPLS World Congress, 2008.

Layer 3 Virtual Private Networks (l3vpn) IETF Working Group.
www.ietf.org/html.charters/l3vpn-charter.html

Multicast in MPLS/BGP IP VPNs, draft-ietf-l3vpn-2547bis-mcast-07.txt.
www.ietf.org/internet-drafts/draft-ietf-l3vpn-2547bis-mcast-07.txt

Multicast in MPLS/BGP VPNs, draft-rosen-vpn-mcast-07.txt.
http://tools.ietf.org/html/draft-rosen-vpn-mcast-07

Multicast in MPLS/VPLS Networks, tutorial. Matthew Bocci and Yakov Rekhter. MPLS World Congress, 2008.

Next-Generation Solution for Multicast in 2547 VPNs and VPLS. Rahul Aggarwal. MPLS World Congress, 2006.

PIM/GRE Based MVPN Deployment Experience and Recommendations, draft-rekhter-mboned-mvpn-deploy-00.txt.
www.ietf.org/internet-drafts/draft-rekhter-mboned-mvpn-deploy-00.txt

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 And Sales Headquarters
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

APAC Headquarters
Juniper Networks (Hong Kong)
26/F, Cityplaza One
1111 King's Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803

EMEA Headquarters
Juniper Networks Ireland
Airside Business Park
Swords, County Dublin,
Ireland
Phone: 35.31.8903.600
Fax: 35.31.8903.601

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

2000291-002-EN Mar 2009

Printed on recycled paper.