WHITE PAPER

DSLAM SELECTION FOR SINGLE-EDGE IPTV NETWORKS

Matching DSLAM capabilities with network requirements

Table of Contents

Table of Figures

Executive Summary

Providing IPTV services using xDSL introduces new considerations. These are primarily focused around how DSLAMs support delivery of broadcast television traffic. Although IGMP itself is a well-known standard, the implementation and use of IGMP in a broadband access network is not governed by any best practices and thus not implemented in a single consistent manner.

Also, introducing IGMP into an existing network may create additional complexity. This document discusses the often-overlooked complexities and considerations for implementing IPTV on an xDSL network. The emphasis is on the capabilities required in a DSLAM to support IPTV service.

Introduction

Selecting a DSLAM that supports IPTV service introduces several new considerations revolving around how the DSLAM supports "broadcast" television channel—in particular, how the DSLAM handles incoming IGMP requests.

IGMP serves three functions in IPTV networks. The best-known function is **channel replication**, in which the DSLAM begins (or stops) sending a broadcast television channel to an additional subscriber. Not all DSLAMs replicate multicast traffic, and those that do have varying levels of IGMP capabilities.

In addition, IGMP requests allow the network to perform **quality of service (QoS) adjustment**. For example, understanding that the subscriber has switched from a high-definition (HD) channel to a standard definition (SD) channel allows the network to make the extra bandwidth available to other applications. Finally, **usage tracking** allows the broadband operator to understand subscriber preferences, enabling informed decisions about what programming to carry.

There are numerous considerations for selecting a DSLAM. This document will attempt to answer the three most common questions:

- Which type of DSLAM— Layer 2, IP-Aware or IP-based—should be deployed?
- How can DHCP-based IPTV service be integrated with the existing PPPoE-based services?
- What are the additional complexities when implementing a separate video edge device?

Baseline IPTV Model

Figure 1 depicts the baseline hardware topology model for IPTV over DSL. There are up to four network elements involved. The television set—or more precisely, a set-top box (STB)—is the IGMP client. The routing gateway (RG) at the subscriber's site and DSLAM are intermediate devices that aggregate traffic. Some networks include an Ethernet switch to provide an additional layer of aggregation.

Although targeted at DSLAMs, many of the topics covered here are equally applicable to aggregation switches in the broadband network.

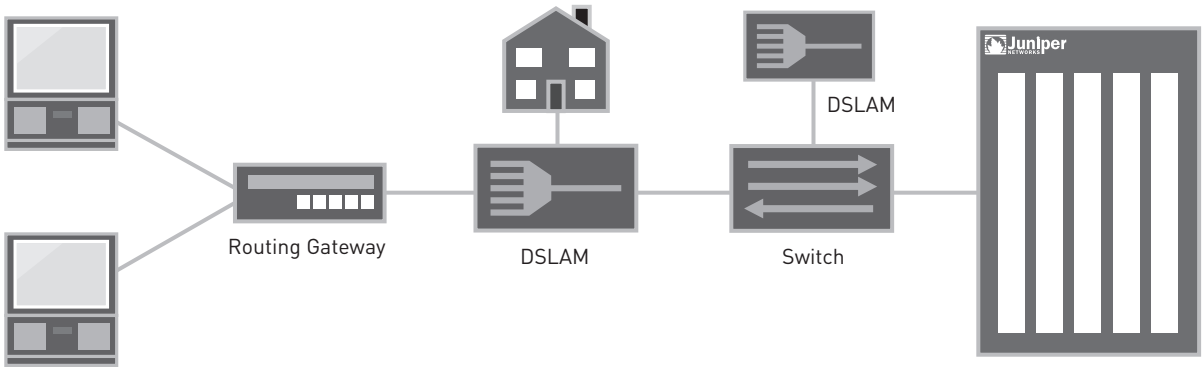

Figure 1: Baseline IPTV topology

Terminology

There are a few terms that are critical to understand:

- **IPTV** refers to traditional "broadcast television" channels sent across an IP network. IPTV excludes video on demand (VOD) content, which is unicast to each requesting subscriber.
- **Channel zapping** indicates a change in the IPTV channel being viewed. It most frequently refers to changing from one channel to another, but also includes starting and stopping of IPTV channel viewing (for example, when switching from a unicast VOD session to an IPTV channel or turning off an STB).
- **Replication** is the ability of a network element to forward an incoming television channel to multiple downstream devices on multiple ports. For example, an aggregation switch can often forward incoming IPTV streams to all downstream DSLAMs, and the DSLAM can then selectively forward channels to individual subscribers.
- **IP video** collectively refers to IPTV, video on demand and IGMP.
- **Broadband Services Router** (BSR) is an edge router that supports IP video, high-speed data (HSD) and voice over IP (VoIP) using either DHCP or PPPoE logical session flows. The functionality provided by this device varies dramatically, based upon operator requirements as well as the capabilities of the installed DSLAMs.

Basic DSLAM Selection

DSLAMs are organized into three broad categories:

- **Layer 2 DSLAMs** primarily convert xDSL to Ethernet (or ATM), providing little additional function. This type of DSLAM is essentially a Layer 2 switch, with some relevant enhancements. For example, all subscriber traffic should flow upstream to the edge router, preventing direct subscriber-to-subscriber communication. Layer 2 DSLAMs cannot interpret IGMP, so they cannot selectively replicate IPTV channels.
- **IP-aware DSLAMs** understand and can respond to IGMP traffic. To support IPTV applications, these DSLAMs interpret IGMP requests and replicate channels based on received requests. IP-aware DSLAMs can be further subdivided into two categories:
  - **Static** IP-aware DSLAMs always receive all multicast television channels. They do not have the ability to request specific channels be forwarded to the DSLAM.
  - **Dynamic** IP-aware DSLAMs can inform the network to begin (or discontinue) sending individual channels to the DSLAM. Configuring IGMP proxy or IGMP snooping on the DSLAM does this.
- **IP DSLAMs** add extensive IP intelligence, possibly including IP routing and advanced queuing based on DiffServ markings.

Within each category, there is a wide range of possible capabilities that can be implemented. The current debate is around how much intelligence to put into the DSLAM. This document provides some guidance for answering this question.

In selecting which type of DSLAM to use, there are three key decisions:

- Does the DSLAM need to replicate IPTV traffic?
- Should only those IPTV channels being viewed be sent to each DSLAM?
- Does the broadband edge router need per-subscriber visibility into what is happening, either to adjust network parameters or to track network usage?

Channel Replication

The first decision is whether the DSLAM should perform channel replication. While most early IPTV networks require this, it is not an absolute requirement. Equipment further back in the network—the aggregation switch or BSR—could provide this function.

The most common implementations support DSLAM replication of IPTV traffic. The commonly quoted advantages are:

- Reduced bandwidth requirements and associated bandwidth costs - Replicating further back in the network (for example, at the BSR) means that the same information (television channel) is sent to the DSLAM multiple times.
- Fast "channel zapping" - Proponents argue that doing this in the DSLAM allows zaps to be processed more quickly.

An emerging model eliminates replication in the DSLAM, instead providing this capability in the BSR (edge router). The arguments in favor of this model are:

- Simplified network operations - It is easier to manage and support fewer centralized BSRs supporting multicast than it is to support more smaller distributed DSLAMs providing this function. Typically one BSR aggregates dozens to hundreds of DSLAMs.
- DSLAM ubiquity - Minimizing the intelligence in the DSLAM allows selecting the lowest cost DSLAM available. In addition, previously deployed DSLAMs may be able to support IPTV.

These arguments are about cost and operations. Of course, selecting this alternative requires negating the earlier arguments about needing replication in the DSLAM.

- Bandwidth: For many DSLAMs, there is no additional cost to sending duplicate IPTV channels to the DSLAM since there is sufficient bandwidth available. In addition, many telcos are expecting heavy Video on Demand (VOD) penetration. Since VOD provides a unique stream to every subscriber, the bandwidth advantage is no longer relevant.
- Channel zapping: The reality is that the network adds no perceivable delay. Assuming a 1 Gbps connection, the network adds well under 1/10 of one second of delay. In addition, current BSRs can support thousands of channel zaps per second.

Bottom line: Most deployed IPTV networks support replication in the DSLAM. In this case, an IP-aware DSLAM is required. Some operators are deploying Layer 2 DSLAMs to reduce total network cost.

IPTV Delivery

If replication is desired in the DSLAM, numerous other considerations come into play. This first is how IPTV is delivered to the DSLAM. There are two models for delivering IPTV—multicast push and multicast pull.

In the **multicast push** model, all television channels are always sent to every DSLAM. Multicast push is implemented by notifying upstream equipment to automatically forward all channels to each DSLAM. This is done using static IGMP commands at the upstream equipment. This was the initial implementation in IP-aware DSLAMs.

In contrast, **multicast pull** delivers a multicast channel to the DSLAM only when a subscriber is viewing the channel. IGMP proxy or IGMP snooping is configured on the DSLAM, and instructs upstream network equipment when to send each IPTV channel to the requesting DSLAM. Many currently available DSLAMs now support this model.

These techniques trade off channel change performance against bandwidth requirements.

- With multicast push, the DSLAM can always immediately forward the desired channel to the requester. Multicast pull may cause slightly longer delays since the DSLAM must analyze the IGMP request, determine that it is not receiving the channel, forward the request to the multicast router, and wait for the router to begin forwarding the channel.
- On the other hand, multicast push requires more bandwidth to each DSLAM. A simple network consisting of 100 standard definition channels at 4 Mbps and 10 high-definition channels at 16 Mbps would require 560 Mbps of bandwidth to each DSLAM. In contrast, multicast pull potentially requires less bandwidth to each DSLAM, since less than half of the channels are typically being viewed at any given DSLAM.

Each alternative is appropriate in certain situations. For example, multicast pull may be preferred for large DSLAMs where most channels are being viewed, but less desirable for smaller DSLAMs that support only a few dozen subscribers. In addition, multicast pull more closely emulates what is required to support VOD traffic, where each subscriber establishes a unique "personal" IPTV session.

Both techniques can be deployed concurrently. For example, the top 25 channels can always be pushed to the DSLAM, while less popular channels will only be forwarded when a subscriber requests to view the channel.

Bottom line: As unicast video delivery such as Video on Demand, network-based PVR, Replay TV, and Internet-based video become popular, it becomes more important to minimize the amount of bandwidth consumed by multicast TV. In this case, choose a dynamic IP-aware DSLAM that supports IGMP proxy or IGMP snooping. If bandwidth is not a concern, use static IGMP to deliver all channels to all DSLAMs.

Network Monitoring

If multicast pull is desired for any or all channels, then either IGMP proxy or IGMP snooping can be used.

Until this point, all DSLAM decisions have been based on whether and how the DSLAM replicates channels. At this point, the other IGMP considerations—QoS adjustment and usage tracking—come into play. The decision here is whether and where these functions will occur.

There are two alternatives:

- In the centralized intelligence model, the BSR adjusts QoS parameters or tracks network usage. (The BSR can also optionally provide channel zapping.) This is most common in "single-edge" networks, where a single BSR is the edge router for all services. By having visibility to all IGMP requests, the BSR can throttle unicast traffic so that the access bandwidth to each subscriber is not over-committed.
- In the distributed intelligence model, the DSLAM performs these functions. This is more common in "multi-edge" networks, where there is a separate edge router terminating IP video traffic. In this case, the DSLAMs must manage traffic from two sources (IP video from one edge router, data and VoIP from another) so additional queuing capability is required. However, while DSLAMs attempt to manage traffic through queuing, they cannot provide assured delivery of all traffic.

Bottom line: If using a centralized model, then use IP-aware DSLAM with IGMP snooping to enable all IGMP requests to be seen by the BSR. If the distributed model is preferred, then select an IP-DSLAM that performs these functions. In general, IP DSLAMs support IGMP proxy (for multicast pull) although static IGMP is also supported (for multicast push).

Baseline Recommendation

The recommended architecture is to deploy dynamic IP-aware DSLAMs to support multicast pull. IGMP snooping is recommended since this allows the BSR to throttle unicast traffic.

Figure 2 depicts the baseline decision process, with the recommended path highlighted in green. Of course, specific requirements can lead to a different network design.

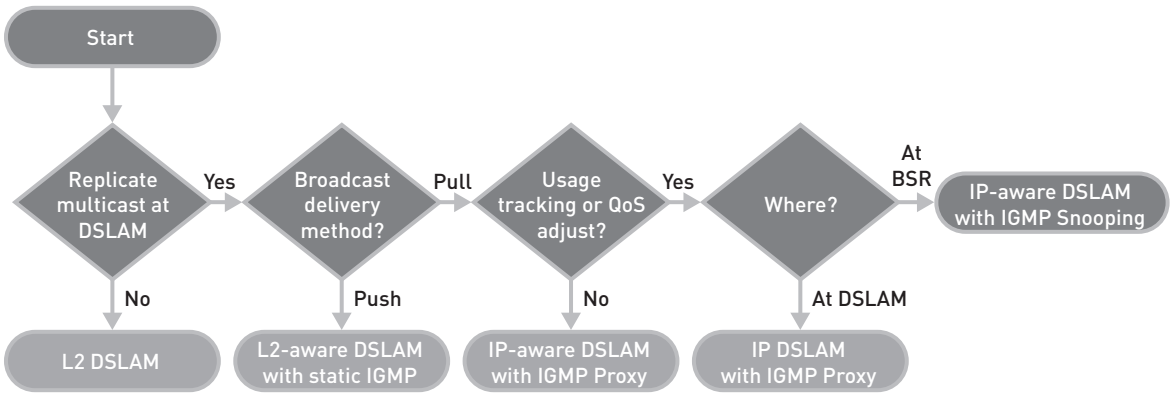

Figure 2: Basic DSLAM selection

Integrating IPTV into the Existing Network

The most common challenge in integrating IPTV into existing networks is that existing networks use PPPoX encapsulation to support data and voice traffic, but PPPoX does not support multicast traffic—such as IPTV—well. This issue is discussed in more detail in the Juniper DHCP/PPPoE document referenced earlier. This section expands upon this issue and discusses how to resolve this challenge.

This discussion only applies when data and VoIP are transported using PPPoX. The DSL Forum has until recently only allowed use of PPPoX, so an overwhelming majority of existing broadband networks use PPPoX to transport data and voice.

An important design decision is whether the customer VLAN (C-VLAN) or the multicast VLAN (MC-VLAN) will be used to transport IGMP requests. On one hand, since IGMP requests initiate channel zaps, they could flow over the same MC-VLAN that delivers IPTV to each subscriber. On the other hand, IGMP requests are typically destined to a single upstream device and should therefore use the unicast (C-VLAN) connection that uses PPPoX.

Figure 3 illustrates these alternatives.

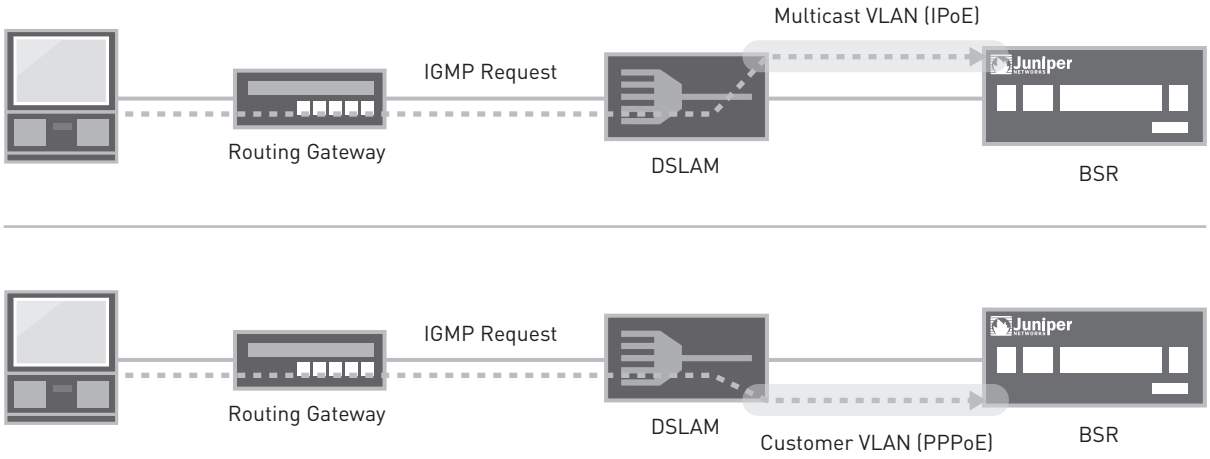

Figure 3: IGMP forwarding options

Both options have certain challenges. Sending an IGMP request on the MC-VLAN does not allow the BSR to throttle unicast traffic being sent down so that the DSL link does not get oversubscribed. On the other hand, sending IGMP requests across a C-VLAN has its own challenges. C-VLAN traffic may be encapsulated in PPPoX, and many DSLAMs cannot inspect IGMP packets encapsulated this way. Therefore, replication cannot occur at the DSLAM. In addition, intermediate switches may require that the same VLAN (the MC-VLAN) be used for both IGMP and the IPTV traffic.

This is resolved by sending IGMP packets across both the C-VLAN and the MC-VLAN. This technique is called IGMP Forking or IGMP Echo, and the DSL Forum describes its use in TR-1011.

Channel Zapping: Using IGMP Forking

IGMP Forking sends channel-zapping requests across both VLANs. Typically it is the residential gateway (RG) that generates two IGMP requests—one encapsulated in PPPoE and one using IPoE. Since the DSLAM does not interpret PPP-encapsulated packets, that request is passed to the edge router. This is used to track network usage and adjust the bandwidth available to other applications.

The DSLAM inspects the IPoE-encapsulated packet (using IGMP proxy or snooping) and performs channel zapping. This request may also be passed upstream, following standard practices for IGMP proxy/snooping. IGMP proxy (as opposed to snooping) is typically enabled on the DSLAM so that the BSR does not receive duplicate copies of every IGMP request.

IGMP Forking is depicted in Figure 4.

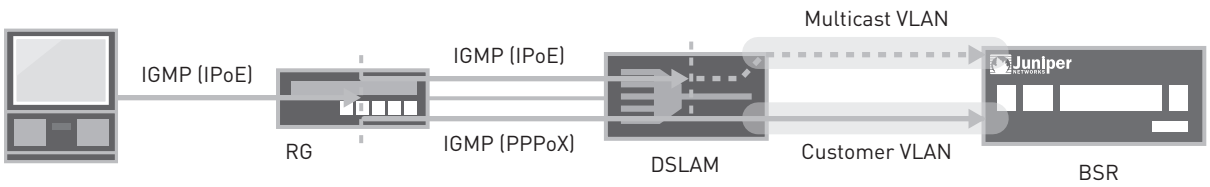

Figure 4: IGMP Forking (IGMP Echo)

In addition to the RG, the DSLAM and BSR must explicitly support this function. For example, TR-1012 requires that the IPoE-encapsulated packet use 0.0.0.0 as the source IP address.

Figure 5 summarizes the decision process for supporting channel zapping on existing PPPoX networks.

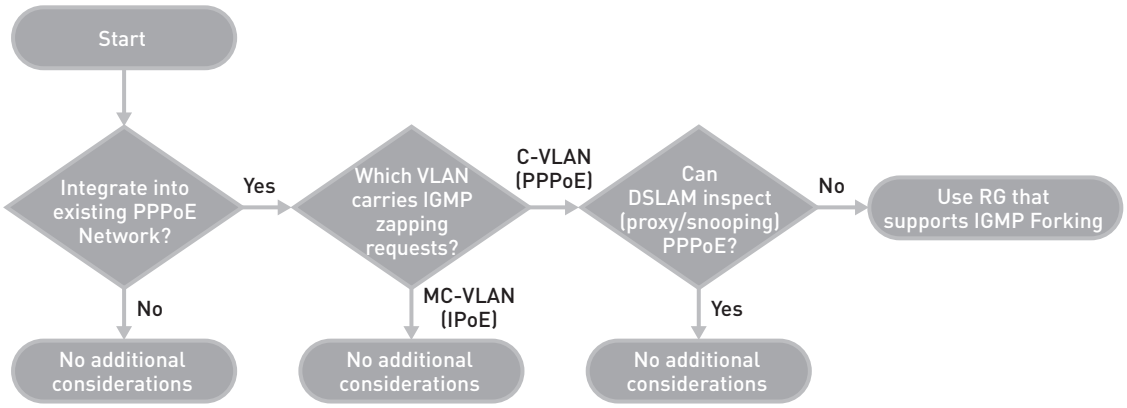

Figure 5: Integrating Channel Zapping in existing PPPoE network

IGMP Queries: Controlling Network Operations

Another IGMP issue concerns the use of IGMP Query messages. IGMP Queries request "what channel everyone is watching," and each active IGMP client responds appropriately. Since the MC-VLAN sends the same packet to all households connected to all DSLAMs on a given physical port, issuing an IGMP Query across an MC-VLAN can generate thousands of responses within a few seconds, potentially slowing network performance. In contrast, sending IGMP Queries across the C-VLAN allows the BSR to throttle the number of IGMP Queries, with each query generating few responses. Since C-VLANs are sent to a single subscriber, there are up to three responses to each query.

If IGMP proxy is being used, then this is not a major issue since the DSLAM aggregates information from downstream devices before responding. In this case, the Query must be sent on the MC-VLAN (which uses IPoE) so that the DSLAM can determine the request.

If IGMP snooping is being used, then the IGMP Query should be sent across the C-VLAN connection to avoid overwhelming the network with responses to the query. The IGMP requests will pass through the DSLAM, and each RG will respond.

Figure 6 depicts the decision process for supporting IGMP queries.

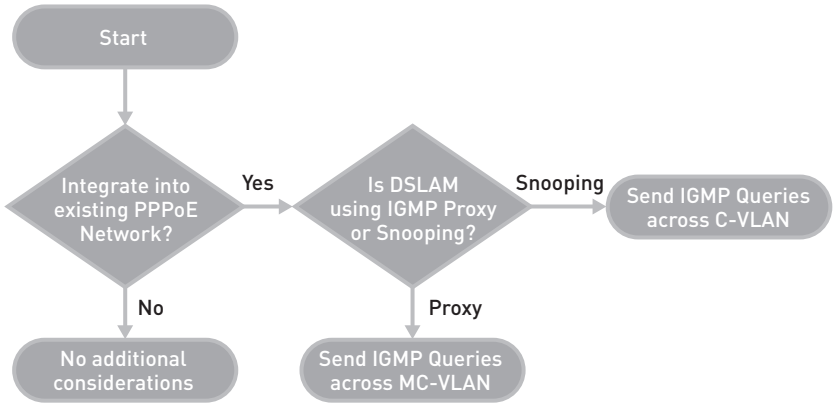

Figure 6: Supporting IGMP queries in PPPoE network

Multicast Support in Juniper Networks E Series Broadband Services Routers

There are several options for building an IGMP-aware xDSL access network for IPTV. The Juniper Networks E Series Broadband Services Routers support all of the models discussed in this document. Juniper Networks JUNOS® Software has the performance to handle all subscriber IGMP messaging when residing behind a snooping DSLAM, and can provide all replication services if no multicast is implemented in the access network. In addition, the E Series supports several features to enhance IGMP and multicast processing at the edge of a broadband network as follows:

- Provides full-fledged multicast router for IPv4 and IPv6
- Offers support for multicast routing protocols such as PIM and MBGP
- Allows proxy support for IGMPv2 and IGMPv3 traffic (for IPv4) as well as MLDv1 and MLDv2 traffic (for IPv6)
- Enables multicast OIF mapping utilizing an out-of-band multicast VC or VLAN
- Supports dynamic QoS adjustment based on IGMP join/leave processing
- Provides IGMP accounting
- Offers multicast call admission control
- Supports L2C protocol

These features allow a service provider to explore and implement any of the models discussed in this document when using the E Series as the BSR.

Contact

For more information, contact Marc Bernstein: mbernstein@juniper.net, 1.978.589.0651

Conclusion

Although IGMP is required in IPTV networks as the de facto method for an STB to perform channel changes for video, there are no well-documented best practices for the implementation of the access network.

Multiple options exist to allow the network to detect channel change requests and forward the appropriate channels. Each architecture has its own benefits and drawbacks based on the multicast optimization required in the network and the ability of various devices to offer IGMP processing.

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

2000189-001-EN Feb 2009