Offering a Multicast Delivery Service in a Programmable Secure IP VPN Environment

Lina ALCHAAL
Netcelo S.A., Echirolles
INRIA Rhône-Alpes, Planète project, France
lina.alchaal@inrialpes.fr

Vincent ROCA
INRIA Rhône-Alpes, Planète project, France
vincent.roca@inrialpes.fr

Michel HABERT
Netcelo S.A., Echirolles
France
michel.habert@netcelo.com

ABSTRACT

The programmable network approach is one possible solution to quickly adapt existing infrastructures to new requirements. This paper shows how programmable networking can be exploited within a VPN environment to offer a secure group communication service across the Internet. We show how the IP VPN approach offloads security, management and administration hassles from the multicast members and we propose a new simple Internet VPN Group Management Protocol, IVGMP, as an alternative to traditional multicast routing protocols.

Keywords: secure group communications, VPN, IPsec

1. INTRODUCTION

Motivations for a Secure Group Communication Service

Group communication is unavoidable when dealing with collaborative work applications and bulk data distribution. If the deployment of native multicast routing is well behind expectation [1], the important activity around application level multicasting proves there is an important need. But one aspect that often lacks is security. In this paper we show how to build a group communication service on top of a fully secure IP VPN (Virtual Private Network) environment, rather than the contrary. It is therefore a non-conventional approach that departs from the work carried out in the MSEC IETF working group for instance. We believe that our approach meets many needs, in particular for the deployment of services in commercial environments requiring a high level of security.

Definition of an IP VPN

An IP VPN [7][12] is an extension of a private network that encompasses links across a shared or public network like the Internet. A secure VPN uses a combination of tunneling and data encryption to securely connect remote clients and remote offices. Thus VPNs can replace troublesome remote-access systems and costly leased lines. There are currently three major tunneling protocols for VPNs [12]: Point-to-Point Tunneling Protocol (PPTP), Internet Protocol Security (IPsec) and Layer 2 Tunneling Protocol (L2TP). IPsec has the advantages of offering advanced cryptographic services and proves to be the best security protocol for LAN-to-LAN VPNs, while other security protocols work better for Host-to-Host connections. We therefore chose IPsec [10].

A Centralized Approach that Meets Secure Group Communication Needs

We assume that a VPN service provider (or VPN SP) is responsible of the VPN deployment and management between the various sites, and that this VPN service provider masters a VPN edge device in each site. The VNOC (Virtual Network Operation Center) is the central point of the service provider that collects the configuration and policy information and that remotely configures the edge device of each site. This centralized but dynamic approach is well suited to our needs. The VPN service provider can easily take in charge the group security management aspects that not only include authentication and access control but also the cryptographic key management, on behalf of the communication group. Finally a VPN topology is dynamic since a site can join or leave a VPN at any time, which fits well with the dynamic nature of a multicast group.

PPVPN Versus VPN SP

The approach discussed in this paper is aimed to VPN SPs who only offer a VPN administrative solution to customers. The assumptions made are completely different from that of the PPVPN IETF working group, where the provider, in addition to providing a VPN solution, also masters the core network that can either be Internet or a private interconnection network. Group communication solutions developed by PPVPN providers can easily take advantage of their own provider equipments (e.g. IP routers or MPLS-enabled infrastructure) to offer multicast-capable VPNs [13]. This is not our case.

The rest of the paper is organized as follows: we discuss in section 2 the limitations of traditional multicast routing protocols in a secure IP VPN environment. In section 3 we detail our proposal, the IVGMP protocol. In section 4 we give some insight on its implementation. Finally we introduce related works and we conclude.

2. LIMITATIONS OF TRADITIONAL MULTICAST ROUTING PROTOCOLS IN SECURE IP VPN ENVIRONMENTS

We first tried to deploy PIM-SM (Sparse Mode) [4], an efficient multicast routing protocol for sparse groups, in an IP VPN environment. We used `pimd` [16], an open-source PIM-SM implementation and tried to deploy it on the VPN edge devices, i.e. the hosts running IPsec in the customer network. We never succeeded for several reasons:

- The Security Associations (SA) management mechanism of the FreeS/Wan [6] IPsec implementation is currently only defined for unicast addresses [14]. This is a problem since a multicast packet, coming either from an application or from PIM-SM (for control purposes), cannot be sent as such in a VPN. For instance a VPN virtual interface built by FreeS/Wan does not set the MULTICAST flag and explicitly checks that no multicast address is used. Yet, there is no fundamental reason for which IPsec could not process multicast packets. There is much activity in the MSEC IETF working group on supporting multicast addresses in a SA, yet this is not finalized and not reflected in the current IPsec/IKE implementations.

- PIM-SM and IPsec ignore each other. For instance the FreeS/Wan [6] IPsec implementation manages its own routing table that cannot be seen by PIM, and there is no way to force PIM-SM packets to go through IPsec virtual interfaces. The same limitations apply to the PIM-DM (Dense Mode) and DVMRP [17] protocols.

An alternative solution could be to deploy a second host in each customer network (in addition to the VPN edge device), controlled by the VPN SP, and to deploy a multicast routing protocol supporting IP-in-IP tunneling on it (e.g. the `mrouted` DVMRP implementation). In that case only unicast tunneled packets are sent through the IPsec tunnel, which solve the problems mentioned above. This solution is not satisfactory though, since (1) it requires that *two hosts*, controlled by the VPN SP, be deployed in each customer network, (2) it generates additional traffic on the customer network (packets cross the LAN twice), and (3) it requires an additional encapsulation.

Two additional more fundamental flaws exist:

- These protocols have been developed to solve the general multicast routing problem over complex networks composed of many subnets, whereas a VPN environment creates a simple overlay network that may even be fully meshed.

- More fundamentally, from a VPN management point of view, the multicast routing protocol approach is not satisfying as it decouples the group communication service from the VPN management service. Therefore the VNOC has limited control and accounting capabilities.

3. THE IVGMP APPROACH

3.1 The Virtual Router Concept

The VPN edge device within each site insures the forwarding of multicast packets issued from the local site over the Internet towards other sites participating in the VPN and vice-versa. Therefore the VPN interconnecting the various edge devices can be modeled as a single virtual router (See figure 1) with several virtual interfaces, one per edge device. Within this virtual router, we introduce the IVGMP protocol to manage the configuration of the virtual interfaces and the forwarding of multicast packets. IVGMP is an alternative to traditional multicast routing protocols that catches the specificities of a programmable IP VPN environment.

3.2 IVGMP Detailed Description

Adding a New Receiving Site to a VPN

We first assume that each site is composed of a single LAN. In order to discover new local group members, the IVGMP protocol running on the edge device relies on IGMP (Internet Group Management Protocol) and its Query/Report mechanism [3]. This mechanism is used both (1) to discover members of new groups for which a new branch must be created in a VPN and (2) to dismantle VPN branches for groups having no member any more in the site. This is done by listening to IGMP traffic on the site's LAN (See figure 2). In order to know if a new VPN branch is needed when an IGMP Report for group G is listened, each VPN edge device maintains the list of multicast groups in which it already participates. In case of a new group, the VPN edge device issues a dedicated VPN command, JOIN_VPN(G), to the VNOC. On receiving a JOIN_VPN request, the VNOC performs some policy checking to determine if this site is authorized to subscribe to this group. A confirmation message is then sent back to the Edge Device and the VNOC automatically distributes the new management policies to all the VPN edge devices. Some accounting operations may also be performed during this process (e.g. to do per group subscription billing). Finally the edge device who joined the VPN asks the VNOC for some additional information (e.g. the list of sites participating in this VPN).

Adding a New Sending Site to a VPN

A similar process is used to manage multicast sources. In that case no IGMP message is issued by the application and a sending host will not respond to IGMP Queries either. Thus an edge device has to listen to all multicast packets coming from the local site, check if a new branch

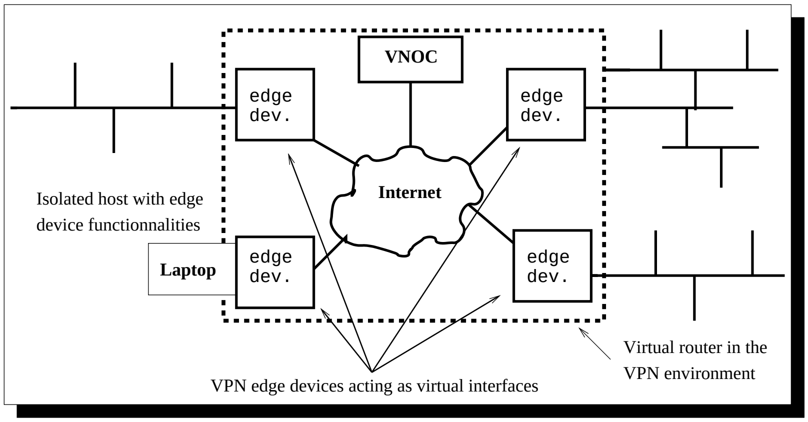
VNOC
edge dev.
edge dev.
Internet
Isolated host with edge device functionnalities
Laptop
edge dev.
edge dev.
Virtual router in the VPN environment
VPN edge devices acting as virtual interfaces

Figure 1: The virtual router concept.

is needed for that multicast group, G, and finally issue a JOIN_VPN(G) to the VNOC as described above.

IVGMP and Multicast Routing Protocols Interoperability

When a site is composed of several subnets, a multicast routing protocol is needed. In that case [7]:

receiver problem: the edge device will not receive IGMP messages sent by group members located in inner subnets not directly attached to him.

source problem: likewise the "top" multicast router (i.e. located on the same subnet as the edge device) will not forward any traffic coming from an inner source to this edge device if there is no receiver on the edge device's subnet.

Several solutions are possible to solve these problems:

- one can use IGMP-proxying [5] to let IGMP Report messages be forwarded by routers up to the edge device. This solution has the drawback of requiring some administration work in the site which is not always possible and desirable. Besides this solution only solves the "receiver problem".
- when there is a *small number of pre-defined multicast addresses* that can be used between the VPN sites, IVGMP can pro-actively subscribe to these groups (i.e. send IGMP reports) each time the top multicast router performs an IGMP query. Therefore the multicast traffic from inner sources, if any, will flow up to him. This solution only solves the "source problem" but can be used along with the IGMP-proxying solution. A major drawback is the useless state information created in routers and the increased IGMP signaling.
- another possible solution is based on a dedicated application used by users that start a sending or receiving application to inform the local IVGMP of the presence of new multicast groups ("source problem") and/or receivers ("receiver problem"). Then IVGMP can then contact the VNOC accordingly and subscribe to this group (to receive multicast traffic) if required. This solutions does not require any modification to the internal site but puts some burden on users. To avoid problems, the announcement is only valid for a limited span of time (and should be re-issued when required).

3.3 IVGMP Critical Appraisal

The IVGMP approach brings several distinctive benefits compared to traditional group communication approaches:

- Simple but efficient security management: security is done on a point-to-point basis, using the well known and operational IPsec/IKE framework. The security protocols dedicated to group communications currently being defined in the IETF MSEC working group are not needed.
- Centralized approach: the presence of a VNOC simplifies the configuration, management, and possibly billing aspects. Note that this VNOC is anyway needed for point-to-point IP VPN management.
- Many communication topologies are possible: so far we assumed that a star topology, centered on the sending site, was created. This is not compulsory and other schemes are possible. For instance the sending site(s) may forward traffic to a security certified node (the VNOC should have such an accreditation) that acts as a reflector to other receiving sites. This solution enables a sending site having limited upstream bandwidth to disseminate data to a large number of receivers without sacrificing security. More elaborated topologies can be envisioned, leading to the notion of VPRN, or Virtual Private Routed Networks [7] (see below).
- No dependency on inter-domain multicast routing: our solution only assumes the presence of a unicast routing service in the core network. This is an advantage in front of the slow deployment of inter-domain multicast routing [1].

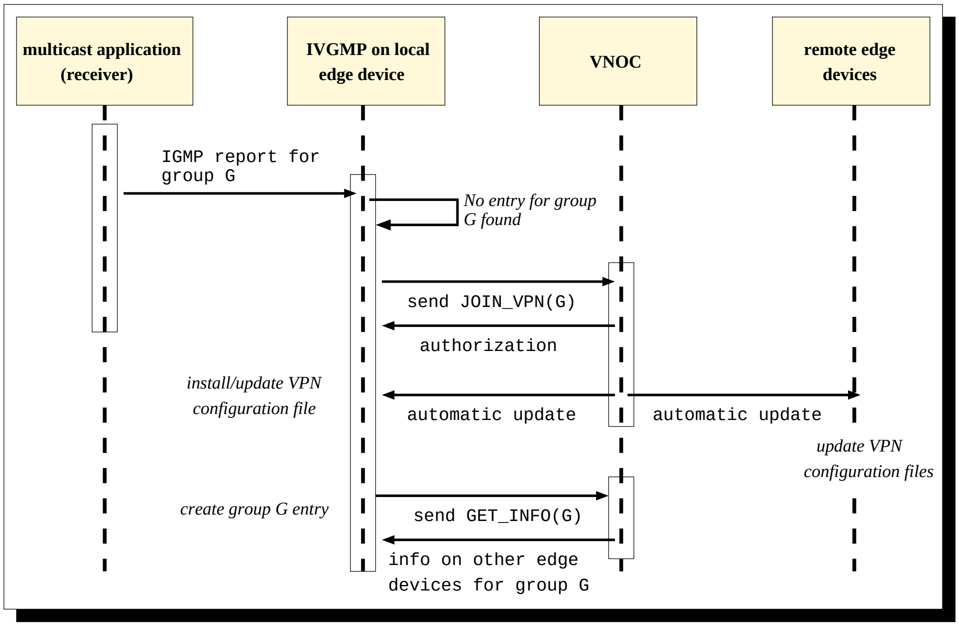
multicast application (receiver)
IVGMP on local edge device
VNOC
remote edge devices
IGMP report for group G
No entry for group G found
send JOIN_VPN(G)
authorization
install/update VPN configuration file
automatic update
automatic update
update VPN configuration files
create group G entry
send GET_INFO(G)
info on other edge devices for group G

Figure 2: Joining a Multicast Group as a Client.

- Compatible with other VPN types: Our approach, since it is based on edge technology, can be used with other VPN types such as MPLS VPN environment when QoS guarantees are needed.

IVGMP also has some limitations:

- Lower communication efficiency than native multicast routing: any solution based on traffic duplication at the edge is non-surprisingly less efficient than a solution based on traffic duplication in the core network as native multicast routing protocols do. The same remark can be done to application level group communication schemes that share some similarities with our approach.

- Scalability problem: This is a direct consequence of the lower communication efficiency.

The scalability problem can be addressed by provisioning some sites (either a subset of the receiving sites or trusted third-party sites) as VPRN nodes, i.e. nodes that can perform traffic forwarding. Each VPRN node belongs to two VPNs for the same multicast group, one including the sending site and another one including one or more receiving sites. A hierarchy of interconnected VPNs is thus created for each multicast group, leading to a distribution of the networking load among the set of VPRN nodes. The scalability of the solution is increased while keeping a centralized (around the VNOC) management.

4. IMPLEMENTATION ASPECTS

We have implemented the IVGMP protocol and integrated it in a PC/Linux edge device (See figure 3). We use a modified version of the FreeS/Wan IPsec implementation [6], the `isakmpd` implementation of IKE for OpenBSD [8], and Netcelo's VPN administration tools / VNOC. The IVGMP-VNOC communication, required for instance to maintain the multicast enabled VPN list, is based on the SOAP (Simple Object Application Protocol) web service technology which offers major advantages in terms of interoperability support and firewall/proxy friendliness.

Packets are processed as follows in each edge device: a multicast packet coming from an active source in the local site is first intercepted by the BPF packet filter [9] running in the edge device and sent to the IVGMP daemon. IVGMP looks for a VPN entry that matches the destination multicast address to decide whether or not the packet should be sent to the other sites. If an entry is found, a copy of the packet is encapsulated in a unicast UDP/IP datagram for each remote site (we use a UDP encapsulation in the current prototype for simplicity), given to IPsec and sent through a tunnel to the remote VPN edge device. In the other direction a packet coming from a remote VPN site is successively processed by IPsec, IP, UDP, and then IVGMP. This latter finally injects the original multicast packet in the local site through a raw socket.

This implementation enabled us to validate the IVGMP concept. Performance aspects and possible optimizations are left for future work, the present paper focusing essentially on architectural aspects.

5. RELATED WORKS

Many efforts have been done in several related domains like programmable networks, secure multicast, VPN technology, etc. For instance, the goal of the MSEC IETF working group is to standardize protocols for securing group communications within (potentially) very large groups, when

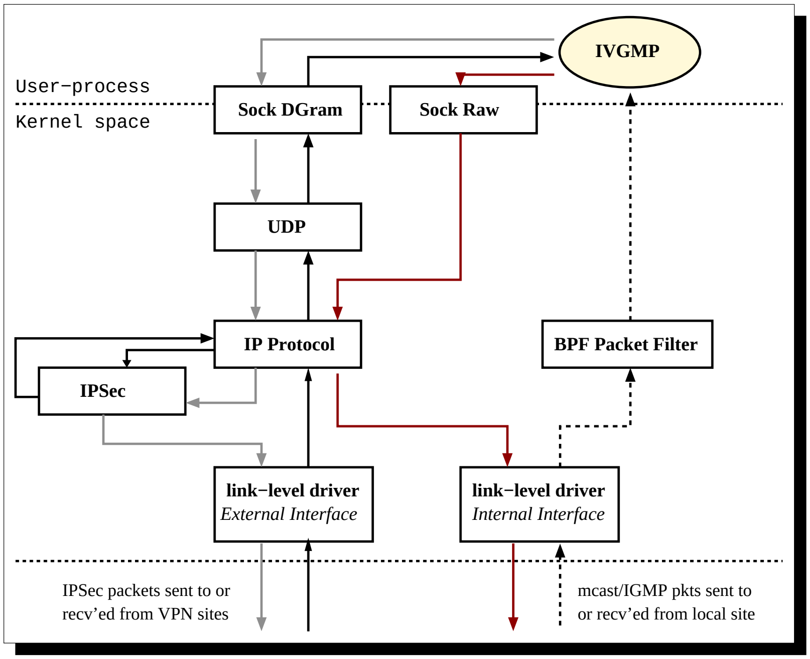

Figure 3: An architectural view of packet processing in IVGMP.

IP multicast routing protocols are used. This is therefore complementary to our approach.

[15] describes a solution to offer a multicast service over MPLS/BGP VPNs, using PIM within the VPN and customer routers at different sites. This solution differs from ours by the fact (1) it is aimed to be used in the service-provider backbones specified in [2], while our approach is based on edge technology, and (2) it does not address the problem of using PIM along with IPsec.

As mentioned in section 1, a PPVPN can easily and efficiently offer a group communication service. [13] describes, at a high level, how the VPN can exploit either a multicast routing service in the provider's network, or an MPLS-enabled infrastructure.

Finally, our approach shares some similarities with the Centralized Multicast (CM) approach [11]. In CM, the data forwarding and control functions are kept separated, and the control part is centralized in distinct control elements. The control elements are arranged in a two-level hierarchy within autonomous systems and are used to set up multicast trees. In our approach too, the control part is centralized in the VNOC. The major difference yet it that CM does not address security.

6. CONCLUSIONS

The approach discussed in this paper fuses miscellaneous sparse technologies like IP VPNs, web services (SOAP), programmable networks, in the same melting pot to get out with a simple flexible way to offer fully secure group communication services over the Internet. We have described IVGMP, a new solution for offering a group communication service based on the programmable IP VPN technology. The nature of this solution ensures its robustness, flexibility and full security. This solution departs from the traditional approach since it brings a group communication service on top of a fully secure infrastructure rather than the contrary. This work is still under progress and is part of a larger effort to master secure communications across the Internet and to manage them in the most transparent way to the end user.

7. REFERENCES

[1] C. Diot, B. Levine, B. Lyles, H. Kassem, and D. Balensiefen. Deployment issues for the ip multicast service and architecture. *IEEE Network*, pages 78–88, Jan. 2000.

[2] E.Rosen and Y. Rekhter. *BGP/MPLS VPNs*, Mar. 1999. IETF Request for Comments, RFC 2547.

[3] B. Fenner. *Internet Group Management Protocol, Version 2*, Nov. 1997. IETF Request for Comments, RFC 2236.

[4] B. Fenner, M. Handley, H. Holbrook, and I. Kouvelas. *Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)*, Nov. 2001. IETF PIM working group, work in progress,<draft-ietf-pim-sm-v2-new-04.txt>.

[5] B. Fenner, H. He, B. Haberman, and H. Sandick. *IGMP-based Multicast Forwarding (IGMP Proxying)*, Nov. 2001. Work in Progress, <draft-ietf-magma-igmp-proxy-00.txt>.

[6] FreeS/Wan org. *FreeS/Wan project home page: an open-source implementation of IPSEC and IKE for Linux*. http://www.xs4all.nl/ freeswan/.

[7] B. Gleeson, A. Lin, J. Heinanen, G. Armitage, and A. Malis. *A Framework for IP based Virtual Private*

Networks, Feb. 2000. IETF Request for Comments, RFC 2764.

[8] N. Hallqvist and A. Keromytis. Implementing internet key exchange, ike. In *Usenix Annual Technical Conference*, June 2000.

[9] V. Jacobson and S. McCanne. A bsd packet filter: A new architecture for user-level packet capture. In *Usenix Winter Conference, San Diego, California*, pages 259–269, Jan. 1993.

[10] S. Kent and R. Atkinson. *Security Architecture for the Internet Protocol*, Nov. 1998. IETF Request for Comments, RFC 2401.

[11] S. Keshav and S. Paul. Centralized muticast. In *7th Int. Conference on Network Protocols (ICNP'99), Toronto, Canada*, Oct. 1999.

[12] D. Kosiur. *Building and Managing Virtual Private Networks*. John Wiley & Sons Inc., ISBN 0-471-29526-4, 1998.

[13] D. Ooms and J. D. Clercq. *Overview of Multicast in VPNs*, Feb. 2002. work in progress, <draft-ooms-ppvpn-mcast-overview-00.txt>.

[14] A. Rodriguez, J. Gatrell, J. Karas, and R. Peschke. *TCP/IP Tutorial and Technical Overview*, 2001. IBM corp. Document Number GG24-3376-06.

[15] E. Rosen, Y. Cai, D. Tapan, I. Wijnands, Y. Rekhter, and D. Farinacci. *Multicast in MPLS/BGP VPNs*, February 2002. work in progress, $< draft - rosen - vpn - mcast - 03.txt >$.

[16] University of South California. *PIM-SM implementation home page, University of South California*. http://netweb.usc.edu/pim/pimd/.

[17] D. Waitzman, C. Partridge, and S. Deering. *Distance Vector Multicast Routing Protocol*, Nov. 1988. RFC 1075, BBN STC, Stanford University.