PIM-SM SECURITY: INTERDOMAIN ISSUES AND SOLUTIONS

Thomas Hardjono and Brad Cain
Bay Architecture Laboratory
Nortel Networks
3 Federal Street, BL3-03
Billerica, MA 01821, USA
thardjon/bcain@nortelnetworks.com

Abstract IP multicast is growing to be the future vehicle of delivery for multimedia voice/video/text in the Internet to its millions of connected users. With PIM emerging as the multicast routing protocol standard in the networking industry, and more specifically PIM-SM (Sparse Mode) for multicasting to sparse groups, the security of PIM represents a crucial factor for the successful wide deployment of IP multicast in the Internet. The current work argues that the authentication-key arrangement for PIMv2 [1] from the PIM WG is insufficient for interdomain authentication of PIM control-messages. The paper analyses some of the deficiencies of the PIM WG proposal, and offers some solutions to these shortcomings, whilst maintaining the key arrangement proposed by the PIM WG.

1. INTRODUCTION

IP multicast is emerging to be the future vehicle of delivery for multimedia in the Internet, with the promise of reaching the millions of users on the Internet. One crucial architecture component to this future vision is the multicast routing protocol that delivers multicast data packets (data stream) to group members, following the basic IP multicast model proposed in [2].

A number of multicast routing protocols have been proposed in the last few years (eg. [3, 4, 5, 6]). However, one protocol that has the promise of emerging as the industry standard in the *Protocol Independent Multicast* (PIM) multicast routing protocol, of which a *dense* mode (PIM-DM) and a *sparse* mode (PIM-SM) have been defined[6]. Currently, PIM has gone through almost two years of development and ex-

B. Preneel (ed.), *Secure Information Networks*

perimentations, and a number of large *Internet Service Providers* (ISP) have began to enable PIM in their routers.

In this paper we analyze the security features proposed for PIM, and in particular for PIM-SM. We argue that the proposed authentication-key arrangement for PIMv2 is insufficient for interdomain authentication of control-messages, and may even threaten the entire multicast infrastructure of PIM-domains. The paper offers a solution to these shortcomings of PIM-SM authentication, whilst maintaining the key arrangement proposed by the PIM Working Group in the IETF.

In Section 2. by a brief description of PIM and the protocol (MSDP) that interconnects PIM-domains. Section 3. presents the limitations and problems with the current security features of PIM and MSDP. A possible solution is presented in Section 4.. The paper is then closed with some remarks and conclusion.

The notations employed in the current work is as follows. Public key pairs are denoted as (SK, PK) representing the Secret-Key and the Public-Key, whilst private/symmetric keys are denoted by the symbol K. The symmetric/private keys are known only to a limited number entities. Two types of public key pairs are distinguished. The first type is best described as "semi-public" (or "closed mode") where the public-key half of the pair is known to a limited number of entities (not globally known nor globally available). As will become evident later, this is a non-traditional manner of using public key cryptography. The second type is the traditional public (or "global mode") usage of the public-key half, where anyone can obtain a (certified) copy of the owner's public-key. To simplify discussion, the first type will be referred to as a closed public-key, whilst the second type as a global public key. It follows that symmetric/private keys are "closed" by nature. All closed public-keys exist only within one domain and is unknown by other entities in other domains. Unless otherwise stated the term "domain" refers to a PIM-domain, the term "routers" refers to PIM-routers, whilst "border routers" refers to PIM-domain border routers.

2. INTERDOMAIN IP MULTICAST

The *Protocol Independent Multicast* (PIM) protocol is a multicast routing protocol designed with a number of motivations, one being the scalability of the protocol. To this end, two types of PIM protocols have been defined, namely PIM dense mode (PIM-DM) and sparse mode (PIM-SM). We briefly describe these in the following (while the interested reader is directed to [6]).

2.1 BACKGROUND: PIM-DM AND PIM-SM

The first, PIM-DM, is aimed at domains or regions whose group-population is dense and thus warrants the use of flood-and-prune techniques similar to that in DVMRP [3]. However, unlike DVMRP and MOSPF [5], PIM is "protocol independent", meaning that PIM does not depend on any specific unicast routing protocol/table (as do DVMRP and MOSPF).

PIM-SM, on the other hand, was designed for network with a sparse group-population in which techniques such as flooding could not be justified from a bandwidth point of view. Thus, in PIM-SM a number of special "meeting-point" routers, called *Rendezvous Point* (RP), are designated to which receivers must send an explicit join request. Following this meeting-point philosophy, a sender (Source) wishing to multicast data to the group initially sends data messages for the group via unicast (encapsulated) to the RP. The RP forwards the data to the receivers using a unidirectional shared tree. In order to avoid the RP becoming a bottleneck for high data-volume groups, when the sender's traffic exceeds a per-defined threshold the sender must create a shortest path tree (SPT) between itself and the RP. Thus, at this point, the intermediate routers (between the RP and the receivers) must also "switch" to the shortest path tree rooted at the Source.

Another entity involved in PIM-SM is the *BootStrap Router* (BSR), whose task is among others to advertise the set of RPs available in the PIM domain to the PIM- routers in the domain. The selection of which RP to use for a group can be based on a number of fair/balanced techniques (eg. hashing the group address).

Similar to other protocols running within a well-defined domain, what remained to be addressed was the issue of interconnecting PIM-domains in order to allow wider scalability across the Internet. Although a number of solutions have been proposed, a promising candidate protocol in this area is the *Multicast Source Discovery Protocol* (MSDP) [7]. In brief, MSDP interconnects RPs in differing PIM-domains to allow the existence of groups in one domain to be known to other potential receivers in other domain and to facilitate the receivers joining groups in other domains. Initially, an RP in a foreign domain acts as an intermediary for the receivers in its domain to join a foreign group. However, after certain conditions are met these receivers join the group directly to the Source located in the foreign domain. Hence, the "source discovery" of MSDP. The reader is directed to [7] for further details on MSDP.

2.2 KEY ARRANGEMENT IN PIM

The PIM Working Group (IETF) has recently put forward a proposal in [1] for the arrangement of cryptographic keys within a given PIM domain, with the aim of deploying the keys for control-packet authentication in the PIM domain. Following [1], when security is enabled, all PIM version 2 messages will carry an IPSEC authentication header (AH) [8]. The authentication mechanism must support HMAC-MD5-96 [9, 10] and HMAC-SHA-1-96 [11] security transformations.

The PIM key arrangement of [1] identifies the following entities in a PIM domain that require keys: the *BootStrap Router* (BSR), the *Rendezvous Point* (RP), the *Designated Router* (DR) and other PIM routers. All keys are relevant and recognized only within one PIM domain. The document of [1] purposely did not specify any key management approach to be adopted for the management of these keys. The three keys defined in [1] are as follows.

- *BSR Public Key* (closed): All BSRs own an identical RSA [12] key pair and uses the private key to sign an entire bootstrap message, whilst other PIM-routers only have the public key to verify the signature. This RSA secret-public key pair is denoted here as (SK_{bsr}, PK_{bsr}). This allows only authorized candidate BSRs to become a bootstrap router.

- *"Equal Opportunity Key"* (closed): All PIM routers in the same domain share a single private key used to compute digests or MACs for PIM control messages. This key is denoted here as K_{eq}. This key is used for "per-hop" authentication of control messages by PIM-routers in a given PIM-domain.

- *RP-Key* (closed): All RPs and BSRs share another private key, known as the "RP-key" and denoted here as K_{rp}. No other routers have this key. For candidate RP advertisement the digest is only calculated with the RP-key K_{rp} (instead of the equal opportunity key K_{eq}). This achieves the effect that only the authorized candidate RPs can advertise their candidacy to the BSR.

It is worth noting that key arrangement specified in [1] is perhaps best referred to as a "closed" or "semi-public" arrangement since the public key pair (namely the BSR public key) is known only to a limited number of entities in the network (namely the BSRs, RPs and PIM-routers only). In this sense, the usage of public key technology departs from its traditional mode of use where the public key of an entity is globally known and certified via a trusted Certification Authority (CA). The

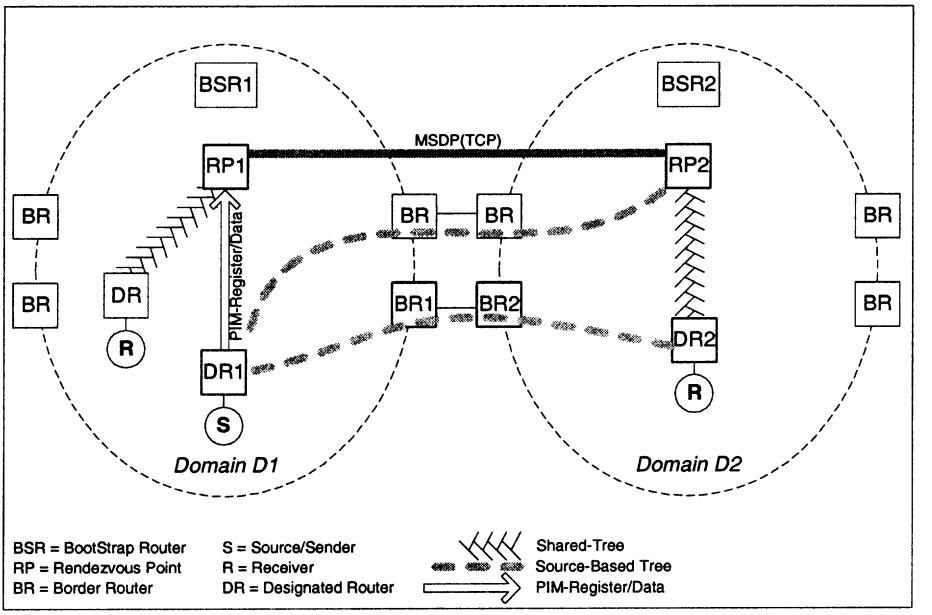

Figure 1 PIM-SM and MSDP

BSR public key is defined as "closed" since in general the advertisement of available RPs within a domain (ie. the RP-set) must be contained within that domain only. Other specific mechanisms must be employed to make known foreign groups and their RPs to a domain (eg. MSDP).

3. INTERDOMAIN AUTHENTICATION

In this section we argue that the key arrangement defined in [1] is sufficient only for per-hop authentication of control messages within a single PIM domain. Its limitations are highlighted when Receivers in a different domain than the Source join the multicast group.

We then propose the use of *Message Authentication Code (MAC) Translation* to alleviate the problem, which is largely caused by the constraints imposed by the key arrangement of [1]. We then point-out remaining deficiencies with MAC translations alone and propose additional mechanisms to overcome the limitations of MAC translations.

3.1 PROBLEM I: LIMITATIONS OF KEY ARRANGEMENT

We illustrate the limitations of the key arrangement of [1] in in Figure 1 consisting of two PIM-domains, namely domain D1 and D2. Each domain has its corresponding BootStrap Router (BSR1 and BSR2), Rendezvous Point (RP1 and RP2) and a number of Border Routers (denoted simply as BR1 and BR2). For simplicity we do not distinguish among the border routers within a PIM domain.

In the normal interdomain operation of PIM-SM, when a Receiver in domain D1 wishes to join a group whose source in also in domain D1 it uses IGMP to inform its Designated Router (DR1) and DR1 sends a Join/Prune message towards RP1. This message is authenticated on a per-hop basis by the PIM-routers using the shared private key K_{eq1}.

In the case of an interdomain membership within domain D2, the RP within domain D2 (namely RP2) must first join the source-based tree from the Source located in domain D1. MSDP is deployed by the RP-peers (RP1 and RP2) in order for RP2 to obtain "Source Active" messages containing the source address in domain D1, the group address and the IP address of RP1. Data received from the Source is then forwarded down the shared-tree rooted at RP2. In this manner, the receivers (more specifically DR2) within domain D2 obtain data originating from domain D2 via RP2. Once the Source's data rate reaches a given threshold, the Receiver's DR2 initiates the switch from the shared-tree to the shortest-path tree rooted at the Source's DR1.

3.1.1 The Need for MAC Translations. The first and foremost limitation of the key arrangement of [1] concerns the scalability of the shared private key (the "equal opportunity key") approach for all routers in a PIM-domain. Although a shared private key together with a keyed hash function represents a suitable approach from the perspective of performance, the use of a shared private key places limits on the scalability of the approach.

In Figure 1 a Join/Prune message issued by DR2 in domain D2 towards the Source (ie. DR1) in domain D1 can only be authenticated (using K_{eq2}) by the PIM-routers up to (and including) the border router BR2 of domain D2. Since the border routers BR1 in domain D1 do not posses the key K_{eq2}, they will not be able to authenticate the message.

To alleviate this shortcoming while remaining faithful to the key arrangement, we propose that some form of a "translation" of the MAC (attached to the Join/Prune message) be conducted between BR2 and BR1.

3.1.2 MAC Translation. In order to solve the problem above of interdomain control-message authentication, some form of "translation" of the MACs attached to the control messages must occur at the Border Routers (BR) of the domain. Assuming that the domain next to D1 (with key K_{eq1}) is domain D2 (with key K_{eq2}), then the MAC attached to the Join/Prune message going from domain D1 to domain D2 must be translated at the border of the two domains.

Two possible configurations may be deployed with respect to the border router(s) of the two domains:

- If the two domains share a common border router, then that border router can perform the translation internally. It must verify the MAC originating from domain D2 (using K_{eq2}), remove the MAC, and compute (and attach) a new MAC for domain D1 (using K_{eq1}). The common Border Router must thus be in possession of both keys.

- If domain D1 has a Border Router BR1 directly connected to the Border Router BR2 of domain D2, then these two Border Routers can share a common (intermediary) key known only to them (say $K_{br1-br2}$). When a control message traverses the boundary of the two domains from D2 to D1, Border Router BR2 in domain D2 must verify and replace the MAC attached to the message with an "intermediate" MAC computed with key $K_{br1-br2}$. The control message with the intermediate MAC is then delivered to the Border Router BR1 of domain D1. The Border Router BR1 must then verify and replace the intermediate MAC with a new MAC computed with its own key K_{eq1}.

Note that the approach of the two Border Routers sharing a common "intermediate" authentication key $K_{br1-br2}$ reflects the need of each domain to be autonomous from one another. From the key autonomy perspective, a re-keying of the authentication key in one domain should not affect another domain.

3.2 PROBLEM II: THE INSUFFICIENCY OF MAC TRANSLATION ALONE

Although MAC translation is required if we are to keep faithful to the key arrangement of [1], on its own MAC translation is insufficient to counter the effects due to the capture of the equal-opportunity-key in a given domain. Since all PIM-routers in the domain have a copy of the the equal-opportunity-key, these routers represents multiple points

of attack [13]. The larger the number of routers, the higher the chances of one of them being compromised.

The function of MAC translation is essentially to "transfer" trust across domain boundaries. However, in the current key configuration of [1], the effects of the capture of a copy of the equal-opportunity-key is also transferred across. When a router in a PIM-domain is compromised and the equal-opportunity-key is captured, any bogus interdomain control messages will be propagated by the (unaware) border routers to other domains. This means that a security breach in one domain has a direct impact on the other domains, an eventually the entire PIM-based multicast infrastructure.

Again, Figure 1 illustrates the situation. A border router BR2 upon receiving a control message (eg. Join) from the Receiver (ie. DR2) is only able to perform "group-authentication" on the message, meaning that it is only assured that whoever issued the message is in possession of the correct key K_{eq2}. The border router BR2 is never assured of the actual identity of the issuer of the message (namely DR2). Thus assuming that K_{eq2} has been captured and is used to issue a bogus control message, BR2 will simply translate the MAC for BR1, and BR1 will propagate the message into domain D1. In the case of a bogus Join message, the Source (ie. DR1) will simply add the Receiver (allegedly DR2) to the SPT rooted the Source (ie. DR1).

A more devastating attack is when a bogus Prune message (signed by a captured equal-opportunity-key) is injected within a transit or intermediary domain, where other receivers downstream exist within the same domain on within another domain. Assuming three adjacent domains D1-D2-D3, with the source in D1 and receivers in domain D3, if a bogus Prune message is injected in the intermediary domain D2 (with a MAC computed using a captured key K_{eq2}) towards the source (ie. DR1) in D1, then the Receivers in D3 will stop receiving data. This is because these Receivers would have been pruned-off the SPT rooted at the Source/DR1 in domain D1.

Hence, in itself MAC translations only solves the scalability problem without solving the problem of the potential spreading of bogus control-messages across domains.

3.3 PROBLEM III: MSDP SECURITY AND SCALABILITY

Currently, no specific security measures are incorporated into the Multicast Source Discovery Protocol (MSDP) [7] that interconnects two Rendezvous Points (RP) within PIM-SM. The most likely approach to be

adopted would be that of deploying authentication based on a keyed-hash function, similar to the (minimal) security solution of [14] adopted for the Border Gateway Protocol (BGP) [15]. It is worth noting that policy issues may also influence the solutions for MSDP, particularly if each RP is under a different administrative control (eg. different ISPs).

In contrast to BGP, in which the BGP-speakers are pre-defined and static, in the case of MSDP which connects two RPs, the relationship between two RPs are more dynamic since they interact based on the multicast groups which they serve. Even if the RP-to-RP communications have been protected using some form of symmetric-key based authentication (eg. keyed-hash as in BGP), further problems remain since the RPs (unlike BGP) communicate also with the Source within another domain.

Consider again the previous case in Figure 1. Here, RP2 in domain D2 discovers the multicast groups (whose Source's are in domain D1) through MSDP running between RP1 and RP2. RP2 then advertises the groups of the foreign domain D1 to its own domain D2. When a host in domain D2 wishes to join a foreign group, it indicates this request to its own RP2. The RP2 then joins the shortest-path tree to the Source/DR1 in domain D1 in order to forward the group's traffic down its own shared-tree.

From a security perspective, the current key arrangement of [1] is insufficient since it does not allow the Source (ie. DR1) in domain D1 to authenticate the control messages (ie. Joins) originating from RP2 in domain D2. Similarly, any control-messages sent by DR1 in domain D1 to the RP2 in domain D2 cannot be authenticated by RP2.

A possible solution that remains faithful to the key arrangement of [1] would be for the RP1 in domain D1 to "vouch" for the Source/DR1, and temporarily become an intermediary between RP2 and the Source/DR1. The thinking behind this is as follows. Since RP2 has already established some level of trust with RP1 when both parties communicated with each other through MSDP (protected using some shared symmetric key, assuming that a BGP-like approach is to be adopted), RP2 should also accept entities vouched by RP1. Thus, from this minimal level of trust, RP2 in domain D2 moves a step further by trusting the Source/DR1 since that Source is already communicating with RP1 well before the Source's group was made known to RP2 through MSDP.

The disadvantage of this approach is that RP2 in domain D2 must establish a shared key with each Source in a foreign domain (ie. domain D1). This per- group based key sharing may become intolerable for RP2 if the number of sources in domain D1 increases. In other words, although the approach of [1] is sufficient for a well-defined PIM domain,

inherently that approach is not scalable, as seen from this inter-domain behavior via MSDP.

4. DR-TO-BR AUTHENTICATION

In this section we present a solution for interdomain PIM control message authentication based on the key arrangement of [1]. The current work seeks to maintain the original key arrangement, particularly that of the equal-opportunity-key using symmetric cryptography due to performance needs within routers. Stemming from the use of such a symmetric key on a per-domain basis is the need to perform MAC translations (as previously described).

We introduce a security entity called the *Domain Key Distributor* (DKD) which has a number of roles in the domain, one of which is to be a key-server for the domain. The DKD thus acts as a public key certificate for open public keys (namely public keys known outside the PIM domain). The DKD also plays an important role in the key management within the domain. Although functionally it is referred to as a single entity, in practice the DKD can be implemented by several servers for reliability and availability reasons. Any such server must be implemented with the strongest security protection available due to their sensitivity.

In the following we present a new arrangement of keys for interdomain authentication that augments the arrangement of [1]. We illustrate the purpose of each key and its usage.

4.1 INTERDOMAIN KEYS

For interdomain authentication purposes the following keys are introduced to the entities that are concerned with trans-domain control messages. Although all of these keys are public keys, some are used in a *closed* manner (in the sense of the BSR public key). Public keys that are employed in the traditional manner are denoted as *global.* The method to generate and manage the keys are beyond the scope of the current work.

- *BR public key* (global): each border router in a PIM-domain is assigned a unique public key pair, designated as (SK_{br}, PK_{br}).

  The BR public key is "global" in that it is used in the traditional sense of public key cryptography. The public half PK_{br} associated with a given border router is known also outside the domain of the border router via some certification and advertising mechanism.

- *RP public key* (global): each RP (including Candidate-RPs [6]) are assigned a unique (non-shared) public key pair, designated as (SK_{rp}, PK_{rp}).

  Similar to the public key of the border routers, the RP public key is made available to entities outside the domain through a certificate signed by a trusted authority.

- *DKD public key* (global): the Domain Key Distributor (DKD) has the public key pair (SK_{dkd}, PK_{dkd}). The public-key PK_{dkd} is global in the traditional sense of public key cryptography. If multiple DKDs are deployed, then each should have a unique pair.

- *DR public key* (closed): each DR that is PIM-capable within the PIM-domain is assigned a unique (non-shared) public key pair, designated as (SK_{dr}, PK_{dr}).

  Since a DR is essentially the subnet router, and since the population of members can vary throughout regions of the Internet, the DR public key can be made either "closed" or "global" depending on the availability of a global public key infrastructure spanning the Internet.

  For the moment we assume that the DR public key is known only to within its own domain (eg. closed, and certified by the DKD of that domain). The DR public key can later be made global without impact to the security of the domain within which a DR resides.

So far only a limited number of entities have been assigned global public keys (ie. the BSR, RPs and the DKD). This is because the current work does not assume the existence of a global public key infrastructure to support the keys. Since the DRs (subnet routers) may far outnumber these special PIM entities, the current work has assumed the DR public keys to be closed. Should routers in the future be shipped with unique public keys, such routers can be suitable Designated Routers with global public keys.

4.2 PURPOSE AND USAGE OF KEYS

The aim of introducing a closed public key pair for the DRs within a PIM-domain is to alleviate the problem that arises from depending solely on the per-hop equal-opportunity-key K_{eq}. Note that the per-hop authentication is not removed or replaced, but rather augmented by the DR public key.

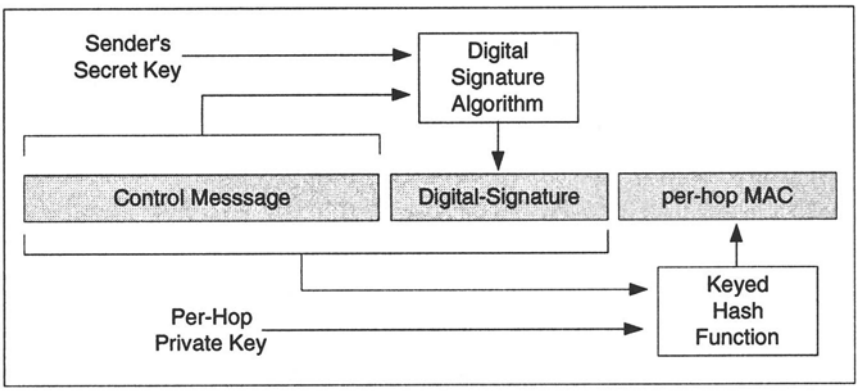

Figure 2 Augmenting with DR Digital Signatures

One reason for introducing only a limited usage of public key cryptography while maintaining the per-hop MAC is performance. Verification of public key digital signatures at every hop (ie. router) would be too burdensome. Thus, our current approach is to verify public key digital signatures at the domain-borders and at the destination.

When a DR (or RP) issues a control-message (ie. Join/Prune) towards the Source, it must digitally-sign the message using its secret key SK_{dr} (or SK_{rp}) before both the message and digital-signature is passed through the keyed-hash function (using K_{eq}) to produce the usual MAC. This is shown in Figure 2. This digital-signature is not verified on a per-hop basis. Only certain entities – such as the border routers within the domain – will have interest in verifying the digital-signature.

Referring to Figure 1 and Figure 3, when a DR2 (or RP2) in domain D2 sends a control-message towards the Source in another domain D1, its border router(s) BR2 will detect the control-message (from the header) and must then verify the digital-signature of the message allegedly coming from the DR. The border router BR2 must also verified the per-hop MAC as usual.

Once verified, two approaches are possible. Either the original digital-signature is left intact (which implies that the Source in D1 must have available the public key certificate of DR2 in D2), or it is replaced by a new one. In the first approach, a public key infastructure may be needed since public key certificates of all DRs must be shared across domains. Since the first case is straightforward, we focus on the second approach which only requires a limited number of public key certificates to be available across domains (namely that of the Border Routers).

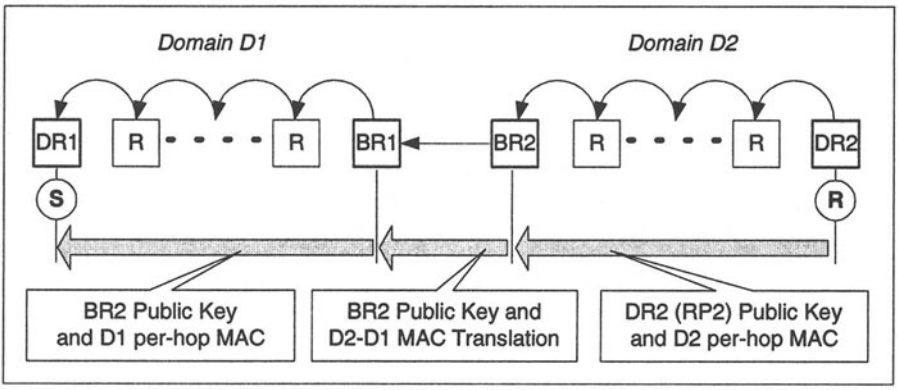

Figure 3 DR-to-BR Authentication

Thus, once BR2 verifies the per-hop MAC, both the digital-signature and the per-hop MAC must be replaced by a new digital-signature produced by BR2. The digital-signature from the DR2 is replaced with a digital-signature from the border router BR2 using its secret key (ie. SK_{br2}), whose public key half PK_{br2} is openly available to entities outside its domain. In effect, border router BR2 vouches for the Receiver (DR2) in domain D2. The per-hop MAC can either be translated (eg. using some intermediate key $K_{br1-br2}$ shared only between the two border routers BR1 and BR2 of domains D1 and D2 respectively), or be simply removed.

When border router BR1 receives the control-message from border router BR2 (originating from DR2), it must verify the digital-signature (from BR2) attached to the message and – if it exists – the per-hop MAC generated using the shared intermediary key (ie. $K_{br1-br2}$). The border router BR1 in domain D1 can obtain the public key certificate of border router BR2 in domain D2 through a number of ways. Among others, through the MSDP-peering between the RPs, through the announcements via the All-PIM Multicast group in the domain, through domain-wide broadcasts by the trusted entity or through other mechanisms.

Upon successfully verifying the digital-signature (from BR2) and the per-hop MAC (using the intermediary key $K_{br1-br2}$) the border router replaces the per-hop MAC with one that is recognized by other PIM-routers in its domain D1 (ie. computed using K_{eq1}). Border router BR1 does not alter the digital-signature created by BR2. Rather, it creates the per-hop MAC over both the control-message and its attached BR2's

digital-signature, using the equal-opportunity-key in domain D1, namely K_{eq1} (Figure 2).

All PIM-routers in domain D1 between the border router BR1 and the Source/DR1 only verifies the per-hop MAC using their copy of the key K_{eq1}. When the Source/DR1 receives the control-message, it verifies the per-hop MAC. Only after a successful MAC verification does it proceed to verify the digital-signature attached to the message (ie. BR2's signature).

The same process applies to the RP2 in domain D2 when it joins the source-based tree (towards DR1) via its border router(s).

4.3 ADVERTISING CERTIFICATES

In this current work we also propose an extension to the *Source Active* message within MSDP [7]. The current Source Active message consists of tuples of the form *[Source IP Address, Group Address, RP IP Address].*

One possible extension consists of the addition of the certificates containing:

- the (global) public key of the RP (certified by a global CA)
- the (closed) public of the Source/DR (certified by the DKD of the originating domain)
- the (global) public key of the DKD (certified by a global CA)

These can be sent as part of an extended tuple, or sent in the form of a distinct tuple (ie. newly added message definition and format to MSDP messages).

Another mechanism to advertise public keys (both closed and global) would be through the All-PIM multicast group.

5. REMARKS AND CONCLUSIONS

In this paper we have analyzed some of the main security problems with the *Protocol Independent Multicast - Sparse Mode* (PIM-SM) protocol, which is emerging to be the industry standard for multicast routing. The PIM Working Group (IETF) has recently proposed the use of the key arrangement defined in [1] for authenticating control-messages exchanged among PIM-routers. No plans are underway to authenticate data messages.

In the current paper we have argued (with enough basis and examples) that the key arrangement of [1] is insufficient for interdomain authentication, which is necessary for interdomain interaction since PIM-SM is designed to cover large areas with a sparse groups of population. We

then propose improvements to the key arrangement together with a combined usage of digital-signatures and MACs.

References

[1] L. Wei, "Authenticating PIM version 2 messages," Nov 1998. `draft-ietf-pim-v2-auth-00.txt` (`http://www.ietf.org`).

[2] S. Deering, "Host extensions for IP multicasting," RFC 1112, IETF, 1989.

[3] D. Waitzman, C. Partridge, and S. Deering, "Distance vector multicast routing protocol," RFC 1075, IETF, 1988.

[4] T. Ballardie, P. Francis, and J. Crowcroft, "Core based trees: An architecture for scalable inter-domain multicast routing," in *Proceedings of ACM SIGCOMM'93*, (San Francisco), pp. 85–95, ACM, 1993.

[5] J. Moy, "Multicast extensions to OSPF," RFC 1584, IETF, 1994.

[6] S. Deering, D. Estrin, D. Farinacci, M. Handley, A. Helmy, V. Jacobson, C. Liu, P. Sharma, D. Thaler, and L. Wei, "Protocol Independent Multicast – Sparse Mode: Motivations and architecture," Aug 1998. `draft-ietf-pim-arch-05.txt` available at `http://www.ietf.org`.

[7] D. Farinacci, Y. Rekhter, P. Lothberg, H. Kilmer, and J. Hall, "Multicast Source Discovery Protocol (MSDP)," tech. rep., IETF, June 1998. `draft-farinacci-msdp-00.txt`.

[8] S. Kent and R. Atkinson, "IP authentication header," RFC 2402, IETF, Nov 1998.

[9] C. Madsen and R. Glenn, "The use of HMAC-MD5-96 within ESP and AH," RFC 2403, IETF, Nov 1998.

[10] R. L. Rivest, "The MD5 message digest algorithm," RFC 1321, IETF, Apr 1992.

[11] C. Madsen and R. Glenn, "The use of HMAC-SHA-1-96 within ESP and AH," RFC 2404, IETF, Nov 1998.

[12] RSA Laboratories, "PKCS1: RSA encryption standard," 1993.

[13] L. Gong, "Increasing availability and security of an authentication service," *IEEE Journal on Selected Areas in Communications*, vol. 11, no. 5, pp. 657–662, 1993.

[14] A. Heffernan, "Protection of BGP sessions via the TCP MD5 signature option," Mar 1998. `draft-ietf-idr-bgp-tcp-md5-00.txt` available at `http://www.ietf.org`.

[15] Y. Rekhter and T. Li, "A Border Gateway Protocol 4 (BGP-4)," RFC 1771, IETF, 1995.