RP-Based Multicast Receiver Access Control in PIM-SM

Thomas Hardjono

Nortel Networks
600 Technology Park Drive
Billerica, MA 01821, USA,
thardjono@baynetworks.com

Abstract. The current work focuses on the issue of *receiver access control* in the context of the *Protocol Independent Multicast* (PIM) protocol. Currently, a host within a subnet can request the multicast router to join any multicast group without that host being authenticated and authorized to join. This (unauthorized) join-request results in the multicast distribution tree being extended towards that subnet, which opens the possibility of attacks. In such an attack, the malicious user/host intentionally extends or "pulls" the tree towards its subnet, effecting a wastage in resources and state within all the affected routers. In this case, the end-to-end encryption of the multicast data does not provide any help, since the (encrypted) packets still flows down the distribution tree to the malicious host. The current work analyzes this problem closer in the context of PIM Sparse Mode (PIM-SM) and offers a solution. The proposed approach also complements the recent developments in IGMPv3 [1] and the Express multicast model of [2].

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 [3].

A number of multicast routing protocols have been proposed in the last few years (eg. [4–7]). 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 [7]. Currently, PIM has gone through over two years of development and experiments, and a number of large *Internet Service Providers* (ISP) have began to enable PIM in their routers.

In this paper we focus on the issue of receiver access control to the multicast distribution tree within the context of PIM-SM. Currently no mechanism exists within PIM-SM to control the joining of hosts to the multicast distribution tree emanating from the *Rendezvous Point* (RP). A host can execute a group

G. Pujolle et al. (Eds.): NETWORKING 2000, LNCS 1815, pp. 120-131, 2000

membership protocol and direct its subnet router to join a multicast group, regardless of whether or not the host is an actual member of the multicast group. Encryption of data – typically end-to-end – does not protect against the malicious host from joining the group (and discarding packets). In itself, the result of the distribution tree being extended to the host's subnet has already consumed bandwidth, which in this case is simply wasted and thus leads to the potential of a denial-of-service attack on other hosts in the domain.

In the following section we briefly discuss the PIM-SM protocol (Section 2), followed by a problem description in Section 3. Section 4 reviews existing work directly relevant to the issue at hand, while Section 5 suggests a solution to the problem based on the use of the RP and the group-key management protocol. The paper is then closed with some remarks and pointers for future work. 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 reader is assumed to be familiar with the PIM-SM protocol, and is directed to [7] for more information.

2 Context: The PIM-SM Protocol

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 related modes of the PIM protocol have been defined, namely the PIM dense mode (PIM-DM) protocol and the PIM sparse mode (PIM-SM) protocol.

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 [4]. However, unlike DVMRP and MOSPF [6], 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 the receivers' join requests are directed (Figure 1). Before any groups can function, a *BootStrap Router* (BSR), must advertise the set of (candidate) RPs which are available in the domain. 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 pre-defined threshold a receiver (ie. its DR) must establish a shortest path tree (SPT) between itself and the Source. 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.

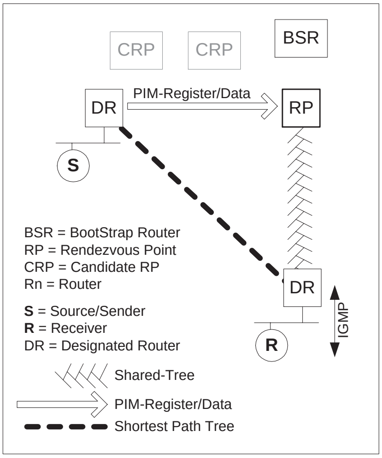

Fig. 1. PIM Sparse Mode

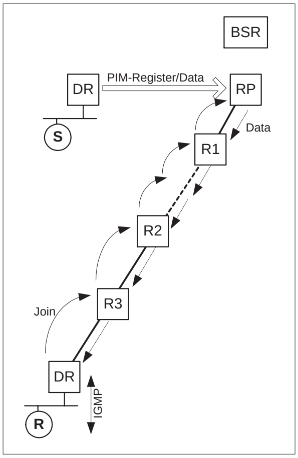

Fig. 2. Unauthorized "pulling"

3 Access Control for Receivers: Problem Statement

One of the main factors that makes IP multicast attractive from the perspective of scalability is the anonymous-receiver model underlying it. Any host in a subnet can join a multicast group without its subnet multicast router passing identification information about the host to other routers upstream in the distribution tree. This allows IP multicast to scale to a large number of participating hosts. However, from the perspective of security, this lack of host-identification information represents a problem for access control. This situation is true for most multicast routing protocols in the industry today, including PIM.

A possible attack that exploits the anonymous-receiver underpinnings of IP multicast is one in which a host simply joins a multicast group, without any intention of using the data being delivered to it. In such an attack, the user/host essentially extends or "pulls" the tree towards the subnet effecting a wastage in resources and state within all the affected routers. In this case, the encryption of the multicast data does not provide any help, since the (encrypted) packets still flows down the distribution tree to the malicious host.

The work of [8] has attempted to increase the security of the PIM protocol by requiring all control-messages to be authenticated via IPsec AH and a symmetric key shared by all routers in the PIM domain (namely, the "equal-opportunity-key", K_{eq}). However, this approach does not solve the main concern here of controlling access to the multicast distribution tree.

More specifically, although the control messages between a designated router (DR) and the RP are per-hop authenticated, the interaction between the Receiver and its DR remain open to attack. Thus, using the example in Figure 2, an unauthorized Receiver R is still able to join a multicast group, even though the control-message that flow upstream (from its DR to router R3, upwards to R1 and finally the RP) is protected from any malicious modifications via the equal-opportunity-key. The RP will still respond in the usual manner by sending out data along the appropriate interfaces towards the subnet of the unauthorized Receiver R.

4 Related Work

The main effort to protect the PIM distribution tree has been put forward in [8], with a key management proposal given in [9]. Other related work is the recent IGMPv3 [1] which allows a filtering of Sources, and the *Express* multicast proposal of [2] which proposes joins using specific (Source, Group) pairs. These are reviewed briefly in the following.

4.1 Control Message Authentication in PIMv2

The PIM Working Group (IETF) has recently put forward a proposal in [8] 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 [8], when security is enabled, all PIM version 2 messages will carry an IPsec authentication header (AH) [10]. The authentication mechanism must support, among others, HMAC-MD5-96 [11,12] and HMAC-SHA1-96 [13] security transformations.

The PIM key arrangement of [8] 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 keys are as follows:

- *BSR Public Key*: All BSRs own an identical RSA [14] key pair[1] 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*: All PIM routers in the same domain share a single private (symmetric) key used to compute digests or MACs for the protection of 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*: All RPs and BSRs share another private (symmetric) 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.

[1] Although the first version of the proposal in 1998 required one RSA key pair to be shared by all BSRs, the second revision of the proposal (IETF-45, July 1999) recognizes the need for each BSR to have a unique RSA key pair.

4.2 IGMPv3 and Express

The most recent version of IGMP (ie. Version 3 [1]) provides a receiver with the ability to filter based on specified sources. Thus, in the interface provided to the receiver, the receiver must not only specify the multicast-address (ie. group) which it wishes to join, but also specify the source-list and the filter-mode. From a multicast perspective, this addition of source-based filtering represents a major step forward as it allows the controlling of the senders to a multicast group. The work of [1] recognizes a number of possible control-message forgeries (eg. the IGMP Query message and Report messages), thereby recognizing the need of user/host authentication in IGMP and in IP multicast.

Along similar lines, a departure from the basic IP multicast service model was recently proposed in the form of the *Express* model [2]. Here, a "multicast channel" is defined to consists of the tuple (S, E) where S is the sender's source address and E is the channel destination address. The receiver or subscriber then requests reception of the data sent to the channel by explicitly specifying both S and E. The work of [2] defines a number of interfaces to the source and subscriber, where a source uses $channelKey(channel, K_{(S,E)})$ while a subscriber uses $newSubscription(channel, K_{(S,E)})$.

As the effort of [2] does not specify authentication methods or key delivery and management approaches to be deployed, we see the current work as complementing the works of [2] and [1]. Indeed, the proposal to be put forward in the following section provides a way for the key $K_{(S,E)}$ of [2] to be delivered to both the source and subscribers of a channel.

5 Coupling with Group Key Management

One major advantage of the PIM protocol from the point of view of access control lies in the fact that the Rendezvous Point (RP) is initially the point of departure for data packets destined for the multicast group members. The Source (Sender) unicasts data packets to the RP encapsulated within the PIM "Register" messages. Thus, the RP presents itself as a suitable point at which a decision is made as to whether data is sent towards a (candidate) Receiver, effecting the creation of a branch within the multicast distribution tree, linking the Receiver and the RP. (Note that the RP also represent a good location to conduct Source/Sender access control, since sources within the PIM protocol must initially unicast data to the RP).

In this section we propose the use of the group key management (GKM) event to aid the RP is deciding access control. 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 [15]. The DKD also acts as a public key certificate for open public keys (namely public keys known outside the PIM domain) and "closed" public keys (only known within the 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.

5.1 Group key management: background

Since data related to a multicast group traverses the public Internet and is therefore subject to tapping or copying by non-members of the group, encryption is the method commonly used to provided control to the data. In the simplest case, shared-key (symmetric) cryptography is used by the Sender/Source and the Receivers, where the data is encrypted by the Sender and decrypted by the Receivers. This shared-key is commonly referred to as the *group-key*, since only members of the multicast group are in possession of the key.

The use of cryptography necessitates the delivery or dissemination of keys, which in this is case is the group-key. Thus, an additional facet to the general problem of multicast security is the method of distributing keys to the appropriate entities involved in a multicast instance and the management of the keys of over given period of time. A *Group-Key Management* (GKM) protocol must not only issue a group-key for a new multicast group, but also update (re-key) the existing group-key under certain conditions and following the prescribed policies, be those general security policies or multicast-specific policies. A GKM protocol must be scalable, secure and must be independent from the underlying unicast and multicast routing protocols [16]. Examples of GKM protocols are [17–19].

It is important to note that the use of a group-key on multicast data is carried-out *end-to-end*, independent of the multicast routing protocol. The routing protocol – in this case PIM – is unaware of the group-key, and it treats the (encrypted) data simply as payload. Neither the RP nor the other PIM entities (eg. routers) understand the notion of the group-key. The encryption of data by a valid Sender (destined for other valid group members that hold a copy of the group-key) does not protect against the multicast distribution tree being attacked by unauthorized receivers.

5.2 GKM event to facilitate receiver access control

Since a new member of a multicast group must use the GKM protocol (designated for that group) in order to obtain a copy of the current group-key, we propose to use that event to also perform receiver access control with the aid of the RP and the DKD. This is further illustrated in Figure 3.

In Figure 3 (Step 1), the receiver R joins the multicast group by first obtaining a copy of the group-key (and other parameters) from the Domain Key Distributor (DKD). User/host authentication is carried-out at this time (eg. based on certificates), and group membership is also verified. This behaviour represents the general approach taken by most – if not all – GKM protocols. In the current context, we introduce an additional cryptographic key called the "DR Key" (*DR-key*), denoted as K_{dr} for the purpose of allowing the RP to

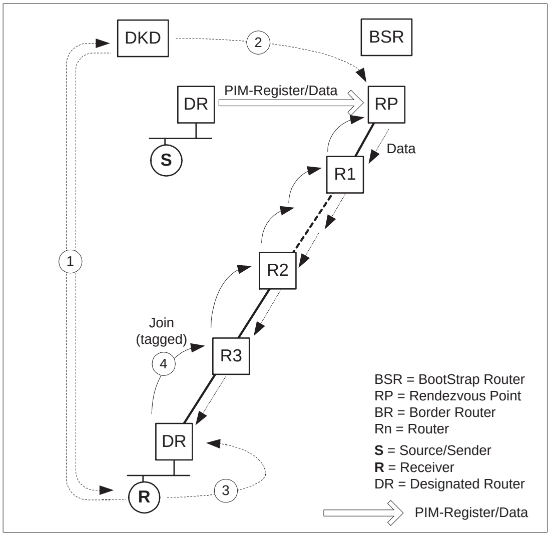

Fig. 3. GKM Event for Receiver Access Control

determine whether a DR in a given subnet is catering for a valid member of a group (as opposed to an unknown member or an attacker).

Here we propose that when a host-member is authenticated/verified by the DKD and is given a copy of the group-key, that host-member is also given a new DR-key (in addition to other parameters supporting its usage). The DR-key is only ever used once, and the DKD provides a list of the current DR-keys to the the RP through a secure channel (Step 2). The host-member then delivers the DR-key to its subnet-router (ie. the DR) through a secure channel, such as a secure unicast protected using IPsec ESP (Step 3). Extensions to IGMP can also be created to cater for this function.

When the DR issues the Join message towards the RP, it must create a "tag", which is a digest/MAC computed using a keyed-hash function and the DR-key. In addition, a nonce is also added to the Join message to prevent replay attacks. Other replay-prevention techniques can also be deployed. The triple *(Join, Tag, Nonce)* is then treated as payload, to be protected using the mechanisms specified in [8], namely using IPsec AH and the equal-opportunity-key (K_{eq}). The tagged-join message is then sent towards the RP by the DR (Step 4).

Upon receiving the tagged-join message, the RP uses its copy of the DR-key to verify the received join message. If the verification fails, the message is dropped and the relevant routers along the path between the DR and RP will time-out due to the fact that no data packets are received from the RP. If the verification is successful, the RP will deliver data to the appropriate interface, following the usual PIM protocol.

Note that in this proposed solution, the routers within the multicast distribution tree do not maintain any host-identification information. Hence, the solution still promotes the “anonymous-receiver” approach underlying the IP multicast model of [3].

5.3 Features and advantages of the basic solution

The solution described above has a number of features which makes it attractive:

- *Independence of group key management.* The solution maintains the important principle of the independence of group key management from the underlying multicast routing [16]. Multicast routing is unaffected by the introduction of the DR-key, since the key is introduced at the RP and the DR.
- *Adherence to the basic IP multicast model.* The solution adheres to the basic IP multicast model of [3], where the receivers are still anonymous from the point of view of the multicast distribution tree. The proposed approach also complements the recent developments in IGMPv3 [1] and the Express multicast model of [2].
- *Deploys existing GKM protocols.* The solution does not require a new GKM protocol, but only the introduction of additional security parameters to the existing deployed GKM protocol.
- *Adherence to the key arrangement of [8].* The solution fits into the key arrangement proposal for PIM and strengthens the security of PIM-SM as it currently stands.

5.4 Join towards on-tree nodes

The approach summarized above is deployable as it stands for new multicast distribution subtrees that emanate directly from the RP. Further refinement is required to cater for the situation where the join (tagged) command hits upon an on-tree node/router, representing the root of a sub-tree.

To refine the proposed basic solution in Section 5.2, consider the scenario in Figure 4. In Figure 4 (a) a Receiver R2 requests its designated router DR2 to join a multicast group through IGMP. Data is already flowing down from the RP to DR1 at the left-hand subtree (via Routers R1, R2, R3 and finally to DR1). In this scenario, Router R2 is the branching-router.

When a router receives a tagged-join message is must verify whether or not it is a branch-router (eg. by checking whether or not it already has an outgoing-interface for the same group G). If it does not have any outgoing-interfaces, then the router concludes that it is not on the distribution tree for that multicast group G.

However, if the router discovers that it is already an on-tree node within the multicast distribution tree for group G, it must perform additional tasks beyond those defined for PIM-SM. Rather than simply forwarding data to another interface, it must wait for an explicit acknowledgement from the RP regarding the new subtree leading to DR2.

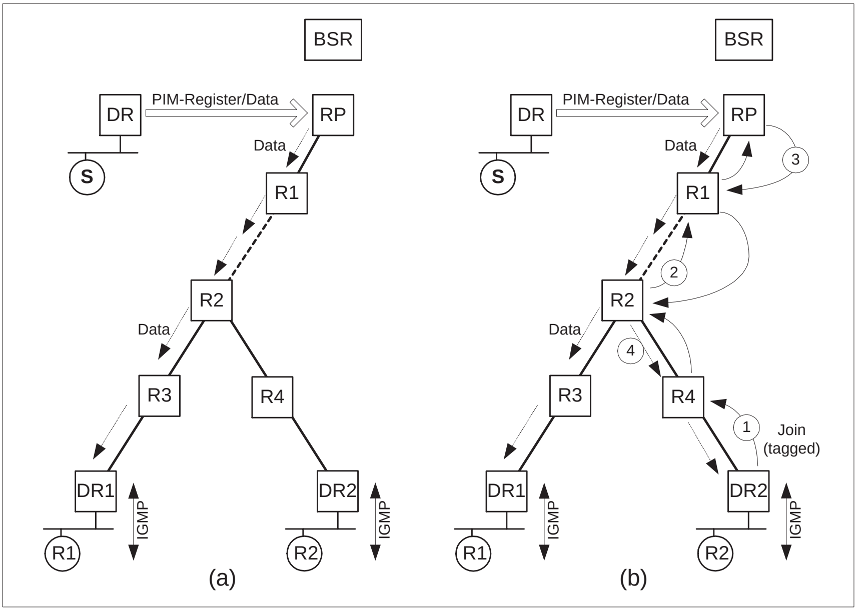

Fig. 4. Receiver Joining Sub-Tree

This is shown in Figure 4 (b), where the Receiver R2 is assumed to have already been authenticated by the DKD and is in possession of the group-key and the unique DR-key (ie. K_{dr2}). The Receiver R2 is also assumed to have securely delivered the DR-key (K_{dr2}) to its designated router DR2. In Step 1 of Figure 4 (b), the DR2 uses K_{dr2} to create the tagged-join (as previously proposed in Section 5.2). This tagged-join travels upstream, first to Router R4 and to Router R2. Since Router R2 will become a branching-router for the same distribution tree (ie. same group), it must not as yet deliver data down to Router R4. The Router R2 must first forward the tagged-join towards the RP and wait for an explicit acknowledgement from RP. This is shown as Step 2. In order to be able to recognize the response from the RP, Router R2 must maintain a copy of the tagged-join message obtained from DR2 via Router R4 (more specifically, the nonce and tag of the tagged-join message).

The explicit acknowledgement from the RP takes the form of a *Signed Join-Ack* message sent by the RP towards Router R2 (Step 3). The Signed Join-Ack message consists among others of the elements *(Ack, Nonce, Tag, (Source, Group))*, where the *Nonce* and *Tag* are those originally sent within the tagged-join message in Step 1 and Step 2. The acknowledgement message is digitally-signed by the RP using its public key pair.

When Router R2 receives the Signed Join-Ack message from the RP, it will use it to compare against the tagged-join messages which it received earlier (and which it has kept whilst awaiting the explicit acknowledgement from the RP).

A match indicates that the RP has given authorization for the new subtree to be created. That is, in Step 4 of Figure 4 (b), Router R2 allows data to flow through its interface towards Router R4 and finally to DR2. Once data flows, Router R2 can discard or erase both the tagged-join message (issued by DR2 earlier) and the Signed Join-Ack message (from the RP).

If the branching Router R2 does not receive the Signed Join-Ack message from the RP after a given period of time (eg. several heart-beats), it will time-out and discard the tagged-join. In this case data will not flow to the DR2. At this point the behaviour of Router R2 is the same as in the PIM protocol.

Again, similar to other routers, the branching Router R2 does not maintain any host identification information regarding the hosts downstream (in this case, the Receiver R2). This is in conformance with the IP multicast model of [3].

6 Remarks and Future Work

In the current work we have analyzed the problem of *receiver access control* in the context of the PIM-SM multicast routing protocol and the multicast distribution tree created by PIM-SM.

The problem of access control for receivers was then explained. Here, one of the possible attacks that exploits the anonymous-receiver underpinnings of IP multicast is one in which a host simply joins a multicast group, without any intention of using the data being delivered to it. In such an attack, the host essentially extends or “pulls” the tree towards the subnet effecting a wastage in resources and state within all the affected routers. In this case, the encryption of the multicast data does not provide any help, since the (encrypted) packets still flows down the distribution tree to the subnet and the (malicious) host.

The solution proposed to counter this deficiency makes use of the fact that the Rendezvous Point (RP) is a good decision-point for access control and that the group-key management event can be extended to facilitate receiver access control.

There are a number of issues that remain to be addressed in the general context of PIM security and in the specific area of member access control:

- *Interdomain sender/receiver access control.* Here, the issue concerns receiver access control when the source and the receiver/DR are located in different PIM domains. The PIM authentication key arrangement of [8] refers only to single PIM domains. The notion of a symmetric key shared by all routers (ie. the equal-opportunity-key) does not scale to multiple PIM domains whose RPs are connected via MSDP [20]. Security issues and possible solutions have been described in [21].
- *The RP as a bottleneck.* Although the RP represents a suitable access control decision point, the advantages from the perspective of security can also become a disadvantage from the perspective of routing performance. Some possible solutions to the bottleneck problem are as follows:
  - Alternative security architectures are deployed, in which the RP (a router) is aided by other network entities (eg. servers) in its task of verifiying

join-requests and deciding access control. The notion of other entities – such as servers – aiding routers is not new and can be found in a number of reliable multicast protocols based on the router-assitance concept (eg. GRA [22]) and in some Tree-based ACK scheme (eg. RMTPII [23,24]). In addition, different key arrangements (over and above the DR-key) can be deployed to create a balance between state held within the routers and the load experienced by the RP.

  - A "distributed" access control approach, where the DR-keys are pushed down the multicast distribution tree from the RP and are maintained by the DRs. In this approach, a DR becomes the access control decision point for potential members in its subnet since it will be in possession of the valid DR-keys. Besides requiring the DRs to be trusted, the approach requires the DR-keys to be available at the DRs via transmission through the multicast distribution tree (eg. the special "all multicast routers" group). That is, in itself this approach requires a low-delay and reliable transmission of the DR-keys to the DR. Futher research is currently being conducted into this approach.

- *IGMP control message authentication.* Assuming DRs can be trusted, one remaining requirement is that all IGMP messages exchanged between a host and the DR (ie. the multicast router) be source-authentic and integrity-protected. To this end, the DR-key – or more generally an IGMP-key – shared between a host and the multicast router can be used to kickstart a secure channel between the two entities. Again, the GKM event may represent a suitable solution towards distributing the IGMP-key to the host (ie. senders/receivers).
- *The problem of "oscillating switching" attacks in PIM.* The "switch" refers to the basic notion in PIM, where after a given transmission threshold has been reached, a Receiver would switch from the shared-tree (emanating from the RP) to the shortest-path tree emanating from the Source. A possible sophisticated attack in this case consists of the attacker throttling and releasing a critical link in the multicast distribution tree (without severing it) with the aim of effecting a switch back-and-forth by the targeted receivers between the shared-tree and the shortest-path tree. In itself this represents a denial-of-service attack.

These and other issues will be the focus for our future research.

Acknowledgements: We would like to thank Brad Cain for his help (and patience) with IGMP and PIM.

References

1. B. Cain, S. Deering, and A. Thyagarajan, "Internet group management protocol version 3," Nov 1999. `draft-ietf-idmr-igmp-v3-02.txt` (Work in Progress).
2. H. Holbrook and D. Cheriton, "IP multicast channels: EXPRESS support for large-scale single-source applications," in *Proceedings of ACM SIGCOMM'99*, (Cambridge, MA), pp. 65–78, ACM, 1999.

3. S. Deering, "Host extensions for IP multicasting," RFC 1112, IETF, 1989.
4. D. Waitzman, C. Partridge, and S. Deering, "Distance vector multicast routing protocol," RFC 1075, IETF, 1988.
5. 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.
6. J. Moy, "Multicast extensions to OSPF," RFC 1584, IETF, 1994.
7. 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` (Work in Progress).
8. L. Wei, "Authenticating PIM version 2 messages," July 1999. `draft-ietf-pim-v2-auth-00.txt` (Work in Progress).
9. T. Hardjono and B. Cain, "Simple key management protocol for PIM," Mar 1999. `draft-ietf-pim-simplekmp-00.txt` (Work in Progress).
10. S. Kent and R. Atkinson, "IP authentication header," RFC 2402, IETF, Nov 1998.
11. C. Madsen and R. Glenn, "The use of HMAC-MD5-96 within ESP and AH," RFC 2403, IETF, Nov 1998.
12. R. L. Rivest, "The MD5 message digest algorithm," RFC 1321, IETF, Apr 1992.
13. C. Madsen and R. Glenn, "The use of HMAC-SHA-1-96 within ESP and AH," RFC 2404, IETF, Nov 1998.
14. RSA Laboratories, "PKCS1: RSA encryption standard," 1993.
15. T. Hardjono, R. Canetti, M. Baugher, and P. Dinsmore, "Secure IP multicast: Problem areas, framework and building blocks," Nov 1999. `draft-irtf-smug-framework-00.txt` (Work in Progress).
16. T. Hardjono, B. Cain, and N. Doraswamy, "A framework for group key management for multicast security," Feb 1999. `draft-ietf-ipsec-gkmframework-01.txt` (Work in Progress).
17. H. Harney and E. Harder, "Group security association key management protocol," Apr 1999. `draft-harney-sparta-gsakmp-sec-00.txt` (Work in Progress).
18. T. Hardjono, B. Cain, and I. Monga, "Intra-domain group key management protocol," Jul 1999. `draft-ietf-ipsec-intragkm-01.txt` (Work in Progress).
19. C. K. Wong, M. Gouda, and S. Lam, "Secure group communications using key graphs," in *Proceedings of ACM SIGCOMM'98*, ACM, 1998.
20. D. Farinacci, Y. Rekhter, D. Meyer, P. Lothberg, H. Kilmer, and J. Hall, "Multicast Source Discovery Protocol (MSDP)," Jan 2000. `draft-ietf-msdp-spec-03.txt` (Work in Progress).
21. T. Hardjono and B. Cain, "PIM-SM security: Interdomain issues and solutions," in *Communications and Multimedia Security (CMS'99)* (B. Preneel, ed.), (Leuven, Belgium), Kluwer, 1999.
22. B. Cain, T. Speakman, and D. Towsley, "Generic router assist (GRA) building block: Motivation and architecture," Oct 1999. `draft-ietf-rmt-gra-arch-00.txt` (Work in Progress).
23. B. Whetten, M. Basavaiah, S. Paul, and T. Montgomery, "RMTP-II specification," Apr 1998. `draft-whetten-rmtp-ii-00.txt` (Work in Progress).
24. T. Hardjono and B. Whetten, "Security requirements for RMTP-II," Nov 1999. `draft-ietf-rmtp-ii-sec-00.txt` (Work in Progress).