Telecommun Syst (2008) 37: 145–156
DOI 10.1007/s11235-008-9108-y

Multicast deflector

Secure video distribution system

Antonio Pinto · Manuel Ricardo

Published online: 19 April 2008

Abstract Technological evolution is leading telecommunications toward all-IP scenarios, where multiple services are transported as IP packets. Among these services is the broadcast of video. A possible mechanism for broadcasting multiple video channels over IP is to use IP multicast, and let each client decide about the reception of a channel. The secure IP multicast specified by the IETF MSEC working group is a candidate solution for securing these broadcast services. In this paper we propose a new solution for supporting the broadcast of multiple video channels which can be accessed only by authorized users; besides, when a video channel is not visualized in the last mile its transmission is temporarily suspended, so that the cable can be used for other services such as standard Internet access.

Keywords Group communications · Video · Multicast · Security

1 Introduction

The technological evolution is leading telecommunications toward all-IP scenarios, where multiple services are transported as IP packets. Solutions for transmitting multimedia contents in IP packets already exist, and the RTP [1] is a key component of these solutions. In networks not over-provisioned there might be a need for traffic prioritization in order to provide QoS to the real-time services. Among these services is the broadcast of video that is expected to reach users with quality and security better than those provided today by cable TV. The latter is mainly based on coaxial cable and enables the transmission of multiple video channels, independently of the channels being visualized. The all-channel, all-time transmission leads to bandwidth waste, but it enables short response times when clients switch between channels.

When considering an all-IP network, a possible solution for transmitting multiple video channels could be to use multicast groups, and to enable each client to decide if he wants to receive a video channel. This solution consists in one-to-many transmission, where each receiver shows its interest in receiving the data by joining a multicast group, which can be represented by an IP multicast address. Each receiver must know previously the multicast address associated to the channel and, by issuing a multicast group join request to its gateway router, it forces the modification of the multicast routing tables, instructing routers to route data toward the receiver.

Security is a major concern of the new generation of IP networks, key issues being authentication and encryption. The first enables access control and its main objective is to guarantee the identity of each element involved in the communications; the second enables confidentiality by transforming plain data into unintelligible data, through the use of computer algorithms. Together, they enable the elimination of problems such as eavesdropping, unauthorized access, and data integrity violation. There are security solutions at multiple communication layers and each application can support security by using its own mechanisms.

A. Pinto (✉) · M. Ricardo
INESC, Porto, Portugal
e-mail: apinto@inescporto.pt

A. Pinto
Escola Superior de Tecnologia e Gestão de Felgueiras,
Politécnico do Porto, Porto, Portugal

M. Ricardo
Faculdade de Engenharia, Universidade do Porto, Porto, Portugal
e-mail: mricardo@inescporto.pt

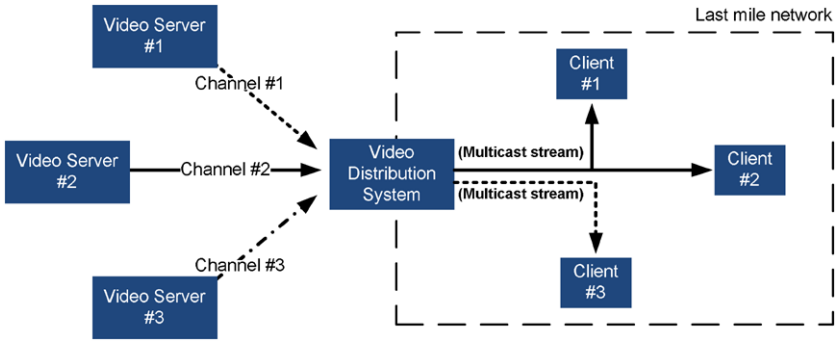

Fig. 1 Reference scenario

Secure IP multicast [2] may be used for the secure broadcasting of multiple channels over IP, but additional aspects need to be considered in scenarios of multiple video channel transmission, where some clients are authorized to view only a subset of the channels transmitted. On the other hand, the network resources consumed by multiple video transmissions, some of which not visioned, should be reduced. The first aspect can be solved by adopting a secure multicast group management protocol that provides each user with a set of cryptographic keys, one for each subscribed channel, that are unique and not shared with the other members of the group. The second aspect can be solved based on the idea that if a channel is not viewed, it needs not to be transmitted; this solution can be implemented by filtering the channel at some network element in the edge of the network, freeing the bandwidth that becomes available for other services.

The reference scenario adopted in this work is presented in Fig. 1. It consists of a set of video servers, a video distribution system, and a set of clients. Each server generates a video channel; the video is distributed as a stream of bytes, whose format is non-specified but must be transported by IP packets. The video distribution system retransmits a received video to the last mile network only if there are authorized clients interested in receiving it; the video distribution system must also be capable of processing video requests sent by clients, authenticate clients, and implement access control mechanisms; besides, it distributes video decryption keys to authorized clients. The client interacts with the video distribution system in order to authenticate himself, request videos, receive video as IP packets, and display them.

From the security point of view some assumptions are made. At the video server there is no need for security besides that required for any Internet connected computer (e.g. firewall), because it resides in the service provider network, assumed to be reasonably secure. The stream sent by the video server is encrypted using a key pre-shared with the video distribution system; it implies that the video distribution system assumes as authentic any video server holding the correct pre-shared key. The stream sent by the video distribution system to the last mile network is encrypted with a new key, different from the original pre-shared key. The client will assume as authentic any message signed with the private key of its serving video distribution system.

The following section introduces secure multicast transmission. The Video Distribution System, that is the object of our work, is presented in Sect. 3, starting by the presentation of the system requirements, followed by the system specification and by its design. The results of the functional and performance tests, obtained both in our test bed and in simulations, are presented and discussed in Sects. 5 and 6. Concluding this paper, Sect. 7 characterizes the results and discusses future work.

2 Secure multicast

Secure multicast is a group transmission technique that enforces confidentiality by adopting cryptographic and access control techniques. A secure multicast architecture is influenced by the group size. It must handle group memberships and security contexts such as the cryptographic decryption keys. The group management functionality of a secure multicast solution must also be scalable. The architecture defined in [3], for instance, is considered efficient in communications of small ad-hoc groups, whereas the architecture defined in [4] is being developed for large groups [5, 6].

2.1 Key management

Usually the data transmitted to a group is encrypted at the source and, then, the key used to decrypt the data is sent to all receivers, thus imposing confidentiality and access control. The entity responsible for key distribution to group members is the Key Server (KS) or Group Controller. A new receiver willing to participate in a secure group communication must signal its interest to join the group by sending an IGMP join request and contact the KS in order to obtain the decryption key. The source also obtains the cryptographic key from the KS and uses it to encrypt the data.

A group change, either by arrival of new members or by departure of members, requires that the (sub)group shared key is renewed; if not, the new members would be able to access previous communications, and the old members could

access future communications. Confidentiality requirements can then be mapped into four key distribution rules [7]: 1) non-group confidentiality; 2) forward secrecy; 3) backward secrecy; 4) collusion resistance. The first rule imposes that users that never participated in the group should not access any cryptographic material. The second rule imposes that members that depart a group should stop accessing cryptographic material, ensuring that these member's are unable to decrypt group communications after leaving the group. The third rule imposes that a receiver new to the group should not access previous cryptographic material, ensuring that this member is unable to decrypt past group communications. The fourth rule imposes that current cryptographic materials should not be inferable by non-members.

Group management in secure multicast communications demands scalable cryptographic key distribution techniques. In [8] the main approaches for key distribution were identified, and they consist of: 1) centralized group key management protocols; 2) distributed key management protocols; 3) decentralized architectures. The first approach is characterized by the existence of a unique entity with the responsibility of managing the entire group; it is focused in bandwidth usage optimization, key storage minimization, and reduction of computational power. The second approach assumes that any member can be a key distribution server and perform access control; all members can contribute to the generation of the group key. The third approach (decentralized architecture) achieves scalability through the division of the global group into a set of subgroups, each one having its subgroup manager.

Some proposals arose for the decentralized architecture approach; they include the Iolus framework (Iolus) [9], the Dual-Encryption Protocol (DEP) [10], the Zero Side Effect Multicast Key Management Using Arbitrarily Revealed Key Sequences (MARKS) [11], the Cipher Sequences (CS) [12], the Kronos approach (Kronos) [13], and the Hydra decentralized group key management architecture [14]. A more recent solution is the one proposed by the MSEC IETF working group [4].

2.2 Multicast security group

MSEC is an IETF working group which addresses secure group communications. This group considers three main problems: data confidentiality, group management and security, and group policy management and enforcement [4]. The MSEC solution consists of a set of blocks, each addressing a problem. Secure group applications can be built by combining these blocks.

The data confidentiality block addresses data security transforms and provides authentication and confidentiality to the group communications. While confidentiality is easily obtained through the use of symmetric encryption, data authentication requires more complex approaches since data encryption key needs to be shared among all the group members. Public-key encryption solves this problem but it is slower than symmetric encryption; hence, it may be inadequate to real-time traffic. The Timed Efficient Stream Loss-tolerant Authentication Protocol (TESLA) [15] is the authentication block proposed by MSEC for that purpose. It maintains the benefits of symmetric encryption, but it is unable to cope with immediate authentication. Senders, in TESLA, adopt time relaxed authentication by using a chain of keys to sign data and by sending the signing key after the time interval it was used to sign the data. The first key is the unique key that is signed digitally.

In order to encrypt, decrypt, sign or authenticate messages, the group elements need to previously agree on a set of parameters such as encryption key and algorithms; this negotiation is carried out under the responsibility of the group management and security block. The central element in group management and key distribution is the Group Controller / Key Server (GCKS) described in [6], which provides key and data encryption keys to new members, after their authentication. The MSEC is currently developing three of these key management blocks: the Group Domain of Interpretation (GDOI) [16] that enables the creation and the management of Security Associations (SAs) for IPSec and other network or application layer protocols, the Multimedia Internet Keying (MIKEY) [17] that addresses the specifics of real-time multimedia applications and can be tunneled over SIP [18], and the Group Secure Association Key Management Protocol (GSAKMP) [19] that enables group policy specification and dissemination.

The third block addresses group policy management and enforcement, and defines policies based on the cryptographic key information holders, encryption algorithms, and on the authorization of the policy creator. This information is gathered into a policy token that, totally or partially, is sent to all group members using the GSAKMP protocol defined in [20].

3 Secure video distribution system

A secure video distribution system is proposed in this section. We start by introducing the system requirements. Then we present the system specification. The design section follows and, there, we describe the architecture of all the elements and their functionality. The application protocols are presented next. Concluding this section, we discuss the innovation of the solution proposed.

3.1 System requirements

In order to describe the system requirements, three scenarios were considered: 1) an unauthorized user tries to access

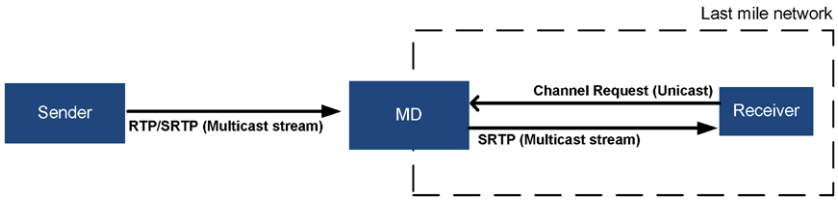

Fig. 2 Secure video distribution system elements

a video channel; 2) an authorized user tries to access a video channel for which he has no permission; 3) an authorized user tries to access a video channel for which he has permission. The 3 scenarios assume the transmission of multiple video channels, and each channel can be viewed only by authorized users; when the channel is not visualized by any user in some last mile network area, it shall not be transmitted. Based on that the following requirements were identified:

1. Encrypted data—video channel streams must be encrypted prior to their transmission;
2. Unique user identification and authentication—each user must be identified individually in order to authenticate himself and enabling access control;
3. Unique channel identification—it must be possible to identify uniquely every channel;
4. Access control by user and channel—it must be possible to prevent an authenticated user from accessing channels for which he has no authorization;
5. Channel source authentication—it must be possible for a client to be sure that he is receiving data from the correct server;
6. Group transmission—each video channel should be transmitted only once for all the interested clients;
7. Transmit only if required—if there are no clients interested in a video channel it shall not be transmitted;
8. Restring user mobility—a set of user credentials should only be valid in the user's network area;
9. Scalability—is should be possible to support as many users as required.

The usage of encrypted data transmissions prevents unauthorized video channel access; it depends on the authentication of receivers and on the encryption of the transmitted video channels. Requirement 6 may enable network bandwidth reduction since the video channels are sent to groups of users, and not to individual users. Both requirements can be fulfilled through the usage of secure multicast transmission of simultaneous video channels. The requirement 4 can also be helped by secure multicast transmission, since it may enable individual user identification, authentication and access control. This works also when each transmitted video channel is translated into a new secure multicast session that uses a multicast group access control algorithm, as those described in [9] and [21]. In these forms of group authentication, each user receives periodically a cryptographic key which enables him to access the video channel cryptographic key; this channel key is encrypted with another key (session key) that was initially negotiated with the KS.

The individual identification of users and channels can be met by providing to each user and to each channel a unique reference or identifier. In order to support user authentication, conventional encryption can be adopted and a cryptographic key can be pre-shared and associated with the user's individual reference. Hence, if a message is received and it carries out a user reference in clear text, then the authentication process can select the user's pre-shared key and decode the encrypted portion of the message that also contains the user's reference; if both references match, the user is authenticated. Authenticating the video channel transmission source will protect users from receiving video channels which are not original; either partially tampered or totally modified video channels. This requirement may suggest the adoption of asymmetric cryptography techniques. Although adequate to this case, asymmetric cryptography introduces significant delays in the video channel transmission when compared to symmetric encryption. This may lead to the selection of the latter.

The solution proposed addresses scalability using the same solution of the MSEC reference framework that approaches scalability through KS elements distribution, forcing group management information and group policy synchronizations and introducing replication mechanisms. As a consequence, we gain the support of user mobility. User mobility is not demanded and will not be provided by the proposed solution; the synchronization/replication mechanism foreseen by the MSEC group will not be used also. The local equipment of the video distribution system will be responsible for the client authentication and it will be the element having or accessing the client credentials. The result of this authentication model is a reduction of local group size, which improves the scalability of the proposed solution [7].

3.2 Video distribution elements

Figure 2 depicts the elements involved in the proposed video distribution system: the Multicast Deflector (MD), the Sender and the Receiver. Our Receiver element is analogous to the Receiver element of the MSEC framework; our Sender element is similar to the MSEC Sender. The group key management functional area of the MSEC framework is also implemented in our MD element; MD has the same

responsibilities of the MSEC Group Controller / Key Server (GCKS) that include user authentication, user authorization, key distribution and generation. But it also includes some functions hold by the MSEC Policy Server, such as rekey interval definition.

The Sender element is responsible for generating the IP packet stream of one or more video channels. If this element resides in the service provider core network, then Real-time Transfer Protocol (RTP) can be used to transport the video channel to the MD element; on the other hand, if the Sender is placed anywhere on the Internet, then Secure Real-time Transfer Protocol (SRTP) must be used to transmit the video stream in order to prevent unauthorized access. In a scenario of multiple video channel distribution, several Sender elements are expected to arise and, despite the original video channel streams being RTP or SRTP, the streams sent in the last mile network must always be transported in SRTP streams.

The main function of the Receiver element is to display video channels. In order to do this it must decode the SRTP streams sent in the last mile network. The Receiver element shares a cryptographic key with the MD that enables it to acquire Session Encryption Keys (SEKs) and to authenticate itself before the MD element. The SEKs are then used by the Receiver to decrypt messages containing the Channel Encryption Key (CEK), which will enable the Receiver to access each video channel SRTP streams. In a more realistic scenario, this module would be included in a Set-Top-Box which would be connected to a TV set, thus maintaining the same remote control based interaction between users and the service of the existing digital cable TV services.

The MD is the key element of the proposed architecture; it is responsible for (1) group management, (2) cryptographic key generation and distribution, (3) Receiver authentication, (4) (re)encryption of video channel streams with SRTP, (5) transmission of requested video streams, and (6) reception of all the video channel streams coming from the senders. Group management is achieved through our Implicitly Managed Group Protocol (IMGP) (not IGMP) and it consists on the creation of a layer four multicast group per video channel. When a Receiver wants to view a channel, it requests it to the MD and, on a successful reply, it becomes part of that channel group for a limited period of time. The MD authenticates Receiver elements by means of a pre-shared cryptographic key. In other words, if the MD can decrypt correctly the Receiver requests, it means that the correct pre-shared key was used and that the Receiver identity was verified. All video channels must be received by the MD; this decision enables fast responses to Receivers requests as well as filtering of video channel transmissions in the last mile network. The MD will only transmit the channels corresponding to the groups having one or more Receivers subscribed. Last mile network area access control is achieved by encrypting the video streams with SRTP.

Multiple levels of network optimizations are accomplished with our video distribution system. At the service provider core network the optimization is achieved by adopting multicast transmission. In the last mile network, considering broadcast media, multicast may not perform the same level of optimization. In this case, optimization is obtained by limiting the channels transmitted in the last mile, to the video channels requested, in conjunction with standard multicast transmission. All the streams are sent to a unique multicast address, simplifying other network functions such as the definition of QoS. The video channels differentiation is made by port numbers, what makes our solution a layer 4 video distribution solution. Multicast traffic filtering at the MD element is based on IP routing configuration. The two multicast addresses used, one for the core network, and the other for the last mile network, are associated to different network interfaces. Filtering is made by not forwarding IP traffic. With this configuration, the core network multicast streams are forced to end in the MD, the multicast streams generated by the MD are sent to the last mile network interface.

User access control is made through the use of cryptographic techniques, where each multicast video stream transmitted to the last mile network is encrypted with a unique cryptographic key that is defined as a CEK. The Receiver element needs to send a channel request to the MD in order to obtain a new SEK that is used to decrypt the multicast messages containing the CEKs. Considering the possibility of existing more than one Receiver interested in the same video channel, the CEK is transmitted to all Receivers simultaneously. This is done by appending several copies of the CEK and by encrypting each copy with a different SEK. When a Receiver gets a CEK message it searches for the part that is encrypted with its SEK and decrypts it.

3.3 System design

The secure video distribution system architecture is shown in Fig. 3, where all the modules of all elements are depicted. The key element MD is explained next, followed by the Receiver element.

3.3.1 Multicast deflector

The MD Architecture consists of 4 main modules: Multicast Deflection Engine (MDE), SRTP, KDP, and IMGP. The MDE is the controller module. It generates new cryptographic keys (SEK or CEK) upon IMGP requests, instructs KDP to distribute CEK to Receivers, configures the IP routing of the MD, and launches the execution of the SRTP, KDP and IMGP modules. The MDE is also responsible for the authorization of Receiver channel requests, for instructing the SRTP module to start or stop transmitting video channel

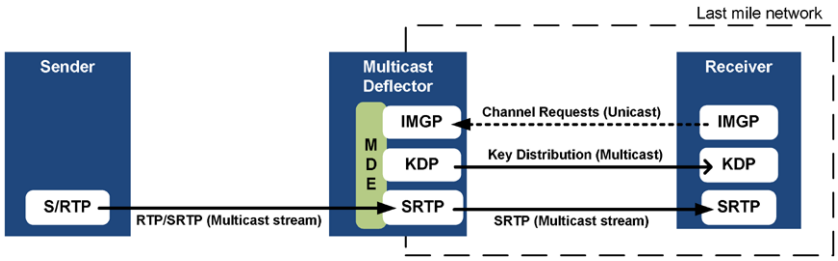

Fig. 3 Secure video distribution system architecture

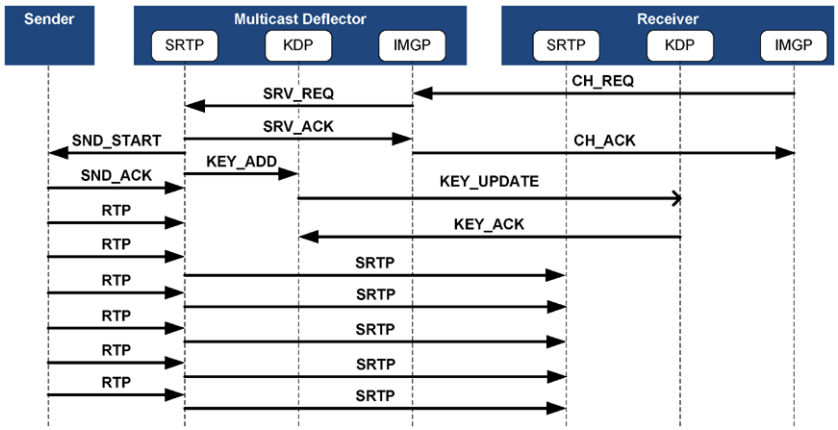

Fig. 4 Successful channel request

SRTP streams, and for distributing CEKs. Group management is another responsibility of the MDE, and it is done by maintaining lists of channels, each one associated to a list of the currently subscribed Receivers. Group management is the main function of the MDE.

The IMGP module implements the IMGP protocol. This protocol is responsible for transporting the SEK of each receiver, when it requests the reception of a video channel. Receivers must send an encrypted channel request message (CH-REQ message) to the MD, which will answer with an acknowledge or not acknowledge message, depending of the request acceptance of the request. The acknowledge message includes the new SEK, that will be used to encrypt the messages exchanged in the Key Distribution Protocol (KDP). The IMGP messages are encrypted with the Receivers pre-shared encryption key, enabling data confidentiality and Receiver authentication.

The KDP module implements the KDP protocol that allows CEK distribution to all SEK owners. It sends UPDATE-KEY messages containing the CEK, in multicast to all local Receivers. Each secure video channel stream has a single CEK at time, which must be delivered to all authorized Receivers interested in that video channel. CEK keys are generated at the MD upon the first IMGP request for that video channel, and subsequent requests will receive the same CEK. When the CEK Time-To-Live (TTL) is reached, a new rekey process will occur with the first channel request reception, and a new CEK will be generated and sent in multicast. The UPDATE-KEY message contains several copies of the CEK, one for each interested Receiver and encrypted with each Receiver's SEK.

3.3.2 Message sequence chart

The messages exchanged between all the elements, including the modules of the MD and the Receiver elements, for a scenario of successful video channel requests, is shown in the Message Sequence Chart (MSC) of Fig. 4. Communication is always started by the Receiver, which uses the IMGP protocol to request a video channel using a TCP connection to the port 3770 of the IMGP module at the MD. When the IMGP module at the MD receives a new TCP connection, it generates a new instance of itself to process Receiver requests, until that connection is closed. The Receiver sends the channel request (CH-REQ) message to the MD, which is then processed by the newly created instance of the IMGP module.

The IMGP module, in turn, sends the SRVREQ message (service request) to the MDE module. Then, the MDE module validates the receiver identity and request and, if all is correct, it generates the CEK and the Receiver SEK. The information required by the MDE to validate the receiver entity may be previously or on-demand downloaded from a central authentication server. The MDE module also instructs the SRTP module to start transmitting the channel SRTP stream encrypted with the CEK and instructs the KDP module to start the key update process. At the end of the Receivers request process, the TCP connection to the MD's

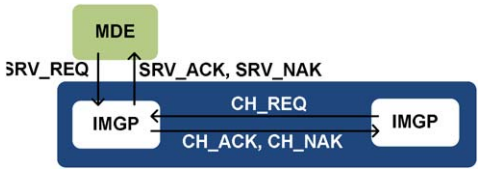

Fig. 5 IMGP subsystem

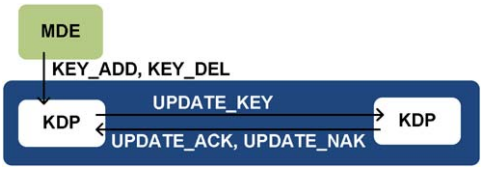

Fig. 6 KDP subsystem

IMGP module is closed, thus terminating that IMGP instance. The key update process allows the CEK multicast transmission to all the Receivers that requested that video channel. By receiving the CEK, the Receivers become able to decrypt the video channel SRTP stream.

The CEK generated by the MDE has a limited TTL. The main purpose of this TTL is to ease the management of interested Receivers; if no new CH-REQ is received during the last TTL then the Receiver is removed from that channel group. If the Receiver intends to continue receiving a video channel, then it must renew the channel request by sending a new CH-REQ message to the MD element.

3.3.3 Implicitly Managed Group Protocol

The IMGP is a simple protocol which manages groups of users interested in a channel and accepts user channel requests. The group join operation is made by the CH-REQ message (Figs. 4 and 5), and it is confirmed positively through an acknowledge message (CH-ACK), or negatively (CH-NAK). Group leave operations are implicit, that is, they are based in timeouts and, therefore, use no protocol messages.

The IMGP module at the MD interacts with its peers at the Receivers. It accepts channel requests, forwards them to the MDE using the SRVREQ message, and waits for the confirmation, which can be received through an acknowledge message (SRV-ACK), or a not acknowledge message (SRV-NAK). The IMGP module at the MD first waits for a video channel request from a Receiver; upon its reception it sends a service request to the MDE and waits for the MDE response. On a positive response, the IMGP sends a channel request message confirmation; on a negative response it sends a not acknowledged channel request. The video channel request message sent by the Receiver to the MD IMGP module is encrypted with a symmetric pre-shared encryption key that enables the confirmation of the Receiver identification.

The channel request message (CH-REQ) sent by Receivers consists of three fields: 1) a message type identification string; 2) the Receiver identification code; and 3) and encrypted field. The encrypted field contains a message identification sequential number, the requested channel identification code, and the Receiver identification code again. This part is encrypted with the Receiver's symmetric cryptographic key that is pre-shared with the MD element. The clear text part of the channel request message enables the MD to understand that it is a channel request and to decide what cryptographic key to use. The encrypted part of the channel request enables the control of the message sequence and the confirmation of the Receiver identification, by comparing the clear text Receiver identification code with the decoded Receiver identification code.

The channel request message confirmation (CH-ACK) sent by the MD to Receiver is composed of four parts: 1) a message type identification string; 2) the message sequence number (plus one that the sequence number in the CH-REQ message); 3) a digital signature; and 4) an encrypted part. The encrypted part is composed by the new SEK that the Receiver must use to decrypt future messages sent from the MD element, and a time to live for that SEK. The signed parts of the messages enable Receivers to validate the MD identity; Receivers must be able to generate the same signature using the MD public key.

3.3.4 Key Distribution Protocol

The KDP is also a simple protocol that allows the distribution of CEK to authorized receivers. It is based in the key distribution messages adopted in Iolus framework [9] and in the protocol described in [21]. The KDP subsystem is shown in Fig. 6, where four message types can be identified: 1) KEY-ADD; 2) UPDATE-KEY; 3) UPDATE-ACK; and 4) UPDATE-NAK. The KEY-ADD message has significant role in the operation of the KDP module at the MD; this message instructs the KDP module to append a new SEK to the list of SEK of a channel, and to initiate the key update procedures towards the Receivers. The KEY-ADD message is sent by the MDE module to inform the KDP module about the cryptographic key of a channel (CEK) or about the cryptographic key of a Receiver (SEK). This message structure consists of four fields: 1) a message type identification string; 2) a numeric code representing the type of key (0 for CEK, other for SEK); 3) an identification code used for channel identification or Receiver identification, depending on the value of the previous field; and 4) the cryptographic key.

The UPDATE-KEY message has a variable size, and depends on the number of Receivers viewing the channel. This

Fig. 7 Test bed #1

message is composed of five fields: 1) a message type identification string; 2) a message sequence number; 3) the channel identification code; 4) a digital signature; and 5) a variable set of pairs of Receiver identification codes and CEK encrypted with Receiver SEK. The number of pairs included in UPDATE-KEY messages depends on the scenario; if the UPDATE-KEY is generated by a KEY-ADD with a Receiver SEK or by an UPDATE-NAK message, then there will be only one pair in the UPDATE-KEY message; if it is generated as a consequence of a KEY-ADD message containing a CEK, then there will be a pair for each Receiver requesting the channel.

4 The innovation of the system

We claim that the MD is a new network element. Its new properties include the capability of optimizing last mile network usage according to Receivers request, and on filtering undemanded video channel transmissions. It enables the control of individual Receiver access to the video channels transmitted in the last mile area network through the use of secure multicast. It is scalable, and offers fast response times in video channel zapping situations.

The solution proposed is based on the architecture proposed by the IETF's MSEC working group. This influence can be observed specially in the decentralized group management by splitting it into several subgroups, each one having a local group manager and a key distributor. The major difference between these two architectures is on the data transmission path. In our solution the MD is part of the data path; in MSEC's this does not happen. The solution proposed requires traffic filtering at the edge of the core network, whereas the MSEC architecture does not consider the data path, which is assumed to be a multicast transmission. Another difference is the importance given to source authentication. In MSEC this is required and supplied; in our solution it is not considered a key issue.

5 Evaluation of the secure video distribution system

The methodology adopted for evaluating our system is based manly on passive testing; the system is used according well-known scenarios and traffic and signaling messages are captured using appropriate tools. The application used to capture the network traffic was the Ethereal, for which a dissector was developed. The latter was used to analyze the network protocols developed (IMGP, KDP), as well as the other messages exchanged between the MD modules.

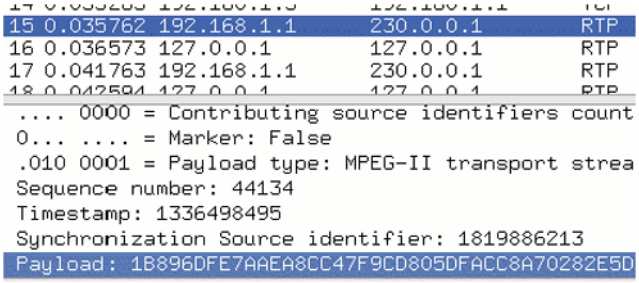

Fig. 8 SRTP packet highlight

The section starts by presenting the functional tests used to verify the system requirements. For that purpose, we present some captures of live network traffic. The performance tests follow, with the aim of demonstrating the behavior of the protocols proposed, from a timed perspective. Concluding, we present a set of comparisons between the proposed solution and conventional systems.

5.1 Functional tests

Our first test bed, shown in Fig. 7, consists of 3 computers each one representing a system element; they are Intel-based architectures, and run Fedora Core 3 Linux. IPv4 was used over 100 Mbit/s full-duplex Ethernet, which fits well to the reference scenario, particularly for the last mile network area where a broadcast layer 2 network is expected. Each Ethernet segment has its own IPv4 subnet. The first represents the core network and the second represents a last mile local area. The network segments are interconnected by the MD, which represents the edge element between the core and the last mile networks.

Requirement 1 says that encrypted data transmission must be used. Considering that the SRTP stream received by the Receiver is decrypted and forwarded to its loop back interface, we can identify the same portion of the stream by verifying that both the RTP and the SRTP packets have the same sequence number. Figure 8 highlights the SRTP packet having the sequence number 44134 that is sent in multicast by the MD. Figure 9 shows the RTP packet having the same sequence number, timestamp and synchronization source identifier, that is, the same information, but now with a different payload since it is decrypted.

Individual and unique identification of video channel streams and Receivers is made by means of simple sequential number generation and assignment. It fulfills Requirement 3, and partially Requirement 2, which also demands

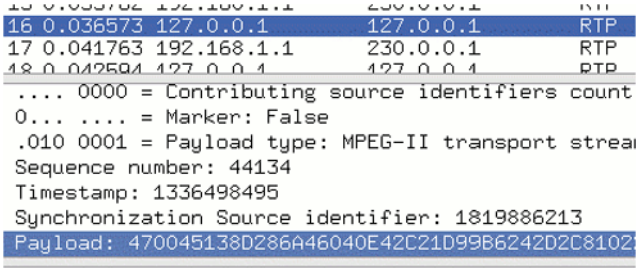

Fig. 9 RTP packet highlight

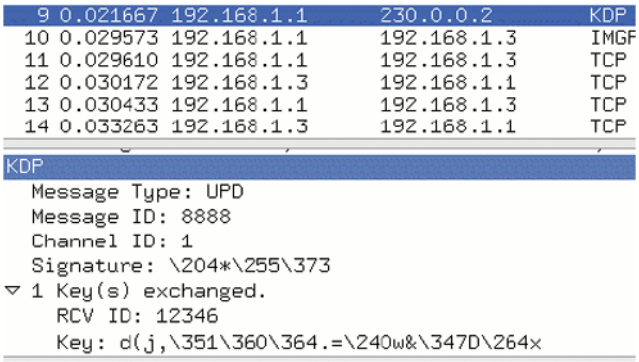

Fig. 10 IMGP UPDATE-KEY message highlight

a mechanism to authenticate each Receiver. Authentication is verified by the IMGP module at the MD and it is initially based in a pre-shared symmetric cryptographic key that is used by the Receiver for encrypting the channel request (CH-REQ) messages. If the IMGP module at the MD can decrypt correctly the channel request, then the user is considered authenticated.

Requirement 4 considers access control by user and by channel. If an authenticated user requests a channel for which it does not have authorization, this request must be rejected. This is accomplished by first authenticating the Receiver and then comparing the channel requested with the list of authorized (i.e. subscribed) channels for that Receiver. If there is a match, the Receiver is authorized and the channel is streamed; if not, the request is not considered and the connection is reset. The access control by video channel is achieved by distributing the CEK to the authorized Receivers (1 copy per Receiver) and encrypted with the Receivers session key. The UPDATE-KEY message is used for that purpose (Fig. 10). The channel source is authenticated through cryptographic techniques. For instance, all the messages sent by the server are digitally signed, and each video SRTP stream is encrypted so the source can be authenticated. Two techniques are used to guarantee the source authentication of the SRTP stream; the first consists on the decryption of the stream by the Receivers, what means that the stream was encrypted with the correct key; the second uses the capability of the libSRTP to sign the messages.

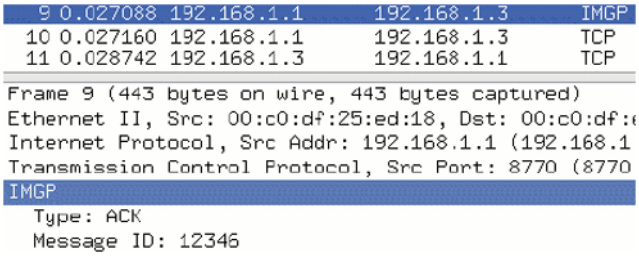

Fig. 11 IMGP CH-ACK message highlight

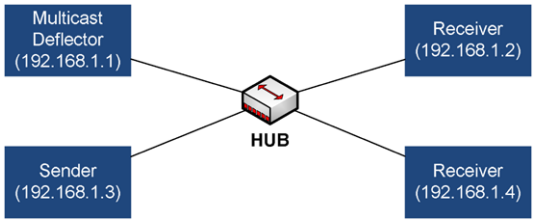

Fig. 12 Test bed #2

Requirement 6 refers to the transmission of video channels, and can be identified in Figs. 7, 8 and 10. The video channel streams in the last mile network are sent to the group address 230.0.0.1, which is the address configured for the transmissions of the encrypted channels in the last mile network. The video streams in the core network are sent with no encryption and in multicast to the group 229.0.0.1. Requirement 7 says that channels must be transmitted in the last mile network only when they are required. This means that a video channel stream shall appear in the network capture only if there was a previous channel request (CH-REQ), which was accepted (CH-ACK) by the MD. This behavior is shown in Fig. 10. Figure 11 shows a request which was rejected and no video channel is transmitted.

The last two requirements, Requirements 8 and 9 are related with Requirement 4. The user mobility restriction is achieved through access control mechanisms, which make the authorization process fail if the credentials presented do not belong to the correct local group, that is, are not associated to the correct MD element. As a consequence, if a failure occurs in a MD element only the Receivers assigned to that MD cease to operate, which improves the scalability of the proposed solution.

5.2 Performance tests

The test bed used for performance tests is shown in Fig. 12. It consists of a MD, 2 Receivers and 1 Monitor. These tests aim to show the limits of the MD when processing the signaling sent by the Receivers. In these tests, Receivers will generate signaling, MD will process it, and the Monitor will capture traffic and make the measurements.

The main parameter to evaluate is the channel request processing time, because it significantly influences the

global performance of the system. The test made to evaluate this parameter consists in sending to the MD channel requests whose rate will increase until the MD's CPU becomes 100% used. For this reason, the values obtained depend on the MD CPU power and implies the existence of multiple Receivers.

The main objective of this test is to calculate the average rate of channel requests that the MD is able to process and the bandwidth needed for this number. In order to carry out the test the Receivers code was modified by removing the code not required for channel request. During a 59 second capture, with an average signaling rate of 0,34 Mbit/s was observed, and 2291 TCP/IP connections were identified between the Receivers and the MD, meaning that 2291 channel request were processed in that time interval, that gives 39 video channel requests processed per second by the MD. During the processing of the video channel requests, the MD reached a CPU utilization of 100%. The machine used as Multicast Deflector was an AMD Athlon 64 Processor 3000+ with 512 MB of RAM.

6 Performance comparison

The performance of the proposed solution was compared with the performance of a conventional video over IP system. The conventional system is assumed to use RTP over IP multicast to transmit a channel. Two additional assumptions were made: 1) a video channel is represented by a multicast address; 2) the well known multicast protocol IGMP is used by the user to select and renew its interest in receiving a channel, and the version 2 of IGMP was selected since it is the version commonly used. The last assumption implies that group leaves are based on timeouts; that is, if a renew is not observed for the IGMP default value of 260 seconds, the channel ceases to be transmitted.

Figure 13 compares the bandwidth required by the proposed solution (MD) with the bandwidth required by the IGMP based conventional system (STD). The bandwidth considered for both systems includes the video streams and signaling. For the case when, for instance, 20 channels are transmitted in both systems, the MD system shows an increase of 5.8% of used bandwidth. This is caused mainly by the additional information required to cipher the RTP streams.

Despite of such overhead, it is still possible to obtain a significant global bandwidth reduction. Figure 14 represents de bandwidth used by a system during a 30 second interval for both the proposed solution (MD) and the conventional solution (STD). Either one (STD-1, MD-1) or two clients (STD-2, MD-2) are in a zapping scenario, changing the channel every 3 seconds. The proposed solution maintains the bandwidth constant since the number of channels

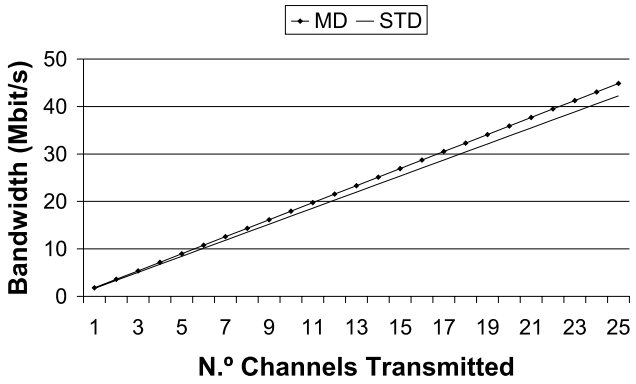

Fig. 13 Bandwidth required to transmit video channels

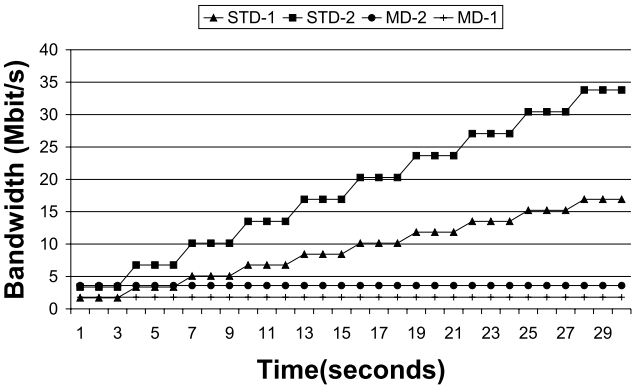

Fig. 14 Bandwidth usage in zapping scenarios

being transmitted is constant: one (MD-1) or two (MD-2). In the IGMPv2 based solutions (STD), a channel is stopped to be transmitted only after the IGMP leave timeout occurs.

7 Conclusion

The work carried out lead us to four main results: 1) video transmission system architecture; 2) multicast filtering network element; 3) group management and access control mechanism; and 4) video system prototype. To the best of our knowledge the second result the multicast filtering network element is new and it constitutes an original contribution.

The video transmission system architecture enables the transmission of simultaneous video channels over secure IP multicast. The video channels are encrypted using symmetric encryption but maintaining access control per user and per channel. This architecture is composed of four modules: MDE, SRTP, IMGP, and KDP. The MDE is the controller module, and it generates new cryptographic keys (SEK or CEK) upon requests. The SRTP module receives streams sent by the Sender, either RTP or SRTP, and encrypts them with the corresponding CEK into a new SRTP stream, which

is sent to a network area associated to a multicast address. The IMGP module is responsible for Receiver authentication and SEK generation at the time of reception of the channel request. The KDP module implements the KDP that allows CEK distribution to all SEK owners.

The multicast filtering network element transmits to the last mile network only the video channel streams requested by clients. It is achieved through a combination of multicast transmission, multicast routing tables, and management of client requests. The MD receives all the video channels, either in RTP or in SRTP. It has pre-configured a multicast routing table that forces the transmission of the channel to a multicast address through the last mile network interface card. These multicast addresses are those expected by the Receivers and they are associated to areas. The video channel sent by the MD can only be triggered by an authorized Receiver, through the use of the IMGP protocol.

The group management and the access control mechanism enables the formation of groups interested in the same data, but still controlling the access of individual clients to the data. The IMGP and the KDP protocols are used for that purpose. IMGP is a two message protocol that enables clients to join groups; it assumes the client leaves after a period of inactivity. The IMGP also supports client authentication to prevent abuse or misusing, through conventional encryption and pre-sharing of cryptographic keys. It forces each user to have an individual key that will be used to obtain each SEK that is also unique. The KDP protocol uses unique SEK to create a new CEK distribution message, and concatenates multiple CEKs, one for each interested Receiver, encrypted with these Receivers SEKs.

Future work opportunities arose during the development of this work. They can be described from two perspectives: improvement of the current scenario and adaptation of the current solution to different scenarios. Several improvements can be made: 1) use only one instance of the IMGP module at the MD per Receiver; 2) obtain information about the users profile; 3) include support for Authentication, Authorization and Accounting (AAA); 4) support for pre-paid user access. The first can be achieved by maintaining open the unicast TCP connection between the Receiver and the MD. The second can be obtained by logging successful video channel requests and the end of their transmissions. The third can be achieved be adapting the IMGP module so that it starts to validate Receiver credentials in an AAA server; it would also enable accounting possibilities. The last can also be achieved using the IMGP module but, now, interacting with a credit control server. The user profile characterization could lead to the optimization of core network, by allowing the MD to not subscribe the multicast groups associated to the channels which are never viewed by its local Receivers.

The adaptation of the solution proposed to different types of last mile media, simultaneously or not, arises also as interesting, particularly when considering group management and key distribution. An example could be the adaptation of the solution to non-broadcast media having bandwidth limits, such as ADSL.

References

1. Schulzrinne, H., Casner, S., Frederick, R., & Jacobson, V. (2003). *RTP: A transport protocol for real-time applications*. IETF: RFC 3550.
2. Pinto, A., & Ricardo, M. (2005). Multiple video channel transmission using secure IP multicast. In *Proceedings of Conftele2005–5th conference on telecommunications*, Portugal.
3. Steiner, M., Tsudik, G., & Waidner, M. (1998). CLIQUES: A new approach to group key agreement. In *Proceeding of IEEE ICDCS'98*.
4. Multicast Security (MSEC), IETF Working Group. http://www.ietf.org/html.charters/msec-charter.html [1 August 2007].
5. Hardjono, T., & Weis, B. (2004). *The multicast group security architecture*. IETF: RFC 3740.
6. Baugher, M., Canetti, R., & Dondeti, L. (2005). *Multicast security (MSEC) group key management architecture*. IETF: RFC 4046.
7. Challal, Y., Bettahar, H., & Bouabdallah, A. (2004). SAKM: a scalable and adaptive key management approach for multicast communications. *ACM SIGCOMM Computer Communication Review*, *34*(2), 55–70.
8. Rafaeli, S., & Hutchison, D. (2003). A survey of key management for secure group communication. *ACM Computing Surveys (CSUR)*, *35*(3), 309–329.
9. Mittra, S. (1997). Iolus: a framework for scalable secure multicast. In *Proceeding of ACM SIGCOMM'97*, Cannes, France.
10. Dondeti, L., Mukherjee, S., & Samal, A. (2000). Scalable secure one-to-many group communication using dual encryption. *Computer Communications*, *23*(17), 1681–1701.
11. Briscoe, B. (1999). MARKS: zero side effect multicast key management using arbitrarily revealed key sequences. In *Proceeding of first international COST264 workshop on networked communication*.
12. Molva, R., & Pannetrat, A. (1999). Scalable multicast security in dynamic groups. In *Proceedings of 6th ACM conference on computer and communications security*.
13. Setia, S., Koussih, S., Jajodia, S., & Harder, E. (2000). Kronos: a scalable group re-keying approach for secure multicast. In *Proceeding of 2000 IEEE symposium on security and privacy*.
14. Rafaeli, S., & Hutchison, D. (2002). Hydra: A decentralized group key management. In *Proceeding of 11th IEEE international workshops on enabling technologies: infrastructure for collaborative enterprises (WETICE'02)*. Los Alamitos, California.
15. Perrig, A., Song, D., Canetti, R., Tygar, J., & Briscoe, B. (2005). *Timed efficient stream loss-tolerant authentication (TESLA): multicast source authentication transform introduction*. IETF: RFC 4062.
16. Baugher, M., Weis, B., Hardjono, T., & Harney, H. (2003). *The group domain of interpretation*. IETF: RFC 3547.
17. Arkko, J., Carrara, E., Lindholm, F., Naslund, M., & Norrman, K. (2004). *MIKEY: multimedia internet keying*. IETF: RFC 3820.
18. Handley, M., Schulzrinne, H., Schooler, E., & Rosenberg, J. (1999). *SIP: Session initiation protocol*. IETF: RFC 2543.
19. Harney, H., Meth, U., Colegrove, A., & Gross, G. (2006). *GSAKMP: Group secure association key management protocol*. IETF: RFC 4535.

20. Colegrove, A., & Harney, H. (1999). *Group security policy token v1*. IETF: RFC 4534.
21. Chu, H., Qiao, L., Nahrstedt, K., & Wang, H. (2002). A secure multicast protocol with copyright protection. *ACM Computer Communication Review Journal (ACM CCR)*, *XXXII*, 42–60.

Antonio Pinto received his diploma (2000) in computer science from Polytechnic of Porto, and M.Sc. (2005) degrees in communication networks and services from Porto University. Currently, he is an assistant professor at Escola Superior de Tecnologia e Gestão de Felgueiras of the Polythecnic of Porto, where he gives courses in computer networks. He is also currently involved in a PhD program on secure multicast services at Porto University.

Manuel Ricardo received a Licenticatura (1988), M.Sc. (1992), and Ph.D. (2000) degrees in Electrical and Computer Engineering from Porto University. Currently, he is an associate professor at the Faculty of Engineering, Porto University, where he gives courses in mobile communications and computer networks. He also leads Wireless and Mobile Networks area of INESC Porto.