Dynamic Exposure Control in P2PSIP Networks

Oussema Dabbebi, Badonnel Rémi, Festor Olivier

Dynamic Exposure Control in P2PSIP Networks

O. Dabbebi, R. Badonnel, O. Festor
INRIA Nancy Grand Est

***Abstract*—Voice over IP services have undergone a large-scale deployment thanks to the development of high-speed broadband access and the standardization of dedicated signaling protocols. They offer new opportunities, in particular in the context of peer-to-peer networks. However they are exposed to multiple security attacks due to a lower confinement in comparison to traditional networks. Protection mechanisms are available, but may significantly impact the service performance. We propose in this paper a risk management strategy for dynamically adapting the exposure of P2PSIP networks. We describe the underlying mechanisms for mitigating risks based on a portfolio of countermeasures. We also detail the mathematical modeling which supports our solution based on the analysis of a case study. Finally we quantify the benefits and limits of this approach through an extensive set of experiments performed with the OMNET++ simulator.**

I. Introduction

IP telephony is among the fastest increasing service in the Internet today. This development was supported by the standardization of dedicated protocols enabling the interoperability among VoIP devices. In particular, the SIP[1] signaling protocol has become the de facto open standard for supporting the creation, modification and termination of VoIP call sessions. A VoIP enterprise architecture is typically composed of a SIP server (providing proxy, registration and location services) and a set of SIP clients (also called user agents). The registration server (logical component) plays a central role because it permits to register the bindings between the public address (address-of-record SIP-URI) of a SIP client and its location (typically, IP address and port included in a contact SIP-URI). When a SIP client wants to establish a session with a SIP client of another domain, the SIP server first requests the SIP server of that domain to obtain the IP address of the destination. The source client can then directly contact the destination client, or can establish the connection through the SIP servers playing the role of proxies. In the meantime, peer-to-peer overlay networks provide new perspectives in terms of robustness and scalability. This has already been illustrated by the success of proprietary solutions such as Skype, while the latter uses a centralized global index server. Research efforts are currently spent for extending the SIP protocol for peer-to-peer environments [1]. In particular, the P2PSIP protocol [2] aims at providing an open decentralized solution where the registration and location servers are implemented by a distributed hash table (DHT), which stores the bindings between the address-of-record SIP-URI (sip:dumont@sip.example.com) and the contact SIP-URI (sip:dumont@1.2.3.4:5060). The underlying P2P architecture is typically based on a two-level hierarchy composed of ordinary peers and super peers responsible for maintaining a Chord DHT [3] containing these bindings.

The emergence of such an open protocol is quite promising for VoIP telephony. It does not require any centralized server, and provides increased performance in terms of fault tolerance and scalability. However it poses new security issues: VoIP communications are even more exposed to security threats than in traditional SIP environments. These threats are inherited from IP telephony such as SPIT[2] attacks, and from the P2P area such as sybil attacks [4]. Many protection mechanisms are available for countering these attacks, but their deployment may significantly decrease the performance of such a critical service. A major challenge is to provide an open P2P VoIP service which is (a) secure, (b) efficient and (c) truly decentralized (without requiring a central security enrollment).

In order to address this issue, we propose a new approach for automatically adapting the exposure of P2PSIP networks. The objective is to minimize the exposure to risks while maintaining network performance. The strategy relies on a set of security countermeasures which are dynamically activated or deactivated with respect to the threat potentiality and the induced cost. The main contributions of this paper are: (a) a dynamic strategy for controlling and adapting the exposure of P2PSIP networks, (b) the instantiation of this solution based on a risk modeling, (c) the identification of attack scenarios as well as the specification of a portfolio of countermeasures, and (d) the evaluation of our solution through an extensive set of experiments performed with the OMNET++ simulator.

The remainder of this paper is organized as follows. In Section II we give an overview of P2PSIP networks and identify attack scenarios. Section III describes existing work in the area and their limits. Section IV presents our dynamic exposure control strategy by detailing the underlying mechanisms and the portfolio of countermeasures. In Section V we depict the risk modeling supporting our solution. Section VI quantifies the benefits and limits of our solution based on an extensive set of experimental results obtained with the OMNET++ simulator. Section VII concludes the paper and presents future research efforts.

II. P2PSIP Network and Attack Scenarios

Using SIP in P2P networks is a key challenge for avoiding the need of any centralized servers. This new area of research has led to different extensions for this signaling protocol, in particular the P2PSIP extension. A P2PSIP network, as defined in [1], is a P2P overlay network for SIP communications, which exploits a distributed hash table (Chord DHT) for

[1]Session Initiation Protocol

[2]Spam over Internet Telephony

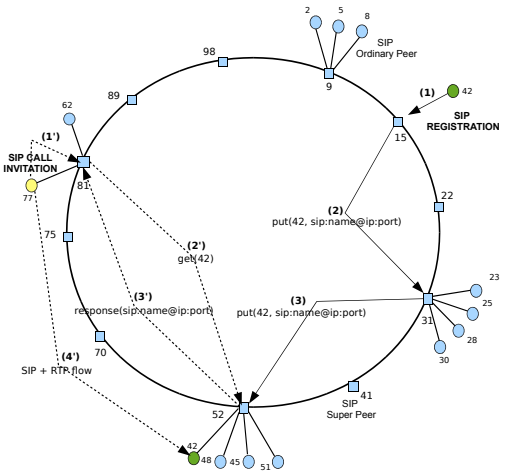

Fig. 1. Registration and Session Initiation in a P2PSIP network

registering and locating SIP users. This network is based on a two-level hierarchical architecture composed of ordinary peers and super peers. Ordinary peers, with lower resources and/or capacities, interact as SIP clients, while super peers, with higher resources and/or capacities, interact as SIP servers and maintain the DHT table. Figure 1 illustrates respectively the registration of a SIP user as well as the establishment of a SIP call session in a P2PSIP network. The peers located on the plotted ring are super peers, while ordinary peers are linked (logical view) to these super peers. During the registration phase (regular black arrows), in order to join the P2P network, a new peer has first to find a set of super peers using multicast or service location protocols. It then selects two of these super peers for redundancy purpose, and sends a SIP REGISTER message to both of them. For readability purpose, only one is represented on the figure (step 1). The super peers then store the identity and location of the peer into the DHT (steps 2 and 3), based on the hashed value of the address-of-record (AoR) SIP-URI. This DHT is maintained by super peers in a dynamic manner. During a call establishment (dashed black arrows), an ordinary peer first sends a SIP INVITE message to one of its super peers (step 1'). This latter requests the DHT serving as a registrar server (step 2'), in order to obtain the location (IP address and port) where to join the destination peer (step 3'). The call session can then be established between the initial peer and its destination (step 4').

A. Attack sources and scenarios

Based on this architecture, we consider three main classes of attackers. The first class noted A_1 corresponds to attackers implementing the P2P stack only. They do not understand SIP messages, but can modify the content of the DHT. In particular, they can modify the table entries by sending malicious messages, as well as injecting sybil peers close to a target super peer. The second class A_2 is attackers implementing P2P and SIP stacks. They typically correspond to SIP super peers. They can generate and modify SIP entries in the DHT. They can also interact as SIP proxy servers in order to perform man-in-the-middle attacks.The third class A_3 is attackers implementing the SIP stack only. These SIP ordinary peers can typically send malicious registration messages to SIP super peers in order to pollute the DHT. As highlighted in [5], the two

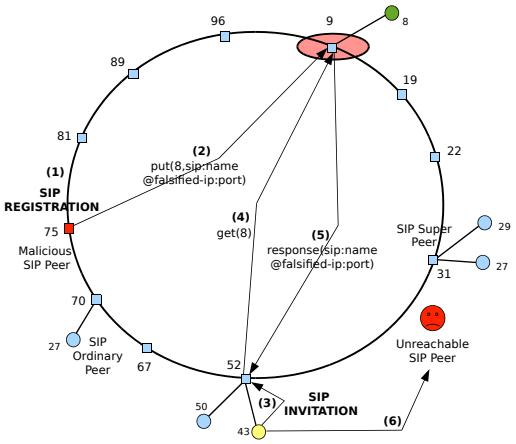

Fig. 2. SIP REGISTER attack in a P2PSIP network

major risks in a P2PSIP network are denial of service (DoS) and call hijacking. In the context of our work, we will focus on two attack scenarios corresponding to the first category. The purpose is to exploit these two scenarios as a basis for experimenting our dynamic exposure control strategy. The first scenario, noted S_1 and depicted in Figure 2, consists in performing a SIP REGISTER tampering attack in the DHT providing the registration service. The objective is to modify a registration entry (steps 1 and 2) in order to make a SIP peer unreachable (steps 3 to 6). The attacker injects in the DHT, a SIP registration entry corresponding to the AoR SIP-URI of the attacked SIP peer associated with a falsified contact SIP-URI. When a SIP peer wants to establish a call session with the attacked SIP peer, it obtains in return from the local super peer the falsified contact SIP-URI, and then the falsified IP address. This attack can be performed by the three classes of attackers. An attacker playing the role of a SIP super peer (class A_2) is more qualified to execute such an attack. A SIP ordinary peer (class A_3) can also execute this attack, but the registration has to be forwarded by its local super peer. While it does not implement the SIP stack, an attacker of class A_1 is also capable to modify an entry in the DHT. A second scenario noted S'_1, targeting also DoS, consists in inserting sybil peers in order to modify the results returned by the DHT. The objective is to insert malicious SIP super peers that precede the target SIP peer in the Chord ring. In that manner, the malicious peers can isolate the target peer, and also its SIP registration entries, by refusing to forward messages to this SIP peer. The consequences of this attack are more significant than the previous one, as it permits to directly control a segment of the network. This attack can be performed by the classes A_1 and A_2, but not by the last class A_3 which does not implement the P2P stack required for generating and inserting sybils.

It is obvious that these security attacks can be avoided

through the introduction of authentication and certification techniques. However the key challenge here is to provide an open P2P VoIP service which is secure, efficient and really decentralized. These techniques are typically in contradiction with the last constraint (when they rely on a centralized certification authority, such as defined in the RELOAD enrollment mechanism [2]) or pose serious issues with respect to the second constraint (in terms of performance and deployability). We therefore argue in favor of an adaptive solution capable of selecting different protection techniques without relying on a central entity, in order to address these three constraints.

III. Related Work

P2PSIP infrastructures are exposed to multiple security attacks. We can typically consider (1) attacks targeting the P2P overlay network, such as sybil attacks, routing attacks and eclipse attacks, (2) attacks related to the signaling protocol, such as caller ID spoofing, call hijacking, and SPIT attacks, and (3) attacks targeting the media transport protocols, such as eavesdropping attacks. A large variety of techniques has already been proposed for dealing with them [5]. In the area of P2P overlay networks, several work target identity assignment attacks in DHTs. In particular, the authors of [6] evaluate different countermeasures against sybil attacks, based on packet tracking, IP address restriction and identity verification. Securing the routing process is also a challenging issue. In [7], the authors consider three major security requirements (peer ID assignment, table maintenance, message forwarding) with respect to the routing and propose protection techniques for each of them. Trust and reward methods are also discussed in [4] with respect to this issue, and self-certification techniques were proposed in [8] to avoid the need of a centralized authority server. Several secure routing techniques for P2PSIP have already been described and evaluated in an emulative study in [9]. RELOAD integrates security features [2]. In particular, it leverages a central enrollment server to provide credentials for each peer in the P2PSIP network. We consider the decentralization as a strong constraint. Our approach is therefore complementary and could serve as a support for dynamically selecting these features based on the environment.

In the area of VoIP telephony, security threats are typically classified into five categories based on their objectives: service disruption and annoyance attacks, eavesdropping and traffic analysis, masquerading and impersonation attacks, unauthorized access attacks, and fraud attacks [10]. Significant efforts have been spent for preventing unwanted communications, in particular SPIT attacks. VoIP SEAL [11] implements a two-stage decision process: the first stage contains modules which analyze a call only by looking at the information which is available before actually answering the call. The second stage consists of modules which actually interacts with the caller or the callee to refine the detection. A survey of protection techniques against SPIT is given in [12]. The authors argue in favor of using and combining complementary techniques, which is fully in coherence with our adaptive strategy. Finally, a game theoretical model, which specifically targets the SPIT threat in P2PSIP, is described and evaluated through Monte Carlo simulations in [13]. The integration of game theory into risk management is an important challenge for capturing the behavioral dimension. A few approaches really address risk management and its dynamic instantiation in the area of VoIP networks and services [14]. Existing work related to risk assessment in VoIP infrastructures includes approaches for assessing threats (defender viewpoint) such as honeypot architectures and intrusion detection systems based on signatures or anomalies [15]. They also include approaches for assessing vulnerabilities (attacker side) such as fuzzing-based discovery and auditing/benchmarking tools. Risk models supporting this assessment may be qualitative (based on linguistic scales), quantitative (based on probabilities) or mixed [16]. Existing work on risk treatments permit to eliminate risks (risk avoidance) by applying best practices, to reduce and mitigate them (risk optimization) by deploying protection and prevention systems [17], to ensure against them (risk transfert) by subscribing an insurance contract or to accept them (risk retention) [18]. When we look further at these approaches proposed for VoIP networks and services, we can clearly observe that most of them do not really address risk management. They partially cover the risk management process, and do not integrate any risk model, or at least not explicitly, which is not the case in our solution.

IV. Dynamic Exposure Control Strategy

In this work, we propose a risk management strategy for adapting the exposure of P2PSIP networks based on a portfolio of countermeasures. Applying runtime risk management techniques in VoIP environments is a key challenge, and we have already showed the benefits of such an approach for VoIP centralized enterprise networks. This need is particularly increased in P2PSIP networks. As it was pointed out by Brian et. al. in [19], **"any P2PSIP protocol must offer a range of security models that can be selected according to the needs of the overlay."**. Risk management provide new perspectives for enabling this selection in an automatic manner, or at least for suggesting to administrators the most convenient security models with respect to the current network context in a semi-automatic manner. The runtime selection of available countermeasures permits to dynamically control the network exposure to risks based on a risk modeling. Risk is the combination of the probability that a given threat exercises a vulnerability on a given system and the resulting impact of that adverse event on this system. It can be mathematically defined by Equation 1 where a stands for an attack and A stands for the set of potential security attacks.

$$\mathcal{R} = \sum_{a \in A} \mathcal{P}(a) \times \mathcal{E}(a) \times \mathcal{C}(a) \quad (1)$$

The $\mathcal{P}(a)$ parameter represents the threat potentiality, $\mathcal{E}(a)$ stands for the exposure of the infrastructure with respect to this threat (based on the existing vulnerabilities), and $\mathcal{C}(a)$ quantifies the consequences of this attack on the infrastructure. Our exposure control strategy summarized in Figure 3 aims

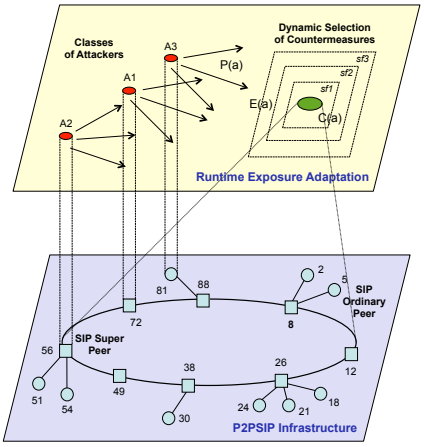

Fig. 3. Dynamic exposure control in P2PSIP networks

at adapting the exposure $\mathcal{E}(a)$ in order to maintain the risk level to an acceptable value. The exposure is controlled through the activation or deactivation of countermeasures. Another important parameter to be taken into account during the selection of countermeasures is the induced costs. Let $Sf = \{sf_1, \ldots, sf_m\}$ be the set of available countermeasures. The objective of the exposure control, as defined by Equation 2, is to maintain the calculated risk level less than a threshold value, noted R_{th}, while minimizing the costs induced by activated countermeasures.

$$maintain(R < R_{th}) \; and \; minimize(\sum_i cost(sf_i)) \qquad (2)$$

In that context, two major mechanisms [14] drive the exposure control strategy. The restriction mechanism reduces the exposure $\mathcal{E}(a)$ by activating countermeasures when the risk level is high. The relaxation mechanism reduces the cost of countermeasures when the risk level is low, typically when the threat potentiality decreases. In order to apply this exposure control in P2PSIP infrastructures, we need first to identify a portfolio of countermeasures. The purpose is not to establish an exhaustive list, but to focus on the variety of P2PSIP countermeasures, which is required for allowing a fine-grained and progressive exposure control. We briefly describe each of these countermeasures and their properties below.

1) Correction by SIP registration: this countermeasure consists in periodically sending SIP REGISTER messages with an adaptive frequency to correct a falsified entry. It checks the contact SIP-URI associated to a given AoR SIP-URI, and corrects this entry in the DHT, as soon as falsification has been detected. The detection is performed by a trusted node querying the value associated to a known AoR SIP-URI. If the value differs from the expected contact SIP-URI, the peer sends a new SIP REGISTER message in order to overwrite the falsified value. The frequency of checks and corrections is progressively adapted to the frequency of the pollution. This countermeasure is not efficient in case of sybil attacks, and may generate a non negligible signaling cost.

2) Replication of SIP identifiers: this countermeasure consists in registering replicated SIP identifiers in the DHT in order to prevent falsified entries. This replication can typically be done by applying several times the hash function to the initial AoR SIP-URI. These replicated SIP identifiers are located in different DHT segments, which complicate the execution and cost of sybil attacks. This countermeasure supposes that, during the call session initiation, the super peer performs several DHT requests in order to obtain the values associated to different replicates of the same peer, and compare them to detect inconsistencies.

3) Adaptation of the P2PSIP routing: the routing strategy in the DHT is typically performed in a recursive manner [1]. A recursive strategy generates a lower routing overhead in comparison to an iterative strategy. While the latter is more expensive, the iterative routing constitutes an interesting approach for increasing the control over the routes taken by P2PSIP messages, in particular registration messages. The objective is to avoid intermediate peers (class A_2) that may alter the P2PSIP messages during the registration phase or the call session initiation phase.

4) Restriction to trusted SIP peers: this countermeasure consists in improving the previous ones by restricting the interactions to a subset of trusted SIP peers. Many trust and reward algorithms have already been proposed in the framework of P2P networks, and may be easily transferred into P2PSIP infrastructures. The objective is to minimize the probability to interact with a malicious SIP peer. Several trust metrics can be envisioned in that context. A natural metric is the probability of SIP peers to properly resolve SIP registration entries they are in charge of. Typically, the checks performed by the first countermeasure to detect falsified entries permits to quantify this probability.

5) Authentication and self-certification: the last countermeasure aims at authenticating the SIP peers, without disgressing the decentralization constraint, i.e. without using a centralized certification authority. Simple authentication mechanisms are possible through the e-mail service but create a strong dependency with respect to that service. We rather consider, for this last countermeasure, a self-certification mechanism such as described in [20]. This solution poses deployability and peer capacity constraints, but does not require any centralized server.

V. Risk Modeling

Based on this portfolio of countermeasures we describe in this section the risk modeling that supports our exposure control strategy. This modeling is derived from Equation 1 of Section IV which quantifies the risk level. The activation of countermeasures permit to reduce the P2PSIP network exposure when the risk level is high, but it also implies additional overhead which may impact the network performance depending on its configuration. This compromise drives our control strategy.

A. Risk potentiality

The potentiality quantifies the intensity of threats in the network. In our attack scenarios, this measure is directly related to the detection of falsified entries in the DHT. Let $V = \{v_1, \ldots, v_n\}$ be the set of n SIP peers. Each v_i is associated to at least one AoR SIP-URI, noted $AoR(v_i)$, and a contact SIP-URI, noted $Contact(v_i)$, containing the IP address. Let h be the hash function which provides the key of a SIP registration entry in the DHT. The key of the SIP peer v_i is therefore given by $h(AoR(v_i))$ and is associated to the value $Contact(v_i)$. The detection of falsified entries is performed by periodically querying the DHT. Let T be the time space divided into non homogeneous time intervals $[t_i,t_{i+1}]$ where t_i indicates the instant time at which the i^{th} check is performed in the network. During this detection, the DHT query is typically performed by a SIP peer v_j different from the SIP peer v_i whose SIP registration entry is checked. The objective is to prevent the attacker to change its behavior when it observes that the SIP peer performing the query is the owner of the requested identity. The frequency of these DHT queries is adaptive over time and is increased when a falsified entry is detected. The checks are performed as follows, with an initial instant time t_0 set to 0 and an initial time interval $[t_0,t_1]$ is set to ϕ.

- Initially, the SIP peer v_j knows both the AoR SIP-URI $AoR(v_i)$ of peer v_i and the expected contact URI $Contact(v_i)$. It calculates the key $h(AoR(v_i))$ based on the hash function h. It then executes a FIND query for this key into the DHT. It obtains in return the value, noted $Contact'(v_i)$, associated to this key.
- If the observed $Contact'(v_i)$ is different from the expected $Contact(v_i)$, the frequency of next checks is increased as soon as the time interval $[t_{i-1},t_i]$ is less than a minimal value ϕ_{min}, and the next check instant time t_{i+1} is set as follows $t_{i+1} = t_i + \frac{t_i - t_{i-1}}{2}$.
- If the observed $Contact'(v_i)$ is equal to the expected $Contact(v_i)$, the SIP registration entry in the DHT is not classified as falsified, and the checks retrieve its initial frequency, which means the instant time of the next check is set as $t_{i+1} = t_i + \phi$.

This detection scheme can be refined by introducing multi-states automata. In the context of this detection, the potentiality is directly calculated based on the frequency of checks, which expresses the intensity of the attack over time. When the potentiality increases, our exposure control solution determines whether new countermeasures have to be activated.

B. Risk consequence

As defined in Equation 1, another important parameter is the consequence of a successful attack. In the context of our attack scenarios related to denial of service, the objective is to determine the damages that will occur in the P2PSIP network, when the attacker succeeds to make the targeted SIP peer unreachable. A natural manner of quantifying these consequences is to consider the SIP incoming call sessions that are lost because the source SIP peer obtains a falsified Contact URI. Let S and O be the subsets of V representing respectively the set of SIP super peers and the set of ordinary SIP peers. Let $A(v_i)$ be the set of SIP ordinary peers logically attached to a super peer $v_i \in S$. Let $InAvg$ be the function that provides the average number of incoming call sessions for a given SIP peer $v_i \in V$ during a regular time period. In order to quantify the consequences, it is important to distinguish the case where the targeted SIP peer is an ordinary peer, from the case where it plays the role of a super peer.

- When the peer v_i is a SIP ordinary peer ($v_i \in O$), the consequences of a successful attack can be directly calculated based on the $InAvg(v_i)$ value, which permits to estimate the average number of lost call sessions.
- When the peer v_i is a SIP super peer ($v_i \in S$), the consequences should also take into account the SIP ordinary peers that are logically attached to the SIP super peer. We therefore add to the $InAvg(v_i)$ value another parameter $\sum_{v_j \in A(v_i)} InAvg(v_j)/r$ with the r factor specifying the level of replication.

The purpose of this quantification is to determine the importance of a given SIP peer, and therefore the impact on the P2PSIP network if this SIP peer is unreachable due to such a denial of service attack. It is possible to refine this parameter by integrating additional factors, by considering, for instance, that SIP call sessions are not equally important.

C. Risk exposure

The last parameter of Equation 1 corresponds to the exposure of the P2PSIP network. We dynamically control this exposure through the activation or deactivation of security countermeasures. The activation of a countermeasure permits to restrict the exposure of the P2PSIP network, while its deactivation increases its exposure. Considering the set $Sf = \{sf_1, \ldots, sf_i\}$ already defined in Section IV, let $Active(sf_i)$ be a function that indicates if the countermeasure sf_i is activated. We define the exposure $\mathcal{E}(a)$ of the P2PSIP network with respect to a given attack based on Equation 3.

$$\mathcal{E}(a) = 1 - \frac{\sum_i \sigma_i(a) Active(sf_i)}{\sum_i \sigma_i(a)} \tag{3}$$

In this equation, the σ_i value quantifies the impact of the countermeasure sf_i on the network exposure. The exposure is maximal when none of the available countermeasures is activated ($\forall\ sf_i \in Sf, Active(sf_i) = 0$). In our attack scenarios, we consider three classes of attackers A_1, A_2 and A_3, and define the impact value as specified by Equation 4.

$$\sigma_i = \sum_j p(A_j) \times p(sf_i|A_j) = \sum_j p(A_j) \times \alpha_{ij} \times \delta_i \tag{4}$$

The $p(A_j)$ term stands for the probability of existence of the attack source A_j while $p(sf_i|A_j)$ is the probability that the countermeasure sf_i counters the source A_j. $p(A_1)$, $p(A_2)$ and $p(A_3)$ were considered as equiprobable in our work, but these probabilities could be refined based on statistical analysis. The $p(sf_i|A_j)$ term can be decomposed as the

product of two elementary terms. The first term, noted α_{ij}, quantifies to what extent the countermeasure sf_i impacts on the source A_j, while the second term, noted δ_i corresponds to the intrinsic characteristics of the countermeasure sf_i, such as the frequency of corrections by SIP registrations, or the number of replicas of SIP identities. The decomposition of the impact value σ_i for a given countermeasure sf_i can easily be represented as a probability tree. The activation of countermeasures can be exclusive or cumulative. The impact of cumulative countermeasures may overlap with respect to a given security attack. We consider this specific issue as out of the scope of this paper. In case of cumulation, an additional term must be integrated into this formulation, in order to quantify and subtract the intersection of the two overlapping countermeasures.

D. Cost of countermeasures

The cost of countermeasures is not required for assessing the risk level and therefore does not appear in Equation 1. However this parameter plays a central role for the selection of countermeasures, as mentioned in Equation 2. The exposure control permanently tries to mitigate the risk level while minimizing the cost of countermeasures. We mean, by cost of countermeasures, the additional overhead that has to be supported by the infrastructure and its non malicious users. In the context of this work, we decompose the cost of countermeasures into three elementary terms. The first term corresponds to the traffic overhead. The objective is to quantify the number of additional signaling messages required for executing the countermeasure. For instance with the first countermeasure, this cost includes the SIP REGISTER messages required for correcting the falsified entries. The second term corresponds to the time overhead during the establishment of a SIP call session. It expresses the additional time a regular user has to wait before effectively establishing the call session. For instance with the second countermeasure, it corresponds to the time required for obtaining and comparing the values of the replicated SIP identifiers. The last term expresses the cost of deployment for the countermeasures. This cost is often not negligible, in particular for the last considered countermeasures.

VI. Experimental Results

In order to evaluate the performance of our exposure control solution, we conducted an extensive set of experiments using the OMNET++ simulator, combined with the OverSim package [21]. During these experiments, we considered a P2PSIP network composed of up to 300 SIP peers based on the Chord DHT. We focused on the first attack scenario described in Section II, with an attacker altering the SIP registration entries in the DHT, in order to make one or several SIP peers unreachable. We assume that 5% of SIP peers are likely to be attacked. We implemented three of the five countermeasures, respectively the correction by SIP registration, the replication of SIP identifiers, and the replication combined with the restriction to trusted SIP peers. We measured the risk level as well as the cost induced by countermeasures, and compared

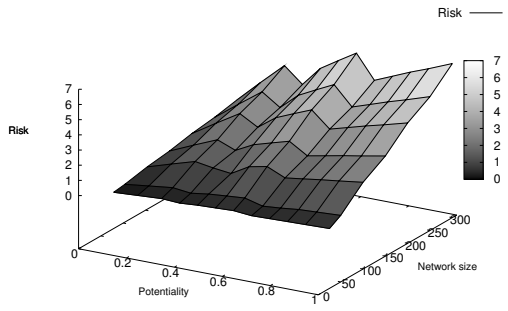

Fig. 4. Risk level in the P2PSIP network

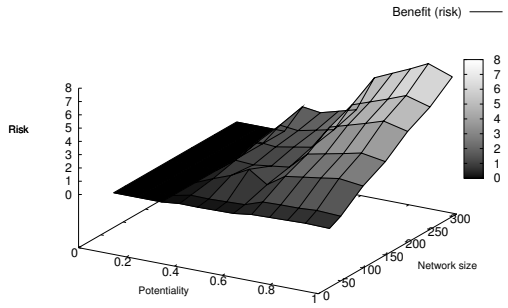

Fig. 5. Comparison to the strategy ψ_1

our approach to other two strategies: an open strategy, noted ψ_1, which consists in systematically minimizing the countermeasure cost, and a closed strategy, noted ψ_2, which consists in systematically minimizing the risk level.

A. Risk analysis

In a first series of experiments we were interested in analysing the behavior of our exposure control strategy with respect to the risk level. We quantified the risk level in function of the threat potentiality for differents sizes of P2PSIP networks. These experimental results are synthesized in Figure 4. The two first axes stand for respectively the threat potentiality and the network size, while the third axis indicates the measured risk level. The threat potentiality is a normalized value (from 0 to 1). We modify this potentiality by varying the frequency at which the attacker injects falsified SIP registration entries in the DHT. The network size varies from 0 (theoretical value) to 300 SIP peers. The risk level corresponds to the effective risk level, i.e. after the application of active countermeasures. We can clearly observe on this graph how the threat potentiality impacts on the risk level. The exposure is initially equal to zero, as no countermeasure is activated. When the potentiality increases over time, the restriction algorithm progressively activates the countermeasures in order to degrade the exposure. Let consider the case of a network composed of 100 SIP peers. The first activated countermeasure is the correction by SIP registration. The objective is to compensate the pollution generated by the attacker, by performing updates of corrected values at a higher frequency. This countermeasure reduces the exposure, and thus partially reduces the risk level.

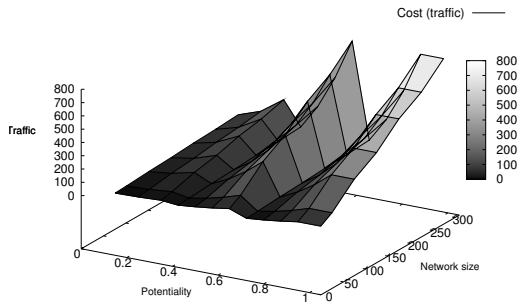

Fig. 6. Traffic overhead

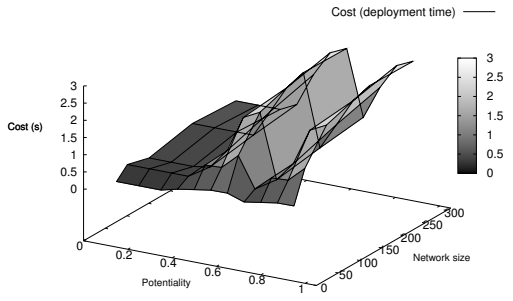

Fig. 7. Deployment time of countermeasures

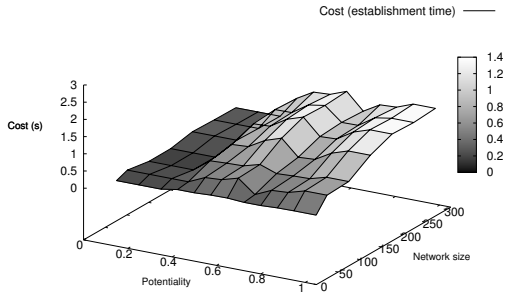

Fig. 8. Time overhead during the session establishment

However the potentiality continues to grow. At a potentiality of 0.3, the replication countermeasure is activated. The number of replicas is not constant, but goes from two to five replicas. The replicated SIP identities are obtained by successively applying the hash function on the initial SIP-URI: $h(AoR(v_i)), h(h(AoR(v_i))), h(h(h(AoR(v_i))))$. The potentiality keeps growing and reaches a value of 0.5. At this value, the algorithm activates the third countermeasure, which permits to restrict the interactions with trusted SIP peers. This leads to a new reduction of the exposure. The risk level falls to a value close to 1. As no other countermeasure is available, the risk level continues to grow until the potentiality reaches its maximal value. The same phenomenon is observed for the different network sizes. However the higher the network size is, the higher is the risk level. This can be easily explained by the fact that we consider 5% of SIP peers likely to be attacked in the network. The highest risk level in these experiments is of 7 attacked SIP peers. We can also infer the risk level expressed in terms of lost call sessions, as defined in Section V. We also compared our solution to strategies ψ_1 and ψ_2. We expected that our strategy outperforms ψ_1, which was the case. We plotted in Figure 5 the benefit of our strategy in comparison to ψ_1. We mean by benefit the difference between the risk level obtained with ψ_1 and the one obtained with our strategy. We quantified a benefit of up to 8 potential attacked SIP peers prevented in the best cases. We also compared our solution to ψ_2. As expected, there is no benefit. However ψ_2 is more expensive, as we will see in the next section.

B. Cost analysis

Another interesting aspect is to analyse the costs of our approach, i.e. the costs induced by selected countermeasures. Our strategy aims at adapting the exposure to the potentiality while minimizing these costs, as defined in Equation 2. During theses experiments, we varied the potentiality from 0 to 1 and observed three different criteria: the traffic overhead generated by countermeasures, the deployment time required by countermeasures, and the time overhead introduced by countermeasures at the establishment of new call sessions.

1) Traffic overhead: the countermeasures generate additional signaling traffic. This traffic is required for correcting the falsified entries in the case of the first countermeasure, and for multiplying the SIP identities in the case of the two other countermeasures. We also integrated in this cost the traffic due to the detection phase, as this one is strongly correlated to the treatment. The traffic overhead depends on the activated countermeasure and also of its intrinsic characteristics, such as the frequency of corrections and the number of replicas. We can observe in Figure 6 that the first countermeasure is globally less expensive than the two other ones. Indeed, this countermeasure requires to query the DHT to check the SIP registration entries (GET messages and their responses), and then to execute the correction by updating the falsified entries (PUT messages and their responses). In comparison, the two other countermeasures require to generate multiple registrations (PUT messages and their responses) corresponding to the different replicated SIP identities, the number of replicas varying from two to five. These experimental results are consistent with our approach: when the potentiality is growing, the strategy progressively reduces the exposure, requiring more expensive countermeasures.

2) Time overhead: the costs can also be quantified in terms of time overhead. We can include in this category the time required for deploying the countermeasures, and the additional time needed for establishing a SIP call session. The deployment time corresponds to the time for applying the selected countermeasures. We measured the duration between the emission of the first message part of the countermeasure, and the reception of the last message part of this countermeasure. As depicted in Figure 7, the first countermeasure is globally less expensive than the two other countermeasures which show a deployment time of up to almost 3 seconds. This can be explained by the time required for propagating

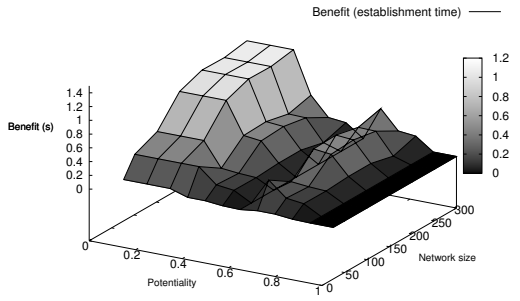

Fig. 9. Comparison to the strategy ψ_2

the replicated SIP identities. We also measured the additional time experienced by end users at the establishment phase (see Figure 8). The two last countermeasures are clearly more expensive with an additional time of up to 1.4 seconds. Indeed, they require to obtain the values of different replicated SIP entities. These values are then compared before contacting the destination SIP peer. We also compared our strategy to ψ_1 and ψ_2. ψ_1 introduces a lower time overhead at the session establishment, however it provides worst performance with respect to the risk level, as depicted in Figure 5. The comparison of our strategy with ψ_2 is detailed in Figure 9. Our strategy globally outperforms ψ_2, the best benefits being observed with a low threat potentiality. These different results illustrate the benefits and limits of our exposure control strategy. This one permits to dynamically adapt the exposure with respect to the potentiality, this adaptation being drived by a compromise between risk level and countermeasure costs.

VII. Conclusions and Future Work

P2PSIP infrastructures offer new perspectives for providing VoIP services in a fully decentralized manner, but are exposed to multiple security attacks. While a large variety of protection mechanisms is available, this security may be in conflict with decentralization, or may impact the service performance. Our objective is to enable both security and performance for an open and fully decentralized P2PSIP service. In order to address this compromise, we have defined a strategy for dynamically controlling the exposure of a P2PSIP network based on a set of countermeasures. We have described the underlying mechanisms and detailed the instantiation of this strategy based on a risk modeling. In that context we have identified attack scenarios regarding denial of service, and established a portfolio of countermeasures, from the simple correction by SIP registrations to the certification of SIP peers. We have then evaluated the performance of our solution through an extensive set of experiments performed with OMNET++. In particular, we have shown how this strategy is capable of maintaining the risk level while minimizing the costs induced by countermeasures. We have quantified the benefits and limits with respect to two alternative strategies. We have also observed how the selection of countermeasures is influenced by the threat potentiality and the network size. This work is consistent with the requirements previously expressed in [19] and [12] regarding the plurality of protections and their selection in these environments.

As future work, we are interested in defining and comparing elaborated metrics for quantifying the intersection between two cumulative countermeasures. We want also to correlate our simulation results with prototyping experiments and to cover collaborative attacks. We have voluntary focused in this work on protections that respect the decentralization constraint. We are planning to evaluate to what extent our solution can integrate and coexist with centralized protections such as the enrollment mechanism defined in RELOAD [2].

References

[1] K. Singh and H. Schulzrinne, "Peer-to-peer Internet Telephony Using SIP," in *Proc. of the International Workshop on Network and Operating Systems Support for Digital Audio and Video*, ser. NOSSDAV '05. New York, NY, USA: ACM, 2005, pp. 63–68.

[2] REsource LOcation And Discovery (RELOAD), IETF Internet Draft, draft-ietf-p2psip-base-15, May 2011.

[3] E. K. Lua, J. Crowcroft, M. Pias, R. Sharma, and S. Lim, "A Survey and Comparison of Peer-to-Peer Overlay Network Schemes," *IEEE Communications Surveys and Tutorials*, pp. 72–93.

[4] J. Seedorf, "Security Challenges for Peer-to-Peer SIP," *IEEE Network*, vol. 20, no. 5, pp. 38–45, 2006.

[5] D. Chopra, H. Schulzrinne, E. Marocco, and E. Ivov, "Peer-to-peer Overlays for Real-Time Communications: Issues and Solutions," *IEEE Communications Surveys and Tutorials*, vol. 1, 2009.

[6] T. Cholez and I. Chrisment, "Evaluation of Sybil Attacks Protection Schemes in KAD," in *Proc. of the 3rd International Conference on Autonomous Infrastructure, Management and Security (AIMS'09)*. Berlin, Heidelberg: Springer-Verlag, 2009, pp. 70–82.

[7] C. Miguel, P. Druschel, A. Ganesh, A. Rowstron, and D. S. Wallach, "Secure Routing for Structured Peer-to-peer Overlay Networks," *SIGOPS Oper. Syst. Rev.*, vol. 36, pp. 299–314, December 2002.

[8] A. Keromytis, "A Survey of Voice over IP Security Research," in *Proc. of the 5th International Conference on Information Systems Security (ICISS'09)*. Berlin, Heidelberg: Springer-Verlag, 2009, pp. 1–17.

[9] J. Seedorf, F. Ruwolt, M. Stiemerling, and S. Niccolini, "Evaluating P2PSIP under Attack: An Emulative Study," in *Proc. of the IEEE Global Telecommunications Conference (GLOBECOM'08)*, December 2008.

[10] Voice over IP Security Alliance, "VoIP Security and Privacy Threat Taxonomy," *www.voipsa.org/Activities/taxonomy.php*, October 2005.

[11] R. Schlegel, S. Niccolini, S. Tartarelli, and M. Brunner, "Spam over Internet Telephony (SPIT) Prevention Framework," in *Proc. of the IEEE Global Communications Conference (GLOBECOM'06), USA*, 2006.

[12] V. M. Quinten, R. van de Meent, and A. Pras, "Analysis of Techniques for Protection Against Spam over Internet Telephony ," in *Proc. of 13th Open European Summer School EUNICE 2007*, July 2007.

[13] S. Becker, R. State, and T. Engel, "Using Game Theory to Configure P2P SIP," in *Proc. of the 3rd International Conference on Principles, Systems and Applications of IP Telecommunications (IPTComm'09)*. New York, NY, USA: ACM, 2009, pp. 6:1–6:9.

[14] A. Gehani and G. Kedem, "RheoStat: Real Time Risk Management," in *Proc. of the 7th International Symposium on Recent Advances in Intrusion Detection (RAID'04)*, Springer, 2004.

[15] D. Shin and C. Shim, "Progressive Multi Gray-Leveling: A Voice Spam Protection Algorithm," *IEEE Network Magazine*, vol. 20, 2006.

[16] T. Bedford and R. Cooke, *Probabilistic Risk Analysis: Foundations and Methods*. Cambridge University Press, 2001.

[17] N. d'Heureuse, J. Seedorf, S. Niccolini, and T. Ewald, "Protecting SIP-based Networks and Services from Unwanted Communications," in *Proc. of the GLOBECOM'08 Conference*, 2008.

[18] ISO/IEC 27005, Information Security Risk Management, http://www.iso.org.

[19] A P2P Approach to SIP Registration, IETF Internet Draft, http://www.p2psip.org/drafts/draft-bryan-p2psip-dsip-00.html.

[20] J. Seedorf, "Using Cryptographically Generated SIP-URIs to protect the Integrity of Content in P2P-SIP," in *Proc. of the 3rd Annual VoIP Security Workshop (VSW'06)*, Berlin, June 2006.

[21] OMNeT++ Network Simulation Framework, http://www.omnetpp.org.