Computer Networks 52 (2008) 2432–2446

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

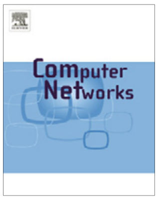

Assessing the risk of intercepting VoIP calls

M. Benini *, S. Sicari

Università degli Studi dell'Insubria, Dipartimento di Informatica e Comunicazione, via Mazzini 5, IT-21100 Varese, Italy

ARTICLE INFO

Article history:
Received 6 February 2007
Received in revised form 29 January 2008
Accepted 5 May 2008
Available online 8 May 2008

Responsible Editor: Chuang Lin

Keywords:
VoIP security
Risk assessment

ABSTRACT

Voice over-IP (VoIP) solutions and services for corporate telephony are usually marketed as 'cost-free' and 'secure': this paper shows that both statements are false in general. Though being no doubt about the economical benefits resulting from the adoption of VoIP products instead of the standard telephony, hidden costs related to VoIP services security arise whenever a company intends to assure the privacy of its phone conversations. This conclusion is extensively justified in the literature and this article aims at reasserting it by analysing the risk that a VoIP phone call may be intercepted when travelling across the Internet. The purpose of deriving a well-known conclusion consists in proving that a general and formal risk assessment method can be used in place of ad-hoc methods not only without losing the strength in the results but also adding up a sound mathematical and engineering foundation.

1. Introduction

Voice over IP (VoIP) services have seen a great raise of interest and popularity in recent years, probably because simple yet effective products, i.e. *Skype* [1], have appeared in the market, promising high-quality and low-cost substitutes for the traditional telephony. However, they are beginning to cause a new set of problems, despite being mature enough to partly fulfil these expectations. In this respect, security is undoubtedly the most questionable aspect of VoIP: in the world of traditional telephony, the privacy and security of conversations are guaranteed up to the physical layer of a network; a phone call can be heard by an intruder either by directly listening to the call, i.e. being in the same room, or by violating the physical security of the phone network itself or its devices, i.e. by putting a phone in parallel on the same line.

The problem of VoIP security has been addressed by many researchers in the telecommunication and in the Internet security fields [2–7] (see also Section 6 for further discussion), as well as by newspapers, i.e. see [8–10]. What emerges so far is that VoIP security is more than just Internet security because of the service's distinctive features: a phone call is in fact a real-time communication, thus any external action causing a delay does actually interfere with its normal flow, disturbing what is meant to be observed.

The intrinsic security problems are also increased by VoIP technology marketing strategies, which have engendered a number of misbeliefs and wrong expectations even impairing the evaluation of the risks connected with the adoption of VoIP-based solutions. In particular, marketing slogans such as 'a cost-free solution'[1] and 'a solution as secure as your network' convey the misleading information according to which VoIP services are both secure and (almost) cost-free, thus evidently underestimating security-related costs and efforts. Hence, despite the presence of mature, stable and solid VoIP products offering important economical benefits, it is to be pointed out how a knowledge of the security and privacy risks associated to their use unfortunately is still lacking.

In the light of the above sketched considerations, the present paper discusses the application of a simple yet

---

* Tel.: +39 0332 218933; fax: +39 0332 218919.
E-mail addresses: marco.benini@uninsubria.it (M. Benini), sabrina.sicari@uninsubria.it (S. Sicari).

[1] This slogan represents an extreme situation, since lots of VoIP-related services actually require the payment of small fees, i.e. a Skype call to a PSTN number.

1389-1286/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.comnet.2008.05.001

effective formal risk assessment methodology to analyse the risk of intercepting a VoIP phone call traversing the Internet. This situation is perceived as a major threat by those companies moving from traditional telephony to VoIP services as for their internal phone system: the risk analysis will prove that the 'cost-free solution' and 'as secure as your network' slogans are both false, since the adoption of VoIP solutions can involve a factual risk and, all the more, the natural and effective countermeasures aimed at mitigating it, are not cost-free and can even strongly impact on the overall security system of the company network itself.

In particular, this article takes into consideration the case of a multi-branched company internally communicating and exchanging information through the Internet: the potential attacker operates in the Internet and her/his goal consists in capturing a live conversation between two phones within the private networks. Besides, here it is going to be considered a more specific situation involving four diverse scenarios: an isolated hacker, a malicious Internet Service Provider (ISP) lying somewhere in the Internet, a malicious ISP on the route of the phone call and, finally, the case of the VoIP traffic travelling in a virtual private network (VPN). The above listed cases are quite common in those geographically distributed organisations planning to move from traditional telephony to VoIP.

A distinctive and unusual aspect of the present analysis and of the scenarios taken into consideration lies in the assumption that the attacks cannot compromise the private networks. Such a postulation has to be regarded as a limitation allowing to confute the slogan 'as secure as your network'; it will in fact be demonstrated that there is a real risk of intercepting phone calls even presupposing 'your network' as being perfectly secure. Moreover, since the attack vectors to break the security of VoIP services inside a private network are well studied, see Section 6, 'internal' threats are widely covered by the existing literature.

However, since it is here taken into account only a subset of the possible threats, aggregating the others (i.e. DNS poisoning, WiFi interception, etc.) as subclasses of general vulnerabilities, the related countermeasures will be general when addressing a class of specific vulnerabilities.

Therefore, by expanding the results in [11], where only the scenario of an isolated attacker has been accounted for, this essay evaluates in depth the risk of the call interception coming from the Internet. It will be concluded that VoIP solutions are cost-effective and their security can be ensured up to a reasonably high level; however they are definitely not cost-free and have a significant impact on the overall security of the networks hosting them. It is to be pointed out that these conclusions are well-known when the literature dealing with the same problem is considered. Hence, the novelty of this contribution lies in the way the results are derived.

The method allowing us to draw these conclusions on a strong scientific basis is in fact used to analyse a general scenario rather than a specific case. Moreover, this work shows how to draw conclusions from a risk analysis not strictly depending on the analysts' expertise,since diverse experts will achieve equivalent (in a strict mathematical sense, see [15] and Sections 3 and 5) results.

Therefore, although the above reported weaknesses are well-known due to a wide number of empirical studies, see Section 6, their structured analysis has been thus far conducted only by means of ad-hoc methods: this paper intends to convey the idea that general and formal risk assessment methodologies are as suitable as ad-hoc methods, since they lead to the same results, though being simpler to apply because of their standardisation. They even produce sounder results because their reliability is certified by a supporting mathematical theory.

2. The VoIP architecture

The standard VoIP architecture, see Fig. 1, is based on a set of hardware or software *IP phones* over an IP network; moreover, the IP network, usually the Internet, can be connected to a traditional phone system (PSTN) by means of a *VoIP gateway* transforming VoIP calls and conversations into phone calls to/from a PBX. In addition, the IP phones may benefit from a *voice server* providing auxiliary support to the VoIP services, i.e. translation from user names to IP addresses and vice versa.

The main components of the architecture are:

- *IP phone*: a terminal (*A* and *B* in the figure) with native VoIP support and the possibility to directly connect to an IP network;
- *VoIP gateway*[2]: a network device (*VG* in the figure) converting signals from/to the telephony interfaces (POTS, T1/E1, ISDN, E&M trunks) and the VoIP protocols;
- *Voice server*: a network server providing the management and administrative functions with the necessary support to the routing of the calls across the network; in a system based on H.323, the server is known as the *gatekeeper*; in SIP/SDP, the server is called *SIP server*; in a system based on MGCP or MEGACO, the server is named *call agent*;
- *IP network*: an interconnection structure based on the TCP/IP protocol family; the IP network can be a private wide-area network, an intranet, or the Internet.

As it has been stated in the Introduction, this work aims at evaluating the *wire tapping* risk in a VoIP system, i.e. the risk of successfully intercepting a live conversation between two IP phones. The core principle lying at the basis of the approach selected for the present risk analysis [12] consists in taking into consideration the dependencies among the system vulnerabilities: evidently, these dependencies are strictly related to the system architecture.

The most direct way to perform a wire tapping attack is to break the security of the private networks hosting the two communication end-points: if an intruder is allowed to enter them, s/he can dispose of a wide range of techniques to listen to VoIP conversations. These threats have been analysed at length in the literature [13,14] as

---

[2] As usual, the term 'gateway' refers to a device connecting different networks; from now on the term 'gateway' alone is reserved to router gateways, the devices connecting networks on the Internet, while 'VoIP gateway' is used when referring to the device translating VoIP into switched telephony and vice versa.

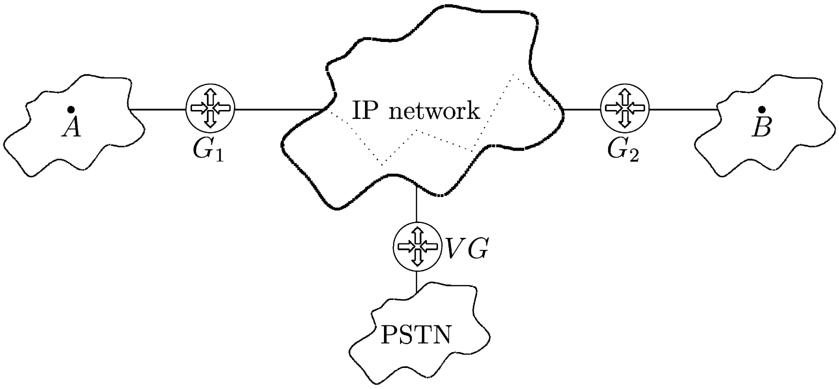

Fig. 1. The VoIP architecture.

discussed in Section 6. However, this form of attack is usually seen as unrelated to the VoIP traffic; on the contrary, it is usually believed – as far as the commercialisation of VoIP services is concerned – that 'a secure network gives a secure VoIP system'. The absolute security of private networks is thus assumed in this paper so as to confute the false beliefs according to which VoIP services security is reduced to private network security. Not only will the analysis finally prove, see Section 5, the importance of private networks security – which constitutes the major weakness as for the security of VoIP conversations – but it will also highlight further ways to break the security of VoIP systems, whose protection will involve significant economical costs. As a matter of fact, the balance between VoIP technology security and its economical advantages seem not as clear as the market typically promises.

In this respect, four scenarios can be identified, where it seems possible to carry out a VoIP phone call interception without breaking the private networks' security:

- *Scenario I: an isolated attacker in the Internet.* In this scenario, represented in Fig. 2, the IP phones A and B lie in two private networks delimited by the G_1 and G_2 gateways (the *border gateways*) connecting them to the Internet. Here a hacker is supposed to be in the public Internet with the scope of intercepting a conversation crossing the Internet from A to B.
- *Scenario II: a malicious ISP outside the route of the conversation.* The difference between this scenario, see Fig. 3, and the previous one lies in the presence of a malicious ISP outside the route of the conversation instead of the isolated hacker: it is to be pointed out how an ISP's knowledge and 'status', availability of devices as well as possibility of managing a piece of the Internet are usually deemed as a major advantage when security attacks are at issue, especially as opposed to a malicious individual's more modest opportunities.
- *Scenario III: a malicious ISP on the route of the conversation.* In this context, Fig. 4, the point where the wire tapping attack is performed is located in a malicious ISP lying on the route of the conversation. ISPs are usually reliable companies providing their clients with a secure transport of data and communications: however, a few recent cases, i.e. see [9,10], have revealed that even

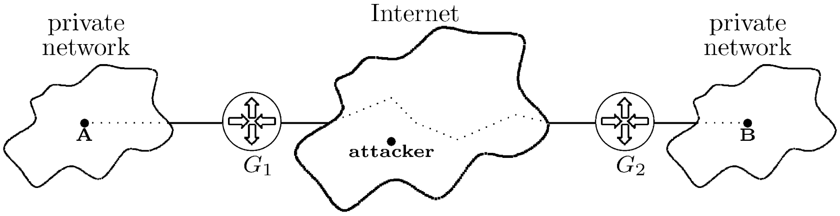

Fig. 2. The first scenario: an isolated attacker in the Internet.

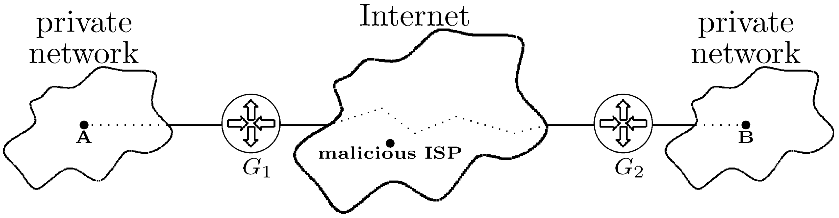

Fig. 3. The second scenario: a malicious ISP outside the route of the conversation.

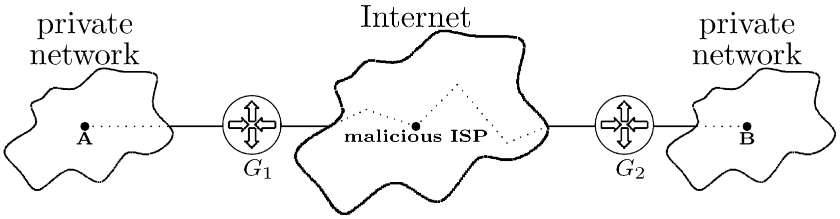

Fig. 4. The third scenario: a malicious ISP on the route of the conversation.

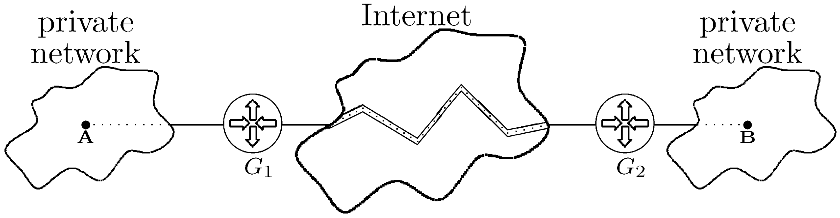

Fig. 5. The fourth scenario: the conversation travels in a VPN.

major telecommunication companies have sometimes been involved in security incidents where they acted as attackers. It seems thus worth considering what may happen when VoIP conversations are exposed to the action of a malicious ISP.

- *Scenario IV: the conversation travels in a VPN.* In Fig. 5 a VPN is adopted to improve the level of security in the architecture. The VPN links together the private networks where the IP phones are located. In this context, the conversation between *A* and *B* takes place as a communication between the private networks embedded in the VPN channel: the VPN traffic is usually encrypted by the border gateways before being transmitted through the Internet. This is the reason why it is interesting to evaluate the possibility of intercepting a VoIP phone call in this situation.

The above outlined scenarios are exhaustive covering as they do any possible position of a potential attacker operating in the Internet both in the case the VoIP traffic is inspectable and it is not (scenario IV). These scenarios can and should quite obviously be dealt with more specifically when analysing a concrete situation: for instance, if countermeasures have been taken to protect a system traffic such as BGP, some of the attacks considered in this paper cannot be launched. This is the reason why the scenarios should be regarded as general frameworks where detailed analyses of concrete situations should be conducted: interestingly enough, it should be noted that the detailed analyses are direct extensions of the scenarios taken into account.

3. Measuring the risk

Risk assessment aims at quantitatively evaluating the danger of an undesired event occurring in a given environment. As far as this paper is concerned, the environment has been described in Section 2 as one of the reference architectures represented in Figs. 1–4; furthermore, the undesired event is evident, that is to say the interception of a VoIP phone call.

Therefore, this section intends to define a specific notion of risk as well as illustrating the methodology employed to evaluate it. The risk assessment procedure is in fact based on a general engineering methodology described in other publications; some introductory information could be found in [12] while, as for the related mathematical treatment, the reader is referred to [15]. This section offers a concise overview of the risk assessment procedure in order to allow a better understanding of its application to the VoIP phone call interception.

As for the present paper's approach, the risk is a function on two variables: the damage potential, that is to say the average loss caused by an attack, and the level of exploitability measuring the easiness to break a system component, as defined in [16]. In this specific case, the risk assessment procedure intends to determine the exploitability levels.

In brief, the risk assessment procedure consists of five steps:

(1) The possible threats to the system are modelled by means of an attack tree [17]: the root node represents the attack goal and, recursively, the children can be alternative subgoals, each one satisfying the parent goal (`or` subtree) or partial subgoals, whose composition satisfies the parent goal (`and` subtree). The tree's leaves stand for the vulnerabilities of the system enabling the attacks modelled by the subtrees.

(2) The dependencies among the identified vulnerabilities are determined: a vulnerability v depends on a vulnerability w if and only if v may become easier to utilise to attain the attack goal when w has already been compromised.

(3) To each vulnerability v in the attack tree is associated a numerical index $E_0(v)$, called its *initial exploitability*, measuring the chances that v may be successfully

used to break the security of the system. Similarly, the dependencies between pairs of vulnerabilities are weighted on the same metric: a value $E(v|w)$ is assigned to each pair (w,v) of dependent vulnerabilities, meaning that the exploitability of v becomes $E(v|w)$ when w has been compromised.

(4) The exploitability $E_i(v)$ of each single vulnerability v is updated to a new value $E_{i+1}(v)$ to take into account its dependencies, until the values reach a fixed point, that is to say when the effects of the dependencies have been fully considered. As proved in [15], the iteration process converges in finite, bounded time, ensuring the termination of the process.

(5) The risk associated to the threat under examination is finally computed by recursively aggregating the exploitabilities along the attack tree. The exploitability of an `or` subtree is the easiest (maximum value) of its children, and the exploitability of an `and` subtree is the most difficult (minimum value) of its children. Finally, the aggregated exploitability of the root node, which measures the level of feasibility of the attack, is combined with the damage potential to assess the risk of the threat.

The first step generates an attack tree, whose leaves form set V of the system vulnerabilities. Likewise, Step 2 produces the *dependency graph* $G = \langle V,D \rangle$, whose nodes are the system vulnerabilities and whose edges are the dependencies: an edge $(v, w) \in D$ means that the exposition of w seems easier when v has been compromised.

In Step 3, the evaluations $E_0(v)$ of the initial exploitability of every vulnerability and the weightings $E(v|w)$ of the identified dependencies are produced. The values $E(v|w)$ obey the constraint $E(v|w) > E_0(v)$, meaning that the exploitation of w eases the abuse of v. These numerical values lie in the range [0,10] where 0 means *impossible to exploit* and 10 means *immediate*. The exploitability values are chosen by security experts conducting the risk analysis by taking into account the relative difficulty in making use of the various vulnerabilities. Although this evaluation is subjective depending on to the experts, it is remarkable the fact that the whole risk assessment procedure depends just on the ordering of the exploitability values, as mathematically proved in [15]. As a matter of fact, since different metrics with the same ordering structure are equivalent, it follows that most of the seemingly different evaluations are actually the same, despite using different values.

Then, in Step 4, the notion of exploitability is generalised by means of the function family $E : \mathrm{N} \times V \to [0; 10]$ mapping the vulnerabilities to [0,10]; thus the values $E_i(v)$, with i varying over natural numbers, are associated to the vulnerability v. The initial value $E_0(v)$ has been fixed in Step 3, while the other values are calculated by means of

$$E_{i+1}(v) = \max(E_i(v); \{\min(E(v|w); E_i(w)) : (w; v) \in D\}) \quad (1)$$

whose rationale is to include the potential influence of the dependencies in evaluating the exploitability of a vulnerability v. This influence manifests itself when it is easier to attack a connected vulnerability w both because $E_i(w)$ is higher, that is to say that w is easier to exploit, and, when w is compromised, the misuse of v is simplified, that is its exploitability becomes $E(v|w)$.

Finally, during Step 5, the outcome of Step 4, that is the fixed point values in the iteration of the formula (1), is distributed along the nodes of the attack tree. The result is an attack tree where every node is decorated by an exploitability value: then, by applying the risk function, one may calculate the risk of the root node and, if needed, the risk of every subtree as the risk of the root of the subtree.

4. The risk analysis of the VoIP scenarios

4.1. Step 1: construction of the attack tree

An attack tree [17] describes how an attacker may break the security of a system: the attacker's goal is the root of the tree. In the present case, the goal consists in intercepting a VoIP phone call crossing the Internet. Intercepting a call means that the attacker is able to listen to the communication and understand its content: in particular, it is to be pointed out that the exact copy of an encrypted VoIP communication is not considered as an interception, since its content, that is the conversation, is not disclosed. On the contrary, a real-time listening is not different from making a copy. Therefore, the main goal of the attack tree generates two subgoals whose aim is to get a copy of the communication and to understand its content. The latter goal, although difficult, is quite standard: given an encrypted communication and being able to know how the data is encoded, the attacker would have to break or guess the encryption key and decode the data to retrieve the original conversation. The goal of copying the communication is more interesting: in fact, the scenarios of Section 2 play a crucial role in the way an attacker may act.

The attacker is assumed to know how to identify the communication s/he is interested in intercepting: this hypothesis implies that the attacker knows something about the structure of the private networks at the endpoints of the communication of interest. This knowledge usually allows to determine the IP addresses of the gateways on the frontiers of the private networks: according to the initial information, the attacker may either use the networks registration data the Domain Name System, the `whois` service or trace the routes to some known internal point within the private networks. Alone or combined, these information disclosure techniques enable the attacker to learn the IP addresses of the frontier gateways; this will be therefore taken for granted from now on.

Moreover, because of the scenarios taken into consideration, and because this article focuses on confuting the misleading thesis according to which VoIP services can be added to existing private networks without altering their security posture and economical costs, the risk analysis will begin by presupposing that private networks cannot be directly attacked. As already hinted at in the Introduction, this attitude is quite common during the transition period from traditional telephony to VoIP services employment: this is the reason why the assumption inevitably confines the risk analysis to those scenarios recognised as risky. The most dangerous scenarios will be eventually result to be those usually considered as trustworthy, that is to say the attack from a 'reliable' ISP and the one inside the private networks.

The attack tree is shown in Fig. 6: it has been constructed by expanding the main goal in two subgoals, as already described. The second one (case 2) has been decomposed in two subgoals: the first one exploits the fact (case 2.1.2) that the information about the voice encoding is written in the control channel. In fact, a VoIP call operates on a double connection [18]: a control channel, utilized to determine the parameters of the communication, start and stop voice transfers, identify the connection of the media channel, etc.; and a media channel, whose function consists in transporting the voice from one end-point to the other.

The vulnerabilities, i.e. the leaves of the attack tree, marked with a (*), are considered to be immediate since, when the attack reaches that point in the tree, the difficulties in exploiting its vulnerabilities will already be overcome. On the contrary, the (+) marks on the vulnerabilities mean that they cannot be evaluated in isolation: for instance, in (case 2.2), if the voice is not encrypted, as in most cases, it is possible to immediately exploit the vulnerability; however, if the media channel makes use of a strong encryption, the same vulnerability becomes almost

Goal: To intercept a VoIP phone call
AND 1. To copy the communication
OR 1.1. To access a gateway on the path
AND 1.1.1. To identify a gateway on the path
OR 1.1.1.1. It is a border gateway (V_3)
1.1.1.2. To identify an intermediate gateway on the path
AND 1.1.1.2.1. To trace the route between the communication end-points (V_4)
1.1.1.2.2. To choose a weak gateway on the detected route (*)
1.1.2. To control the identified gateway
AND 1.1.2.1. To connect to the administration channel of the gateway (telnet, ...) (V_5)
1.1.2.2. To force the administrator's password
OR 1.1.2.2.1. Default or weak password (V_1)
1.1.2.2.2. To sniff the password (V_2)
1.1.3. To identify the communication in the traffic crossing the identified and controlled gateway
XOR 1.1.3.1. The traffic lies in a VPN
1.1.3.1.1. To decode the VPN traffic (V_8)
1.1.3.2. The traffic is inspectable
AND 1.1.3.2.1. To copy the control channel (*)
1.1.3.2.2. To identify the media channels (*)
1.1.3.2.3. To copy the media channels (*)
1.2. To divert the traffic through a malicious gateway
AND 1.2.1. To identify a gateway on the path (see case 1.1.1)
1.2.2. To poison the route between a border gateway and the identified gateway
OR 1.2.2.1. It is a intra-autonomous system gateway
1.2.2.1.1. To announce a false OSPF bandwidth (V_6)
1.2.2.2. It is a inter-autonomous system gateway
1.2.2.2.1. To announce a false BGP route (V_7)
1.2.3. To identify the communication in the traffic (see 1.1.3)
2. To decode the content of the communication
AND 2.1. To understand the coding algorithm
OR 2.1.1. To guess the coding algorithm (*)
2.1.2. To read the algorithm in the control channel (*)
2.2. To determine the encryption key (+)

Fig. 6. The combined attack tree.

impossible to attain, thus the whole (case 2) subgoal results impracticable.

Going into further detail, it is to be highlighted that the first subgoal (case 1 in the attack tree) may be reached either by gaining control of a gateway on the route followed by the communication to be copied, or by diverting the route. In the first case, the attacker can access the gateway as system manager (case 1.1.2) and then single out the communication of interest in the crossing traffic (case 1.1.3): if the communication does not travel in a VPN tunnel (case 1.1.3.2 in the attack tree, corresponding to the scenarios I, II and III), the attacker can easily obtain the RTP ports of the media channels involved in the communication by inspecting the control channel and consequently copy them; if the communication travels in a VPN tunnel (case 1.1.3.1 and the fourth scenario), the control channel cannot be directly inspected, thus the VPN traffic has to be decoded.

In the second case, if the attacker chooses to divert the traffic (case 1.2), s/he will consequently poison the route in such a way that the communication will flow through a malicious gateway under her/his control (case 1.2.2): then, he may go on listening to the communication as already described (cases 1.1.1, 1.1.2 and 1.1.3 in the attack tree). In both cases, the first step consists in individuating a suitable gateway in the route followed by the communication (cases 1.1.1 and 1.2.1): the gateway may be either a border gateway, i.e. a gateway on the frontier of one of the private networks, or an intermediate gateway; the selection will depend on its vulnerability to attacks, a feature which can be easily tested for every gateway on the identified route.

As a matter of fact, VoIP protocols peculiarities limit the possibilities of an attacker and, as a consequence, the shape of the attack tree: in fact, the admissible attacks must not interfere with the existing connections on the gateways, otherwise the VoIP call will be influenced and therefore drop. This is due to the fact that VoIP calls are real-time, streaming connections, hence any loss or detour will highly probably bringing the communication to conclusion, thus destroying what an attacker intended to observe.

Moreover, an attack to a gateway involves a stricter subset of the techniques the Internet attacker has generally at her/his disposal: the attacker can reach his/her goals only by avoiding influencing any existing connection. For instance, the category of *denial of service* attacks is banned, since they aim at substituting a device after its collapse due to resource shortage: as for the VoIP traffic, a gateway collapse can delay the voice call due to a network congestion, thus causing the call to drop. Similarly, since the goal is to copy an existing connection, the attacks to access a gateway limit to trying to login as system manager: most examples of threats involving the exploitation of software bugs in the operative system are either not deep enough to permit to copy the desired connections, or even too invasive, making the existing connections die or be delayed. Consequently, the attack tree in Fig. 6 can be deemed quite exhaustive in the development of the scenarios taken into account.

To sum up, the identified vulnerabilities are listed in Table 1. A few remarks are all the more worth reporting as follows:

Table 1
Detected vulnerabilities

Vulnerability	Description
V_1	The identified gateway has a weak authentication mechanism
V_2	A link connected to the identified gateway can be sniffed
V_3	Information disclosure on the private networks
V_4	The source routing option is enabled in one of the gateways on the frontier of the private networks
V_5	The identified gateway can be remotely controlled from the attacker's position in the Internet
V_6	The identified gateway exchanges OSPF announces with its neighbours
V_7	The identified gateway exchanges BGP announces with its neighbours
V_8	The encryption algorithm or the encryption key of the VPN channel is weak

- Identification of a gateway's address is fundamental in order to attack it, either by accessing it or diverting its traffic. V_3 vulnerability implies that, from the information regarding the conversation target of the attack, which has been assumed to be learnt, the potential intruder may reconstruct the IP addresses of the gateways on the frontiers of the private networks.
- The case 1.1.1.2.1 in the attack tree requires to trace the route between the conversation end-points and, in particular, between the two gateways on the frontiers of the private networks. This goal can be accomplished by means of the `traceroute` service, calculating the route between two nodes in the Internet as well as reporting an estimate of the round trip time. By using the source routing option of the IP protocol [19], one can make a `traceroute` from a malicious host to one of the border gateway follow a route crossing the other border gateway. In this manner, it is possible to see the optimal route between the two border gateways as well as estimating the round trip time between every pair of nodes in the route. This is the reason why the source routing option enabled in one of the border gateways has been listed as V_4 vulnerability in the table.
- V_6 and V_7 vulnerabilities have been introduced to model the fact that, in order to poison the route between the border gateways, one has to announce a false route to

Table 2
The difficulty in exploiting the vulnerabilities in the scenarios

Vulnerability	Isolated hacker	Off-path malicious ISP	On-path malicious ISP	VPN
V_1	easy	easy	very easy	?
V_2	very difficult	very difficult	very easy	?
V_3	on average	on average	very easy	?
V_4	difficult	difficult	easy	?
V_5	difficult	difficult	very easy	?
V_6	very difficult	difficult	very easy	?
V_7	very difficult	difficult	very easy	?
V_8	?	?	?	very difficult

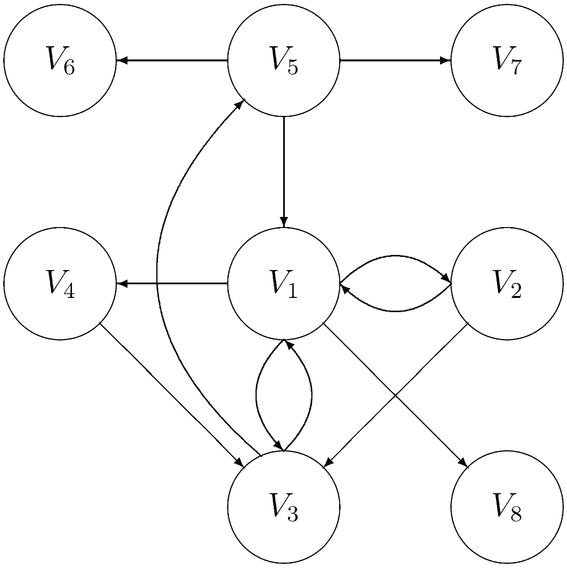

Fig. 7. The dependency graph.

a neighbour that is a legitimate gateway on the legal route. This is hardly ever possible, since only 'important' gateways are used to announce long-range routes, though it is still likely to construct false announces if one has the control of a malicious gateway credited as a legal OSPF or BGP gateway by its neighbours.

- V_8 vulnerability has been introduced because decrypting a VPN, see case 1.1.3.1.1, depends on the adoption of a weak algorithm or a weak set of encryption keys.

The identified vulnerabilities are differently important in the light of the diverse scenarios. Table 2 shows a qualitative evaluation of the difficulty in exploiting the vulnerabilities in the scenarios taken into consideration. V_8 vulnerability makes sense only within the fourth scenario, while the vulnerabilities ranging from V_1 to V_7 influence the possible attacks only within the other scenarios, hence the '?' signs.

4.2. Step 2: the dependency graph

The identified vulnerabilities are not independent: in fact, it suffices to break one of them to easier exploit the others as well. The overall framework, encoded as a dependency graph, see Section 3, is represented in Fig. 7.

Its edges can be explained as follows:

- exploiting a weak authentication in the identified gateway, i.e. V_1 vulnerability means having control of the gateway, thus V_2 vulnerability is immediately achieved; moreover, if the identified gateway is a border gateway, V_3, V_4 and V_8 vulnerabilities are achieved as well.
- misusing V_2 vulnerability means that the traffic on a link connected to the identified gateway can be observed by the attacker; if the administrator of the gateway connects via the sniffed link, V_1 is attained; moreover, the content of the traffic allows the attacker to acquire information about the private networks when the gateway forwards the traffic originated from or directed to a private net, thus simplifying the exploitation of V_3 vulnerability.
- exploiting V_3 vulnerability means collecting useful information about the private networks; if the identified gateway is a border gateway, then the collected information may reveal that the gateway is controlled also from outside, simplifying V_5, and may even give suggestions to guess the password of the gateway, thus simplifying V_1.
- It is evident that achieving V_4 means discovering the route between the two border gateways, thus implying an information disclosure, i.e. V_3.
- abusing V_5 means being aware that the identified gateway can be remotely controlled, which simplies V_1; the way to acquire this knowledge usually reveal some suggestions of the system traffic originating from the gateway, in particular the enabled routing protocols, thus allowing the exploitation of V_6 and V_7 vulnerabilities.

Hence, from a different viewpoint, the difficulty in exploiting a vulnerability – given the successful misuse

Table 3
The difficulty in exploiting the dependencies

	V_1	V_2	V_3	V_4	V_5	V_6	V_7	V_8
V_1	–	very easy	difficult	difficult	–	–	–	difficult
V_2	difficult	–	difficult	–	–	–	–	–
V_3	difficult	–	–	–	difficult	–	–	–
V_4	–	–	easy	–	–	–	–	–
V_5	on average	–	–	–	–	difficult	difficult	–
V_6	–	–	–	–	–	–	–	–
V_7	–	–	–	–	–	–	–	–
V_8	–	–	–	–	–	–	–	–

Table 4
Conversion of the qualitative evaluations into quantitative ones

very easy	easy	on average	difficult	very difficult	impossible
9	7	5	3	1	0

of a depending one – is summarised in Table 3: every entry in the table qualitatively measures the difficulty in attaining the vulnerability in the column, taking into account the previous exploitation of the vulnerability in the row: for instance, in the case the column is V_3 and the row is V_4, the table cell will measure $E(V_3|V_4)$.

As a matter of fact, the dependencies change neither in their presence nor in their evaluation in the scenarios introduced in Section 2: in fact, the scenarios act by changing the degree of exploitability of the single vulnerabilities, making dependencies useful or useless according to the initial exploitability assessment.

4.3. Step 3: evaluating exploitabilities

In the previous steps, a qualitative evaluation of the ability to exploit various vulnerabilities has been accounted for. The qualitative judgement is debatable to the extent that it has been conceived by security experts, basing their evaluation on their experience and knowledge. The reader could either agree on the evaluations provided or continue applying the method starting with a different viewpoint: further on, see Section 5, it will be highlighted that the initial assessment will have a weak influence on the conclusions of the present work. Nevertheless, the application of the risk assessment methodology, as described in Section 3, is needed so as to justify the conclusions themselves, as it will appear in the end.

Therefore, qualitative evaluations are converted into numbers, following the metric shown in Table 4: the exploitability values are in the range 0–10. The resulting dependency graphs in the various scenarios are depicted in Figs. 8–11.

In the first scenario, the initial assessment reveals that most vulnerabilities are difficult to be exploited due to the attacker's low status: an hacker does not have direct access to a trusted gateway, and thus s/he cannot poison the Internet routes (V_6 and V_7); s/he cannot either sniff a link directly connected to a gateway on the path followed by the conversation (V_2); he may use the source routing option of a gateway (V_4) or the control channel of a gateway (V_5), i.e. by trying to connect via the SSH or `telnet` protocols; these vulnerabilities are nonetheless difficult to misuse in her/his position. On the contrary, a weak authentication on the gateway (V_1) or, to a less extent, collecting information about the private networks can be successfully used to harm.

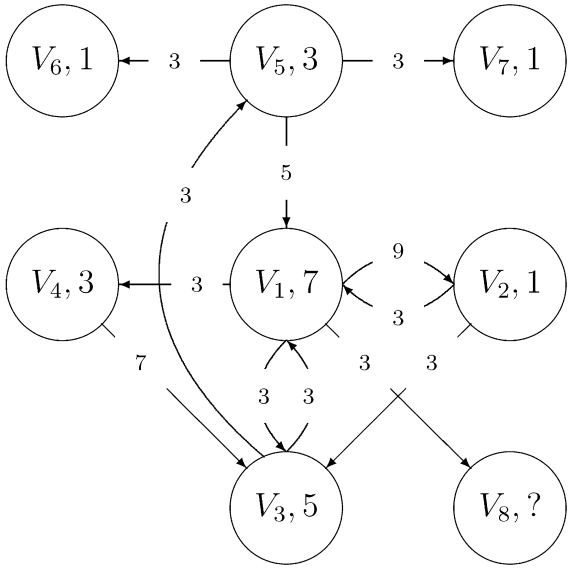

Fig. 8. The weighted dependency graph in the hacker scenario.

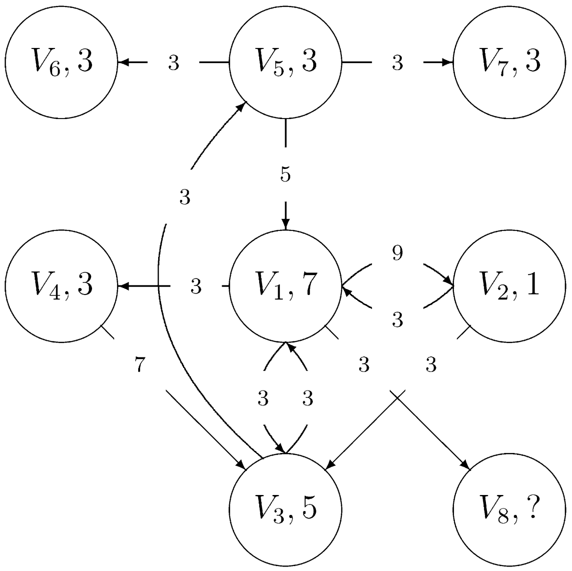

Fig. 9. The weighted dependency graph in the off-path ISP scenario.

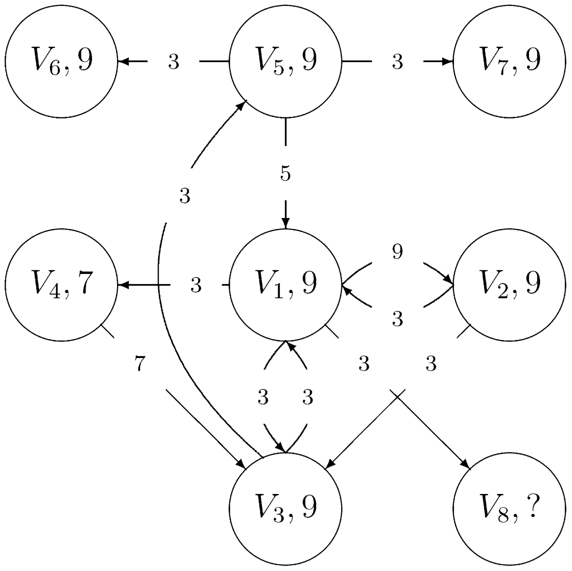

Fig. 10. The weighted dependency graph in the on-path malicious ISP scenario.

Differently, in the second scenario, see Fig. 9, the attacker holds a higher status in the Internet, that is the hacker is

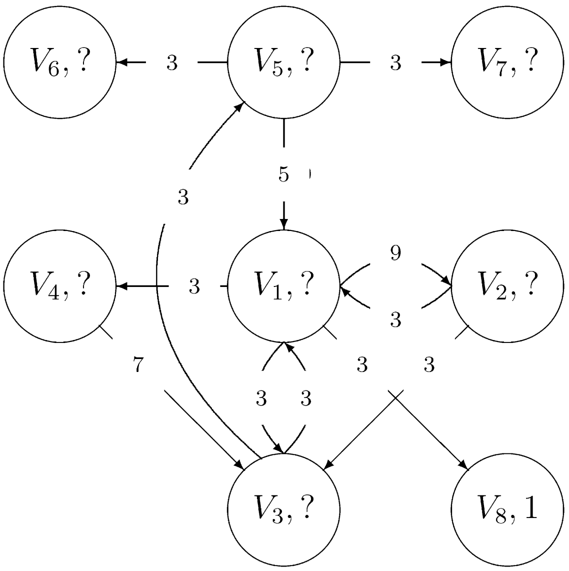

Fig. 11. The weighted dependency graph in the VPN scenario.

an ISP with a trusted gateway exchanging routing information with its neighbours. As for the first scenario, V_6 and V_7 vulnerabilities are easier to exploit since, if the attacker's gateway directly exchanges routing information with a gateway located on the conversation path, it is easier to poison its routes. Of course, this is a relative judgement: a combination of proximity and clever misuse of the attacker's gateway is required to successfully mount this kind of attack, hence the corresponding exploitability value is still 'difficult'.

The third scenario, in Fig. 10, illustrates what happens when the attacker is an ISP lying on the route followed by the conversation. In this case, all vulnerabilities can be very easily exploited, since the gateway through which the conversation flows is controlled by the attacker: it is just a matter of identifying the conversation among the many connections.

The fourth and last scenario, see Fig. 11, describes the situation where the VoIP call does not merely travel in the Internet, being embedded as it is in a VPN channel; the conversation is thus encrypted and not easily separable from the other connections in the channel. In this case, it is usually very difficult to decrypt the VPN channel; moreover, the exploitation of the other vulnerabilities should not influence the final difficulty in launching a successful attack to the system. In Section 4.5 it will be proved that, in fact, the aggregated exploitability of the root node of the attack tree depends mainly on V_8, as expected.

4.4. Step 4: propagating the dependencies

As it has already be pointed out in Section 3, the propagation of the dependencies is repeatedly calculated by applying formula (1): the results are displayed in Table 5.

In particular, the unknown ('?') values on V_8 vulnerability have been dealt with by establishing a lower bound for the corresponding exploitability value. In this way, there emerges the amount of influence induced by the dependencies on the exploitability of V_8 in the scenarios I, II and III. Instead, in the fourth scenario, the unknown values have not been lower-bounded since, as it will be highlighted in the following section, their influence on the overall risk assessment is limited to their upper bounding of V_8 vulnerability.

4.5. Step 5: aggregation and risk assessment

In order to determine the exploitability of the root node in the attack tree, i.e. the feasibility of intercepting a VoIP phone call, the exploitability values of the leaves are aggregated on every subtree as described in Section 3. The attack tree obtained as the result of the aggregation process is shown in Fig. 12: a label, indicating the exploitability values in the four scenarios under examination, is attached to every node; as already said, the nodes marked with (*) are immediate, i.e. their label is [I: 10, II: 10, III: 10, IV: 10]. The case 2.2 is marked with the [I: *x*, II: *x*, III: *x*, IV: *x*] label, whose exact value depends on the employed instance of the VoIP protocol; in particular, the value *x* represents the feasibility to decrypt the VoIP media channel.

Slightly surprisingly, the first and second scenarios get the same results, which means that the different status held by an hacker and an off-path ISP does not affect the risk under analysis. The final exploitability value in the root node can be easily lowered by encrypting the media channel, to the detriment of a wider bandwidth

Table 5
Propagation of dependencies

iteration	$E(V_1)$	$E(V_2)$	$E(V_3)$	$E(V_4)$	$E(V_5)$	$E(V_6)$	$E(V_7)$	$E(V_8)$	scenario
0	7	1	5	3	3	1	1	?	I
1	7	7	5	3	3	3	3	P 3	
2	7	7	5	3	3	3	3	P 3	
0	7	1	5	3	3	3	3	?	II
1	7	7	5	3	3	3	3	P 3	
2	7	7	5	3	3	3	3	P 3	
0	9	9	7	9	9	9	9	?	III
1	9	9	7	9	9	9	9	P 3	
2	9	9	7	9	9	9	9	P 3	
0	?	?	?	?	?	?	?	1	IV
1	?	?	?	?	?	?	?	63	
2	?	?	?	?	?	?	?	63	