VoIP SEAL: A Research Prototype for Protecting Voice-over-IP Networks and Users

Jan Seedorf, Nico d'Heureuse, Saverio Niccolini, Thilo Ewald
NEC Laboratories Europe
Kurfürstenanlage 36, 69115 Heidelberg
{firstname.lastname}@nw.neclab.eu

Abstract: While deployment of Voice-over-IP (*VoIP*) systems is gaining momentum, security remains a major challenge for this new technology. Due to the inherited flaws of IP-networks, lacking authentication mechanisms, and complexity in terminals (among others) many threats to VoIP systems exist. Thus, there is a need for protecting VoIP systems and its users against attacks. To address these challenges, we present a modular protection framework for VoIP networks and a corresponding prototype implementation. Our solution, VoIP SEAL, can protect against various kinds of VoIP threats at different stages of a call-setup process. Using a modular approach, different protection mechanism can be combined and interact with each other. This allows for flexible protection against diverse types of VoIP attacks. We describe our overall architecture as well as some recent enhancements in detail.

1 Introduction

Voice-over-IP (*VoIP*) deployment has increased in recent years. VoIP promises to save cost for providers and customers through convergence of voice and data networks. However, the technology introduces new security threats to the network. For instance, due to the inherent complexity, many VoIP terminals are vulnerable to attacks [WLS04]. Also, new threats like Spam over Internet Telephony (*SPIT*) are expected to emerge in the future. Motivated by these threats we have developed a framework to enable protection against different kinds of attacks on VoIP networks. In this paper we present our prototype implementation of this modular protection approach against different VoIP threats. Our solution, *VoIP SEAL (VoIP SEcure Application Layer Firewall)*, can detect unsolicited and malicious signalling messages at different stages of the call-setup process. Since it is advisable to block malicious or unwanted communication attempts as early as possible during session establishment, protection should start at the signalling level. The Session Initiation Protocol (SIP) [RSC+02] has emerged as the predominant signalling protocol for setting up and managing VoIP sessions. Our efforts and prototype implementation therefore concentrate on SIP, while in principle our solution could be applied to other signalling protocols (e.g., H.323) as well.

2 VoIP Security Threats

Voice-over-IP networks face a broad range of attacks against which protection is necessary. For instance, the availability of the infrastructure and terminals can be attacked. By sending a large amount of unsolicited messages or single malformed messages a VoIP server or terminal can become inoperative. With eavesdropping of packets or man-in-the-middle attacks conversations may be intercepted or modified. In addition, private information can be obtained by attackers (e.g., passwords or account-related information). Unauthorized or unaccountable resource utilisation can lead to impersonation and fraud (e.g., with replay attacks). Because most VoIP providers offer some services with cost (e.g., gateway services to the classical telephone network), billing attacks impose a severe threat to VoIP systems and its customers. Furthermore, social threats like Spam-over-Internet-Telephony (SPIT) can annoy users and directly harm user acceptance of the new technology. Because potentially a telephone rings with every attack, SPIT is much more obtrusive than e-mail spam.

3 A modular Protection Approach

Current VoIP systems and signalling standards like SIP offer only weak protection against the threats mentioned in the previous section [PS05]. Because of the variety of different VoIP attacks, there is no general solution against all kinds of threats. A key challenge therefore is to come up with a flexible architecture where single protection measures can be combined for overall protection. To address this challenge, we have developed a protection framework based on user interaction [SNTB06]. Figure 1 shows this 5-stage protection framework for VoIP signalling. This architecture enables a flexible way to protect against VoIP attacks because single protection modules at different stages can interact which each other. In the framework, protection modules can be integrated at different stages of the call-setup process:

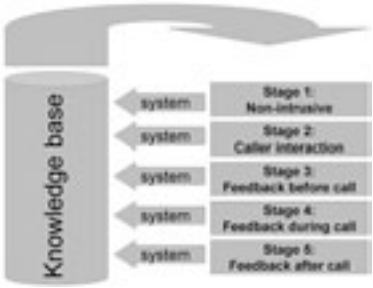

Figure 1: Protection Framework based on Caller Interaction

Stage 1: *Non-intrusive* methods can filter SIP messages without any user-interaction. Examples for such protection modules are: blacklists (calls from identities on this list are automatically blocked), whitelists (identities on this list are explicitly allowed to communicate), statistical analysis, or pattern-based detection of malformed messages.

Stage 2: *Caller interaction* protects the callee of a call by testing the trustworthiness of the caller in an interactive way. An example for such a stage-2 module is a CAPTCHA-test, where the caller needs to prove that she/he is not a machine by solving some test (to protect against botnets and automatically generated unsolicited calls) [QNT+07].

Stage 3: Prior to accepting a call a user can be asked to provide *feedback* regarding the identity of the caller. This enables the callee to either accept or reject a call based on some information about the caller (*consent-based communication*).

Stage 4: With *feedback during the call* a user can indicate to the system that the current call is unsolicited (e.g., by pressing a special button on his/her phone). Such feedback is especially helpful to calibrate statistical detection modules.

Stage 5: A callee can also mark a call as unsolicited once the session has terminated to provide *feedback after the call*.

4 Prototype Implementation and Recent Enhancements

We implemented a prototype of the framework described in Section 3. The prototype currently focuses on stages 1 and 2 but also includes basic implementations for stages 4 and 5. Figure 2 shows how protection modules at different stages can be combined: In the first stage (*non-intrusive*) multiple modules can be deployed. Each of them rates the likelihood of an incoming message to be unsolicited and contributes to an overall score s. Depending on this score – and on the system's configuration – a message is further processed: a) If the score is above a threshold θ_h, i.e., $s > \theta_h$, the message is considered malicious and either blocked with an error message or forwarded to a mailbox system for logging, b) if s is below a threshold θ_l the message is considered non-malicious and directly forwarded to the next signalling hop, c) if $\theta_l \leq s \leq \theta_h$ the message is forwarded to the second stage for further tests (*caller interaction*). Feedback from stages 3-5 calibrate some of the protection modules in stage 1.

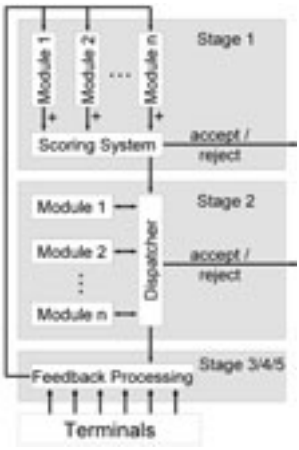

Figure 2: Modular Protection Architecture

Protection Modules with Caller Interaction Currently, three different protection modules which interact with the caller (stage 2) have been implemented:

Voice analysis: This test analyses the voice energy level of the incoming media stream. The test fails if the caller's behavior does not fit into a predefined pattern, e.g., the caller is not silent while a short welcoming message is being played, or the voice energy level does not increase for a short time after the caller is asked to tell his name [QNT+07].

DTMF test / Simple IQ test: In this test, the caller is asked to enter some digits (e.g., "Please enter 1, 2, 3", or "Please enter the sum of 2 and 3").

Greylisting: Greylisting (known from the email world) has been adapted for VoIP: the caller is asked to call again within a certain time interval. The test is passed only if the caller calls back within the specified interval.

Figure 3: Selection of Stage 2 Tests on Hardphone User Interface

Figure 4: Black-/Whitelist Administration on Hardphone User Interface

Personalisation From the caller's point-of-view, each stage 2 tests creates a different level of inconvenience. While the voice analysis test imposes the lowest inconvenience for the caller, greylisting might even cause a level of inconvenience which is not acceptable. To address this issue, we have implemented a per-callee personalisation which allows users to specify preferences for the tests applied for incoming calls. The preferences can depend on a multitude of variables, such as company guidelines or even the time of day. Furthermore, also the settings of stage 1 modules are configurable by users, e.g., by allowing users to define personal black- and whitelists. In our prototype a simple web-interface allows users to edit their personal black- and whitelists and select a preferred stage 2 test. The same functionality has also been implemented using the mini-browser functionality of a Snom 360 SIP hard phone (see figure 3 and 4). This allows the user to directly configure his/her personal protection profile directly from the phone.

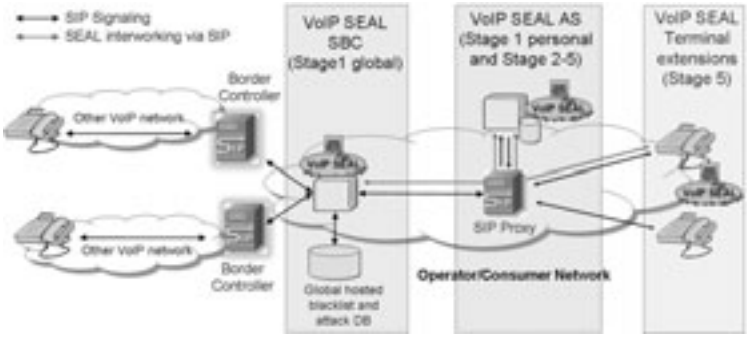

Figure 5: Distributed VoIP Security Detection

Distributed Detection Scheme Figure 5 shows a possible deployment scenario of VoIP SEAL. In this scenario VoIP SEAL is located at two different points of the network: one instance is deployed at the border of the network (SEAL-SBC[1]), a second instance is deployed next to the SIP proxy and acts as an Application Server (AS). The two instances cooperate with each other in order to provide a scalable, flexible, and personalised protec-

[1]The name "SEAL-SBC" was chosen to indicate that this instance of VoIP SEAL is placed before, next to, or on top of a Session Border Controller (SBC); the SEAL-SBC itself does *not* implement the full feature set of an SBC.

tion of the SIP infrastructure. Each instance can implement one or multiple stages of the protection framework. In the current prototype implementation, each VoIP SEAL instance adds its results, i.e., the scores it calculated, as additional SIP headers to the incoming SIP messages. These headers are then evaluated by subsequent VoIP SEAL instances. The prototype uses proprietary SIP headers for transmitting the scores from one instance to the next. Currently, there is no standard for transmitting security scores between SIP entities. However, there is ongoing work in that direction in the IETF [WNTS07]. In the current setting, SEAL-SBC is used for detection of attacks on the network infrastructure (e.g., (D)DoS attacks), as well as for basic filtering and signaling analysis (stage 1, e.g., global black-/whitelists). The second VoIP SEAL instance, the AS, allows the application of advanced tests. If multiple ASs are deployed, load balancing techniques can be used to enhance scalability. The AS applies personalised stage 1 tests (e.g., personal black- and whitelists) as well as resource consuming stage 2 tests (with caller interaction).

5 Conclusion

In this paper, we have presented a prototype implementation for protecting VoIP networks and its users. We have described our overall architecture, based on a modular protection framework, and discussed recently added protection modules. Currently, VoIP SEAL is being evaluated at several large European telecommunication providers to test its protection modules and performance. In the future, we plan to integrate distributed monitoring and automatic signature generation for detecting malicious messages into our prototype.

References

[PS05] J. Posegga and J. Seedorf. Voice over IP - Unsafe at Any Bandwidth? In *Proc. Eurescom Summit 2005 - Ubiquitous Services and Applications*, pages 305–314, Heidelberg, April 27-29, 2005. VDE Verlag.

[QNT+07] J. Quittek, S. Niccolini, S. Tartarelli, M. Stiemerling, M. Brunner, and T. Ewald. Detecting SPIT Calls by Checking Human Communication Patterns. In *Proc. IEEE International Conference on Communications ICC '07*, pages 1979–1984, 2007.

[RSC+02] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. Johnston, J. Peterson, R. Sparks, M. Handley, and E. Schooler. SIP: Session initiation protocol, RFC 3261, 2002.

[SNTB06] R. Schlegel, S. Niccolini, S. Tartarelli, and M. Brunner. ISE03-2: SPam over Internet Telephony (SPIT) Prevention Framework. In Saverio Niccolini, editor, *Proc. IEEE Global Telecommunications Conference GLOBECOM '06*, pages 1–6, 2006.

[WLS04] C Wieser, M. Laakso, and H. Schulzrinne. SIP Robustness Testing for Large-Scale Use. In S. Beydeda, V. Gruhn, J. Mayer, R. Reussner, and F. Schweiggert, editors, *SOQUA/TECOS*, volume 58 of *LNI*, pages 165–178. GI, 2004.

[WNTS07] D. Wing, S. Niccolini, H. Tschofenig, and M. Stiemerling. Spam Score for SIP, internet draft (work in progress), 2007.