Analysis of threats on a VoIP Based PBX Honeypot

N. McInnes, E.J. Zaluska, G. Wills
School of Electronics and Computer Science
Faculty of Engineering and Physical Sciences
University of Southampton
SO17 1BJ
United Kingdom

Abstract— **Many organisations are moving over from legacy telecommunications to Voice over IP (VoIP), enabling greater flexibility, resilience and an overall cost reduction. Session Initiated Protocol (SIP) is now considered to be the main VoIP protocol in the business–to-business market, but the correct implementation and configuration is not always well-understood. The failure to configure SIP systems correctly has led to significant fraud exploiting a range of vulnerabilities and billions of dollars every year being stolen from companies of all sizes through PBX Hacking via the medium of Toll Fraud. Previous research into this area is now dated but suggests a fast-changing approach by the attackers. Industry organisations such as the Communications Fraud Control Association (CFCA) acknowledge that this is a fast-growing problem. To quantify the size of the current problem, a Honeypot experiment was undertaken using a popular phone system used by businesses. The Honeypot ran for 10 days and recorded just under 19 million SIP messages. This research has identified the rate of attack is approximately 30 times more aggressive than previous reported research.**

Keywords-component; VoIP; SIP; Toll Fraud; PBX Hacking; Honeypot; Misuse; Attack

I. Introduction

Voice over Internet Protocol (VoIP) is a group of protocols that enables voice communication over an IP network.

Current telephone networks are in the process of moving from a closed legacy approach in network design (traditionally copper cable infrastructure using Signaling System No. 7 (SS7) technologies) to an open approach based on other mediums (for instance IP). This transition enables telephone costs to be significantly reduced.

However, as the market for VoIP paid calls grows, it also increases the possibility of misuse and fraud. Billions of dollars a year is being lost through various activities such as toll fraud, where calls are made to phone numbers (through a hacked component) where the calling party receives revenue from such calls at the expense of the victim – completely unknown and undetected until they receive their telephone bill [1].

Annual fraud figures published by the Communications Fraud Control Association (CFCA) suggest that around $38 billion are lost to fraud [2] and PBX hacking itself has increased by over 60% resulting in losses of over $7 billion in 2015 [3].

Previous research using VoIP Honeynets to assess the impact of attackers is now over 5 years old. One of the results of these early studies was that the attackers changed their behavior over time [4]. This paper discusses the methodology and initial results of a limited Honeypot experiment that ran for just 10 days to provide an up-to-date view on the attack methods currently being used by attackers.

II. Background Literature

Ahmed et al. states many consumer based services are value-added and peer-to-peer in nature, often using their own proprietorial protocols [5] (Skype, Viber and Whatsapp are examples). In these cases, protocols are closed, custom built and have dedicated mobile phone applications.

In comparison, business services typically focus on interoperability, because most businesses want to use existing equipment, which requires the use of a common protocol. As suggested by Abdelnur et al, SIP is the-IETF recommended protocol for VoIP. [6].

A. SIP

SIP is the industry recommended VoIP Protocol for interconnectivity and an open standard (IETF RFC 3261) [7]. This has enabled a number of manufactures, namely Cisco, Avaya, Yealink and software vendors such as Asterisk, Freeswitch to build products and platforms that work together.

SIP works by splitting a phone call into signalling and media elements [8]. The signalling, which carries messages containing the initial call message (called an 'Invite') with the call details. This is usually on User Diagram Protocol (UDP) port 5060 [7]. The audio channels are set up on 2 random UDP ports (for example in the case of Asterisk these ports are between 10,000-20,000 for bi-directional audio).

SIP uses Uniform Resource Identifiers (URI) adopting the following format: SIP/Extension@IP [9]. This approach is similar to other web protocols, for instance a web URL or File Transfer Protocol (FTP) address.

SIP is regularly used in one of 2 modes depending on the context it is to be used within: IP Authentication or Registration (to be able to identify a user). IP Authentication is regularly used for SIP trunking (equivalent to a trunk line) between a provider’s switch and a customer’s Private Branch Exchange (PBX), where registration is most often used to register a handset to a PBX.

IP authentication requires a PBX to have a dedicated public IP where the PBX is public facing. For the scope of this paper, it is usually at this interface that attacks from the Internet will occur.

B. Private Branch Exchange (PBX)

A PBX is a phone system that is usually located inside a business’s office and has desk phones connected to it. PBXs originally had complex command line interfaces, however as web technologies have improved more PBXs have entered the market with easy-to-use Graphical User Interfaces (GUI).

User devices can regularly change IP, therefore a device needs to periodically register with a SIP server. This can be to a PBX to reconfirm where to send invites (for receiving or making a call) [10]. Registration is performed by username and password authentication and is used between a PBX and a device (e.g. handset), although both types can be used in other use cases (i.e. a trunk can also be username/password based). Fig. 1 shows a normal use case of how the SIP protocol is used in business communications.

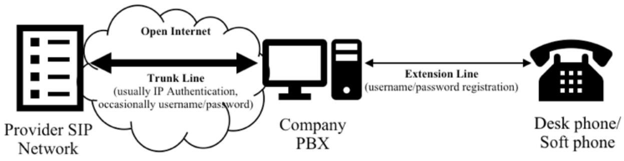
Open Internet
Trunk Line
(usually IP Authentication, occasionally username/password)
Provider SIP Network
Company PBX
Extension Line
(username/password registration)
Desk phone/ Soft phone

Figure 1. Normal Business SIP Use Case

Many workers work remotely (home workers). This requires a firewall to accept incoming traffic from the general internet. This introduces new sets of challenges that PBX owners need to consider.

C. VoIP Vulnerabilities and Types of Attacks

Many potent VoIP vulnerabilities have been identified. Rebahi et al. studied more than 220 different vulnerabilities [11]. Proprietorial software, services and open source all appear to have their own vulnerabilities.

Skype users have reported account breaches via a login vulnerability (weak credentials) or SIP not encrypting the username, password, registration details, messages or media.

In contrast, Sengar suggests that most SIP vulnerabilities are not due to the weaknesses in the protocol itself, but result from poor SIP credentials and misconfigured systems [12]. Other researchers such as Hoffstadt et al. [13] and Gruber et al. [14] support this assessment.

D. PBX Penetration Studies

Hoffstadt et al. [13] from the University of Duisburg, Essen, Germany have become leading researchers in VoIP Attacks for fraudulent purposes. Their paper presented at the IEEE 2012 Conference on Trust, Security and Privacy in Computing and Communications began with them building a Honeynet (a collection of Honeypots) setup as an Intrusion Detection System (IDS), located in Germany and the USA.

This Honeynet collected over 47.5 million messages over approximately a 2-year period. Hoffstadt et al. used a novel method to collect not only messages contacting their systems, but also the entire subnet (monitoring the Level 3 Switch). On building the Honeynet, Hoffstadt et al. analysed previous work and established low level interaction, Honeypots have weaknesses by not being able to provide full overview and only enable basic "fingerprinting". Hoffstadt et al. identified that considerable amounts of data will be available and could use Packet Capture (PCAP) and UDP Sockets for SIP Traffic analysis. They did this by building 2 networks, one with SIP components, the other without.

Hoffstadt et al. [13] discovered that to determine if a device is SIP enabled, attackers would send out Option messages to probe whether a device is SIP enabled or not (where the device would reply if it was) [13]. An Option is a message sent out to a SIP server which replies with a list of features it supports. This led Hoffstadt et al. categorising an attack into 4 stages:

1. Initial SIP Server Scan – Scan IP with Option messages looking for replies
2. Extension Scan – Scan for extensions looking at differences in error messages (404 not found, 403 Forbidden, 401 Unauthorised)
3. Extension Hijacking – Using dictionary attacks on extensions
4. Toll Fraud – Making successful calls.

In Essen’s analysis of the signalling details, Hoffstadt et al. noted that tools such as SIPVicious were being used to automate an attack. This confirms claim by Ronniger et al. that there are tools available to attack VoIP Infrastructures [15].

Hoffstadt et al. established that once a non-SIP component was open to the public internet, it appeared to be continually under Option attacks [13].

In contrast, when a SIP component replied, little to no Options were further received, but moved onto stage 2. Stage 3 of the attack used typically 10,000 various usernames with different password combinations (over 55,000 attempts) which took a little over a minute to complete. On successfully registering (stage 4) it was observed that various prefixes (numbers) were used to dial out (i.e. 011 for an international line the US and 00 for an international line in Germany) [13].

Although most attacks were automated when scanning and brute forcing, the author discovered that stage 4 would happen several months after successfully registering with an extension [13]. The victim may have difficulties in researching the hack as evidence may already have been destroyed by natural log cycles in order to save storage space.

In further papers, Hoffstadt et al. extend their work to introduce logic. This allowed them to dynamically create extensions where that extension was being probed by giving the impression the extension is valid.

Furthermore, a system was introduced to answer calls for random periods to simulate a call. This enabled Hoffstadt et al. to follow attackers from stage 1, where multiple IP addresses may be involved [16]. Later on, Hoffstadt et al. created a Generic Attack Replay Tool (GART) allowing the replaying of attacks by capturing key information to assist in building tools that can detect and prevent attacks at a later date [17].

Repeated studies found that once an attacker has gained access, they attempt to call premium-rate numbers or high-cost numbers in various countries, suggesting attackers earn money for doing so [13,14]. Gambia, Palestine and Somalia appear to be regularly attempted. Gruber et al. goes further suggesting that most calls go to African countries [4].

Since 2011, researchers at Vienna University of Technology have also been running a Honeynet. There findings coincide with Essen's in that many calls were to African countries (Ethiopia and Egypt) and when attempts are made, most of the time the attempt is made from an Egyptian IP address[14].

Researchers at the University of Duisburg partnered with Researchers at the Network Systems Group to develop novel methods for implementing and improving monitoring nodes around distributed Honeynets with monitoring points in China, Norway and Germany. This reconfirmed their initial conclusions that attackers scan large segments of the internet. Moreover, they determined not all attackers were involved in the different stages of attacks, concluding that attackers share information about potential victims with other attackers [18].

III. Methodology

Researchers at Essen and Vienna, both developed Honeynets (the collection of Honeypots) to monitor SIP messaging at different geographical locations. These studies did not facilitate the use of a PBX, but only certain software engines used by PBX software.

At the SIP Signalling level, messages can contain metadata regarding the software configurations and versions to replies in Invites, Options and Registration requests.

To meet the requirements of having a high interactive VoIP engine, the decision was taken to use a popular open source PBX software package known as FreePBX.

Freepbx is a feature rich platform. which claims over 1 million installs. This means that, if hackers are actively seeking out VoIP Systems to attack, then this would be a good candidate because of its wide installation base.

As FreePBX is lightweight in nature, a Virtual Machine is sufficient to run the Honeypot which is located in the United Kingdom. In this experiment, the Honeypot would reply with the Server type as FPBX-VERSION where VERSION is the version of FreePBX being used in the SIP Messaging. FreePBX is built on top of the Open Source Asterisk Engine.

A. PBX Configuration

For ease of setup, monitoring and collecting data, it is important that the PBX is setup in a way which maximises opportunities to collect and study how a PBX hack attempt has occurred while minimising and reducing risk.

The overall goal is to study a hack attempt, not to allow the system itself to get infected by malware or similar. Certain flows will be allowed to allow toll fraud, but overall the system is to be locked down and no actual calls can be made as it is not connected to a phone provider.

The PBX is installed via the ISO disk image provided by FreePBX which allows a straightforward installation. The ISO is built on the Linux Centos 7 Operating System.

FreePBX comes with several built-in security features such as Fail2ban which will block an IP for a period of time if the incorrect credentials are provided. However, for the purpose of this experiment to assist in obtaining a full understanding of the scale of the current problem, Fail2ban will be disabled. (Leaving Fail2ban enabled could limit the number of attacks on the PBX.)

To secure all ports, the Virtual Private Service provider provides a free virtual firewall which provides the ability to block all connections except a select few TCP and UDP ports.

For the purpose of this experiment the following ports were monitored:

- TCP: 5060-5070
- UDP: 5060-5070, 10,000-20,000

As seen by previous researchers, it is expected that attackers would scan extensions looking at the replies to determine if an extension exists. Therefore, to trigger different responses, various extensions have been created to provide an illusion of a live production system (reinforcing meta data replies to SIP requests). These extensions can be seen in Table 1.

To record SIP interactions on the PBX, different methods could be used. SIP interactions could be recorded in the Asterisk log or via a Packet Analysis. Packet Analysis is preferred over the Asterisk Log as it allows for detailed inspection of all events that occur over the network interface of the system which may not be recorded via the Asterisk Log.

Using a utility such as Wireshark, knowledge can be gained to determine whether previous findings in previous studies are still valid and what new techniques are being used.

TABLE I. SAMPLE EXTENSIONS TO BE CREATED ON HONEYPOT

Username (Extension)	*Password*
1001	fdfAS243%32
1002	1002
1003	1003
1125486	Dgfg35DGS24g
10000	10000
50000	50000
100000	100000
5001	5001
5003	dfdfSDG3435s

In addition to the above, using the Virtual Machine Provider's Tools, Bandwidth Data and Machine Resources will be monitored to see what impact if any an attack has on the VM.

The specification of the virtual machine was 1 x 2.4Ghz Intel virtual core with 1024gb of RAM and SSD storage.

IV. Results

The data analysed was collected over a 10-day period (24th September 00:00 BST – 3rd October 23:59 BST 2018). During this period, just under 19 million SIP messages were received from malicious actors. These were split into the following SIP Message Types as seen in Table 2. For the purpose of this experiment, Register, Invite and Option has the meaning set out by IETF RFC 3261 [7].

TABLE II. DAILY BREAKDOWN OF SIP MESSAGE TYPES RECEIVED

	SIP Message Type Received		
Date	***Register***	***Invite***	***Option***
24/09/2018	1,494,872	1,488	78
25/09/2018	45,247	1,667	91
26/09/2018	2,014	2,266	84
27/09/2018	478,208	1,153	66
28/09/2018	12,037	1,636	121
29/09/2018	3,030,372	1,667	114
30/09/2018	2,770,527	2,774	91
01/10/2018	1,914,163	315	34
02/10/2018	1,921,432	34,778	102
03/10/2018	7,204,257	10,471	95
Total	18,873,129	58,215	876

The IP of the Honeypot has not been publicly advertised and real devices are not connected to the SIP system. Therefore, it can be presumed that all inward SIP traffic is malicious.

During the period the Honeypot ran, no successful registration was made to the server. When analysing Packet analysis to investigate a selection of credential combinations, some attempts were with the same detail combination. The attack on the 3rd October were all with the same credential combination.

A. System Resources

During the 10-day period the Honeypot ran, approximately 20GB of inward SIP traffic was recorded which averaged 2GB per day.

The highest average bandwidth utilisation recorded by the Virtual Machine provider was 600Kbps. The highest average CPU utilisation recorded by the Virtual Machine was 30% and the average SIP dictionary registration attack was for 12 hours (excluding 3rd October). During these times the above average results of bandwidth and CPU utilisation were observed during the period.

B. User Agents

During the 10 days the Honeypot ran, various user agents appear to have been used by third parties to send messages to the system were observed.

Some of these were a combination of random letters and numbers, others were known user agents, desk phones or phone systems. The breakdown based on SIP Message type can be seen in table 3.

TABLE III. USER AGENTS DETECTED

Register	***Invite***	***Option***
• Vaxsipuseragent/ 3.1 • MGKsip release 1110 • VoIPSIP V11.0.0 • Eyebeam • FPBX	• Linksys-SPA924 • SIPCLI/V1.8 (some were V1.9) • Pplsip • voipxx • Various random characters: o zazann, o zxcvfdf11	• Friendly-scanner • Avaya • Cisco-sipgateway/ IOS-12.X • sipvicious

C. 3^{rd} October Registration Attack

The attack witness on the 3rd October 2018 was multiple times bigger than any other attack. Instead of the 12 hour attacks witnessed on other days, 3rd October had a continual 24 hour period attack where 7.2 million registration attempts were made. The user agent used in the SIP messaging appeared to be from MGKsip release 1110.

V. Discussion of Results

The date period within this experiment is 10 days. The experiment performed by Hoffstadt et al. (Essen) between late 2009 and early 2012 experienced 47.5 million SIP messages.

In contrast our experiment has experienced just under 19 million SIP messages in only 10 days compared to just over 2 years in the Essen experiment. This gives a mean average of 1.9 million messages per day. In 1 day alone (3rd October), over 7.2 million SIP messages were received which is over 15% of what Essen witnessed across multiple locations combined over a 2-year period.

Based on our results, if this experiment ran for the same time period to that of the Essen experiment, it is expected to result in over 1.4 billion SIP messages.

This represents a rate of attack 30 times more aggressive than in the Essen experiment. This view reinforces the trend observed by the CFCA that PBX hacking has increased significantly in recent years [3].

When a dictionary attack was being performed, the resources being consumed by the PBX were significant which would question whether a negative impact of service would be experienced by a business user or owner of a PBX that was being attacked. A large amount of upload bandwidth was consumed by a rate of approximately 600Kbps.

For businesses that do not have a leased line and rely on ADSL, this could have significant consequences by reducing the quality of calls due to bandwidth constraints imposed by technologies and providers.

In addition, 30% of the CPU was being constrained which could limit the call capacity of the PBX where both would have the symptoms of audio breaks in the calls, also known as choppy calls.

A final consequence on the bandwidth being used in these attacks are where businesses pay per GB of data transferred or have a monthly bandwidth quota. It is still common among business ADSL lines to have data quotas applied to a broadband line with additional fees applying for exceeding this threshold.

The large CPU resources being consumed by these attacks can be explained by the multiple processes that occur when a SIP message is received. Asterisk would receive a message and if a register or invite was received would have to perform an SQL lookup in a database and provide a reply while at the same time writing events to a log. Many PBX systems do not have an SSD, but a hard disk. If a basic PBX had to deal with this volume of connection attempts, then the PBX would most likely become overloaded or not be usable by a business due to IO constraints.

Unlike previous similar experiments, little to no Options were received. This could be because we have not specifically checked whether the research conclusion from Essen are still correct (i.e the PBX is continually being sent Option SIP messages until the PBX responds and after responding, few further Options will be received.)

As from previous studies from over 5 years ago, several user agents are still being used such as friendly-scanner and user agents with various random characters. Of interest though are user agents which give the impression of real, widey-used devices such as Avaya, Cisco-SIPGateway and Linksys-SPA942. This would suggest that attackers are either spoofing these user agent names or hardware has been compromised into some form of botnet. This hardware is usually expensive at the higher end of the market. The significant attack on the 3rd October appeared to originate from an attacker using MGKsip Release.

Unlike other documented experiments, an actual PBX was used, although the SIP process would remain the same, the metadata exchanged would contain information regarding the PBX which would give information about the software being used. This would lead an attacker to believe the system is genuine and a production system. This may partially explain why this experiment witnessed significantly more attacks.

Hoffstadt et al. concluded that not all attackers were involved in various stages of attack and potentially shared information with other attackers. Based on anecdotal evidence physical hardware appearing to be involved in attacks, this could reinforce the idea that data is shared but via a botnet of infected machines.

During the Honeypot period, no attack was able to successfully register to any of the PBX extensions. Although on most days, when an attack was under way, different credentials were being attempted. On the 3rd October, all 7.2 million attempts used the same details. This could be explained by a botnet which is automated and, in this case, misconfigured, but more concerningly due to the volume could be seen as a Denial of Service attack as the result it could have on a business.

VI. Conclusions and Future Work

The data gathered in this 10-day Honeypot experiment has yielded interesting findings and established that VoIP Attacks are not only still occurring, but are significantly greater in numbers (up to 30 times more aggressive than previous studies) and new software and possibly compromised hardware is being used to generate these attacks which could suggest a wider problem from a PBX Hacking botnet.

Previous documented user agents still occur in some occasions, but this is outweighed by new appearing user agents. In addition, attacks are of the size that could congest business broadband connections resulting in further indirect charges.

Further work could include investigating indirect ways (i.e non-SIP methods) that hackers use to infiltrate a PBX and using the data gathered to assist in development towards a filter that could be used to attempt to detect, limit and prevent PBX hacking.

References

[1] New York Times, " Phone Hackers Dial and Redial to Steal Billions" October 2014. [Online]. Available: https://www.nytimes.com/2014/10/20/technology/dial-and-redial-phone-hackers-stealing-billions-.html (Access Date: 1 October 2018)

[2] M. Sahin, A. Francillon, P. Gupta and M. Ahamad, "SoK: Fraud in Telephony Networks" *Proceedings - 2nd IEEE European Symposium on Security and Privacy, EuroS and P 2017,* pp. 235-240, 2017.

[3] Marketwired, "Argyle Data Recommendations From CFCA's 2015 Fraud Survey Analysis: Think Globally, Act Locally" March 2016. [Online]. Available: https://finance.yahoo.com/news/argyle-data-recommendations-cfcas-2015-100000320.html. (Access Date: 1 October 2018)

[4] M. Gruber, D. Hoffstadt, A. Aziz, F. Fankhauser, C. Schanes, E. Rathgeb and T. Grechenig, "Global VoIP security threats - Large scale validation based on independent honeynets" *Proceedings of 2015 14th IFIP Networking Conference, IFIP Networking 2015,* 2015.

[5] Ahmed, A. S. Shaon and R. Hasan, "Evaluation of popular VoIP services" ICAST 2009 - 2nd International Conference on Adaptive Science and Technology, pp. 53-63, 2009.

[6] A. Humberto, A. Tigran, R. Michael and R. State, "Abusing SIP authentication" *Proceedings - The 4th International Symposium on Information Assurance and Security, IAS 2008,* pp. 237-242, 2008.

[7] Internet Engineering Task Force, "SIP: Session Initiation Protocol" June 2002. [Online]. Available: https://www.ietf.org/rfc/rfc3261.txt. (Access Date: 1 October 2018)

[8] P. Park, Voice over IP Security, IN, USA: Cisco Press, 2008.

[9] T. Z. Jyh-Cheng Chen, IP-Based Next-Generation Wireless Networks: Systems, Architectures, and Protocols, New Jersey: John Wiley & Sons Inc, 2004.

[10] J. Davidson, J. F. Peters and M. Bhatia, Voice Over IP Fundamentals, Indianapolis: Cisco Press, 2006.

[11] Y. N. M. Rebahi, T. Magedanz and O. Festor, "A survey on fraud and service misuse in voice over IP (VoIP) networks" *Information Security Technical Report,* vol. 1, no. 1, pp. 12-19, 2011.

[12] H. Sengar, "VoIP Fraud : Identifying a Wolf in Sheep ' s Clothing" Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security - CCS '14, pp. 334-345, 2014.

[13] D. Hoffstadt, A. Marold and E. P. Rathgeb, "Analysis of SIP-based threats using a VoIP Honeynet System" Proc. of the 11th IEEE Int. Conference on Trust, Security and Privacy in Computing and Communications, TrustCom-2012 - 11th IEEE Int. Conference on Ubiquitous Computing and Communications, IUCC-2012, pp. 541-548, 2012.

[14] M. Gruber, C. Schanes, F. Fankhauser and T. Grechenig, "Voice calls for free: How the black market establishes free phone calls-Trapped and uncovered by a VoIP honeynet" *2013 11th Annual Conference on Privacy, Security and Trust, PST 2013,* pp. 205-212, 2013.

[15] M. Ronniger, F. Fankhauser, C. Schanes and G. Thomas, "A robust and flexible test environment for voip security tests" *010 International Conference for Internet Technology and Secured Transactions (ICITST),* 2010.

[16] D. Hoffstadt, N. Wolff, S. Monhof and E. Rathgeb, "Improved detection and correlation of multi-stage VoIP attack patterns by using a Dynamic Honeynet System" *IEEE International Conference on Communications,* pp. 1968-1973, 2013.

[17] A. Aziz, D. Hoffstadt, S. Ganz and E. Rathgeb, "Development and analysis of generic voip attack sequences based on analysis of real attack traffic" *Proceedings - 12th IEEE International Conference on Trust, Security and Privacy in Computing and Communications, TrustCom 2013,* pp. 675-682, 2013.

[18] A. Aziz, D. Hoffstadt, E. Rathgeb and T. Dreibholz, "A distributed infrastructure to analyse SIP attacks in the Internet" *2014 IFIP Networking Conference, IFIP Networking 2014,* 2014.