Anatomy of SIP Attacks

JOÃO M. CERON, KLAUS STEDING-JESSEN, AND CRISTINE HOEPERS

João Marcelo Ceron is a Security Analyst at CERT.br/NIC.br. He holds a master's degree from Federal University of Rio Grande do Sul, Brazil, where he worked with honeypots for botnet detection. He currently works with incident handling, and his research interests include honeypot data analysis.
ceron@cert.br

Klaus Steding-Jessen is CERT.br/NIC.br Technical Manager. He holds a PhD in applied computing from the Brazilian National Institute for Space Research, where he worked with honeypots for spam detection. His research area is the use of honeypots for detecting Internet infrastructure abuse by attackers and spammers. He is also one of the authors of the chkrootkit tool.
jessen@cert.br

Cristine Hoepers is CERT.br/NIC.br General Manager. She holds a PhD in applied computing from the Brazilian National Institute for Space Research in the area of honeypot data analysis. Her research interests include the use of honeypots for incident detection and trend analysis. She has spoken at several security conferences, including FIRST, APWG, MAAWG, LACNIC, and AusCERT. cristine@cert.br

In the past few years we have seen a steady increase in the popularity of VoIP (Voice over IP) services. Scans for SIP (Session Initiation Protocol [4]) servers have been reported for many years, and to gather more details about these activities we emulated SIP servers in a network of 50 low-interaction honeypots, and collected data about these attacks for 358 days. What will follow is a description of our observations and advice on how to prevent these attacks from being successful.

Tracking SIP Servers Abuse

For quite some time, the security community has been reporting an increase in scans for the SIP default port 5060/UDP, as well as some anecdotal evidence of other types of abuse. Similarly, at the CERT.br honeyTARG Honeynet Project [3] (a chapter of The Honeynet Project), port 5060/UDP was consistently among the top-10 targeted ports. Bearing that in mind, we have been tracking the abuse of SIP servers more closely since last year.

This project consists of 50 low-interaction honeypots, based on Honeyd [7], deployed in the Brazilian Internet space. In order to enable Honeyd to collect SIP attack information, we implemented a listener that emulates Asterisk Server [2] and allows the definition of which extensions are available, as well as their default responses and passwords. This software allows us to collect the initial stages of a SIP session, logging information such as attack origins and the phone numbers the attackers attempted to call. For privacy reasons, we chose not to record audio sessions, limiting the implementation only to the SIP signaling.

Figure 1 presents a SIP conversation fragment logged by our listener. There are two SIP methods: `REGISTER` and `INVITE`. The first part is a `REGISTER` request. This is used by a user agent (UA) for registering contact information, such as its current IP address. The second part illustrates the `INVITE` method, which is used to establish a media session between UAs. In this log, a UA places a call from the extension 100 to the external phone number "201*****274" (sanitized number). Additionally, the "user-agent" field shows that this UA has provided the identification string "X-Lite release 1006e stamp 34025", a common softphone.

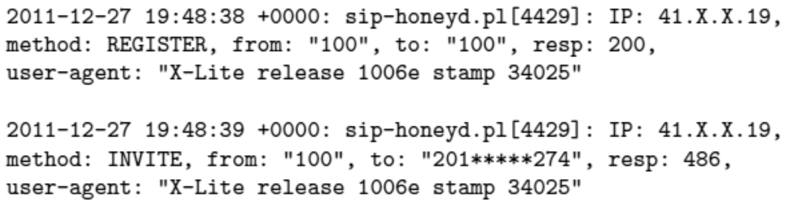

```
2011-12-27 19:48:38 +0000: sip-honeyd.pl[4429]: IP: 41.X.X.19,
method: REGISTER, from: "100", to: "100", resp: 200,
user-agent: "X-Lite release 1006e stamp 34025"

2011-12-27 19:48:39 +0000: sip-honeyd.pl[4429]: IP: 41.X.X.19,
method: INVITE, from: "100", to: "201*****274", resp: 486,
user-agent: "X-Lite release 1006e stamp 34025"
```

Figure 1: Honeypot log showing the attacker's IP, the phone number being requested, and the user agent identification string

Making Sense of the Data

The traffic targeted to port 5060/UDP in our honeypots was related to the following attack steps:

1. **Scanning**: searching for SIP servers.
2. **Enumeration**: once a SIP server is identified, the attackers try to enumerate the server configuration, available extensions, and so on.
3. **Brute force:** attackers try to access extensions that are protected with weak passwords.
4. **Abuse:** after gaining access to a PABX extension, the attackers will try to call external PSTN (Public Switched Telephone Network) numbers, usually to place international calls.

In a preliminary analysis of the collected data, we were able to identify that the attackers would try to call a given number by using several prefixes to increase the attack success (see Figure 2). This occurs because a SIP server can be configured in different ways—for example "0" or "9" to access PSTN lines. In some countries, such as Brazil, one must also specify the telecommunication operator to be used for long distance calls.

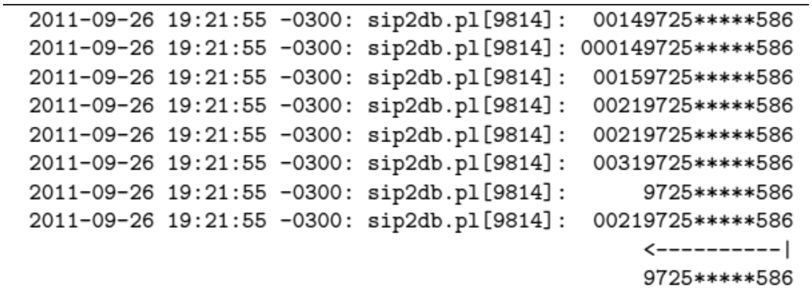

```
2011-09-26 19:21:55 -0300: sip2db.pl[9814]:  00149725*****586
2011-09-26 19:21:55 -0300: sip2db.pl[9814]: 000149725*****586
2011-09-26 19:21:55 -0300: sip2db.pl[9814]:  00159725*****586
2011-09-26 19:21:55 -0300: sip2db.pl[9814]:  00219725*****586
2011-09-26 19:21:55 -0300: sip2db.pl[9814]:  00219725*****586
2011-09-26 19:21:55 -0300: sip2db.pl[9814]:  00319725*****586
2011-09-26 19:21:55 -0300: sip2db.pl[9814]:      9725*****586
2011-09-26 19:21:55 -0300: sip2db.pl[9814]:  00219725*****586
                                                 <----------|
                                                 9725*****586
```

Figure 2: A phone number requested in different ways in order to identify the correct prefix for dialing long distance or international calls

Figure 2 illustrates the many variations one attacker was using for the phone number "9725*****586". To deal with this situation, which we called redundancy, we implemented a heuristic to identify it and to store only a unified SIP session related to this number in the database. Besides reducing the size of the database, this heuristic also helped us to identify the phone number's country code and to correlate calls placed at different times and coming from different sources, to a unique phone number.

Table 1 summarizes the data that reached our honeypot infrastructure from September 2011 to September 2012.

Data	Count
`REGISTER` messages	64,249,923
`INVITE` messages	1,007,697
Unified `INVITE` messages	153,773
Unique IPs	7,752
Unique Autonomous System Numbers - ASNs	858
Total number of days	358
Unique source country codes - CCs	83

Table 1: Summary of the data collected from September 2011 to September 2012

The majority of the `REGISTER` messages are from automated scans. Most of them have the signature of the SipVicious toolkit [5], a collection of tools for auditing SIP-based VoIP systems. The `INVITE` messages are actual abuse attempts directed to the listeners, i.e., phone call attempts. The unified `INVITE` messages are the `INVITE` messages after redundancies were identified. Note that the number of unique ASNs and CCs demonstrate a high dispersion of the origin of the attacks.

In the following sections, we will focus on the analysis of the unified `INVITE` messages, including the phone numbers that were called the most and the abuse sources.

User Agents and IDS Evasion

An important piece of information logged is the user agent identification string provided by the SIP clients that connected to the honeypots and tried to place a call. The most frequent user agents provided are presented in Figure 3.

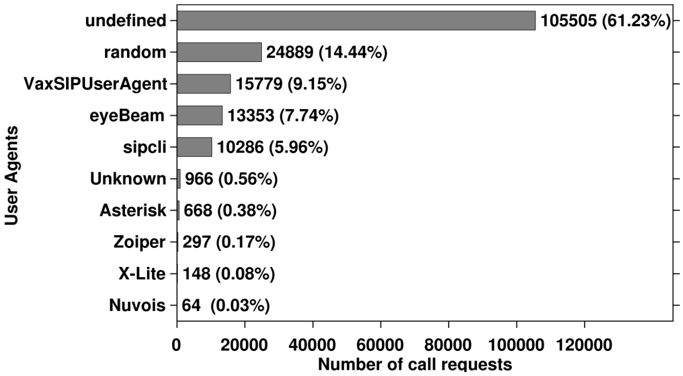

Figure 3: Top user agent identification strings provided by the SIP clients that tried to place calls

Note that 61.23% of the connections came from SIP clients that didn't provide any user agent string. We grouped all these clients under an identification string we called "undefined." This behavior is not expected from SIP clients and may suggest that many attackers are using customized tools to abuse SIP servers. Also note that the third most frequent user agent is the "VaxSIPUserAgent," which is used by a software development kit, also suggesting customized tools.

Additionally, there was a group of attempts where the user agent almost never repeated. In every new session, the client provided a random 20-character user agent, as shown in Table 2. This behavior was the second most frequent and was observed even in sequential requests coming from the same IP address. Our best guess is that this is being used to hide attack fingerprints or to evade IDS detection.

Timestamp	IP	User Agent String
2012-01-23T04:02:15Z	194.X.X.131	DmQCAsNRKZYayfosaXES
2012-01-23T04:02:17Z	194.X.X.131	yy3BHtWnCBPco3knmRqG
2012-01-23T04:02:19Z	194.X.X.131	KdUhQNVVxaZYfHg0rXFD
2012-01-23T04:02:21Z	194.X.X.131	otYvAff8mpZviS2CfF6M
2012-01-23T04:02:23Z	194.X.X.131	5y5ttWMXPbFIeyHb4l4D
2012-01-23T04:02:25Z	194.X.X.131	YDjb3Q8Wiw6442YCXMnE

Table 2: Examples of random user agent identification strings captured by the honeypots

We have also observed user agents commonly used by SIP servers, such as Asterisk. These user agents could be fake (set by the attacker) or could represent compromised SIP servers used to abuse other servers. The remaining user agents presented in Figure 3 refer to popular softphones.

As we can see, almost 85% of all connections came from customized or potentially malicious software.

Where Is It Coming From?

When looking into the source of the abuse attempts, we can try to identify specific patterns in the geographical origin and try to identify other characteristics that could give some insights about possible motivations.

Based on the source IP addresses of the attempted calls, we were able to estimate the source country code (CC) for the attacks. The country code allocation is based on information provided by the Regional Internet Registries (RIRs). Figure 4 shows these top CCs.

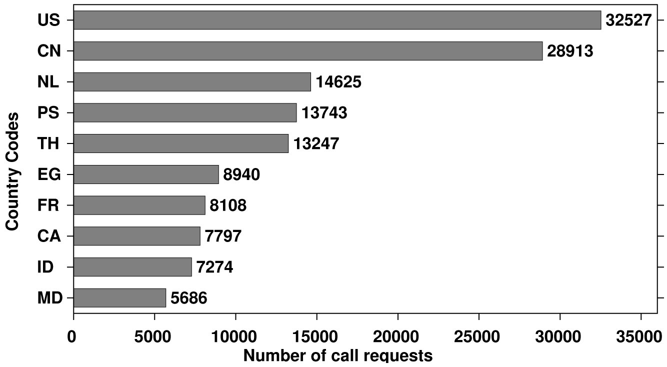

Figure 4: Top country codes of call requests (based on source IPs), aggregated by requested phone numbers

In Table 3 we list the top 10 IP addresses. For each IP, there is also information about how many different IDD (International Direct Dialing) codes it tried to call and the user agent string provided.

#	Count	IP	CC	IDDs	User Agent String
01	19,562	113.X.X.205	CN	142	undefined
02	11,027	83.X.X.16	NL	62	undefined
03	7,553	71.X.X.9	US	67	VaxSIPUserAgent/3.0
04	7,486	50.X.X.99	US	27	undefined
05	6,412	49.X.X.93	TH	38	undefined
06	5,647	85.X.X.212	FR	2	undefined
07	5,343	122.X.X.83	TH	37	undefined
08	4,672	24.X.X.37	CA	6	undefined
09	4,640	194.X.X.36	MD	48	random
10	4,234	202.X.X.204	ID	27	undefined

Table 3: Top source IPs, country codes, number of countries called, and the user agent provided

The most frequent CCs observed are US and CN, which are also the ones for three of the unique IPs that tried to place most of the calls. Note that the top 10 IPs were responsible for 44% of all call attempts. Additionally, the first IP address is responsible for 67% of all attempts coming from Chinese IPs. Likewise, the third and fourth IPs were responsible for 47% of all call attempts coming from IPs allocated to the US.

Another interesting fact is that the user agents provided by the top source IPs are not those of popular softphones but, instead, are possibly from customized attack tools. And, most interestingly, all the user agents provided by the ninth IP were 20-character random strings, as discussed in the previous section.

This combination of few IPs with distinctive user agents points to the possibility of these being rogue VoIP servers or proxies used as hubs to place phone calls.

Considering that one of the expected behaviors of a rogue VoIP server is high geographic dispersion of the destination phone calls, we tried to corroborate this hypothesis with additional analysis. We used AfterGlow [1] to explore the relationship among the top IP (113.X.X.205) and the destination of all calls. Figure 5 presents this IP address and the country codes it attempted to call. The CC was determined using the Perl library Number::Phone::Country, that associates an IDD to a country code. We can see that this IP, a possible VoIP server, places calls to 142 different countries.

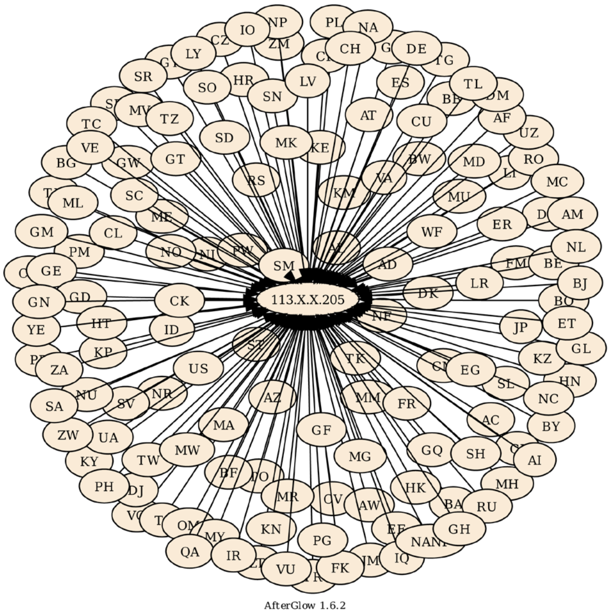

Figure 5: Destination country code for all calls placed by the IP 113.X.X.205

Who They Are Calling, and Why…

To gain more understanding of the abuse, we have also studied the nature of the phone numbers the attackers attempted to call. The most requested phone numbers fall into the following categories:

- **Cell phones:** identified by the number prefix
- **Financial services:** customer services from financial institutions (mainly Bank of America and Citibank)
- **Pre-paid phones:** pre-paid card services for international phone calls (Net2Phone)
- **Customer relations:** e-commerce customer relation services

The most called phone number was Bank of America's Credit Card Customer Service, totaling 5,090 attempts. Only seven IPs requested this phone number, and four of them have the rogue VoIP server behavior that was discussed earlier. Actually, 64% of all calls to Bank of America's Credit Card Customer Service came from the eighth IP listed in Table 3.

Source IP	Destination IDD	Count	(%)
PS	IL	7305	4.23%
EG	EG	6138	3.56%
MD	CZ	5559	3.22%
US	CZ	4535	2.63%
FR	RU	4264	2.47%
CA	800 (Free)	3296	1.91%
US	IL	2088	1.21%
US	ZW	1904	1.10%
CA	CZ	1903	1.10%
DE	CZ	1749	1.01%

Table 4: Most frequent combinations of source IPs and destination IDD country codes

The IP addresses that originated the calls were, for the most part, not the same as the IDD destination country. Table 4 shows the most frequent pairs, consisting of IPs allocated to a country code that are calling numbers in a given destination IDD country code.

Based on the data analyzed so far, we can present some hypotheses about the attackers' motivations:

1. Abusing SIP servers in order to place free phone calls or to gain anonymity;
2. Abusing the premium-rate telephone numbers business model;
3. Reselling VoIP services by abusing poorly configured SIP servers; and
4. Validating personal identifiable information, such as credit cards and bank account details.

Securing Your SIP Server

The types of activities observed reinforce the importance of implementing the current SIP security best practices [6]. Most attacks would have been prevented or mitigated by following one or more of these recommendations:

- **Protect the SIP server from the Internet:** be more restrictive in terms of which extensions can be reached from external IP addresses.
- **Use strong passwords:** use long, hard-to-guess passwords. Most SIP clients require the password to be entered only once, so there is no need to create easy-to-remember passwords. The current recommendation is to use at least 12-character passwords, including numbers, symbols, and lower and uppercase letters.
- **Create usernames different from extensions:** most brute force attempts try usernames that match the extension numbers.
- **Monitor the SIP use in your organization:** monitor your SIP server logs for abuse attempts, but also keep an eye on your PSTN billing information, looking for unusual long distance and international calls.

Conclusion

As the adoption of SIP services grows, being aware of the characteristics of the abuse against them is increasingly important. As our analysis showed, almost 85% of all call requests came from customized or potentially malicious software, and some of the calls may be related to unlawful activities. Also, because there are attackers currently taking advantage of poorly configured servers, the need to increase monitoring is clear.

The good news is that the implementation of basic VoIP security best practices will prevent most of the attacks seen in the wild.

References

[1] AfterGlow: Link Graph Visualization: http://afterglow.sourceforge.net/.

[2] Asterisk: The Open Source Telephony Projects: http://www.asterisk.org/.

[3] CERT.br: honeyTARG Honeynet Project: http://honeytarg.cert.br/.

[4] RFC 3261—SIP: Session Initiation Protocol: http://www.ietf.org/rfc/rfc3261.txt.

[5] SIPVicious—Tools for auditing SIP-based VoIP systems: http://blog.sipvicious.org/.

[6] John Todd, "Seven Steps to Better SIP Security with Asterisk," 2009: http://blogs.digium.com/2009/03/28/sip-security/.

[7] Niels Provos and Thorsten Holz, *Virtual Honeypots: From Botnet Tracking to Intrusion Detection,* Addison-Wesley, 2008.