



Adding salt to pepper a structured security assessment over a humanoid robot A Structured Security Assessment over a Humanoid Robot

Giaretta, Alberto; De Donno, Michele; Dragoni, Nicola

Published in:

Proceedings of 13th International Conference on Availability, Reliability and Security,

Link to article, DOI:

[10.1145/3230833.3232807](https://doi.org/10.1145/3230833.3232807)

Publication date:

2018

Document Version

Peer reviewed version

[Link back to DTU Orbit](#)

Citation (APA):

Giaretta, A., De Donno, M., & Dragoni, N. (2018). Adding salt to pepper a structured security assessment over a humanoid robot: A Structured Security Assessment over a Humanoid Robot. In *Proceedings of 13th International Conference on Availability, Reliability and Security*, Article 3232807 Association for Computing Machinery. <https://doi.org/10.1145/3230833.3232807>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Adding Salt to Pepper

A Structured Security Assessment over a Humanoid Robot

Alberto Giaretta
Centre for Applied Autonomous
Sensor Systems (AASS)
Örebro University
Sweden
alberto.giaretta@oru.se

Michele De Donno
DTU Compute
Technical University of Denmark
Denmark
mido@dtu.dk

Nicola Dragoni
DTU Compute
Technical University of Denmark
Denmark
Centre for Applied Autonomous
Sensor Systems (AASS)
Örebro University
Sweden
ndra@dtu.dk

ABSTRACT

The rise of connectivity, digitalization, robotics, and artificial intelligence (AI) is rapidly changing our society and shaping its future development. During this technological and societal revolution, security has been persistently neglected, yet a hacked robot can act as an insider threat in organizations, industries, public spaces, and private homes. In this paper, we perform a structured security assessment of Pepper, a commercial humanoid robot. Our analysis, composed by an automated and a manual part, points out a relevant number of security flaws that can be used to take over and command the robot. Furthermore, we suggest how these issues could be fixed, thus, avoided in the future. The very final aim of this work is to push the rise of the security level of IoT products before they are sold on the public market.

CCS CONCEPTS

• **Security and privacy** → **Embedded systems security**; **Penetration testing**; *Mobile and wireless security*; *Vulnerability management*; • **Computer systems organization** → **Robotics**;

KEYWORDS

Security, Internet of Things (IoT), Robot, Penetration Testing, Pepper

ACM Reference Format:

Alberto Giaretta, Michele De Donno, and Nicola Dragoni. 2018. Adding Salt to Pepper: A Structured Security Assessment over a Humanoid Robot. In *Proceedings of International Conference on Availability, Reliability and Security (ARES 2018)*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3230833.3232807>

1 INTRODUCTION

In 1990 an Artificial Intelligence (AI) winter violently hit the field and took it to its lowest point in history. Since then, AI popularity and enthusiasm constantly grew, up to the point where AI start-ups were able to raise more than \$6 billions in 2017, according to Venture

Scanner [13]. Even though AI permeates many fields without being acknowledged to be AI (as Nick Bostrom advocated in 2006 [4]), it is most obviously associated with robotics. We are witnessing an explosion of domestic robots which, equipped with Internet of Things (IoT) connectivity, make our lives more convenient and joyful.

But all that glitters is not gold. While these devices capabilities and potentialities are stunning, there seems to be a shallow approach to the security requirements of devices which are meant to closely interact with human beings and populated environments. A hacked robot, used for instance in a private home or even worse in a public space, like an airport, can have tremendous consequences for the safety of human beings, especially when it is easy to remotely turn it into a “cyber and physical weapon”, exposing malicious behaviors.

It is a matter of fact that the Internet of Things security state has been under the microscope for long time [9], and threats like Mirai [7, 8] showed that Distributed Denial of Service (DDoS) attacks can impact critical infrastructures and, subsequently, threaten human beings. Nevertheless, Internet-enabled robots are extremely powerful IoT devices equipped with a substantial number of sensors, motile capabilities, and often robotic arms that can grab and handle things. All in all, these characteristics make robotic devices one of the most critical classes of IoT devices currently sold off the shelf.

In this paper, we perform a thorough security assessment over Pepper, a commercial human-shaped social robot by SoftBank Robotics. The assessment is conducted in two main phases: an automated assessment and a manual one. The results of our experiments are critical. We were able to steal the user credentials, perform a privilege escalation, and steal data. Moreover, we found out that it is possible to physically command the robot without authentication, use it to spy people and, potentially, even directly harm them.

The security flaws highlighted in this work are another evidence of the general trend of shallowness undertaken by IoT manufacturers with regard to the security of their devices. Indeed, it often results in the overlook of a number of well-established security best practices and countermeasures that exposes the final product to basic, yet very dangerous, security flaws which could be easily preventable. We have the feeling that commercial robots get on the market too quickly, evolving from research frameworks to final products without enough prior security investigation. Thus, our

ARES 2018, August 27–30, 2018, Hamburg, Germany

© 2018 Association for Computing Machinery.

This is the author’s version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Proceedings of International Conference on Availability, Reliability and Security (ARES 2018)*, <https://doi.org/10.1145/3230833.3232807>.

final goal is not only to report our findings, but also to highlight that such devices should undergo a prudent security evaluation, before becoming commercial products.

Please note that the findings of this paper are the result of experiments conducted on two samples of Pepper, provided to two different universities for research purposes. Therefore, we cannot assert that all the results are applicable to every Pepper on the market.

Paper Outline. The paper is organized as follows. Section 2 describes the testbed robot. Section 3 discusses related work. Section 4 and Section 5 respectively present the automated and manual security assessments we performed on Pepper, and related outcomes. Section 6 proposes some mitigations to the discovered vulnerabilities, and how similar issues might be prevented in future products. Lastly, Section 7 wraps up our investigation and draws the conclusion.

2 WHAT IS PEPPER?

Manufactured by SoftBank Robotics, formerly Aldebaran Robotics, Pepper is a human-shaped social robot, designed to engage people and interact with them, not to help with practical duties. Its main characteristic is the capability to infer basic human emotions and react accordingly: it will be joyful if the engaged human is happy, comforting if sad, and so on.

Shown in Figure 1, Pepper strong points are perception and expression, and to achieve them the manufacturer equipped it with different sensors. Within its head lie 4 microphones, 2 HD cameras, a 3D depth sensor, and 3 touch sensors. Also, it has 2 more touch sensors, one per each hand, and 1 gyroscope in its torso. Last, the base is equipped with another gyroscope, 2 sonar sensors, 6 laser sensors, and 3 bumper sensors.




Figure 1: Pepper in one of the Örebro University robotic labs.

Currently, Pepper is used to welcome and inform people in various scenarios. In Japan, it is commonly seen in SoftBank Mobile stores and even in private houses. Allegedly, Nestlé is planning to equip more than 1000 Nescafé sale points with a Pepper. Furthermore, early in 2017, a restaurant at Oakland Airport chose Pepper as an information tool and, around mid-February 2018, Lufthansa deployed it in Munich Terminal 2, as a first test phase aimed to evaluate whether passengers might accept it or not.

Moreover, Pepper can be operated through Robot Operating System (ROS) [21], which is one of the most prevalent middlewares in robotics. Even though it is improperly called an Operating System (OS), ROS exposes many characteristics typical of OSes, such as hardware abstraction and low-level device control. Its collection of libraries and tools greatly simplifies the task of operating a large variety of robots that deeply diverge from each other from the hardware point of view.

3 RELATED WORK

To the best of our knowledge, there is no scientific paper in the literature discussing a structured security assessment of Pepper, nor of any other humanoid robot. However, for the sake of completeness, there are some works worth to mention.

First, the security consultancy company IOActive wrote a white paper [2] about a number of cyber-attacks performed over a set of different robots, amongst which Pepper. Their work gives an overview of the security vulnerabilities they found, and lists some generic recommendations that could be put in practice to improve the overall robotics security. Our study corroborates some of their results but, since they did not specifically focus on a single robot, their evaluation of Pepper is incomplete and many things are not analyzed nor reported. Indeed, our structured methodology helped us to depict a clearer overview about the general security state of this particular robot.

Other works related to robotics security can be found in the literature. Clark et al. [3] identify security threats in the robotics field, classify them based on the respective architectural layer, and discuss both their impact and possible countermeasures. Other works [11, 27] highlight and discuss a number of security issues of ROS (Robot Operating System). As already mentioned in Section 2, even though ROS is not necessary for Pepper to operate, it is a widespread middleware in the robotics community and compatible drivers are already available to command Pepper through ROS. This entails that ROS vulnerabilities can easily add up to the operated robot ones, and lead to an even worse situation.

4 AUTOMATED ASSESSMENT

The first step of our security assessment is to perform an automated analysis of the target. This is a fundamental process to check the basic security level of any device and it is described in this Section.

First of all, we performed a port scanning. Secondly, we ran different tools to perform an automatic vulnerabilities scanning. All the information gathered in this process will be extremely useful in the subsequent manual phase. The different steps and consequent results of the automated assessment are described in the rest of the Section.

4.1 Port Scanning (Nmap)

The aim of a port scanning is to discover the open network ports on a target, as well as the services bound to each open port. In our investigation, we used *Nmap* [17] to run a full port scanning, both for TCP and UDP ports. As shown in Table 1, the manufacturer used out-of-date software, released around 2014.

On the one hand, this is a benefit, because older software is usually more stable than the newest version. On the other hand,

Table 1: Nmap port scanning results.

Port	Protocol	Service	Released
9	mDNS	DNS-based service discovery	-
22	SSH 2.0	OpenSSH 6.6	2014
80	HTTP	nginx 1.4.7	2014
8002	HTTP	Tornado httpd 3.1.1	2013
9559	-	naoqi	-

sticking to old software is a bad practice from a security point of view, since often the software team does not provide long term support, nor security updates, for outdated software. Therefore, in case a new vulnerability is discovered, old software has lower chances to receive updates, with respect to newer one. The most obvious example is Heartbleed [10], an implementation flaw found in dozens of versions of OpenSSL, around 2014. This bug exposed, for many years, private encryption keys of users to malicious attackers: a patch and new version of OpenSSL were immediately issued, and everyone was exhorted to immediately update the affected software.

Although a quick investigation on CVEdetails [15] (the main information source for software vulnerabilities) shows that no critical security issues actually exist for the installed software packages, new vulnerabilities are discovered every day. As aforementioned, a system not regularly updated is potentially much more vulnerable to new exploits than a maintained one.

Apart from outdated software listening on standard port services, the port scanning showed an unusual naoqi service, binded to port 9559. We further investigate the matter in Section 5.5.

4.2 Automated Vulnerabilities Scanning

The aim of an automated vulnerabilities scanning is to discover if the robot services are vulnerable to well-known security issues. We used two tools to perform the scanning: *OpenVAS* [19], a complete vulnerabilities scanner framework, and *OWASP ZAP* [18], a web application vulnerabilities scanner. The results obtained by running these tools against Pepper are discussed below.

4.2.1 OpenVAS. The result obtained by executing OpenVAS against Pepper is highlighted in Table 2. The tool raised 3 vulnerabilities and marked them accordingly to the severity: one as a medium and two as low.

Table 2: OpenVAS results. Severity scores are assigned by the tool itself.

Vulnerability	Severity
SSH weak encryption algorithms supported	MEDIUM
SSH weak MAC algorithms supported	LOW
TCP timestamp	LOW

The first warning is the troublesome one. Probably due to compatibility reasons, the manufacturer left enabled some known weak

ciphers which could lead to third party exploitation. As an example, even though arcfour (RC4) is remarkably quick, many issues have been discovered throughout the years. Its usage led to severe problems in the past (e.g. RC4 is the main reason why WEP protocol turned out to be insecure) and it should not be used anymore, nowadays.

The second warning is related to the support of weak Message Authentication Code (MAC) algorithms. Today, some hash algorithms, such as MD5 and the ones producing a 96-bit hash value, are considered weak and thus not suitable to generate MACs [14]. Indeed, due to their proneness to hash functions collision attacks, it is possible to exploit them to conduct SSH downgrade attacks [1]. Therefore, it is recommended to deprecate such algorithms and rely upon collision-resistant algorithms, such as SHA-2.

The third, and last, warning is a minor issue. Potentially, a malicious third party could estimate the target uptime and infer, with time, if the system is regularly patched (which often entails a reboot). Moreover, an attacker could evaluate whether a Denial of Service (DoS) attack was successful or not, again, by gathering the target uptime.

4.2.2 OWASP ZAP. The result obtained by executing OWASP ZAP against Pepper is highlighted in Table 3. The tool raised 3 warnings and marked them accordingly to the severity: one as a medium and the other two as low.

Table 3: OWASP ZAP results. Severity scores are assigned by the tool itself.

Vulnerability	Severity
X-Frame-Options Header Not Set	MEDIUM
Web Browser XSS Protection Not Enabled	LOW
X-Content-Type-Options Header Missing	LOW

The first warning implies that users are potentially exposed to Clickjacking attacks, a category of malicious techniques that trick users into clicking objects disguised as innocuous, through strategic layering of opaque and transparent layers. As an example, a transparent layer can be overlapped to a login web page and used to intercept keystrokes of the users.

The second warning highlights that a simple, yet effective, directive to protect from XSS (Cross-site scripting) attacks is not enabled on the web page, which exposes clients to arbitrary execution of malicious code. This issue, together with the first one, can have huge impact over the general security status of a web service.

The third warning, again, indicates that users are potentially exposed to XSS attacks and malicious code execution. The MIME-Sniffing is a feature that some browsers use to derive the correct data flow, in case that the contacted web server does not declare correctly the MIME type. Unfortunately, if the website accepts uploads from the users without thoroughly checking the data content, this feature can be exploited by a malicious user. First, the attacker disguises a *.html* file as a *.jpg* one by changing the extension, and uploads it to the server (*.html* are usually not allowed, whereas *.jpg* ones are). Then, the attacker tricks the victim into opening the *.jpg* file. Now,

th MIME-Sniffing feature of the victim browser discovers that the *.jpg* file is, as a matter of fact, a *.html* file, therefore the browser overrides the server MIME type and execute the file as a *.html*. In Section 5.3, we will show that one of the Pepper default applications enables the users to upload arbitrary files.

5 MANUAL ASSESSMENT

In the previous Section we have shown how we performed the initial security assessment of Pepper and gathered some basic information. Based on such information, we started looking deeper into the listening services and the related communication flow. As tools, we used *Ettercap* [16] to perform an ARP spoofing attack and then we analyzed the intercepted traffic through *Wireshark* [26]. Furthermore, we used *Hydra* [24] to perform a dictionary brute-force attack and *SSLsplit* [22] as a Man-in-the-Middle (MitM) tool.

For the sake of completeness, we also investigated the system through a “uname -a” terminal command, which resulted in the following information: “Linux Pepper 4.0.4-rt1-aldebaran-rt #1 [...] i686 Intel(R) Atom(TM) CPU E3845 @ 1.91GHz”. After a quick verification of the CPU model, we report that Pepper is prone both to Meltdown and Spectre attacks [12].

5.1 ARP Spoofing and Traffic Analysis (Ettercap & Wireshark)

Pepper offers to users a nice web page where some administrative tasks, such as factory reset and password change, can be easily performed. As soon as we opened the web page, we noticed that the communication is established over HTTP, rather than HTTPS. Therefore, we decided to capture the login communication flow to steal the login credentials. As shown in Figure 2, we were easily able to retrieve in clear text the *user:password* pair, where *user* is always *nao*.

```
- Hypertext Transfer Protocol
  GET /libs/qimessaging/1.0/qimessaging.js?v=1.2.0 HTTP/1.1\r\n
  Host: 10.0.0.6\r\n
  User-Agent: Mozilla/5.0 (X11; Linux x86_64; rv:52.0) Gecko/20100101
  Accept: */*\r\n
  Accept-Language: en-US,en;q=0.5\r\n
  Accept-Encoding: gzip, deflate\r\n
  Referer: http://10.0.0.6/\r\n
  Connection: keep-alive\r\n
  Authorization: Basic bmFvOmNhcmVzc2VzLTlwaWME=\r\n
  Credentials: nao: [REDACTED]\r\n
```

Figure 2: Authentication is performed over an unencrypted HTTP channel, easy to spoof for any attacker.

According to Pepper documentation [23], root login over SSH is disabled and only *nao* user can perform a SSH connection, which is a positive thing from a security perspective. However, end-users can only set the password for *nao* profile and no tool suggests, nor enables, them to change the root password. Besides, *root* password is clearly written in the manual, and even easily guessable (i.e., *root*). Unluckily, according to the documentation [23], the *user:password* pair to log in the web application is the same for SSH. Therefore, after stealing the credentials, transmitted in clear text over HTTP, we logged as a normal user (i.e., *nao*) via SSH and we issued a *su* command to perform a privilege escalation. In order to prevent this privilege escalation is sufficient to modify the root password.

However, since there is no easy tool to change the password, the end-user should know how to issue a password change command through terminal, which is often not the case, thus, the chances that the root password will remain the default one are high.

Moreover, we noticed that the web application does not provide any logoff feature, and this could be a potential problem. Indeed, in case that a user logs in from a public/shared workstation, either he makes sure to clear the browser cache, or he has no other way to logoff from the website.

As a remark, we want to stress that performing authentication over an unencrypted communication channel is a wrong and consistently dangerous security practice which should be always avoided, especially in commercial products.

5.2 SSH Dictionary Brute-force (Hydra)

Knowing that the username *nao* is fixed for this device, we investigated whether a brute-force attack could be performed over SSH or not. In particular, we assessed if any kind of brute-force protections have been deployed, such as IP addresses blacklist, upper bound of allowed parallel connections, and so on.

In order to assess the presence of countermeasures, we used Hydra as a dictionary attack tool and we fed it with a 1000 entries dictionary. To ensure that brute-force operations were carried on without glitches, as last entry of the dictionary we appended the real password, expecting to have a positive result at the end. We started by executing Hydra with 12 concurrent threads and the correct password was found in approximately 7 minutes, without experiencing any limitations. Then, we increased the concurrent threads to 16, and it took approximately 5 minutes to discover the password. Thus, we decided to double the threads and, again, we did not experience any limitations, confirmed by the fact that the password was found in approximately 2 minutes. Last, we used the maximum number of threads allowed by a single Hydra instance, which is 64 threads, and the password was found in exactly 1 minute, as shown in Figure 3.

This experiment clearly shows that no countermeasures to brute-force attacks have been deployed with Pepper.

For the sake of clarity, we chose to perform a dictionary attack to have a time-bounded result and assess whether brute-force attacks are feasible or not. Nevertheless, the lack of countermeasures, and the consequent feasibility of the attack, holds even for pure permutation brute-force attacks.

```
Hydra (http://www.thc.org/thc-hydra) starting at 2018-02-28 15:04:43
[WARNING] Many SSH configurations limit the number of parallel tasks, it is recommended to reduce the tasks: use -t 4
[DATA] max 64 tasks per 1 server, overall 64 tasks, 1001 login tries (l:l/p:1001), ~16 tries per task
[DATA] attacking ssh://10.0.0.6:22/
[22][ssh] host: 10.0.0.6 Login: nao password: [REDACTED]
1 of 1 target successfully completed, 1 valid password found
[WARNING] Writing restore file because 55 final worker threads did not complete until end.
[ERROR] 55 targets did not resolve or could not be connected
[ERROR] 64 targets did not complete
Hydra (http://www.thc.org/thc-hydra) finished at 2018-02-28 15:05:43
```

Figure 3: SSH dictionary brute-force attack with Hydra. Comparing start and end time, Hydra was able to test 1000 entries in 1 minute (16 tries per second).

5.3 Simple Animated Messages (SAM) Application

As previously discussed in Section 4.2.2, the necessary condition to perform a MIME-Sniffing attack is the capability to upload a malicious script by changing its extension (e.g., from *.sh* to *.jpg*). Of all the default applications provided by Pepper, the Simple Animated Messages (SAM) application caught our attention. Once the user accesses the app web page at `<ip-address-pepper>/apps/sam/`, he can design a simple choreography that makes Pepper move, say something through a text-to-speech service, and show a picture on its on-board tablet. Moreover, the user can save the choreography and later edit it.

We were particularly interested in the picture upload feature, therefore we tried to upload a *.jpg* file and we found out that the file is temporarily stored in the directory `/data/home/nao/.local/share/SAMService/uploads`. Once the user saves the choreography, the file is moved to `/data/home/nao/.local/share/SAMService/backups/<ChorName>`, where `<ChorName>` is the name chosen by the user to save the choreography. Instead, if the user does not save the choreography, the directory uploads is emptied as soon as the application is closed.

Then, we tried to upload different files and we quickly figured out that the application performs no control over the file extension. As a matter of fact, we were able to upload images, text files which extensions have been modified to images, and even plain text files without performing extension editing. We did not perform a proof-of-concept attack, but it is evident that a malicious attacker can easily leverage this application to attack the machine of the administrator.

A possible attack scenario is the following. The attacker could create a script, save it as a *.jpeg* file, and substitute the original file with the forged one. As a result, the on-board tablet of Pepper would stop to show the intended picture and, in order to fix this unexpected problem, the administrator could be led to access the SAM web page. At this point, given that no MIME-Sniffing protection is enabled, the browsers running on the administrator machine could interpret the forged file as a script and subsequently execute malicious arbitrary code.

5.4 Man-in-the-Middle (SSLsplit)

When it comes to secure communications, one of the critical points is how certificates are handled during the TLS/SSL handshake. As shown in [5], wrong implementations of the TLS protocol can expose the communication to Man-in-the-Middle (MitM) attacks [6].

Analyzing Pepper communications through Wireshark, we noticed a number of encrypted packets, which led us to assess whether the handshake is carried on correctly or not. In order to do so, we used SSLsplit as a MitM tool and forged a self-signed certificate as similar as possible to the original one; after having ARP-spoofed both Pepper and the local gateway, we tried to feed the fake certificate to Pepper. These attempts were unsuccessful. Pepper seems to reject our fake certificates and keeps on refreshing the connection until it gets the correct certificate. This implies that the certificates validity is correctly assessed by Pepper and that a MitM attack is not easily done.

On a side note, ARP spoofing a target without performing a successful MitM attack usually cuts the target capability to connect to the Internet. However, after attacking Pepper, its standard applications kept on functioning even without Internet connection. Thus, to date, we cannot say exactly what kind of communication Pepper carries on the Internet, but we know that it does not seem to be vital for the basic operations of the robot.

5.5 Remote Control without Authentication

As a strong point, Pepper exposes an API that enables users to program and command it. The APIs are available for a number of different programming languages, such as Python, C++, and Java. Users can operate the robot both through programs stored & run on Pepper itself, and through execution of remote scripts. In the latter case, Pepper APIs enable users to access all the sensors, cameras, and microphones included, as well as, all the moving parts it has been equipped with, no exclusions.

While this functionality has interesting implications, we discovered that it is highly insecure. Indeed, as previously said in Section 4, Pepper exposes a service on port 9559 which accepts TCP messages and reacts accordingly. As long as the messages comply to the API, by-design, Pepper accepts packets from whoever sends them. In other words, we found out that anyone can write a simple script that remotely commands Pepper without providing any credential. Therefore, anyone can use Pepper to do a number of different things which include, but are not limited to, the following:

- Spy on people through Pepper cameras and microphones;
- Interact with people and ask any kind of question, aiming to gather some personal information;
- Hurt people by abruptly moving towards them or grabbing them;
- Shut down and factory reset Pepper, through ALSystem-Proxy API module.

In a research environment, the fact that no authentication is required to remotely control the robot could be a reasonable choice to keep things simple. However, as previously discussed in Section 2, Pepper is currently deployed in many scenarios where human beings are present and actively interact with it, and anyone can buy it off the shelf, thus, this security flaw can result in unpleasant outcomes. Moreover, it confirms once again the alarming trend that many robot manufacturers [2] are undertaking by not considering the security as a key point in the design of their products.

6 RESULTS AND COUNTERMEASURES

In this Section, we recap the steps we followed and the results obtained. In addition, we suggest some ideas to fix the issues we found and, consequently, to avoid the same problems in future products. The section follows the same order used in attacking the robot, from the initial port scanning to the remote control exploit.

6.1 Wrapping Up: Steps and Outcomes

As previously stated, we followed a structured methodology while conducting this security assessment, and we believe that others could benefit from the same kind of approach. Shown below, the steps that we deemed mandatory to have a clear and complete security overview of the analysed robot.

1. Automated assessments:
 - 1a) Run a port scan software;
 - 1b) Execute automated suites (e.g., OpenVAS).
2. Manual assessments:
 - 2a) Analyse network traffic with Wireshark;
 - 2b) Run bruteforce attacks;
 - 2c) Assess SSL certificates handling;
 - 2d) Manually investigate other interesting characteristics, such as uncommon open ports.

Following the aforementioned methodology we found a number of different vulnerabilities, briefly recapped in Table 4. Note that we manually assigned the severity marks in Table 4, according to our experience, whereas the scores in Table 2 and Table 3 were automatically assigned from the respective tools.

The practical implications of Table 4 are critical for a device capable of doing as many things as Pepper. Stealing the admin credentials is really easy and, once in, getting root credentials takes just a command. From there, everything is possible and the only limit is the attacker skills; for instance an attacker can steal sensitive data or alter the communication flow. Even worse, due to design choices, a malicious attacker can remotely access to Pepper hardware without any kind of authorization. This entails that anyone could easily spy through Pepper cameras and microphones, as well as make it perform physical actions that could even harm human beings. Moreover, Pepper can potentially be used to hack other devices through its web services.

6.2 Automated Assessment - Countermeasures

The port scanning phase is fundamental to gather relevant information about the target: the more information, the more chances for the attacker to penetrate the system. One interesting way to counteract this technique is by using a software called Portspooof [20], whose operation mode is counter-intuitive, yet really effective. Indeed, Portspooof opens all the 65535 TCP ports and returns a SYN+ACK reply to any incoming request. Moreover, for each TCP port that does not offer a service, Portspooof emulates a real service through fake replies and banners stored in a local database. As a result of these two combined techniques, port scanning becomes much less informative, since the attacker cannot distinguish between honest and fake services.

Next step, running some automated assessment tools is a good security practice for a new product. Once it has been done, the manufacturer can revise the vulnerabilities list and fix them. In our case, we found out that Pepper instance of OpenSSH supports weak ciphers, yet it takes a little effort to edit the configuration file `/etc/ssh/sshd_config` and remove them. The same goes for the other warnings, which are clearly due to lack of personalization during the installation phases. In addition, what usually comes out from such assessments is that the software that run on similar devices are seriously outdated and this triggers a reflection. In our opinion, manufacturers are not only responsible to deliver a well-functioning product, but also to maintain it through regular security analyses and consequent updates.

6.3 Manual Assessment - Countermeasures

As already mentioned, the administration panel of Pepper runs over HTTP, leaking username and password in clear text. Even if it should be trivial, we consider important to remark, once again, that whenever sensitive information have to be communicated, HTTPS instead of HTTP is a must-have best practice. Communicating sensitive data over an insecure channel, such as HTTP, is strongly deprecated nowadays.

Beside this, the fact that the *root* password cannot be easily changed in Pepper deserves a comment of its own. On the one hand, it is true that remote root login is denied and an attacker has to know *nao* password before escalate privileges. On the other hand, if the attacker finds out *nao* password then root escalation cannot be prevented, since the user is never asked to change the root password. In recent past, we have seen that the same approach with credentials led to catastrophic results for the IoT world, with large-scale malware attacks such as Mirai [8, 8]. Thus, we deem essential to enforce users to change all their devices passwords, the first time they power them on.

On a side note, nowadays one of the main problems with the username/password paradigm is that people have too many services and devices to administrate. This growing complexity leads to adopting weak password and reusing them on different platforms, which entails a low security level. To face this new challenge, we envision a future where the user/password paradigm is abandoned in favor of smarter Access Control (AC) mechanisms, based on blockchain smart contracts. As an example, UniqID is a blockchain access management startup, which aims to enable IoT devices to authenticate and communicate directly with each other, with no credentials needed, by means of a private blockchain [25].

Speaking about passwords, defending from remote bruteforce attacks is of paramount importance. There is a couple of easily applicable countermeasures to bruteforce attacks, as we already mentioned in Section 5. First of all, IP blacklisting can help to cut out a single attacker that tries to infiltrate the device, but this is not sufficient. A motivated attacker, indeed, could use a botnet and coordinate a massive distributed bruteforce attack, in order to struck the attack from a large pool of IPs. Therefore, putting in place an upper bound of simultaneous connections can help, by strongly slowing down the tries per minute rate.

With respect to the SAM application, even though this is a non-critical application, it is useful to comment on this issue in order to address a more general problem. Whenever a user is allowed to upload a specific kind of file, whether it is a picture or any other type, the upload function should carefully analyse if the file adheres to the expected input. Controlling the file extension is not enough, and failing to filter wrong files can lead to arbitrary execution of code.

Another thing, manufacturers need to pay particular attention also to certificates handling, since MitM attacks can be highly disruptive. Even though our investigation showed that certificates are apparently handled correctly by Pepper, this is something that the manufacturer should always verify and ensure for any new product.

Last, but definitely not least, a reflection about the robot remote operations is mandatory. We cannot define this a security bug, since an authentication process was not designed from the beginning, but

Table 4: Discovered Pepper vulnerabilities. We personally assigned the severity scores, according to our experience.

Vulnerability	Practical Consequences	Severity
Port scanning is not hindered	Easy to assess the attack surface	LOW
Software not updated	Higher exposition to exploits	MEDIUM
Unrestricted file upload (in SAM)	Arbitrary code can be executed on the device	MEDIUM
Admin web page over HTTP	Credentials in clear-text	HIGH
No bruteforce countermeasures	High exposure to password cracking	HIGH
Root password cannot be easily changed	High exposure to privilege escalation	HIGH
Unauthenticated remote control	Perform any action without authorization (e.g., take pictures)	HIGH

rather an intrinsically insecure feature. Nevertheless, this finding perfectly reflects the core message that we try to convey in this paper. Security assessments and features should be inherently part of a product life-cycle, not a layer that manufacturers impose on a finite product, in an attempt to make it somehow (in)secure.

To address this critical situation, the whole APIs should be re-designed in order to achieve good security standards, and it is clear to see that this would end up being not only expensive but also time consuming. Once again, the problem is that security is not a property that the average consumer values when buys commercial products, since users are generally not technically skilled enough to evaluate it. Therefore, unless it becomes a characteristic that people ask when they purchase an IoT device, the only way to enforce security is through regulations.

With Pepper, it seems that the manufacturer took a device designed to be used as a research framework and made a commercial product out of it. However, this is not the only case, Cerrudo et al. [2] showed that other products from different manufacturers have similar faults, and the authors express exactly our same concerns.

7 CONCLUSION

In this paper, we performed a security assessment over a human-shaped social robot manufactured by SoftBank Robotics, namely Pepper. Our assessment, conducted through an automated phase and a manual one, pointed out a number of security flaws which are an evidence of the general trend of shallowness that many IoT manufacturers are undertaking with regard to the security status of their final products.

The vulnerabilities we discovered enable an attacker to easily spoof login credentials, steal data stored in the robot, hack other connected devices that interact with it, and even physically harm human beings. In particular, we found out that Pepper APIs are inherently flawed and accept TCP packets from any unauthenticated source, assuming that only legitimate users will call the APIs.

This case is a blatant example of our core message: manufacturers should consider security aspects of their products, before selling them on the market. Until now, traditional IoT devices were very simple, therefore their security flaws did not raise enough awareness about the consequent risks. Now, we are starting to deal with devices that cannot only jeopardize the security of human beings, but also their safety. Thus, overlooking security assessments during the system design phase can lead to very dangerous consequences, as well as expensive remedies.

REFERENCES

- [1] Karthikeyan Bhargavan and Gaëtan Leurent. 2016. Transcript collision attacks: Breaking authentication in TLS, IKE, and SSH. In *Network and Distributed System Security Symposium—NDSS 2016*.
- [2] Cesar Cerrudo and Lucas Apa. 2017. Hacking Robots Before Skynet. IOActive Website. Retrieved March 10th, 2018 from <https://ioactive.com/pdfs/Hacking-Robots-Before-Skynet.pdf>
- [3] George W. Clark, Michael V. Doran, and Todd R. Andel. 2017. Cybersecurity issues in robotics. In *2017 IEEE Conference on Cognitive and Computational Aspects of Situation Management (CogSIMA)*. 1–5. <https://doi.org/10.1109/COGSIMA.2017.7929597>
- [4] CNN. 2006. AI set to exceed human brain power. Retrieved March 10th, 2018 from <http://edition.cnn.com/2006/TECH/science/07/24/ai.bostrom/>
- [5] Mauro Conti, Nicola Dragoni, and Sebastiano Gottardo. 2013. MITHYS: Mind The Hand You Shake - Protecting Mobile Devices from SSL Usage Vulnerabilities. In *Security and Trust Management, Rafael Accorsi and Silvio Ranise (Eds.)*. Springer Berlin Heidelberg, Berlin, Heidelberg, 65–81.
- [6] Mauro Conti, Nicola Dragoni, and Viktor Lesyk. 2016. A Survey of Man In The Middle Attacks. *IEEE Communications Surveys Tutorials* 18, 3 (2016), 2027–2051.
- [7] Michele De Donno, Nicola Dragoni, Alberto Giarretta, and Angelo Spognardi. 2017. Analysis of DDoS-capable IoT Malwares. In *Federated Conference on Computer Science and Information Systems*. IEEE, 807–816.
- [8] Michele De Donno, Nicola Dragoni, Alberto Giarretta, and Angelo Spognardi. 2018. DDoS-Capable IoT Malwares: Comparative Analysis and Mirai Investigation. *Security and Communication Networks* 2018 (2018), 1–30. <https://doi.org/10.1155/2018/7178164>
- [9] Nicola Dragoni, Alberto Giarretta, and Manuel Mazzara. 2017. The Internet of Hackable Things. In *Proceedings of 5th International Conference in Software Engineering for Defence Applications*, Paolo Ciancarini, Stanislav Litvinov, Angelo Messina, Alberto Sillitti, and Giancarlo Succi (Eds.). Springer International Publishing, 129–140.
- [10] Zakir Durumeric, James Kasten, David Adrian, J Alex Halderman, Michael Bailey, Frank Li, Nicolas Weaver, Johanna Amann, Jethro Beekman, Mathias Payer, et al. 2014. The Matter of Heartbleed. In *Proceedings of the 2014 Conference on Internet Measurement Conference*. ACM, 475–488.
- [11] Se-Yeon Jeong, I-Ju Choi, Yeong-Jin Kim, Yong-Min Shin, Jeong-Hun Han, Goo-Hong Jung, and Kyoung-Gon Kim. 2017. A Study on ROS Vulnerabilities and Countermeasure. In *Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction (HRI '17)*. ACM, New York, NY, USA, 147–148. <https://doi.org/10.1145/3029798.3038437>
- [12] Paul Kocher, Daniel Genkin, Daniel Gruss, Werner Haas, Mike Hamburg, Moritz Lipp, Stefan Mangard, Thomas Prescher, Michael Schwarz, and Yuval Yarom. 2018. Spectre Attacks: Exploiting Speculative Execution. *arXiv preprint arXiv:1801.01203* (2018).
- [13] Medium. 2017. Artificial Intelligence Funding Trends - Q3 2017. Retrieved March 10th, 2018 from <https://medium.com/@VentureScanner/artificial-intelligence-funding-trends-q3-2017-4e8a9b80fe53>
- [14] CERT website. 2008. Vulnerability Note VU#836068 - MD5 vulnerable to collision attacks. Retrieved June 25th, 2018 from <https://www.kb.cert.org/vuls/id/836068>
- [15] CVEdetails Website. [n.d.]. CVE Details. Retrieved March 10th, 2018 from <https://www.cvedetails.com/>
- [16] Ettercap Website. [n.d.]. Ettercap Home Page. Retrieved March 10th, 2018 from <http://www.ettercap-project.org/index.html>
- [17] Nmap Website. [n.d.]. Nmap: the Network Mapper. Retrieved March 10th, 2018 from <https://nmap.org/>
- [18] OWASP Website. 2018. OWASP Zed Attack Proxy Project. Retrieved March 10th, 2018 from https://www.owasp.org/index.php/OWASP_Zed_Attack_Proxy_Project

- [19] OpenVAS Website. [n.d.]. OpenVAS. Retrieved March 10th, 2018 from <http://www.openvas.org/>
- [20] Portspooft Website. 2012. About Portspooft. Retrieved March 10th, 2018 from <http://drklwi.github.io/portspooft/>
- [21] ROS Website. [n.d.]. About ROS. Retrieved March 10th, 2018 from <http://www.ros.org/about-ros/>
- [22] SSLsplit Website. 2018. SSLsplit - transparent SSL/TLS interception. Retrieved June 25th, 2018 from <http://www.roe.ch/SSLsplit>
- [23] SoftBank Website. [n.d.]. NAOqi OS - Getting started. Retrieved March 10th, 2018 from <http://doc.aldebaran.com/2-5/dev/tools/opennao.html>
- [24] THC Hydra Website. [n.d.]. THC Hydra. Retrieved June 25th, 2018 from <http://sectools.org/tool/hydra/>
- [25] UniquID Website. [n.d.]. UniquID Inc. - Blockchain Identity Access Management. Retrieved March 10th, 2018 from <http://uniquid.com/>
- [26] Wireshark Website. [n.d.]. Wireshark - Learn Wireshark. Retrieved March 10th, 2018 from <https://www.wireshark.org/#learnWS>
- [27] Ruffin White, I. Henrik Christensen, and Morgan Quigley. 2016. SROS: Securing ROS over the wire, in the graph, and through the kernel. *arXiv preprint arXiv:1611.07060* (2016). <https://arxiv.org/abs/1611.07060>