Automated Dynamic Firmware Analysis at Scale: A Case Study on Embedded Web Interfaces

Andrei Costin
EURECOM
Sophia Antipolis, France
costin@eurecom.fr

Apostolis Zarras
Ruhr-University Bochum
Germany
apostolis.zarras@rub.de

Aurélien Francillon
EURECOM
Sophia Antipolis, France
francill@eurecom.fr

Abstract—**Embedded devices are becoming more widespread, interconnected, and web-enabled than ever. However, recent studies showed that these devices are far from being secure. Moreover, many embedded systems rely on web interfaces for user interaction or administration. Unfortunately, web security is known to be difficult, and therefore the web interfaces of embedded systems represent a considerable attack surface.**

In this paper, we present the *first fully automated framework* that applies dynamic firmware analysis techniques to achieve, in a scalable manner, automated vulnerability discovery within embedded firmware images. We apply our framework to study the security of embedded web interfaces running in Commercial Off-The-Shelf (COTS) embedded devices, such as routers, DSL/cable modems, VoIP phones, IP/CCTV cameras. We introduce a methodology and implement a scalable framework for discovery of vulnerabilities in embedded web interfaces regardless of the vendor, device, or architecture. To achieve this goal, our framework performs full system emulation to achieve the execution of firmware images in a software-only environment, i.e., without involving any physical embedded devices. Then, we analyze the web interfaces within the firmware using both static and dynamic tools. We also present some interesting case-studies, and discuss the main challenges associated with the dynamic analysis of firmware images and their web interfaces and network services. The observations we make in this paper shed light on an important aspect of embedded devices which was not previously studied at a large scale. Insights from this paper can help users, programmers and auditors in efficiently testing and securing their Internet enabled embedded devices.

We validate our framework by testing it on 1925 firmware images from 54 different vendors. We discover important vulnerabilities in 185 firmware images, affecting nearly a quarter of vendors in our dataset. We also perform comprehensive failure analysis. We show that by applying *relatively easy* fixes during corrective maintenance it is possible to remediate at least 61.3% of emulation failures and at least 25.2% of web interface launch failures. These experimental results demonstrate the effectiveness of our approach.

I. INTRODUCTION

Embedded devices are present in many complex systems, like cars, planes, and programmable logic controllers. Such devices also appear massively in customer products such as network gateways and IP cameras. Those devices are becoming more pervasive and "invade" our lives under many different forms (e.g., home automation, smart TVs).

Embedded systems, in particular Small Office/Home Office (SOHO) devices, are often known to be insecure [40, 80]. Their lack of security may be the consequence of the harsh market competition. For instance, the time to market is crucial and the competition puts high pressure on the design and production costs, and enforces short release timelines. Vendors try to provide as many features as possible to differentiate products, while customers do not necessarily look for the most secure products.

Some embedded systems have clear and well-defined security goals, such as the pay-TV smart cards and the Hardware Security Modules (HSM). Therefore, such devices are rather secure. However, many embedded systems are not designed with a clear threat model in mind. This gives little motivation to manufacturers to invest time and money in securing them. This fact motivated several researchers to evaluate the state of security of such embedded devices [11, 17, 23, 25, 71, 79].

Moreover, during the past few years, embedded devices became more connected forming what is called the Internet of Things (IoT). Such devices are often put online by composition; attaching a communication interface to an existing (insecure) device. Most of these devices lack the user interface of desktop computers (e.g., keyboard, video, mouse), but nevertheless need to be configured and maintained. Albeit some devices rely on custom protocols used by "thick" clients or even legacy interfaces (i.e., telnet), the web quickly became the universal administration interface. Thus, the firmware of these devices often embed a web server running web applications, for the rest of this paper, we will refer to these as *embedded web interfaces*.

It is well known that making secure web applications is not a trivial task. In particular, researchers showed that more than 70% of vulnerabilities are hosted in the (web) application layer [73]. Attackers who are familiar with this fact use various techniques to exploit web applications. Well known vulnerabilities, such as SQL injection [18] or Cross Site Scripting (XSS) [81], constitute a significant portion of the vulnerabilities discovered each year [21], and are frequently

used in real-world attacks [38]. Additionally, vulnerabilities such as Cross Site Request Forgery (CSRF) [14], command injection [78], and HTTP response splitting [60] are also often present in web applications.

Given such a track record of security problems in both embedded systems and web applications, it is natural to expect the worse from *embedded web interfaces*. However, as we discuss further, those vulnerabilities are neither easy to discover, analyze, and confirm, nor do the vendors perform the necessary security quality assurance of their firmware images.

Analysis of embedded web interfaces: While there are solutions that can be used during the design phase of the software [46,67,74,75], it is also important to discover and patch existing vulnerabilities before they are found and exploited "in the wild". This is possible to do either by static analysis on their source code [13,29,31,58], or by dynamic analysis where their code or web interface is typically exercised against a number of known attack patterns [15,17].

Unfortunately, these techniques and tools can be inefficient or difficult to use for detecting vulnerabilities inside embedded web interfaces [15, 39]. For instance, performing static analysis on embedded web interfaces seem to be a rather simple task once the firmware has been unpacked. One main limitation of this approach is that the web interfaces often rely on various technologies (e.g., PHP, CGIs, custom server-side languages). However, the static analysis tools are usually designed for a particular technology, and many static tools are often concentrated around some trendy environment (e.g., PHP) leaving the others "uncovered". In addition to this, though sound static analysis tools exist, many other static analysis tools are merely "glorified greps" and have a large number of *false positives (FP)*, which make them problematic to reliably use in an automated large scale study. On the other hand, dynamic analysis tools [37,49] are more generic as they are less sensitive to the server-side language. Nevertheless, they require the system or the web interface to be functional. Unfortunately, it is challenging to create an environment that can perfectly emulate firmware images for a broad range of devices based on a variety of computing architectures and hardware designs.

Scalable dynamic analysis of embedded web interfaces: The easiest way to perform dynamic analysis is to perform it on a live device. However, acquiring devices to dynamically analyze them is expensive and does not scale. At the same time, it is ethically questionable, if not illegal, to test devices one does not own (e.g., devices on the Internet). Another option is to extract the web interface files from a device and load them to a test environment, like an Apache web server. Unfortunately, a large majority of the embedded web interfaces use native CGIs, bindings to local architecture-dependent tools or custom web server features which cannot be *easily* reproduced in a different environment (Section II-D1).

Emulating the firmware is an elegant method to perform dynamic analysis of an embedded system, since it does not require the physical device to be present and can be completely performed in a controlled environment while being easy to scale. But emulation of unknown devices is not easy because an embedded firmware expects specific hardware to be fully present, such as peripherals or memory layouts. Previous at-

tempts were made at improving emulation of firmware images by forwarding hardware I/O or `ioctl` to the hardware [59,85]. These techniques achieve a rather good emulation, but require the presence of the original device and a great deal of manual setup, which does not scale. We noticed that in Linux-based embedded systems the interaction with the hardware is usually performed from the kernel. Moreover, the web interfaces often do not interact with the hardware or this interaction is indirect.

A. Overview of our Approach

To perform scalable security testing of embedded web interfaces we developed a distributed framework for automated analysis (Figure 1) and we tested it in a cloud setup. We started our analysis with a dataset of 1925 unpacked firmware images that contain embedded web interfaces.[1] Then, for each unpacked firmware we identify any potential web document root present inside the firmware. At this point we make a pass with static analysis tools on the modules of the service under test [2]. Next, we propose a partial emulation of firmware images by replacing their kernel with a stock kernel (targeting the same architecture) and emulating the whole userland of the firmware using the QEMU [36] emulator. We currently support the most frequent architectures that are well supported in QEMU and plan to extend later to other architectures. We then `chroot` the unpacked firmware and start the `init` program, the init scripts or sometimes directly the web server. Once (and if) the service under test is up and operational [3], we perform dynamic analysis on it. Finally, we analyze the results, and whenever applicable we perform manual analysis and investigate the failures.

B. Contributions

In this paper we present a completely automated framework to perform scalable dynamic firmware analysis. We demonstrate its effectiveness by testing the security of embedded web interfaces. Our framework mainly relies on the emulation of the firmware images. This allows to test the embedded web interfaces using off-the-shelf dynamic analysis tools.

In summary, we make the following main contributions:

- We present the first framework that achieves *scalable and automated dynamic analysis of firmwares*, and that was precisely developed to discover vulnerabilities in embedded devices using the software-only approach.

- We highlight the challenges in emulating the firmware images and testing the web interfaces of embedded systems, and describe the techniques that can be used for such tasks.

---

[1]We focused mainly on Linux-based firmware images. Linux-based firmware images are in general well structured and documented, therefore they are easier to unpack, analyze and emulate. However, our approach can be easily extended in the future to firmware based on system that are similarly well structured into bootloaders, kernels and filesystems (e.g., VxWorks, QNX). Monolithic firmware is more challenging to fully emulate and in general requires additional frameworks, such as Avatar [85].

[2]In the particular case of this study, the modules are the files within the *web document root*.

[3]In the particular case of this study, the service under test is the *web interface*.
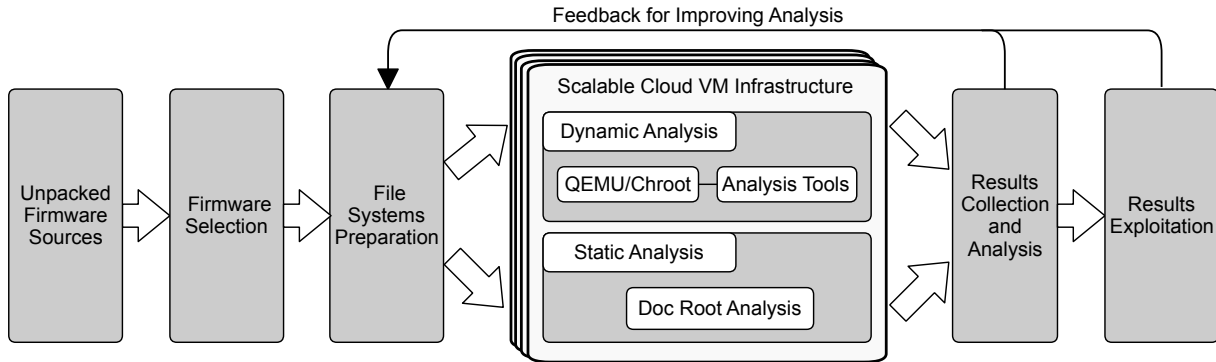
Figure 1: Overview of the analysis framework.

- We describe our framework which leverages multiple techniques and state of the art tools.

- We perform the first large scale, comprehensive, security study on web interfaces of embedded systems.

- We automatically discover 225 previously unknown serious vulnerabilities in 45 firmware images.

C. Outline

The remainder of this paper is organized as follows. In Section II, we explore techniques to emulate and analyze embedded firmware and their web interfaces. In Section III, we present our framework. In Section IV, we describe our dataset. In Section V, we showcase our results and analyze several case-studies. In Section VI, we discuss ethical aspects of our work and its limitations, as well as we propose solutions to them. We summarize the state of the art in Section VII and then conclude with Section VIII.

II. EXPLORING TECHNIQUES TO ANALYZE EMBEDDED WEB INTERFACES

In this section, we summarize the different possibilities for static or dynamic analysis of embedded web interfaces, their limitations, and motivate our final choices.

A. Static Analysis

There are many practical advantages to static analysis tools; they are often automated and do not require setting up too complex test environments. In general, they only need the source code (or application) to be provided to generate an analysis report. It is also relatively easy to plug new static analysis tools for increased coverage or wider support of file formats and source code languages. Finally, as a result of all the above, such tools are scalable and easy to automate.

However, static analysis techniques have well understood limitations. On the one hand, they cannot find all the vulnerabilities, i.e., *false negatives (FN)*, while on the other, they also alert on non-vulnerabilities, i.e., *false positives (FP)*, which becomes increasingly problematic in large scale automated setups. Additionally, we found that embedded devices' firmware often rely on uncommon technologies for which security static analysis tools often do not exist (e.g., `lua`, `haserl`,

`binary CGIs`). Albeit there exist a number of static analysis tools for the PHP language [29,57], in our dataset only 8% of embedded firmware images contain PHP code in their server-side. This is not really a surprise since PHP is not primarily designed for embedded systems. We nevertheless analyze these cases with *RIPS* [29] in Section V-B. Finally, *binary static analysis* can be applied to binary CGIs to find vulnerabilities such as buffer overflow, (remote) code execution, command injection *In this paper we use "command injection" and "command execution" terms interchangeably*. (e.g., `Firmalice` [77] or `Weasel` [76]). Also, new techniques start to appear that are able to cope with the diversity of CPU architectures found in embedded systems [72].

B. Dynamic Analysis

Dynamic analysis—an analysis that relies on testing an application by running it—has many benefits. First, dynamic analysis of web interfaces is mostly independent from the server-side technology that is used. For instance, the very same tool can test web interfaces that are implemented in PHP, native CGIs or custom web scripting engines. Second, it can be used to confirm vulnerabilities found in the static analysis phase. Although there exist many dynamic analysis tools for security testing of web applications [15], unfortunately, they often require significant effort to setup (e.g., environment setup), and sometimes additional customization such as adding new vulnerability modules for scanning, testing or validation.

For this particular study, we selected web penetration tools that are free and open source so that we can easily adapt and integrate them in our framework as well fix their defects when needed. Based on this we selected *Arachni* [1], *Zed Attack Proxy (ZAP)* [2], and *w3af* [3] to be used in our framework.

However, our approach and framework are designed in a way that allows great flexibility. For example, as depicted in the Figure 3 other tools such as Metasploit [4] and Nessus [5] can supplement or replace the web penetration tools mentioned above. In this way, we can achieve additional security and vulnerability testing that can help us increase the surface of vulnerability discovery for both known and unknown vulnerabilities.

---

[4]http://www.metasploit.com/
[5]http://www.tenable.com/products/nessus-vulnerability-scanner

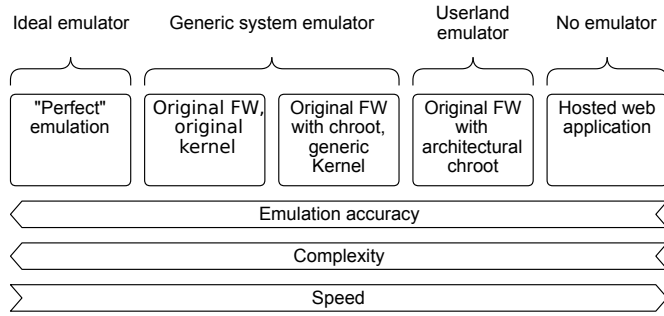| Ideal emulator	Generic system emulator		Userland emulator	No emulator
"Perfect" emulation	Original FW, original kernel	Original FW with chroot, generic Kernel	Original FW with architectural chroot	Hosted web application

Emulation accuracy

Complexity

Speed

Figure 2: Embedded web interfaces emulation possibilities: from perfect emulation of a hardware platform to hosting the web interface. (The arrows show a general increasing trend, actual evolution of the properties may not be linear.)

C. Limitations of Analysis Tools

Our framework relies on existing web analysis tools, which have their own limitations. For instance, the number of FPs and FNs of this study are a direct consequence of the vulnerability finding tools we rely on. An example of their limitations is their ability to detect *command injections* vulnerabilities. Those are frequently missed because such flaws are often hard to discover via automated testing [4,15]. For example, tools try to inject commands, such as `ping <ip>` assuming that the network is functional and that the targeted system supports the `ping` utility. We overcome some of these limitations by taking advantage of our "white box" approach (Section III-D.1).

Additionally, the tools we use were not initially targeting vulnerabilities in embedded web interfaces and were not designed to be integrated in a framework like ours. As a consequence, we found many problems with those tools which were severely impacting the success rate of the vulnerability discovery. We were able to improve or fix many of them at the cost of a significant engineering effort. Nevertheless, fixing these bugs proved necessary to obtain better results.[6] This highlights that better web application analysis tools are needed, especially ones that are in particular adapted for testing embedded web interfaces.

D. Running Web Interfaces

Dynamic analysis of web applications requires a functioning web interface. There are different ways to launch the web interface that is present in the firmware of an embedded system, however, none of them are perfect. Some methods are very accurate but infeasible in our setup, such as emulating the firmware in a perfect emulator (which is not available). Other methods are much less accurate, like extracting the web application files and serving them from a generic web server. Therefore, we evaluated different approaches (Figure 2) and describe their advantages and drawbacks.

1) Hosting Web Interfaces Non-Natively: A straightforward way to launch a web interface from a firmware is to extract and then launch it under a web server on an analysis environment,

without trying to emulate the original web server and firmware. The web application is located (i.e., the document root, as described in Section III-B2), extracted and "transplanted" to the *hosting environment*. The main advantage of this technique is that it does not require emulation, which dramatically simplifies the deployment and thus is easy to automate and scale.

However, this approach has many limitations. For example, it is not possible to handle platform dependent binaries and CGIs. We analyzed the document roots within 1580 firmware candidates for emulation and found that 57% out of these were using binary CGIs or were in some way bound to the platform.[7] In addition to this, the firmware images often use either customized web servers, or versions which are not available on normal systems, thus a generic web server (e.g., Apache) has to be used in the hosting environment. We evaluated this technique and we present results of its evaluation in Section V-D, where we also compare its performance to other techniques we use.

2) Firmware and Web Interface Emulation: A preliminary step to emulate a firmware is to know its architecture. While this may seem straightforward, it is actually often a complicated step to perform automatically at scale. For instance, some firmware packages contain files for various architectures (e.g., ARM and MIPS). Sometimes, vendors package two different firmware blobs into a single firmware update package. The firmware installer then picks the right architecture during the upgrade based on the detected hardware. In such cases, we try to emulate this filesystem with each detected architecture. We detect the architecture of each executable in a firmware using ELF headers or statistical opcode distribution for raw binaries. We then decide on the architecture by counting the number of architecture specific binaries it contains. Once we detect the right architecture, we use the QEMU emulator for that particular architecture. There are different possibilities for emulating the firmware images, which we now compare.

a) Perfect emulation: Ideally, the firmware would be complete (including the bootloader, kernel, etc.) and a QEMU configuration which perfectly emulates the original hardware would be available. However, QEMU only emulates few platforms for each CPU architectures and thus perfectly emulating unknown hardware is impossible in practice, especially considering that hardware devices can be arbitrarily complex. In addition to this, hardware in embedded devices is often custom and its documentation is, in general, not available. It is therefore infeasible to adapt the emulator let alone to apply this at a large scale.

b) Original kernel and filesystem on a generic emulator: Reusing the kernel of the firmware could lead to a more accurate emulation, in particular because it may export interfaces for some custom devices that are needed to properly emulate the system. Unfortunately, kernels for embedded systems are often customized and hence do not support a wide range of peripherals. Therefore, using the original kernel is unlikely to work very well on a generic emulator. Additionally, in our dataset only 5% of the firmware images were containing the kernel making this approach not feasible.

---

[6]We plan to submit the bugfixes to be included in the upstream releases of the tools.

[7]This is a lower bound as we did not count web scripts calling local system utilities, e.g., using the `system()` call.

c) Firmware chroot with a generic kernel and filesystem: Lacking the original kernel, it is is possible to rely on a complete generic system (for the same CPU architecture of the firmware), which is then used as a base for the analysis.[8] From this generic system we chroot to the unpacked firmware and execute the shell (e.g., `/bin/sh`) or the init binary (e.g., `/sbin/init`). Finally, we start the web server's binary along with the web interface document root and web configuration.

Ideally, it should be possible to directly boot the firmware filesystem instead. However, using a generic file system provides a consistent environment to control the virtual machine and perform our analysis and monitoring of the system. The advantage of this approach is that it allows emulation of the web interfaces and web server software in their original file system structure and can execute native programs.

This approach, however, has few drawbacks. First, emulating the system is not very fast.[9] Additionally, the emulator environment setup and cleanup introduces a significant overhead. Furthermore, with this approach we cannot fully emulate the peripherals and specific kernel extensions of the embedded devices. Even so, few firmware images and a limited part of embedded web interfaces actually interact directly with the peripherals. One such example is a web page that performs a firmware upgrade which in turn requires access to `flash` or `NVRAM` memory peripherals.

We found that this approach offers the best trade-off between emulation accuracy, complexity, and speed (see Figure 2). It is also scalable and provided the best results in analyzing dynamically the web interfaces (see Section V-D).

d) Architectural chroot: One way to improve the performance and emulation management aspects of our framework is by using *architectural chroot* [5] (also known as *QEMU static chroot*). This technique uses `chroot` to emulate an environment for architectures other than the architecture of the running host itself. This basically relies on the Linux kernel's ability to call an interpreter to execute an ELF executable for a foreign architecture. Registering the userland QEMU as an interpreter allows to transparently execute `ARM` Linux ELF executables on an `x86_64` Linux system. However, we found that this approach was not very stable, making it impossible to use it at a large scale. Finally, while this approach has the advantage of improving emulation speed, in essence, it is unlikely to improve the number of firmware packages we can finally emulate. Therefore, we did not use this technique in our setup, and we leave this for our future work.

III. ANALYSIS FRAMEWORK DETAILS

In order to perform a large scale and automatic analysis of firmware images we designed a framework to process and analyze them (Figure 1). First, we obtain a set of unpacked firmware images, analyze and filter them (Section III-A). Next, we perform some pre-processing of the selected unpacked files. For instance, some firmware images are incompletely unpacked or the location of the document root is not obvious (Section III-B). We then perform the static and dynamic analyses (Section III-C). Finally, we collect and analyze the reported vulnerabilities (Section III-D) and exploit these results (Section III-E).

A. Firmware Selection

The firmware selection works as follows. First, we select the firmware images that are successfully unpacked and are Linux-based systems which we can natively emulate and chroot (see Section III-B). Second, we choose firmware instances that clearly contain web server binaries (e.g., `httpd`, `lighttpd`) and typical configuration files (e.g., `boa.conf`, `lighttpd.conf`). In addition to these, we select firmware images that include server side or client side code related to web interfaces (e.g., `HTML`, `JavaScript`, `PHP`, `Perl`). Our dataset is detailed in Section IV.

B. Filesystem Preparation

To emulate a firmware the emulator requires its root filesystem. In the simplest case the unpacked firmware directly contains the root filesystem. However, in many cases the firmware images are packed in different and complex ways. For instance, a firmware can contain two root filesystems, one for the upgrade and one for the factory restore, or it can be packed in multiple layers of archives along with other resources. For these reasons, we first need to detect the potential candidates for root filesystems. We achieve this by searching for key directories (e.g., `/bin/`, `/sbin/`, `/etc/`, `/usr/`) and for key files (e.g., `/init`, `/linuxrc`, `/bin/sh`, `/bin/bash`, `/bin/dash`, `/bin/busybox`). Once we discover such files and folders relative to a directory within the unpacked firmware, we select that particular directory as the *root filesystem* point. There are also cases where it is hard or impossible to detect the root filesystem. A possible reason for this is that some firmware updates are just partial and do not provide a complete system. We extract each detected root filesystem and pack it as a standalone root filesystem ready for emulation.

Unpacking firmware images can produce "broken" root filesystems which we attempt to fix. Additionally, in order to start the web server within the root filesystem, we need to detect the web server type, its configuration, and the document root. For these reasons, we have to use heuristics on the candidate root filesystems and apply transformations before we can use them for emulation and analysis.

1) Filesystem Sanitization: Unpacking firmware packages is not a perfect procedure. First, unpacking tools sometimes have defects. Second, some firmware images have to be unpacked using an imperfect "brute force" approach [23]. Finally, some vendors customize archives or filesystem formats. For example, some filesystems have symbolic links that are incorrectly unpacked because they were represented as text files containing the target of the link.[10] All these lead to an incorrect unpacking and thus the unpacked firmware image differs from the filesystem representation intended to be on the device. This reduces the chances of successful emulation and therefore we need a sanitization phase.

---

[8]We use the pre-compiled Debian Squeeze packages from [56].

[9]We measured that emulation is one order of magnitude slower than native execution.

[10]For example, the symbolic link `/usr/bin/boa ->` `/bin/busybox` is represented with a text file named `/usr/bin/boa` that contains the string `/bin/busybox`.

This sanitization phase is performed by scripts that traverse unpacked firmware filesystems and fix such problems. Sometimes, there are multiple ways to fix a single unpacked firmware. This results in multiple root filesystems to be submitted for emulation, increasing our chances of proper emulation of a given firmware. Implementing these heuristics added a 13% processing overhead. At the same time, it allowed us to increase the successful emulations by 2% and the successful web server launches by 11%.

2) Web Server Heuristics: Within the firmware, we locate web server binaries and their related configuration files (e.g., `boa.conf`, `lighttpd.conf`).[11] The path of the web server and its configuration file is sufficient to start the web server using a command such as `/bin/boa -f /etc/boa/boa.conf` (a real example from our dataset). Additionally, we extract important settings from the configuration files (e.g., the document root).

Sometimes, we miss important parameters which are required to properly start the web server, such as the document root path or the CGI path. Often this happens because of a missing configuration file (e.g., partial firmware update) or because the parameters are supplied via the command line from a script which is not available. In these cases, we experiment with all the potential document roots of the firmware. To find a potential document root (within the root filesystem) we first search for index files (e.g., `index.html`, `default.html`) with possible file extensions (`HTML`, `SHTML`, `PHP`, `ASP`, `CGI`). Then, we build a set of *longest common prefix directories* of these files. This can result in multiple document root directories, for example a second document root can be found in a recovery partition. Once we discover the document roots, we prepare the possible commands to start the web server. With this, we increase the chances of bringing the web server up and operational.

We also build an optimized site map for each such document root directory. We use the site maps to hint the dynamic analysis tools which URLs they have to analyze. In general, dynamic analysis tools crawl the web application to discover its site map. However, this is inefficient and can easily miss some pages and even whole sets of vulnerabilities [33]. Thus, we instruct the tools to restrict their analysis to the supplied site map and we do this for multiple reasons. First, it significantly lowers the time required to complete the dynamic analysis. No time is wasted to analyze uninteresting files, such as image files, or to (inefficiently) crawl the web application [33]. Second, it reduces the chances for the web interface or the emulator to crash by limiting the resource load, e.g., number of requested files. Third, it increases the chances that the files that are reported as vulnerable by static analysis will also undergo dynamic analysis.

There are several possible improvements to our tests. Restricting the site map allows to complete tests in reasonable time but may miss URLs when content is dynamically generated or monolithic web server binaries are used. Another improvement would be to use a tool like `ConfigRE` [82] to automatically infer configuration files.[12]

---

[11] Namely: `httpd`, `boa`, `lighttpd`, `thttpd`, `minihttpd`, `webs`, `goahead`.

[12] Unfortunately, `ConfigRE` is not available anymore.

C. Analysis Phase

Once the filesystems are prepared, we emulate each of them in an analysis Virtual Machine (VM) where dynamic testing is performed (Figure 3 and Section II-B). We also submit the document roots to the static analyzers (Section II-A). This phase is completely automated and scales well as each firmware image can be analyzed independently.

D. Results Collection and Analysis

After dynamic and static analysis phases are completed, we obtain the analysis reports from the security analysis tools. We also collect several logs that can help us make further analysis as well as improve our framework. These are typically required to debug the analysis tools or our emulation environment. For instance, we collect SSH communication logs with the emulator host, changes in the firmware filesystem, and capture the network traffic of the interaction with the web interface.

a) File systems changes: We capture a snapshot of the emulated filesystem at several different points in time. We do this (i) before starting the emulation, (ii) after emulation is started, and (iii) after dynamic analysis is completed. Then, we perform a filesystem `diff` among these snapshots. Interesting changes are included in both log files and new files. Log files are interesting to collect in case a manual investigation is needed. New files can be the consequence of a *OS command injection* or more generally of a *Remote Code Execution (RCE)* vulnerability triggered during the dynamic analysis phase. This often occurs when dynamic testing tools try to inject commands (e.g., `touch <filename>`). Sometimes, the command injection can be successful but not detected by the analysis tools. However, it is easy to detect such cases with the filesystem diff.

b) Capturing communications: Performing dynamic analysis involves a lot of input and output data between the (emulated) device and the dynamic analysis tool. Capturing the raw input and output of the communication allows to increase accountability in case of emulation problem.

For instance, a successful *OS command injection* can go undetected by the tools. Also, such vulnerability can be difficult to verify, even in a "white box" testing approach (Section II-C). Once the testing phase is over, it can be discovered that a command injection was, in fact, successful. In such case, we need to rewind through all HTTP transactions to find the input triggering the particular vulnerability and afterward we can look for incriminating inputs and parameters (e.g., a `touch` command).

The testing tools often behave like fuzzers as they try many malformed inputs one after the other. Because of this, a detected vulnerability may not be a direct result of the last input. For example, it can be a result of the combination of several previous inputs. It is therefore important to recover all these previous inputs in order to successfully reproduce the vulnerability.

E. Results Exploitation

After collecting all the details of the analysis phase, we perform several steps to exploit these results. First, we validate the
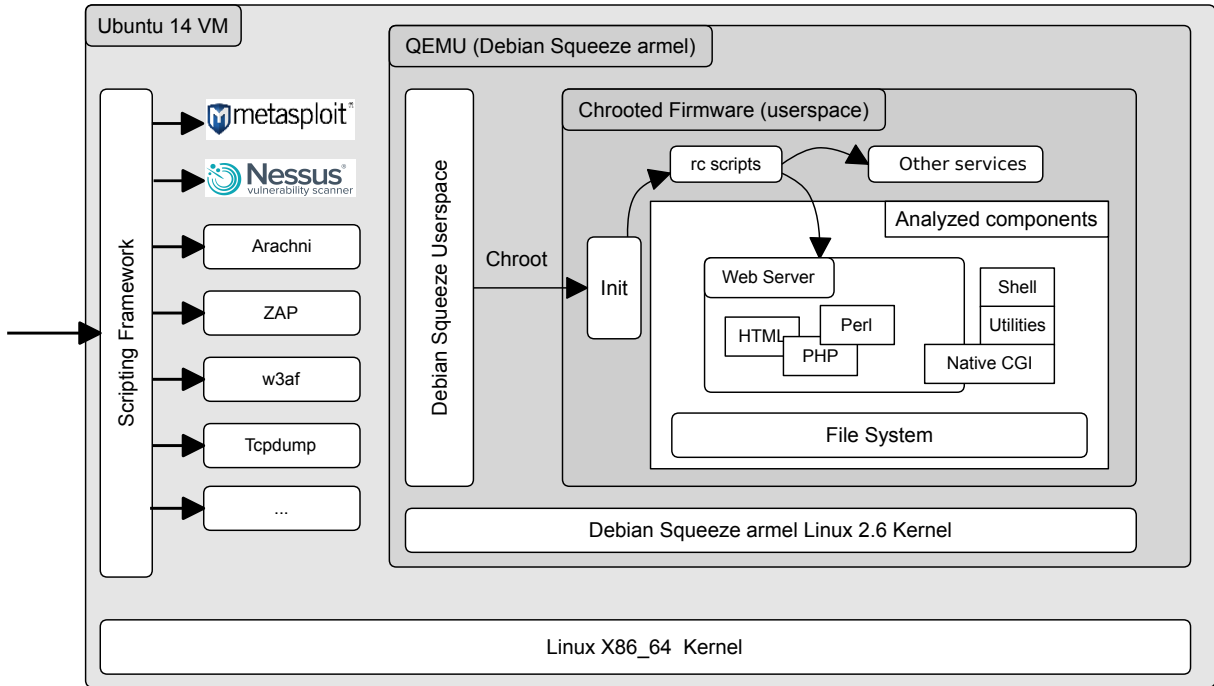
Figure 3: Overview of one analysis environment for Linux armel with a 2.6 kernel.

high impact vulnerabilities by hand and try to create a proof-of-concept exploit. This could be fully automated in the future, as was done for other fields of vulnerability research [12]. Unfortunately, none of the tools we currently use provide such functionality. Additionally, from the static analysis reports we manually select the high impact vulnerabilities (e.g., command injection, XSS, CSRF) and the files they impact. We then use these to explicitly drive the dynamic analysis tools and aim mainly at two things: (i) get the dynamic analysis tools to find the vulnerabilities they missed (if they did) and (ii) find the bugs or limitations that prevented the dynamic tools to discover these vulnerabilities in the first place. Even though manual analysis does not scale, it can help uncover additional nontrivial vulnerabilities (see Table IX). Finally, we summarize all our findings in vulnerability reports to be submitted as CVEs.

IV. DATASET

We started with a set of firmware images that we collected over time from publicly available sources. Table I presents details about the counts of the firmware images each of the phases in our framework.

First, we chose the firmware instances which were successfully unpacked and which were Linux-based embedded systems (1925). These were the systems which seem the easiest to emulate. Then, we selected firmware instances that clearly contained a web server binary (e.g., `httpd`, `lighttpd`) and typical configuration files (e.g., `lighttpd.conf`, `boa.conf`). In addition to these, we also chose firmware images that included server-side or client-side code associated with web interfaces (e.g., `HTML`, `JS`, `PHP`, `CGI`). Once we applied all the heuristics to the firmware candidates (1580),

we tried to chroot to them and start their web interface emulation. Unfortunately, we were able to chroot to only a part of the firmware candidates (488). Then, we were able to start the embedded web interfaces only for a part of the firmware images which successfully chrooted (246). Finally, we were able to discover high impact vulnerabilities only in a part of all the web interfaces that were successfully emulated (185).

Challenges and Limitations: Inevitably, our dataset and the heuristics we apply lead to a bias. First, it almost only contains firmware images that are publicly available online. Second, Linux-based devices only account for a portion of all embedded systems. Third, because we use pre-compiled Debian Squeeze images from [56] we performed our tests mainly on ARM, MIPS and MIPSel firmware images. However, as we present in Table II, adding support for additional architecture should be straightforward and requires mainly engineering effort. For example, when pre-compiling the Debian for architectures in Table II and using mainline QEMU version with additional patches, our framework could ideally support the emulation of $\approx 97\%$ of the firmware images in our dataset. Finally, there exist firmware images running as monolithic software or embedding web servers which we currently do not detect or support. We are aware of this bias and the results herein should be interpreted without generalizing them to all embedded systems. In essence, these choices were needed to perform this study and it will be an interesting future work to extend the study to more diverse firmware images.

TABLE I: Number of firmware images and corresponding vendors at each phase of the experiment.

Dataset phase	# of FWs (unique)	# of root FS	# of vendors (unique)
Original dataset	1925	–	54
Candidates for chroot and web interface emulation	1580	1754	49
Improved by heuristics	1580	1982	49
Chroot OK	488	–	17
Web server OK	246	–	11
High impact vulnerabilities (static + dynamic)	185	–	13

TABLE II: Distribution of CPU architectures, QEMU support of those CPUs, and the success rates of chroot and web launch for each architecture. (The failure analysis is detailed in Section V-F.)

Arch.	QEMU support	Original firmware	Chroot OK	Web server OK
ARM	mainline	35%	53%	55%
MIPS	mainline	19%	21%	17%
MIPSel	mainline	17%	26%	28%
Axis CRIS	patch [52,53]	16%	–	–
bFLT	mainline	5%	–	–
PowerPC	mainline	3%	–	–
Intel 80386	mainline	2%	–	–
DLink Specific	no	$\approx 1\%$	–	–
Unknown	no	$\approx 1\%$	–	–
Altera Nios II	patch [83]	$\ll 1\%$	–	–
ARC Tangent-A5	no	$\ll 1\%$	–	–
Total	–	**1925**	**488**	**246**

V. RESULTS AND CASE STUDIES

A. Summary of Discovered Vulnerabilities

Our automated system performed both static and dynamic analysis of embedded web interfaces inside 1925 firmware images from 54 vendors. We found serious vulnerabilities in at least 45 firmware images out of those 246 for which we were able to emulate the web server. These include 225 *high impact* vulnerabilities found and verified by dynamic analysis. Static analysis reported 145 unique firmware packages to expose 9046 possible vulnerabilities. Aggregating static and dynamic analysis reports, a total of 185 firmware images are responsible for 9271 vulnerabilities, affecting nearly a quarter of vendors in our dataset.[13]

B. Static Analysis Vulnerabilities

PHP is one of the most used server-side web programming languages [34]. Over the past years, many researchers focused on investigating vulnerabilities in PHP applications and creating static analysis tools [29, 57]. However, to the best of our knowledge, we are the first to study the prevalence of PHP in embedded web interfaces and their security. In our dataset the 8% of embedded firmware images contain PHP code in their server-side. We extracted the PHP source code from those firmware packages and analyzed the code using RIPS. RIPS

---

[13]Some firmware images contribute to both static and dynamic firmware counts.

TABLE III: Distribution of web server types among the 246 started web server.

Web server	% among started web servers
minihttpd	37%
lighttpd	30%
boa	4%
thttpd	3%
empty banner	26%

TABLE IV: Web technologies used by the started web servers (combinations possible).

Web interface contains	% of started web servers
HTML	98%
CGI	57%
PHP	2%
Perl	3%
POSIX shell	11%

reported 145 unique firmware packages to contain at least one vulnerability and a total of 9046 reported issues. The detailed breakdown is presented in Table V. We observe that cross-site scripting and file manipulation constitute the majority of the discovered vulnerabilities, while command injection (one of the most serious vulnerability class) ranks third.

C. Dynamic Analysis Vulnerabilities

Our framework was able to perform dynamic security testing on 246 distinct web interfaces, and the general results are presented in Table VI. In particular, we discovered 21 firmware packages which are vulnerable to command injection. The impact of such vulnerabilities can be significant as a large number of devices may be running these firmware images (e.g., Section V-G).

Additionally, we found that 32 firmware packages are affected by XSS and 37 are vulnerable to CSRF. Even though XSS and CSRF are usually not considered to be critical vulnerabilities, they can have a high impact. For example, Bencsáth et al. [16] were able to completely compromise an embedded device only by using XSS and CSRF vulnerabilities.

The above vulnerabilities affect the firmware of multiple type of devices in our dataset, such as SOHO routers, CCTV cameras, smaller WiFi devices (e.g., SD-cards). Leveraging tools such as Shodan [68] or ZMap [35], it is possible to correlate these firmware images to populations of online devices using multiple correlation techniques [23], which we leave for future work.

In summary, we found vulnerabilities in roughly one out of four (24%) of the dynamically tested firmware images, which demonstrates the viability of our approach.

D. Evaluation of Hosting Web Interfaces

"Hosting" embedded web interfaces seems a promising approach as it permits testing a web interface without emulating the complete firmware. Indeed, many firmware images are difficult (or impossible) to emulate. We therefore tested

TABLE V: Distribution of PHP vulnerabilities reported by RIPS static analysis. (The typical error rates of each type of vulnerability can be found in [29].)

Vulnerability type	# of issues	# of affected FWs
Cross-site scripting	5000	143
File manipulation	1129	98
Command execution	938	41
File inclusion	513	40
File disclosure	461	87
SQL injection	442	10
Possible flow control	171	56
Code execution	141	21
HTTP response splitting	127	27
Unserialize	119	15
POP gadgets	4	4
HTTP header injection	1	1
Total	**9046**	**145 (unique)**

TABLE VI: Distribution of dynamic analysis vulnerabilities. (The vulnerabilities followed by a "†" sign have low severity and are known to be reported with a very high false positive rate, therefore they are not mentioned elsewhere in this paper, including when we mention a total number of vulnerabilities found.)

Vulnerability type	# of issues	# of affected FWs
Command execution	*51*	*21*
Cross-site scripting	*90*	*32*
CSRF	*84*	*37*
Sub-total HIGH impact	*225*	*45 (unique)*
Cookies w/o HttpOnly †	9	9
No X-Content-Type-Options †	2938	23
No X-Frame-Options †	2893	23
Backup files †	2	1
Application error info †	1	1
Sub-total low impact †	5843	23 (unique)
Total	**6068**	**58 (unique)**

the "hosting" approach on all the firmware images from our dataset where our web server heuristics tools could extract a document root (Section III-B2). The document roots are then "transplanted" into testing hosts containing a generic web server,[14] and then the dynamic analysis is performed with the same tools as for the emulated firmware images.

Table VII shows the high impact vulnerabilities found in this experiment and also presents a comparison with the emulation approach. We can immediately see that unsurprisingly the "hosted" analysis allows to test web interfaces from many more firmware images, but surprisingly it almost only reports CSRF vulnerabilities. In fact the technique did not allow to detect any new command injection or XSS vulnerabilities. We expect that the lack of results for some categories of vulnerabilities is due to the fact that using the "hosted" approach with a static web server configuration has some undesired side effects. For instance, the headers of the HTTP responses will be different from those of the real device's web server, while these headers may have an important security role (e.g., Cookies w/o HttpOnly, No X-Content-Type-Options,

[14]Ubuntu 14.10, Kernel 3.13, Apache2, PHP 5.5.9, Perl 5.18.2.

No X-Frame-Options). In fact, for command execution and XSS vulnerabilities we had to perform manual interventions into the hosting environment to make the web interface more functional. Then we had to rerun the dynamic analysis to discover a part of vulnerabilities already reported by the fully automated emulation approach. In few instances we had to install additional apache2 modules, while in some others we had to disable .htaccess configuration files which came with the transplanted document roots. Yet in several other cases we had to adjust a variety of shebang (#!) paths in the interpreted scripts' headers to point to the correct interpreter path of the hosting environment. These manual interventions clearly do not scale and limit the "hosted" approach. In the future, we plan to address such limitations with approaches to automatically reconfigure the "hosting" environment based on the semantics of the transplanted document root. Overall, we conclude that both firmware emulation and "hosting" web interfaces are useful and complementary techniques. Moreover, whenever the emulation is possible, it finds more vulnerabilities.

TABLE VII: Comparison of firmware images affected by high impact vulnerabilities found with the *Emulated and Chroot* method and the ones found with the *Hosted* technique. The firmware and vulnerabilities marked with a "†" are found using the *Hosted* technique, which is not yet integrated in the fully automated framework. Therefore, they are not aggregated elsewhere in this paper, including when we mention a total number of vulnerabilities and affected firmware.

	Emulated (unique FWs)	Hosted (unique FWs)	Hosted Contribution (added unique FWs)
Command execution	*21*	*15 †*	*0*
Cross-site scripting	*32*	*13 †*	*0*
CSRF	*37*	*307 †*	*269*
Total tested FWs	**246**	**515 †**	**269**
Total vulnerable FWs	**45**	**307 †**	**262**

E. HTTPS and Other Network Services

We also explored how often embedded devices provide HTTPS support. In our dataset, nearly 19% of the original firmware images contained at least one HTTPS certificate. This provides a lower bound estimate of firmware images that could provide a web server with HTTPS support.[15] Similarly, around 24% of the firmware instances starting an HTTP web server, also started an HTTPS one. We also expect this number to be a lower bound estimate as an HTTPS web server might not start for multiple reasons. It is unfortunate that so few devices provide HTTPS support.

While in this paper we focus on the security of web interfaces, we found it interesting to also report on the other network services that are automatically started during the dynamic analysis. Indeed, these additional network services which we detected using netstat[16] may be vulnerable on their own (e.g., TFTP [7], TR-069 [10], RTSP [9], Debug [8]). For this we compare the netstat output before and after

[15]As some devices may generate certificates dynamically.
[16]Scanning the virtual machine with the NMAP scanner [6] was both too slow and provided too shallow results.

starting the chroot and the init scripts. This provides a very precise information on which program is listening on which port (Table VIII).

TABLE VIII: Distribution of network services opened by 207 firmware instances out of 488 successfully emulated ones. The last entry summarizes the 16 unusual port numbers opened by services such as web, telnetd, ftp or upnp servers.

Port type	Port number	Service name	# of FWs
TCP	554	RTSP	91
TCP	555	RTSP	84
TCP	23	Telnet	60
TCP	53	DNS	23
TCP	22	SSH	15
TCP	Others	Others	58
Total			**207 (unique)**

F. Analysis of the Failures

The failures at various stages limit the coverage of the tested firmware images. For example, Table II shows that chroot failed for around 69% of the original firmware images, and around 50% of the successfully chrooted firmware packages failed to start the embedded web interface. To increase the coverage and hence the chances of finding more vulnerabilities in more firmware images, we have to perform analysis of the failures and improve our framework.

Such a *failure diagnosis* is very time consuming [64], it requires the exploration of the failure symptoms (e.g., message patterns, error codes, unstructured or inconsistent logs). However, this information differs from one system (i.e., device, firmware) to another.

Ideally, such fixes should resolve the failures permanently. However, in practice failures often reoccur [62]. One reason is that the corrective maintenance activities, failure diagnosis and solution development can take a long effort. Another reason is that the deployed solution is not completely effective. Additionally, failures can become more recurrent in older systems (e.g., old devices and firmware). Moreover, once the firmware complexity grows, human analysts become overloaded with the generated logging or failure information. Therefore, scalable failure analysis approaches are required

a) Analysis: As mentioned before, during the experiments our framework encountered 1092 firmware images where *chroot* failed, and 242 firmware emulations where *web interface* launch failed. However, this is too many failures to analyze manually, and the diversity of the systems makes automated log analysis challenging. Therefore, we performed the analysis on a sample of the data and we apply statistical methods and confidence intervals to reason about failures, their root causes and to find ways to improve the system. For this, we consider a 95% confidence level and a $\pm$ 10% confidence interval for the accuracy of estimations. Those parameters allow to provide coarse grained results while remaining within a reasonable number of failures to manually analyze.

We analyzed the log samples of 88 randomly selected (out of 1092) firmware files that *failed to chroot*. Among those

we found actual chroot failures and cases where chroot was successful but which we failed to exploit.

There were 36 cases where chroot was actually the cause of the failure (which we extrapolate to 40.9% $\pm$ 9.8% cases out of 1092). In these sampled logs we found two main reasons of chroot failure:

- Chroot failed for 10 firmware images because of *exec format errors* (extrapolated to 11.3% $\pm$ 6.3% out of 1092 firmware images). Those failures are the consequence of an incorrect guess of the CPU architecture or due to a `/bin/sh` that contain *illegal instructions*.[17] We believe that those error cases should be relatively *easy* to fix, i.e., by changing the QEMU architecture (e.g., because architecture was improperly detected in the first place) or improving QEMU (e.g., to support, or ignore, non standard instructions).

- Second, chroot failed for 26 samples because the firmware images were only *partial firmware updates* (which we extrapolate to 29.5% $\pm$ 9.1% out of 1092). Such firmware images do not contain any shell or busybox binary that our framework could chroot to. Those cases can also be fixed by replacing missing utilities, however, at this point the firmware under analysis will start to diverge from the actual device's firmware.

The remaining 52 chroot failures were found to be *false positives*. In fact, those firmware images did chroot successfully but our framework failed to detect this. This can occur in systems that set the environment in unusual ways or take long time to respond to chroot and environment queries (timeout). We therefore extrapolate that 59.1% $\pm$ 9.8% of 1092 chroot failure cases were in fact successfully chrooted. We estimate that those cases should be relatively easy to fix. The fixes could include more adaptive timeouts, and more robust handling of various shells and environments.

In summary, for the *chroot failed* failures we estimate that 62 samples should be relatively *easy* to fix, meaning that 70.4% ($\pm$ 9.1%) of the failures should be *easy* to fix fix that would allow the emulation to advance one step further.

Similarly, we analyzed log samples of 69 randomly selected (out of 242) firmware files that successfully chrooted but failed to start the web interface. We found 45 instances where *missing device* were the cause of the failure. Some examples of missing devices are `eth1`, `br0`, `/dev/gpio`, `/dev/mtdblock0`. We estimate that fixing the *missing devices* in the emulator is generally *hard*, and sometimes even impossible due to the lack of specification sheets. This also means that 65.2% $\pm$ 9.5% of those 242 web server failures are in general *hard* to fix. We also found 15 firmware samples that failed or hung during the initialization of the emulation. Some of these errors were *"Init is the parent of all processes"* and *"init: must be run as PID 1"*. The reason for such errors could be the chrooted nature of the emulation. However, we expect this not to be too difficult to fix. This translates to 21.8% $\pm$ 8.2% of original 242 web interface failure will be likely *easy* to fix. Finally, we identified 9 firmware samples

---

[17] For example, MIPS processors can have customized opcodes. This is possible by using User Defined Instruction Sets (UDIs) [51].

that reached the web interface launch but failed to launch. We therefore estimate that this is the case for 13.0% ± 6.7% of firmware images that produced *web server* failures. Examples of errors are *"(server.c.621) loading plugins finally failed"* and *"(log.c.118) opening errorlog /tmp/log/lighttpd/error.log failed: No such file or directory"*. However, we estimate this failure category can be relatively *easy* to fix.

In summary, for the *web server* failures we estimate that 24 samples should be relatively *easy* to fix, meaning that 34.8% (± 9.6%) of the failures should have *easy* fixes that would eventually allow the launch of embedded web interfaces.

b) Further improvements: The failure analysis and determination in large-scale deployments [20] can be improved and automated in several ways. One way is to perform log pre-processing [86], log mining [65] and analysis [84]. This approach often uses clustering and machine learning techniques to classify an unknown execution of the system based on its logs and based on previously seen logs of that system [22,66]. Filesystem instrumentation is another approach to automate the failure analysis [48]. Such an approach is generic because it does not assume that the system is based on particular components or existing log files. The failure causes are determined by looking at differences between file accesses (e.g., which file, when) under both normal and abnormal executions. However, these approaches assume that there exist samples of the analyzed system that runs under normal and abnormal conditions. Also, some of these approaches require domain-specific knowledge [65]. These techniques are not trivial to apply in our case. First, we aim at emulating unknown firmware images regardless of the type or application domain of the device for which the firmware is intended. Second, most of the times we do not have samples of a non-failure run of the firmware.

Another way to trace and analyze the failures is to use tracing [30]. For example, `strace` is a debugging tool that provides useful diagnostics by tracing system calls and signals. Unfortunately, `strace` is broken for stock kernels 2.6[18], which also affects the builds for embedded architectures (e.g., ARM, MIPS).

Finally, kernel level instrumentation and analysis could be a reliable approach to monitor [55] and detect the failures and their root cause. For example, `Kprobes` [61,69] can be used to dynamically monitor any kernel routine and collect debugging and performance information non-disruptively.[19] Unfortunately, Kprobes is often not enabled in default kernels we used but more importantly it's support for various CPU architectures is not stable (at least in some old kernel versions that we need to use for emulation).

These limitations do not represent themselves research challenges. However, it takes more than a trivial engineering effort to address them and overcome their effect in a systematic and generic manner. We leave the resolution of such limitations as an engineering challenge for future work.

G. Case Study: Netgear Networking Devices

In the results set we discovered many interesting cases. One of them is the case of at least 8 different device types from *Netgear*, one of the major networking equipment manufacturers. First, our dynamic analysis framework *automatically discovered* several previously unknown yet important *command injection* and *XSS* vulnerabilities that can be exploited by non-authenticated users [20]. Then, we also manually verified and confirmed that the discovered vulnerabilities are indeed exploitable on the emulated interfaces and on some physical devices we acquired for confirmation.

The affected modules are written in PHP and are used to write (i.e., store onto the hardware board) and display back to the user some manufacturer data, such as MAC address and hardware registers values. These values are supplied by the user and therefore can be controlled by the attacker. To write the manufacturer data onto the hardware board, the affected modules use the unsafe `exec` call with unsanitized input which leads to *command injection*. To display the manufacturer data back to the user the affected modules use the unsafe `echo` with unsanitized input which leads to *XSS*. The PHP modules affected by the discovered vulnerabilities do not have a clear purpose because the information they provide can be accessed from other pages of the web interface. Also, unless the user or the attacker knows the affected PHP modules names, the user cannot reach these modules by browsing the web interfaces because no other pages link to them. Therefore, one could speculate that these modules, at best, might be forgotten debugging files used during development or, at worse, could be classified as potential backdoors.

Interestingly enough, since the affected modules were written in PHP, we were also able to discover these vulnerabilities by using the static analysis tools such as RIPS. This confirms that our approach to use a combination of static and dynamic analysis (Section I-A) is sound and efficient, and combining them both can drive the dynamic analysis to focus more on particular modules that are flagged in the static analysis phase. Once again, the disadvantage of the static analysis tools is the high number of false positives and the verboseness of the output. For this reason, if we perform only the static analysis on these modules we could have missed these vulnerabilities easily.

An additional manual analysis (Table IX) revealed that those devices suffer from some more pre-authentication vulnerabilities, such as privilege escalation to web admin, unencrypted configuration storage and unauthorized configuration downloads (e.g., WPAx keys, passwords).

Additionally, by using the emulated interfaces we extracted the web interface keywords that we used to perform searches on Shodan and Google for potentially affected online devices. Shodan reported around 500 affected devices, which seem to be a small population of affected devices connected directly to the Internet via their WAN interface. However, many wireless routers are used mainly in WLANs and cannot be found from

---

TABLE IX: Distribution of vulnerabilities motivating the manual analysis (Section III-E). Firmware images relate to similar products of one vendor. (These vulnerabilities were manually found so we don't consider them when we mention a total number of vulnerabilities found by our automated framework.)

Vulnerability type	# of affected FWs
Privilege escalation (full admin)	19
Unauthorized configuration download	19
Unencrypted configuration storage	19
Total high impact	**19 (unique)**

TABLE X: List of the SHA-256 and the byte size of each PHP module that our framework automatically found vulnerable to the *command injection* and *XSS* in at least 8 device types from Netgear.

SHA-256	Byte size
03bd170b6b284f43168dcf9de905ed33ae2edd721554cebec81894a8d5bcdea5	4847
2311b6a83298833d2cf6f6d02f38b04c8f562f3a1b5eb0092476efd025fd4004	3646
325c7fe9555a62c6ed49358c27881b1f32c26a93f8b9b91214e8d70d595d89bb	4838
33a29622653ef3abc1f178d3f3670f55151137941275f187a7c03ec2acdb5caa	4922
35c60f56ffc79f00bf1322830ecf65c9a8ca8e0f1d68692ee1b5b9df1bdef7c1	4914
40fbb495a60c5ae68d83d3ae69197ac03ac50a8201d2bccd23f296361b0040b9	3582
453658ac170bda80a6539dcb6d42451f30644c7b089308352a0b3422d21bdc01	5039
4679aca17917ab9b074d38217bb5302e33a725ad179f2e4aaf2e7233ec6bc842	3638
56714f750ddb8e2cf8c9c3a8f310ac226b5b0c6b2ab3f93175826a42ea0f4545	4166
70fe0274d6616126e758473b043da37c2635a871e295395e073fb782f955840e	3544
760bde74861b6e48dcbf3e5513aaa721583fbd2e69c93bccb246800e8b9bc1e6	3684
8bf836c5826a1017b339e23411162ef6f6acc34c3df02a8ee9e6df40abe681ff	4964
9f56e5656c137a5ce407eee25bf2405f56b56e69fa89c61cdfd65f07bc6600ef	4256
a5ef01368da8588fc4bc72d3faaa20b21c43c0eaa6ef71866b7aa160e531a5b4	3791
dcefcff36f2825333784c86212e0f1b73b25db9db78476d9c75035f51f135ef6	3552

the Internet. The *WIGLE* project provides access to worldwide scans of wireless networks and can be used to detect devices of this particular vendor. At the date of this writing, the *WIGLE* project reports that several millions [21] of wireless devices from *Netgear* are deployed worldwide. However, lacking detailed information of device types in the *WIGLE* database, we cannot exactly tell how many actual devices worldwide are affected by these particular vulnerabilities.

Finally, the Table X lists the SHA-256 and the byte size of each PHP module that our framework automatically found vulnerable to the *command injection* and *XSS* we detailed above. These PHP modules can be found in more than 30 firmware firmware images that are available in total for the affected 8 device types [22].

H. Case Study: Samsung CCTV Cameras

Another interesting case study is the one of a Samsung CCTV camera model. The affected camera model has Ethernet networking, provides a web interface and multiple advanced functions (e.g., face detection and tracking). These cameras are intended for SOHO and enterprise setups, and cost several hundreds US dollars. Our system *automatically discovered* a *command injection* and multiple *XSS* vulnerabilities in their

---

[21]https://wigle.net/stats#ssidstats

[22] Following a responsible disclosure practice, we intentionally omitted the names of the affected devices and vulnerable modules.

web interfaces as follows. First, our system applied the static analysis. Since a part of the camera's web interface consists of CGI scripts written in PHP, the RIPS tool reported multipled potential vulnerabilities of different types, including *command injection* and *XSS*. interface (i.e., the web interface is not exclusively implemented by native Then, our system took advantage of the partial PHP implementation of the web binaries) and applied the *Hosted* technique (Section II-D) to transplant the camera's web interface onto a Ubuntu Linux host. Finally, our system used the report from the RIPS tool to focus the dynamic analysis phase onto the modules contained in the report. As a result, the dynamic analysis tools were able to confirm that one particular module [23] of the web interface allowed an attacker to successfully perform *command injection* and *XSS*. This once again confirms that our approach to use a combination of static and dynamic analysis (Section I-A) is sound and efficient. Also, it also confirms that "hosting" the embedded web interfaces non-natively is effective. We show it can successfully achieve discovery and confirmation of high-severity vulnerability similar to the emulation-based approaches.

VI. DISCUSSION

A. Limitations of the Emulation Techniques

Although our approach is able to discover vulnerabilities in embedded web interfaces that run inside an emulated environment, setting up such environments is not always straightforward. We discuss several limitations we encountered and outline how they could be handled in the future. In fact, many of these limitations are the results of the failures analyzed previously in Section V-F.

1) Forced Emulation: Even though most of the firmware instances in our dataset are for Linux-based devices, they are quite heterogeneous and their actual binaries vary. Examples include `init` programs that have different set of command parameters or strictly requiring to run as PID 0, which is not the case in a chrooted environment. Ideally, there should be a simple and uniform way to start the firmware, but this is not the case in practice as devices are very heterogeneous. In addition to this, unless we have access to the bootloader of each individual device, there is no reliable way to reproduce the boot sequence. Obtaining and reverse-engineering the bootloaders themselves is not trivial. This usually requires access to the device, use of physical memory dumping techniques, and manual reverse-engineering, which is outside the scope of this paper. We emulate firmware images by forcefully invoking its default initialization scripts, (e.g., `/etc/init`, `/etc/rc`), however, sometimes, these scripts do not exist or fail to execute correctly leading to an incomplete system configuration. For instance, it may fail to mount the `/etc_ro` partition at the `/etc` mount point, and then, the web server is missing some required files (e.g., `/etc/passwd`).

2) Emulated Web Server Environment: Even if the basic emulation was successful, other problems with the emulated web server environment are common. For example, an emulated web interface return for many requests the HTTP response codes `500 Internal Server Error` or `404 Not Found`. Manual inspection of the cases when code

---

[23]The details of the vulnerability can be found at http://firmware.re/vulns/

12

`500` is returned suggests that some scripts or binaries are either missing from the root filesystem or do not have proper permissions. Code `404` was often returned due to the wrong web server configuration file being loaded, which lead to the document root pointing at a wrong directory. To overcome this, we try to emulate the web interface of a firmware using all combinations of the configuration files and document roots we find in this firmware.

3) Imperfect Emulation: The ability to emulate embedded software in QEMU is incredibly valuable, but comes at a price. One big drawback is that some very basic peripheral devices are missing in the emulated environments. A very common emulation failure is related to the lack of non volatile memories (e.g., NVRAM) [28,43]. Such memories are used by embedded devices to store boot and configuration information. Several approaches to overcome such limitations exist. One is to have an universal or on-the-fly NVRAM emulator plugged into QEMU, for example instrumented at kernel-level or implemented using `Avatar` [85]. Another approach is to intercept calls to the commonly used `libnvram` functions (such as `nvram_get` and `nvram_set`) and return fake data [28,43]. While these tools are easy to compile and use, it is not trivial to automatically generate meaningful application data without producing false alerts or breaking the emulation. We plan to integrate these techniques in our future versions.

B. Outdated Firmware Versions

One concern about our approach could be that the firmware files in our experiments were not necessarily the latest available versions. This in turn could imply that the vulnerabilities we automatically discovered are not necessarily applicable to the latest versions of the affected firmware images. Although such a concern is legitimate, in practice there are several caveats to this concern that in our view still make our methodology and findings valuable.

First, it is important to know and understand how many embedded devices that are vulnerable or have outdated firmware will update their firmware in such a case. On the one hand, many embedded devices are SOHO devices which means that the users decide *if and when* they will upgrade their firmware version. On the other hand, researchers showed that even simple improvements, such as changing the default credentials of the embedded devices, are not always applied by the users during long period of times [41]. For example, it was found that 96% of accessible devices having factory default root passwords still remain vulnerable after a period of 4 months [27]. On the other hand, a firmware download and update is a more complex task than a change of the default credentials. Therefore, unless the devices are connected to the Internet and have a firmware auto-update functionality that is effectively *enabled*, it is reasonable to expect that in practice the firmware updates are applied far less than expected/desired, or are applied at best as often and as fast as the credentials are updated.

Second, even though the embedded devices should keep their firmware updated, this in not always feasible, e.g., for field-deployed devices. Such devices often cannot be remotely updated and require the physical access of an operator in the field to do so. However, even in such cases the upgrade of the firmware is not always straightforward. Cerrudo [19] showed that in some cases embedded devices could be buried in the roadway, making firmware updates that require physical access very challenging, if not impossible.

Third, even the latest firmware releases could still contain the very same vulnerabilities as the older versions [24]. Therefore, vulnerabilities discovered in older firmware versions can prove extremely useful as direct input or mutation template for testing the latest firmware versions.

In summary, we believe that a security study performed only on the latest firmware releases could provide important details for securing embedded devices (e.g., critical vulnerability discovery, patching 0-days). At the same time, however, such a study would not be completely accurate as many existing devices run outdated firmware versions. Ultimately, the goal of this work is not to find (all) the vulnerabilities in (all) the latest firmware versions. The main goal is to provide a methodology and insights that can be applied on any firmware version in order to automatically discover vulnerabilities in embedded firmware, and in particular in embedded web interfaces.

C. Manual Interventions

Our framework is designed to be as automated as possible. However, manual interventions are sometimes necessary or even desirable. First, for each newly encountered web server type we need to write a tool, which will then automatically detect, parse, and launch instances of this particular web server. Automation of such a step could be improved, for example, using `ConfigRE` [82]. Second, manual inspection of the results and of the affected software allows to confirm vulnerabilities and sometimes leads to finding new ones. This is part of the power of our methodology, i.e., pointing the finger on likely vulnerable software. In our experience this last phase was very productive as there were only a few false positives left after the dynamic analysis phase.

D. Ethical Aspects

In our study we are particularly careful to work within legal and ethical boundaries. First, we strictly follow the *responsible disclosure* policy. To this end, we try our best to notify vendors, CERTs and Vulnerability Contribution Programs (VCP) for vulnerabilities we discover during our experiments. We also try to assist vendors in reproducing these issues. Second, as previously mentioned, our framework does not operate on live embedded devices, rather on their emulations. This avoids both accessing devices we do not own and breaching the terms of use. Also, there is no risk to interfere unintentionally with devices which are not under our control or to "brick" an actual device. In limited cases when confirmation of an issue requires a physical device, we do perform such validations on devices under our control and in an isolated test environment.

VII. RELATED WORK

Analysis of embedded devices is not a new idea. Costin et al. [23] preformed a large scale analysis but they did it only through a simple static analysis. Bojinov et al. [17] studied the security of embedded management interfaces but performed the analysis manually on only 21 devices. Similar studies

were recently preformed on popular SOHO devices [47, 54] each performing manual analysis on about 10 devices and uncovering flaws in them. In contrast to these, we show that by automating the analysis we can scale to testing hundreds of devices and find thousands of vulnerabilities.

In addition, several projects scanned the Internet, or parts of it, to discover vulnerabilities in embedded systems [11, 26, 27, 44, 68, 70]. In most cases these approaches lead to discovery of devices with known vulnerabilities such as default passwords or keys, and in several notable cases helped the discovery of new flaws [44]. However, such approaches raise serious ethical problems and in general only allow to find devices that are vulnerable to known (manually found) bugs.

Web static analysis is a very active field of research, Huang et al. [50] were the first to statically search for web vulnerabilities in the context of PHP applications. Balzarotti et al. [13] showed that even if the developer performs certain sanitization on input data, XSS attacks are still possible due to the deficiencies in the sanitization routines. Pixy [57] proposed a technique based on data flow analysis for detecting XSS, SQL or command injections. RIPS [29], on the other hand, is a static code analysis tool that detects multiple types of injection vulnerabilities. While in this work we can, in principle use any of those detection mechanisms we only used RIPS which has low false positives and is still openly available.

There are several recent works that rely on emulation in order to discover or verify vulnerabilities in embedded systems. Zaddach et al. [85] proposed `Avatar`, which is a dynamic analysis framework for firmware security testing of embedded devices. `Avatar` executes embedded code inside a QEMU emulator, while the I/O requests to the peripherals of the embedded system are forwarded to the real device attached to the framework. Kammerstetter et al. [59] targeted Linux-based embedded systems that are emulated with a custom kernel which forwards `ioctl` requests to the embedded device that runs the normal kernel. Li et al. [63] proposed a hybrid firmware/hardware emulation framework to achieve confident SoC verification. Authors used a transplanted QEMU at BIOS level to directly emulate devices upon hardware. Unfortunately, those approaches are not possible to use without having access to a physical device, which does not scale as our approach does.

Meanwhile, Shoshitaishvili et al. [77] presented `Firmalice`, a static binary analysis framework to support the analysis of firmware files for embedded devices. It was demonstrated to detect three *known backdoors* in real devices, but it requires manual annotations and is therefore not possible to use in a large scale analysis.

Fong and Okun [39] took a closer look at web application scanners, and their functions and definitions. Bau et al. [15] conducted an evaluation of the state of the art of tools for automated "black box" web application vulnerability testing. While results have shown the promise and effectiveness of such tools, they also uncovered many limitations of existing tools. Similarly, Doupé et al. [33] performed an evaluation of eleven "black box" web pen-testing tools, both open-source and commercial. Authors found that crawling ability is as important and challenging as vulnerability detection techniques and many classes of vulnerabilities are completely overlooked.

They conclude that more research is required to improve the tools. Holm et al. [45] performed a quantitative evaluation of vulnerability scanning. Authors showed that automated scanning is unable to accurately identify all vulnerabilities. They also showed that scans of Linux-based hosts are less accurate than the scans of Windows-based ones. Doupé et al. [32] proposed improvements to the "black box" vulnerability testing tools. Authors observed the web application state *from the outside*. This allowed them to extend their testing coverage. Then they drove the "black box" web application vulnerability scanner. They implemented the technique in a crawler linked to a fuzzing component of an open-source web pen-testing tool. Such improvements to analysis tools will benefit our framework as we can integrate them in our analysis phase.

Finally, Gourdin et al. [42] addressed the challenges of building secure embedded web interfaces by proposing `WebDroid`, which was the first framework specifically dedicated to this purpose.

VIII. CONCLUSION AND FUTURE WORK

In this work, we presented a new methodology to perform large scale security analysis of web interfaces within embedded devices. For this, we designed the first framework that achieves *scalable and automated dynamic analysis of firmwares*, and that was precisely developed to discover vulnerabilities in embedded devices using the software-only approach. Our framework leverages off-the-shelf static and dynamic analysis tools. Because of the limitations in static analysis tools, we created a mechanism for automatic emulation of firmware images. While perfectly emulating unknown hardware will probably remain an open problem, we were able to emulate systems well enough to test the web interfaces of 246 firmware images. Our framework found serious vulnerabilities in at least 24% of the web interfaces we were able to emulate, including 225 *high impact* vulnerabilities found and verified by dynamic analysis. When counting static analysis, 9271 issues were found in 185 firmware images, affecting nearly a quarter of vendors in our dataset. These results show that some embedded systems manufacturers need to start considering security in their software life-cycle, e.g., using off-the-shelf security scanners as part of their product quality assurance.

Our work motivates the need for additional research in several areas. First, there are probably ways to improving emulation quality of unknown hardware. Second, automatically synthesizing web exploits would make vulnerability confirmation easier. Finally, responsibly disclosing vulnerabilities is time consuming and difficult (and in our experience is worse with vendors of SOHO devices). It becomes an open challenge when it needs to be performed at a large scale.

We plan to continue collecting new data and extend our analysis to all the firmware images we can access in the future. Further we want to extend our system with more sophisticated dynamic analysis techniques that allow a more in-depth study of vulnerabilities within each firmware image.

REFERENCES

[1] http://www.arachni-scanner.com/.
[2] https://code.google.com/p/zaproxy/.
[3] http://w3af.org/.

[4] http://owasp.org/index.php/Top_10_2013-A1-Injection.

[5] http://www.darrinhodges.com/chroot-voodoo/.

[6] http://nmap.org.

[7] CVE-2007-1435, CVE-2011-4821.

[8] CVE-2010-2965, CVE-2014-0659.

[9] CVE-2014-4880, CVE-2013-1606.

[10] CVE-2014-9222.

[11] Internet Census 2012 – Port scanning /0 using insecure embedded devices. http://internetcensus2012.bitbucket.org.

[12] T. Avgerinos, S. K. Cha, B. L. T. Hao, and D. Brumley. AEG: Automatic Exploit Generation. In *ISOC Network and Distributed System Security Symposium (NDSS)*, 2011.

[13] D. Balzarotti, M. Cova, V. Felmetsger, N. Jovanovic, E. Kirda, C. Kruegel, and G. Vigna. Saner: Composing Static and Dynamic Analysis to Validate Sanitization in Web Applications. In *IEEE Symposium on Security and Privacy*, 2008.

[14] A. Barth, C. Jackson, and J. C. Mitchell. Robust Defenses for Cross-Site Request Forgery. In *ACM Conference on Computer and Communications Security (CCS)*, 2008.

[15] J. Bau, E. Bursztein, D. Gupta, and J. Mitchell. State of the Art: Automated Black-Box Web Application Vulnerability Testing. In *IEEE Symposium on Security and Privacy*, 2010.

[16] B. Bencsáth, L. Buttyán, and T. Paulik. XCS Based Hidden Firmware Modification on Embedded Dievices. In *International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*, 2011.

[17] H. Bojinov, E. Bursztein, E. Lovett, and D. Boneh. Embedded management interfaces: Emerging massive insecurity. *BlackHat USA*, 2009.

[18] S. W. Boyd and A. D. Keromytis. SQLrand: Preventing SQL Injection Attacks. In *Applied Cryptography and Network Security*, 2004.

[19] C. Cerrudo. Hacking US (and UK, Australia, France, etc.) – Traffic Control Systems. http://blog.ioactive.com/2014/04/hacking-us-and-uk-australia-france-etc.html, Apr 2014.

[20] M. Y. Chen, E. Kiciman, E. Fratkin, A. Fox, and E. Brewer. Pinpoint: Problem determination in large, dynamic internet services. In *Dependable Systems and Networks, 2002. DSN 2002. Proceedings. International Conference on*, pages 595–604. IEEE, 2002.

[21] S. Christey and R. A. Martin. Vulnerability type distributions in CVE. *Mitre Report*, 2007.

[22] I. Cohen, J. S. Chase, M. Goldszmidt, T. Kelly, and J. Symons. Correlating Instrumentation Data to System States: A Building Block for Automated Diagnosis and Control. In *OSDI*, volume 4, pages 16–16, 2004.

[23] A. Costin, J. Zaddach, A. Francillon, and D. Balzarotti. A Large-Scale Analysis of the Security of Embedded Firmwares. In *USENIX Security Symposium*, 2014.

[24] A. Cui, M. Costello, and S. J. Stolfo. When Firmware Modifications Attack: A Case Study of Embedded Exploitation. In *ISOC Network and Distributed System Security Symposium (NDSS)*, 2013.

[25] A. Cui, J. Kataria, and S. J. Stofo. From Prey to Hunter: Transforming Legacy Embedded Devices into Exploitation Sensor Grids. In *Annual Computer Security Applications Conference (ACSAC)*, 2011.

[26] A. Cui, Y. Song, P. V. Prabhu, and S. J. Stolfo. Brave New World: Pervasive Insecurity of Embedded Network Devices. In *International Symposium on Recent Advances in Intrusion Detection (RAID)*, 2009.

[27] A. Cui and S. J. Stolfo. A Quantitative Analysis of the Insecurity of Embedded Network Devices: Results of a Wide-area Scan. In *Annual Computer Security Applications Conference (ACSAC)*, 2010.

[28] Z. Cutlip. Emulating and Debugging Workspace. http://shadow-file.blogspot.fr/2013/12/emulating-and-debugging-workspace.html.

[29] J. Dahse and T. Holz. Simulation of Built-in PHP Features for Precise Static Code Analysis. In *ISOC Network and Distributed System Security Symposium (NDSS)*, 2014.

[30] X. Ding, H. Huang, R. B. Jennings, Y. Roan, S. Sahu, and A. A. Shaikh. Automatic software fault diagnosis by exploiting application signatures, Jan 2011. US Patent 7,877,642.

[31] A. Doupé, B. Boe, C. Kruegel, and G. Vigna. Fear the EAR: Discovering and Mitigating Execution After Redirect Vulnerabilities. In *ACM Conference on Computer and Communications Security (CCS)*, 2011.

[32] A. Doupé, L. Cavedon, C. Kruegel, and G. Vigna. Enemy of the State: A State-Aware Black-Box Web Vulnerability Scanner. In *USENIX Security Symposium*, 2012.

[33] A. Doupé, M. Cova, and G. Vigna. Why Johnny Can't Pentest: An Analysis of Black-box Web Vulnerability Scanners. In *Detection of Intrusions and Malware, and Vulnerability Assessment (DIMVA)*. 2010.

[34] L. Duflot, Y.-A. Perez, and B. Morin. Netcraft. PHP Usage Stats. http://www.php.net/usage.php, June 2007.

[35] Z. Durumeric, E. Wustrow, and J. A. Halderman. ZMap: Fast Internet-wide Scanning and Its Security Applications. In *USENIX Security Symposium*, 2013.

[36] F. B. et al. QEMU – Quick EMULator. http://www.qemu.org.

[37] V. Felmetsger, L. Cavedon, C. Kruegel, and G. Vigna. Toward automated detection of logic vulnerabilities in web applications. In *USENIX Security Symposium*, 2010.

[38] Firehost. The Superfecta Report Special Edition. *Superfecta Report*, 2013.

[39] E. Fong and V. Okun. Web Application Scanners: Definitions and Functions. In *Annual Hawaii International Conference on System Sciences (HICSS)*, 2007.

[40] D. Geer. Cybersecurity as Realpolitik. *BlackHat*, 2014.

[41] B. Ghena, W. Beyer, A. Hillaker, J. Pevarnek, and J. A. Halderman. Green Lights Forever: Analyzing the Security of Traffic Infrastructure. In *USENIX Workshop on Offensive Technologies (WOOT)*, 2014.

[42] B. Gourdin, C. Soman, H. Bojinov, and E. Bursztein. Toward Secure Embedded Web Interfaces. In *USENIX Security Symposium*, 2011.

[43] C. Heffner. Emulating NVRAM in Qemu. http://www.devttys0.com/2012/03/emulating-nvram-in-qemu/.

[44] N. Heninger, Z. Durumeric, E. Wustrow, and J. A. Halderman. Mining Your Ps and Qs: Detection of Widespread Weak Keys in Network Devices. In *USENIX Security Symposium*, 2012.

[45] H. Holm, T. Sommestad, J. Almroth, and M. Persson. A quantitative evaluation of vulnerability scanning. *Information Management & Computer Security*, 19(4):231–247, 2011.

[46] P. Hooimeijer, B. Livshits, D. Molnar, P. Saxena, and M. Veanes. Fast and Precise Sanitizer Analysis with BEK. In *USENIX Security Symposium*, 2011.

[47] HP-Fortify-ShadowLabs. Report: Internet of Things Research Study. http://h20195.www2.hp.com/V2/GetDocument.aspx?docname=4AA5-4759ENW, 2014.

[48] L. Huang and K. Wong. Assisting failure diagnosis through filesystem instrumentation. In *Proceedings of the 2011 Conference of the Center for Advanced Studies on Collaborative Research*, pages 160–174, 2011.

[49] Y.-W. Huang, S.-K. Huang, T.-P. Lin, and C.-H. Tsai. Web Application Security Assessment by Fault Injection and Behavior Monitoring. In *International Conference on World Wide Web (WWW)*, 2003.

[50] Y.-W. Huang, F. Yu, C. Hang, C.-H. Tsai, D.-T. Lee, and S.-Y. Kuo. Securing Web Application Code by Static Analysis and Runtime Protection. In *International Conference on World Wide Web (WWW)*, 2004.

[51] P. Ienne and R. Leupers. *Customizable embedded processors: design technologies and applications*. Academic Press, 2006.

[52] E. Iglesial. CRIS target port of Qemu. http://repo.or.cz/qemu/cris-port.git.

[53] E. Iglesial. Status of CRIS Architecture Support in Linux Kernel. https://lkml.org/lkml/2014/9/15/1082.

[54] Independent Security Evaluators. SOHO Network Equipment (Technical Report), 2013.

[55] T. Jarboui, J. Arlat, Y. Crouzet, and K. Kanoun. Experimental analysis of the errors induced into linux by three fault injection techniques. In *Dependable Systems and Networks, 2002. DSN 2002. Proceedings. International Conference on*, pages 331–336. IEEE, 2002.

[56] A. Jarno. Debian pre-compiled images for QEMU. https://people.debian.org/~aurel32/qemu/.

[57] N. Jovanovic, C. Kruegel, and E. Kirda. Pixy: A Static Analysis Tool for Detecting Web Application Vulnerabilities (Short Paper). In *IEEE Symposium on Security and Privacy*, 2006.

[58] N. Jovanovic, C. Kruegel, and E. Kirda. Static analysis for detecting taint-style vulnerabilities in web applications. *Journal of Computer Security*, 18(5):861–907, 2010.

[59] M. Kammerstetter, C. Platzer, and W. Kastner. PROSPECT – Peripheral Proxying Supported Embedded Code Testing. In *ACM Symposium on Information, Computer and Communications Security (ASIACCS)*, 2014.

[60] A. Klein. Divide and Conquer: HTTP Response Splitting, Web Cache Poisoning Attacks and Related Topics. *Sanctum whitepaper*, 2004.

[61] R. Krishnakumar. Kernel korner: kprobes-a kernel debugger. *Linux Journal*, 2005(133):11, 2005.

[62] I. Lee and R. K. Iyer. Diagnosing rediscovered software problems using symptoms. *Software Engineering, IEEE Transactions on*, 26(2):113–

127, 2000.

[63] H. Li, D. Tong, K. Huang, and X. Cheng. FEMU: A Firmware-Based Emulation Framework for SoC Verification. In *IEEE/ACM/IFIP International Conference on Hardware/Software Codesign and System Synthesis (CODES+ ISSS)*, 2010.

[64] J. Li, T. Stålhane, J. M. Kristiansen, and R. Conradi. Cost drivers of software corrective maintenance: An empirical study in two companies. In *Software Maintenance (ICSM), 2010 IEEE International Conference on*, pages 1–8. IEEE, 2010.

[65] C. Lim, N. Singh, and S. Yajnik. A log mining approach to failure analysis of enterprise telephony systems. In *Dependable Systems and Networks With FTCS and DCC, 2008. DSN 2008. IEEE International Conference on*, pages 398–403. IEEE, 2008.

[66] T.-T. Y. Lin and D. P. Siewiorek. Error log analysis: statistical modeling and heuristic trend analysis. *Reliability, IEEE Transactions on*, 39(4):419–432, 1990.

[67] B. Livshits and S. Chong. Towards Fully Automatic Placement of Security Sanitizers and Declassifiers. In *ACM Symposium on Principles of Programming Languages (POPL)*, 2013.

[68] J. Matherly. SHODAN – Computer Search Engine. http://www.shodan.io.

[69] A. Mavinakayanahalli, P. Panchamukhi, J. Keniston, A. Keshavamurthy, and M. Hiramatsu. Probing the guts of kprobes. In *Linux Symposium*, volume 6, 2006.

[70] H. Moore. Security Flaws in Universal Plug and Play: Unplug, Don't Play. *Rapid7, Ltd.*, 2013.

[71] M. Niemietz and J. Schwenk. Owning your home network: Router security revisited. *CoRR*, abs/1506.04112, 2015.

[72] J. Pewny, B. Garmany, R. Gawlik, C. Rossow, and T. Holz. Cross-Architecture Bug Search in Binary Executables . In *IEEE Symposium on Security and Privacy*, San Jose, CA, May 2015.

[73] J. Prescatore. Gartner, quoted in ComputerWorld, 2005.

[74] M. Samuel, P. Saxena, and D. Song. Context-Sensitive Auto-Sanitization in Web Templating Languages Using Type Qualifiers. In *ACM Conference on Computer and Communications Security (CCS)*, 2011.

[75] P. Saxena, D. Molnar, and B. Livshits. SCRIPTGARD: Automatic Context-Sensitive Sanitization for Large-Scale Legacy Web Applica-

tions. In *ACM Conference on Computer and Communications Security (CCS)*, 2011.

[76] F. Schuster and T. Holz. Towards reducing the attack surface of software backdoors. In *ACM Conference on Computer and Communications Security (CCS)*, 2013.

[77] Y. Shoshitaishvili, R. Wang, C. Hauser, C. Kruegel, and G. Vigna. Firmalice: Automatic Detection of Authentication Bypass Vulnerabilities in Binary Firmware. In *ISOC Network and Distributed System Security Symposium (NDSS)*, 2015.

[78] Z. Su and G. Wassermann. The Essence of Command Injection Attacks in Web Applications. In *ACM Symposium on Principles of Programming Languages (POPL)*, 2006.

[79] Tripwire vulnerability and exposure research team (VERT). SOHO Wireless Router (in)Security. White Paper, 2014.

[80] J. Viega and H. Thompson. The state of embedded-device security (spoiler alert: It's bad). *IEEE Security & Privacy*, 10(5):68–70, 2012.

[81] P. Vogt, F. Nentwich, N. Jovanovic, E. Kirda, C. Kruegel, and G. Vigna. Cross Site Scripting Prevention with Dynamic Data Tainting and Static Analysis. In *ISOC Network and Distributed System Security Symposium (NDSS)*, 2007.

[82] R. Wang, X. Wang, K. Zhang, and Z. Li. Towards automatic reverse engineering of software security configurations. In *ACM Conference on Computer and Communications Security (CCS)*, 2008.

[83] C. Wulff. Altera NiosII Support. https://lists.gnu.org/archive/html/qemu-devel/2012-09/msg01229.html.

[84] W. Xu, L. Huang, A. Fox, D. Patterson, and M. I. Jordan. Detecting large-scale system problems by mining console logs. In *Proceedings of the ACM SIGOPS 22nd symposium on Operating systems principles*, pages 117–132. ACM, 2009.

[85] J. Zaddach, L. Bruno, A. Francillon, and D. Balzarotti. Avatar: A Framework to Support Dynamic Security Analysis of Embedded Systems' Firmwares. In *ISOC Network and Distributed System Security Symposium (NDSS)*, 2014.

[86] Z. Zheng, Z. Lan, B. H. Park, and A. Geist. System log pre-processing to improve failure prediction. In *Dependable Systems & Networks, 2009. DSN'09. IEEE/IFIP International Conference on*, pages 572–577. IEEE, 2009.