Developing a Web-based Training Platform for IT Security Education

Mandy Knöchel[1], Sebastian Karius[1] and Sandro Wefel[1]

Abstract: With the increasing importance of IT security in society and economy, the need to train computer science students in IT security is becoming more and more important. In the process of developing a new IT security undergraduate course, it was found that the current training concepts used in existing master courses were not suitable as it was too much effort for the undergraduate students to set up their own practical exercise environments. To remedy this problem, we have built a web-based exercise platform for IT security training using Docker containers. This enables the students to easily use prebuilt exercise environments specifically designed for the new course. The platform was evaluated during the winter semester 2020/21 through a small survey among the students. The feedback indicated that the platform is convenient to use and helped the students during the course and in exam preparation.

Keywords: IT Security, E-Learning, Education, Virtual Laboratory, Docker

1 Introduction

The field of IT security is becoming increasingly important in our society and the economy. In order to reflect the importance of this field, the Institute for Computer Science at the Martin Luther University Halle-Wittenberg is offering a new undergraduate course "Fundamentals and Practice of IT Security", which is intended to give students a first insight into the subject. The course is designed as an elective course for the 5th semester of the undergraduate computer science program and related programs such as Bioinformatics. Up to this point, training in the field of IT security was only available in the master's degree program.

The introduction of this new course comes with new requirements. Besides teaching the theoretical basics, a large part of the training in IT security consists of practical exercises. Here, students are put in the position of an attacker and carry out attacks on specially designed example applications. The students are thus given an insight into the approach of real attackers so that they can then learn how to defend against these attacks. The exercises in the undergraduate course cover attacks ranging from SQL injection over cross-site scripting to network-based attacks like man-in-the-middle. So far, practical exercises in the master's degree courses have been carried out using virtual machines. These are pre-made training machines which are available on various platforms on the internet [Vu21, Pe21] and can be run by the students on their computers. This comes

[1] Martin Luther University Halle-Wittenberg, Institute for Computer Science, 06099 Halle (Saale), {mandy.knoechel, sebastian.karius, sandro.wefel}@informatik.uni-halle.de

with some disadvantages. Since these exercise machines must be set up by the students independently, knowledge in computer networks, virtual machines and programs such as VirtualBox is necessary. Due to the degree course scheme, networking fundamentals are not taught prior to the IT security course making it necessary that only basic networking knowledge is required. Many students further lack prior knowledge of virtualization technology. To make matters worse, the pre-made training machines are generally not designed to track completion. Often, even the corresponding solution is given directly. Also, many of these machines are not designed for beginners and require knowledge of special attack tools, exceeding the level of this course. Hence, there was a need for a way to create custom practice exercises and provide students easy and efficient access to them, even with multiple students working at the same time such as during a seminar.

To provide easier access for undergraduate students, we developed a new web-based training platform for practical security training. The platform provides each student with their own virtual training containers, requiring only a web browser to complete most of the tasks. This makes the platform not only suitable for local classes but also for remote learning approaches. This paper aims to provide an insight into the design and operation of this new training platform, as well as describing experiences in practical use during the undergraduate course. Furthermore, a survey was conducted at the end of the course in the winter semester 2020/21 in order to assess the experiences of the students.

2 Related Work

Practical training in IT security is handled very differently at various universities around the world. The use of virtual machines is very common in this context [WTM11, SFV19], although the design of the tasks can differ considerably. Besides the approach chosen by us, where the students take over the role of the attacker, there are also concepts where the students have to secure given programs against possible attacks [MC14] or examine entire applications as penetration testers [Ro20]. However, these approaches are difficult for students with little to no prior knowledge.

Besides the use of virtual machines, another approach teaching IT security is to prepare students for participation in hacking competitions or Capture-The-Flag (CTF) contests [CCC15, VSC20]. These CTF contests are competitions where participants compete in teams in which they must either face other teams in an attack/defence scenario or attack given applications to capture hidden flags (Jeopardy-style CTF). However, the high level of difficulty of these competitions and the high knowledge barrier can cause students to quickly become demotivated and possibly give up [CC14]. Therefore, this approach may only be considered for advanced courses.

In addition to courses at universities, there are also a number of online learning websites that aim to teach IT security skills [Pe21, Ha21]. However, a fee is often charged for full access to all functions and exercises. Furthermore, a certain level of basic knowledge is usually required here as well, which makes those platforms not suitable for our case.

3 System Architecture

To get an impression of the use and appearance of the platform, we have set up a demo course[2] with some exercises. The landing page provides an overview of all available exercises for the course. When clicking on an exercise a detailed view is presented with a description on how to perform the task. Starting the exercise causes a virtual container to be launched in the background and the URL to the container web page is shown. To track the progress of each user, we employ the use of a secret key which is revealed upon completing the task. If the user enters the correct key, the exercise is marked as solved.

Our system consists of three components: a web server, a database and a Docker host. User information, i.e. which containers are available for and which container is finished by each user is stored in the database. Information about the container itself, i.e. the content and state of the container is stored by the Docker host and is accessible via an HTTP API provided by the Docker host. Via this API the web server is also able to send commands to the Docker host. We decided to use Docker because it has two main advantages over other solutions. One is the extensive API, as the state of the containers can be easily retrieved and manipulated via HTTP requests. Therefore, it is not necessary to additionally store the state of the containers separately. A second advantage is the better performance of containers, as Docker uses a task-based virtualization, also called containerization, which is more lightweight compared to other virtualizations like KVM used for example in VirtualBox. The difference is, that containerization does not require a hypervisor which means the container does not contain an operating system, hence the difference between container (task-based) and virtual machine (operating system based). This results in faster container creation, faster deletion, and most importantly, the ability to launch a large number of containers simultaneously which is essential for live courses. Besides Docker, LXC also uses containerization similar to Docker. However, in a previous test using LXC via Proxmox[3], it performed significantly worse than Docker and it does not provide an HTTP API as described above.

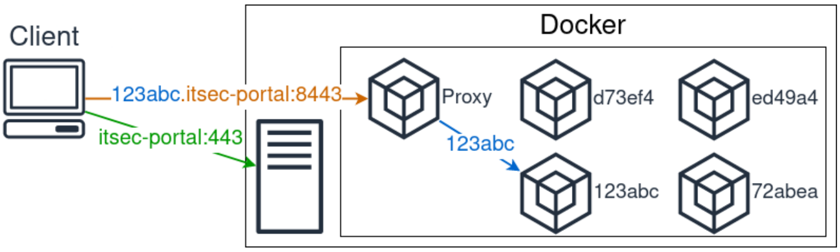

Fig. 1: Path of requests to the web server (green) and inside the Docker host (orange) with the subdomain (blue) to identify the target container.

Most of the exercise containers provide a web server with a vulnerable website. Fig. 1

[2] ITSec-Portal, german demo course: https://itsec-portal.informatik.uni-halle.de (user: demo, pass: ge3rz59gw)
[3] Proxmox is an open-source server virtualization management platform, https://www.proxmox.com

shows the path of a request to the website of the exercise platform (green) and to the individual exercise containers (orange). We set up a proxy container that redirects the traffic to the container created for the exercise which is indicated by a special subdomain. We use subdomains instead of sub-paths because most publicly available vulnerable applications we want to provide for the students are not designed to work on a sub-path. Subdomains, on the other hand, work just fine. To get a unique subdomain we calculate a hash based on the user information which doubles as the name for the container. Since Docker is able to resolve container names to container IP addresses the proxy only needs to extract the subdomain resp. the container name and use it as the URL for the proxy target. Non-web-based exercises such as man-in-the-middle attacks require more than one container and a separated network to allow the containers to communicate with other containers of the same exercise but not with containers of other exercises. Additionally, the students must be able to reach this network from the outside. Fortunately, Docker provides the option to create virtual networks between containers. To allow the students to reach the virtual network we created a VPN container. At creation time the VPN port of the container will be bound to a free port on the host. This way the student can use a VPN tool to connect to the VPN container and be tunnelled inside the virtual network. Alternatively, we provide a container that has the required attack tools already installed accessible via web-CLI.

4 Evaluation

The exercise platform was utilized for the first time with the introduction of the undergraduate course in winter 2018/19 and has since been used annually in the winter semester course. The lecture in winter 2020/21 was special as the course had to be taught fully online for the first time due to the Corona pandemic. The exercise platform played a decisive role in the success of the online course, as the students were able to perform the exercises independently at home without having access to the computer pool otherwise used for the practical seminars.

The undergraduate course consists of a theoretical part in form of a lecture and a practical seminar, both conducted weekly. The seminars demonstrate the attacks in practice and teach the students attack techniques and tools. The platform has proven to be very beneficial in this regard, as the students can experiment with the demonstrated techniques in real-time during the seminar using dedicated practice machines, without having to spend time setting up the exercises. The seminar has been conducted with up to 30 students at a time, with no noticeable impact on the performance of the exercise platform. The time to work on the exercises for the respective subject area was in general one week. Points were awarded for successful completion of the tasks, of which 60 % were required to achieve the course credits. The tasks were worked on independently by the students at home using the exercise platform, whereby questions regarding the content of the exercises or the platform itself could be addressed to the instructor at any time by e-mail or in the forum of the e-learning platform provided by the university.

Question	**0**	**1**	**2**	**3**	**4**
I quickly understood how the platform works.	**82 %**	18 %	0 %	0 %	0 %
The platform made it easier to participate in the seminars.	**59 %**	23 %	6 %	0 %	0 %
The exercises were useful for understanding the topics.	**94 %**	0 %	6 %	0 %	0 %
The platform helped me to prepare for the final exam.	**59 %**	18 %	12 %	0 %	0 %
The exercise platform should also be used in other courses.	**65 %**	12 %	6 %	0 %	0 %

Tab. 1: Numerical feedback about the exercise platforms (N=17). Rows do not always add up to 100 %, because not every question was answered by every student.
(0 = Agree, 1 = Somewhat Agree, 2 = Neutral, 3 = Somewhat Disagree, 4 = Disagree)

To better assess the students' experience with the exercise platform, a survey was conducted at the end of the latest course. Of the 29 students who participated until the end of the course, 17 students completed the survey. The respondents were asked to answer questions about their personal impressions of the course and the exercise platform, assess their own prior knowledge and report technical problems. The participants of the survey consisted of 11 computer science students and 6 bioinformatics students, most of whom were in their 3rd to 7th academic semester, with the median being the 5th semester. As expected, many students estimated their prior knowledge in the field of IT security as rather low with 60 % of the students declaring none or little[4]. Overall, the students' feedback on the training platform was very positive. 88 % reported no technical problems at all, only one student describing issues with an older browser version. Table 1 shows the results of the numerical feedback questions regarding the training platform on a 5-point Likert-scale. The majority of students indicated that the training platform was easy to use and very helpful for the practical exercises. This is also reflected in the open-ended feedback responses, where the students were asked to describe both positive and negative impressions. 12 students gave open-ended feedback with examples of positive feedback being "*The practical exercises were fun and gave me a better understanding of the material*" and "*I liked the game environment for trying out the security vulnerabilities*". Criticism was expressed by some students regarding the graphical interface of the exercise platform, where the overview of the latest tasks was perceived as somewhat unclear. The difficulty regarding some exercises was described as too high by a few students, leading to them being unable to solve the tasks successfully. In our opinion, this is due to the varying levels of prior knowledge of the students, although a reduction in the level of difficulty may lead to other students feeling underchallenged.

[4] Survey results for the question “How much prior knowledge did you have in the subject of IT security before the course?” (N=17): 29.4 % none, 29.4 % little, 35.3 % medium, 5.9 % much, 0 % very much

5 Conclusion and Future Work

In this paper, we introduced our new web-based exercise platform for IT security training, which was used during a new undergraduate course for IT security. The platform uses Docker containers to provide each student with their own training environment, eliminating the need for students to set up exercises themselves and requiring only a web browser for most tasks. The use of Docker has proven to be very beneficial, as container virtualization requires fewer resources, allowing many students to work with the platform simultaneously. The Platform was evaluated in a survey with the students in the winter semester 2020/21, where it was shown that the platform supports the students during the practical exercises and they furthermore enjoy solving them. Future plans include using the platform for other IT security courses, especially in the master's program and exploring concepts of gamification such as high scores or achievements to better motivate the students.

Bibliography

[CC14] Chung, K.; Cohen, J.: Learning Obstacles in the Capture The Flag Model. In: 2014 USENIX Summit on Gaming, Games, and Gamification in Security Education, 2014.

[CCC15] Carlisle, M.; Chiaramonte, M.; Caswell, D: Using CTFs for an Undergraduate Cyber Education. In: 2015 USENIX Summit on Gaming, Games, and Gamification in Security Education, 2015.

[Ha21] Hack The Box: Hacking Training For The Best, https://www.hackthebox.eu/, accessed: 28/06/2021.

[MC14] Martini, B.; Choo, K.-K. R.: Building the Next Generation of Cyber Security Professionals. In: 22nd European Conference on Information Systems (ECIS 2014), pp. 1-13, 2014.

[Pe21] PentesterLab: Learn Web Penetration Testing: The Right Way, https://pentesterlab.com/, accessed: 28/06/2021.

[Ro20] Robles-Gómez, A. et al.: Emulating and Evaluating Virtual Remote Laboratories for Cybersecurity. Sensors, 11/20:3011, 2020.

[SFV19] Sigholm, J.; Falco, G.; Viswanathan, A.: Enhancing Cybersecurity Education through High-Fidelity Live Exercises (HiFLiX). In: Proceedings of the 52nd Hawaii International Conference on System Sciences, pp. 7553-7562, 2019.

[VSC20] Vykopal, J.; Švábenský, V.; Chang, E.-C.: Benefits and Pitfalls of Using Capture The Flag Games in University Courses. In: Proceedings of the 51st ACM Technical Symposium on Computer Science Education, SIGCSE '20, pp. 752-758, 2020.

[Vu21] VulnHub ~ Vulnerable by Design, https://www.vulnhub.com/, accessed: 28/06/2021.

[WTM11] Willems, C.; Tringides, O.; Meinel, C.: Practical IT Security Education with Tele-Lab. UPGRADE: The European Journal for the Informatics Prof., 12(5), pp. 145-152, 2011.