

# Scalable Learning Environments for Teaching Cybersecurity Hands-on

Jan Vykopal\*, Pavel Čeleda\*, Pavel Seda\*, Valdemar Švábenský\*, and Daniel Tovarňák†

\*Faculty of Informatics, Masaryk University, Brno, Czech Republic

†Institute of Computer Science, Masaryk University, Brno, Czech Republic

{vykopal | celeda | seda | svabensky}@fi.muni.cz, tovarnak@ics.muni.cz

## I. INTRODUCTION

**Abstract**—This Innovative Practice full paper describes a technical innovation for scalable teaching of cybersecurity hands-on classes using interactive learning environments. Hands-on experience significantly improves the practical skills of learners. However, the preparation and delivery of hands-on classes usually do not scale. Teaching even small groups of students requires a substantial effort to prepare the class environment and practical assignments. Further issues are associated with teaching large classes, providing feedback, and analyzing learning gains. We present our research effort and practical experience in designing and using learning environments that scale up hands-on cybersecurity classes. The environments support virtual networks with full-fledged operating systems and devices that emulate real-world systems.

The classes are organized as simultaneous training sessions with cybersecurity assignments and learners' assessment. For big classes, with the goal of developing learners' skills and providing formative assessment, we run the environment locally, either in a computer lab or at learners' own desktops or laptops. For classes that exercise the developed skills and feature summative assessment, we use an on-premises cloud environment. Our approach is unique in supporting both types of deployment. The environment is described as code using open and standard formats, defining individual hosts and their networking, configuration of the hosts, and tasks that the students have to solve. The environment can be repeatedly created for different classes on a massive scale or for each student on-demand. Moreover, the approach enables learning analytics and educational data mining of learners' interactions with the environment. These analyses inform the instructor about the student's progress during the class and enable the learner to reflect on a finished training. Thanks to this, we can improve the student class experience and motivation for further learning.

Using the presented environments KYPO Cyber Range Platform and Cyber Sandbox Creator, we delivered the classes on-site or remotely for various target groups of learners (K-12, university students, and professional learners). The learners value the realistic nature of the environments that enable exercising theoretical concepts and tools. The instructors value time-efficiency when preparing and deploying the hands-on activities. Engineering and computing educators can freely use our software, which we have released under an open-source license. We also provide detailed documentation and exemplary hands-on training to help other educators adopt our teaching innovations and enable sharing of reusable components within the community.

**Index Terms**—cybersecurity education, interactive learning environment, sandbox, virtual machines, cyber range, educational data mining, learning analytics, learning technology

Cybersecurity is a complex professional domain involving technology, people, information, and processes to enable assured operations [1]. The dependence of modern society on digital technologies rapidly increased the demand for skilled workers. As a result, cybersecurity has become a focus area of K-12 and higher education, as well as professional learning. Educational institutions of all types expand their program and course offerings in cybersecurity for students and introduce re-skilling opportunities for mid-career professionals.

We support these efforts by designing and providing scalable environments for in-person or remote hands-on learning. These environments emulate real-world systems, applications, and infrastructures using virtual networks with full-fledged operating systems, devices, and applications used in authentic workplace settings. Learners' interaction with the emulated systems is driven by a serious game or a step-by-step tutorial facilitated by the learning environment with or without a human instructor's assistance. Regardless of the session's goal, the environments provide learning analytics. Students are informed about their progress during and after the session, which enables them to reflect on their learning. Instructors receive evidence of students' learning required for summative assessment, either a practical exam or a competition.

The environments can be deployed locally (either at individual hosts in a computer lab or at learners' desktops or laptops) or in a cloud. This flexibility enables many students to learn in a small environment or fewer students to learn in an extensive or complex environment. In both cases, students can practice from their school, workplace, home, or other places connected to the Internet. The environment can be repeatedly created for different classes on a massive scale or for each student on-demand.

From an instructor's perspective, the environments are described as code using open and standard formats, definitions of individual hosts and their networking, configuration of the hosts, and tasks that the students solve. Consequently, these components can be reused in other learning technologies.

Our environments have been iteratively improved and used in practice since 2013; the current, third generation is being developed since 2018. The learners reported they value the realistic nature of the environments that enable exercising theoretical concepts and tools in a safe way. The instructors stated that

our approach saves their time required to prepare and deploy hands-on activities. We have released the presented interactive learning environments as two open-source projects [2], [3] so that other instructors can freely use them. We also provide detailed documentation and exemplary hands-on training to help others adopt our approach. The educators can thus focus on the topic and training itself and do not need to spend extra time setting up the learning environment.

This article is divided into seven sections. Section II summarizes the state of the art of environments for teaching and learning cybersecurity skills. Section III identifies building blocks of a cybersecurity hands-on class. Section IV introduces reusable components of the building blocks, which we use to deliver our classes at scale. Section V details the environments we have been developing using the described building blocks. Section VI reports our experience from using these environments in teaching practice. Section VII concludes the paper.

## II. RELATED WORK

Cybersecurity professionals require specialized education and hands-on training sessions to improve their skills and knowledge to secure the society. To enable the hands-on training sessions, cyber ranges and their lightweight alternatives are a valuable tool and a catalyst in these efforts.

### A. Cyber Ranges

Cyber ranges are interactive learning environments representing the network systems, tools, and applications. The main benefits of cyber ranges are providing (i) performance-based learning and assessment, (ii) real-time feedback, (iii) simulated on-the-job experience, (iv) an authentic environment where teams can collaborate to improve team capabilities [4], [5].

In the last decade, a number of commercial and non-commercial cyber ranges were developed, e.g., Cyber Range Instantiation System, DETER, ELTA Cyber Academy, Airbus Cyber Range, Silensec Cyber Range, and Circadence [6]–[9]. The comprehensive summary of cyber ranges is presented in [10]–[12]. These cyber ranges mainly differ in technologies, integrated capabilities, and functionalities.

Yamin et al. [12] classify the functions and capabilities of cyber ranges into the following taxonomy (i) *scenario*; (ii) *monitoring*; (iii) *scoring*; (iv) *management*; and (v) *teaming*. The scenario defines the execution environment and the execution steps. The monitoring includes the methods and tools supporting real-time data collection for further analysis. Scoring defines a way to evaluate the training participant's behavior. The management defines the roles and resources required in the system. Teaming includes an individual and a group of individuals with specific roles in the cybersecurity training (e.g., red teams or blue teams). This survey found that only a few publicly available cyber ranges and security testbeds support all of these functions.

Pham et al. in [13] propose a Cyber Range Instantiation System for large-scale scenarios with a huge number of participants. They discuss that traditional approaches using

dedicated and isolated physical computer infrastructure are inefficient. The main drawbacks of the traditional approach are expensive creation, maintenance, and low scalability in cases with tens, hundreds, or more participants. Being aware of this, they constructed virtual cyber range environments using the Kernel-based Virtual Machine (KVM) virtualization platform. They present their solution as unique in terms of security and other pre-installed features. However, they conclude their cyber range is not prepared for massive-scale environments that involve hundreds or thousands of virtual machines (VM).

Urias et al. in [14] discuss the limitations and needs for cyber range deployments. They highlight the rising difficulty of the simulation scenarios used in training sessions. Together with the expected functions and capabilities presented in [12], scalability is one of the most critical requirements to consider when deploying cyber ranges [14].

Currently, many cyber ranges allow conducting training sessions with complex scenarios. However, they require sophisticated infrastructure, considerable resources, and workforce to prepare and maintain the cyber range before and during the training. Considering these, the cyber ranges are not suitable for less complex scenarios (e.g., to learn basic Unix commands in introductory university courses) with hundreds of students.

### B. Lightweight Labs

Lightweight cybersecurity labs typically use virtual environments (e.g., VirtualBox) to run training scenarios on participants' personal computers. Using this approach, the instructors can organize training sessions with a huge number of participants without the need to provide and manage many centralized resources (such as in the cloud).

Du et al. in [15] introduce SEED lab environment based on an instructional OS (Minix) and a production OS (Linux). The project developed over 30 labs covering a wide area of topics ranging from network security to security of mobile applications. They discuss that the students' feedback is important to improve the local-based environment. They identified two factors: *efficiency* and *effectiveness*. Efficiency focuses on the appropriateness and design of labs and effectiveness on student's learning output. Although hundreds of students can be served through this lab environment, the VM images are fixed. It is difficult for instructors to modify them, and thus it reduces the scalability.

Irvine et al. [16]–[18] describe fully-provisioned Linux-based lab exercises with an emphasis on cybersecurity. They developed a framework consisting of 50 lab exercises named Labtainers. The solution is distributed as a single VM for either VirtualBox or VMware. The framework uses Docker containers to instantiate networked hosts within a single VM. Labtainers enable conducting training sessions with many participants since it requires fewer resources (hardware, software, workforce, and finances) compared to cyber ranges and dedicated VMs. Further, it can reduce the instructors' effort necessary to prepare the training since it does not require deploying all the scenarios and configurations on the server. The instructors can lack institutional IT equipment and staff to manage and deploy a

fine-tuned lab environment. Although the provided solution applies network virtualization using Docker containers, this solution does not support functions and capabilities typical for cyber ranges presented in [12]. For instance, it does not allow creating teams or does not provide learning analytics and complex scoring visualizations.

To conclude the review of related work, cybersecurity hands-on education currently uses lightweight labs [18] or cyber ranges [10]–[12]. The cyber ranges have its drawbacks in the difficult deployment and heavy resource requirements (despite being offloaded from students to instructors). The lightweight labs can serve a huge number of users using their local hosts. However, the labs do not provide all functions and capabilities of cyber ranges. To the best of our knowledge, there are no published attempts or open-source examples that benefit from both of these areas.

### III. BUILDING BLOCKS OF A CYBERSECURITY HANDS-ON CLASS

Cybersecurity hands-on classes can be held in many different formats. Students follow step-by-step instructions, solve a complex assignment autonomously, play a serious game (such as capture the flag or cyber defense exercise), or work together on a long-term project. Instructions and assignments can be provided by a tutor, which can be a human or a machine. Regardless of the format, students usually interact with a learning environment that features the following building blocks: *sandbox*, *class delivery method*, and *learning analytics*. In this section, we define and detail these blocks.

#### A. Sandbox

Sandbox is an isolated environment for practicing cybersecurity skills. It is the essential component within the interactive learning environment. For the rest of the paper, we assume that the term *sandbox* denotes either a single *virtual machine* (VM), or a network of VMs. Users can run the sandboxes on their own hosts or access them remotely if deployed in a cloud.

VM is an instantiated representation of a *virtual image* using a *virtualization platform* and metadata defining hardware of the VM. *Virtual image* is a snapshot of a computer disk that can be distributed as a file. The images typically carry the installation of an operating system with specific applications, tools, and files.

Once all the VMs are booted, networked, and running, i.e., they constitute a sandbox, users can interact with them. Two basic access methods can be used: *console*<sup>1</sup> and *remote access*. The former is provided by the virtualization platform and is always present, enabling users to interact with input and output devices of VMs. The latter relies on the installed operating system and its applications. The most common remote access protocol is Secure Shell (SSH) for Linux-based systems, and Remote Desktop Protocol (RDP) or PowerShell Remoting Protocol (PSRP) for Windows systems.

<sup>1</sup>Also KVM, Keyboard, Video, and Mouse.

#### B. Class Delivery Method

The class delivery method is a computer-assisted instruction that employs the sandbox. It is provided as a part of the learning environment. It can use teacher- or student-centered approach. For instance, instruction can be provided using static (rich) text *assignments* or *serious games*. Some instructional methods require the assistance of an instructor, who also assesses students' progress or outcome. These methods include labs integrated with lectures, project-based learning, and problem-based learning [19]. Other methods fully rely on a learning environment that presents assignments, provides feedback, or assesses students' outcomes without a human instructor's assistance. These methods include self-directed learning or automated tutoring systems [19].

#### C. Learning Analytics

Learning analytics (LA) [20] and educational data mining (EDM) [21] are two growing disciplines that leverage data from educational contexts. They aim to understand and improve teaching and learning [22]. In this paper, we refer to components analyzing students' interactions in both the sandbox and the class delivery as *learning analytics*. It can be used during and after the class with different goals. During the class, instructors can monitor students' progress and provide formative assessment to students. It can also be used for summative assessment, i.e., student testing and grading. After the class, LA enables students to reflect on the class and plan further learning. Instructors can use LA to improve the class for future runs.

### IV. SCALABLE APPROACH

Preparing a hands-on class is a laborious and complex task. Therefore, instructors strive to use the above-mentioned building blocks repetitively. Less often, they also improve the class based on students' feedback or a technology change. The manual deployment of the sandbox and the class delivery is time-consuming, and it does not scale well for big classes.

This section describes reusable components of the building blocks introduced in Section III. In particular, we introduce:

- *Topology definition* and *provisioning definition*, two parts of *sandbox definition*, which specifies the internal structure of the sandboxes (networks and hosts) and configuration of the hosts.
- *Training definition* specifying the tasks and questions for students in the class.
- *Learning analytics stack*, a key component for providing both formative and summative assessment to students.

These components make delivering cybersecurity hands-on classes scalable. The components are cornerstones of the two learning environments described in Section V.

#### A. Topology Definition

Topology definition specifies a name of the definition, hosts, routers, networks, network mappings, router mappings, and host or router groups [23]. Listing 1 is an example of a topology

definition of a sandbox shown in Figure 1 expressed in a YAML format.

```

1  name: small-sandbox
2
3  hosts:
4    - name: server
5      base_box:
6        image: debian-9-x86_64
7        man_user: debian
8        flavor: tiny1x2
9
10   - name: home
11     base_box:
12       image: debian-9-x86_64
13       man_user: debian
14       flavor: tiny1x2
15
16  routers:
17    - name: server-router
18      cidr: 100.100.100.0/29
19      base_box:
20        image: debian-9-x86_64
21        man_user: debian
22        flavor: tiny1x2
23
24    - name: home-router
25      base_box:
26        image: debian-9-x86_64
27        man_user: debian
28      cidr: 200.100.100.0/29
29      flavor: tiny1x2
30
31  networks:
32    - name: server-switch
33      cidr: 10.10.20.0/24
34
35    - name: home-switch
36      cidr: 10.10.30.0/24
37
38  net_mappings:
39    - host: server
40      network: server-switch
41      ip: 10.10.20.5
42
43    - host: home
44      network: home-switch
45      ip: 10.10.30.5
46
47  router_mappings:
48    - router: server-router
49      network: server-switch
50      ip: 10.10.20.1
51
52    - router: home-router
53      network: home-switch
54      ip: 10.10.30.1
55
56  groups:
57    - name: user-accessible
58      nodes:
59        - home
60        - home-router

```

Listing 1: YAML topology definition of a sandbox.

This topology defines two hosts, two routers, two networks, and one named group. One host is named `server` and the other `home`. The definition specifies host properties, such as virtual image (`base_box`) or the user-name for host provisioning (`man_user`). The hosts are assigned to respective networks with particular IP addresses

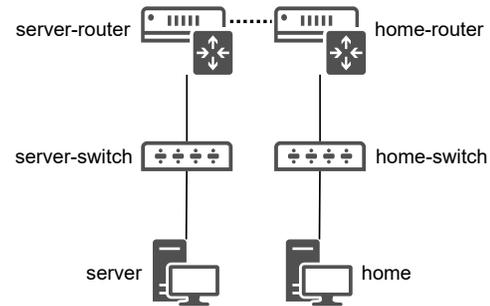


Fig. 1: An example of a network topology of a sandbox.

(`net_mappings`), and the networks are assigned to the routers (`router_mappings`). The definition also creates one named group (`user-accessible`), which will be used for provisioning.

### B. Provisioning Definition

Provisioning definition specifies configuration changes at the hosts defined in the topology definition. It is used to customize the *base boxes*, which contain only a minimal installation of a particular operating system (such as Ubuntu Server 20.04). The provisioning definition provides inputs for software configuration management system, which is responsible for installing and configuring the hosts deployed as base boxes. The provisioning definition contains prescriptions for the installation and configuration of applications. Also, it provides data specific to the particular class (such as installation of Apache web server and provisioning of a web application running at this server).

The common approach is providing hosts as virtual images in binary files. However, these images cannot be easily updated or modified. In contrast, this approach gives instructors more flexibility in preparing the sandbox and allows them to reuse the parts of the definitions for another class.

Our approach uses Ansible for software configuration management. Ansible executes ordered lists of tasks provided as machine-readable YAML files [24]. Listing 2 is a simple example of two tasks installing a web server and provisioning files of a web application at the host named `server`.

```

1  - hosts: server
2    become: true
3    tasks:
4      - name: Install Apache, MySQL and PHP5
5        apt:
6          name: [apache2, mysql-server,
7                php5-mysql, php5]
8          state: present
9          update_cache: yes
10
11     - name: Copy app to the web root
12       copy:
13         src: web-app/
14         dest: /var/www/html
15     ...

```

Listing 2: Ansible configuration file in YAML format.

The provisioning definition can contain multiple files and directories following the Ansible conventions.

### C. Base Boxes

Base boxes can be provided by third parties or can be specifically tailored. Custom-made base boxes enable more flexibility, transparency, and better performance. However, their preparation is laborious and requires expert skills, which instructors may lack. For this reason, we provide open-source custom base boxes of common operating systems fine-tuned for use in our environments (both local and cloud).

The custom base boxes are created using Packer, an open-source tool for creating identical machine images for multiple platforms from a single source configuration [25]. Packer uses machine-readable configuration files, which define the process of creating the images from an ISO file containing the installation image or virtual hard drive.

### D. Training Definition

Training definition specifies consecutive tasks that have to be solved by each student in the class. Figure 2 shows a generic structure of a training. In each training phase (P), the student must complete a task and submit an answer (text) to prove the solution. Then, the student enters the next phase until completing the last phase. Each phase features one or more optional hints and a worked-out solution, which can be displayed if needed. Additionally, the training may include questionnaires (Q) at the beginning (pre-test) and the end (post-test), and information (I) for students, such as introductory instructions.



Fig. 2: Generic structure of training with several phases (P) and optional questionnaires (Q) and informative phases (I).

The training definition is a machine-readable description of the training and its training phases. Listing 3 shows an example of a part of a training definition (metadata and the first training phase with three hints) expressed in a JSON format.

```

1 {
2   "title": "Secret Laboratory",
3   "description": "A cybersecurity game.",
4   "prerequisites": [ "Basic knowledge of Unix",
5                     "Basic networking" ],
6   "outcomes": [ "nmap", "metasploit" ],
7   "phases": [ {
8     "title": "Introduction",
9     "phase_type": "INFO",
10    "order": 0,
11    "estimated_duration": 0,
12    "content": "Place for a background story."
13  }, {
14    "title": "Looking for a vulnerable service.",
15    "max_score": 100,
16    "level_type": "TRAINING",
17    "order": 1,
18    "estimated_duration": 5,
19    "flag": "service-name-1.23",
20    "content": "Now you need to scan the server
21              to find possible vulnerabilities. The IP
22              address of the server is **10.1.26.9** .
23              The name of the vulnerable service starts
24              with \"s\". \n\n
  
```

```

25   As a flag, submit the name of
26   the vulnerable service in the following
27   format: _service-version_. All characters
28   are lowercase. For example: _dvwa-2.050_.",
29   "solution": "`root@attacker:~# nmap -sV
30   -p 10000 10.1.26.9\n`\n\n
31   The flag is: **service-name-1.23**",
32   "hints": [ {
33     "title": "Which port should you scan?",
34     "content": "The vulnerable service is
35               running on port 10000. You can also pass
36               this information to nmap
37               (**-p \"port range\"**).",
38     "hint_penalty": 10,
39     "order": 1
40   }, {
41     "title": "Which option gives you the
42             version?",
43     "content": "To determine the name and
44               version of the service, you need to pass
45               the argument **-sV** to nmap -
46               service/version detection.",
47     "hint_penalty": 15,
48     "order": 2
49   }, {
50     "title": "Which tool should you use?",
51     "content": "You should use **nmap** to
52               scan the server (see **man nmap**).",
53     "hint_penalty": 10,
54     "order": 0
55   } ],
56   "incorrect_flag_limit": 100,
57 }, {
58   "...": "..."
59 }
  
```

Listing 3: Training definition in a JSON format.

Structuring the training to linear phases guides the students by dividing the complex assignment into smaller, more digestible units. In addition, the predefined order of tasks (levels) enables to practice processes or phases in the correct order (e.g., typical cybersecurity attack lifecycle phases). This is a different format than the popular *jeopardy* game [26], where players can choose any task from available tasks at any time.

### E. Learning Analytics Stack

Learning analytics stack is a mechanism of processing events from the learning environment. The events capture interactions of the student with the environment, particularly sandbox and class delivery method (training). The events can capture various actions of the student. These can include: starting or finishing the particular training or its individual phase; submitting a correct or incorrect answer in a training phase; displaying a hint or solution of the phase; logging into a host in a sandbox; typing a command at the command line of a host in the sandbox.

The events are machine-readable strings logged using the standard Syslog protocol [27] in a predefined format, i.e., they can be produced by any application capable of Syslog logging. The learning environment forwards to and stores all events at the central storage, which transforms and further processes them. The central storage uses the ELK stack (Elasticsearch, Logstash, and Kibana) [28], which provides interfaces for both learning analytics applications and instructors.

```

Feb 17 2021 9:17:33  username="root"  client  src="10.10.40.5"
      timestamp                username      hostname  host IP address

      wd="/home"      cmd="ssh alice@server"  cmd_type="bash"  uid="1"
      working directory      command                command type      sandbox ID

```

Fig. 3: A log entry for a command executed on one machine in a sandbox.

For example, Figure 3 and Listing 4 show how a command `ssh alice@server` executed by a student in the Linux terminal at a machine in the sandbox is timestamped and logged into Syslog as a string (Figure 3). Then, it is transferred and stored at the central storage as an entry for further processing (Listing 4).

```

1 {
2   "timestamp": "2021-02-17T09:17:33+02:00",
3   "username": "root",
4   "hostname": "client",
5   "host_ip": "10.10.40.5",
6   "wd": "/home",
7   "cmd": "ssh alice@server",
8   "cmd_type": "bash-command",
9   "sandbox_id": "1"
10 }

```

Listing 4: An entry at the central storage created from the Syslog record in Figure 3.

Another example is an event of submission of an incorrect answer (flag in a capture the flag game) `.invoices2019` to a task in a training phase (a level of the game). This event is timestamped and logged into Syslog (Listing 5) and finally arrives at the central storage.

```

1 {
2   "flag_content": ".invoices2019",
3   "actual_score_in_level": 100,
4   "total_score": 300,
5   "game_time": 3045985,
6   "timestamp": 1610618680221,
7   "type": "events.trainings.WrongFlagSubmitted",
8   "count": 1,
9   "user_ref_id": 19,
10  "phase_id": 36,
11  "training_run_id": 28,
12  "training_instance_id": 12,
13  "training_definition_id": 7,
14  "sandbox_id": 104,
15  "pool_id": 40
16 }

```

Listing 5: An entry created by the training portal.

## V. LEARNING ENVIRONMENTS

This section introduces two learning environments we developed based on the common components described in Section IV. KYPO Cyber Range Platform (KYPO CRP) [2] is a cloud-based platform designed for running multiple classes in parallel or classes requiring sandboxes with many hosts. Cyber Sandbox Creator (CSC) [3] is a lightweight, distributed lab environment using a commercial off-the-shelf computer in the lab or students' own desktop or laptop.

Both environments use the same formats for topology, provisioning, and training definitions, and same formats of events processed by learning analytics stack. The key difference is the used virtualization technology for sandbox instantiation. This is most evident for base boxes, which are almost identical except for the features and limits bound to the different underlying virtualization technologies. Another important difference lies in user roles and access control to sandbox and training instances. Table I highlights the differences between these two environments.

Both environments have been released as open-source software with documentation and examples of training.

### A. Cloud-based Learning Environment

The design of the learning environment KYPO CRP is shown in Figure 4. First, the instructor checks that estimated resources for the training session are available in the cloud. The estimated resources are based on the number of students in the class and the size of the sandbox. After that, the instructor logs in to the web interface of KYPO CRP and allocates a pool of sandboxes for the class. The pool size is usually set to the number of students plus a few more for a reserve. Further, the instructor provides a Gitlab repository with sandbox and provisioning definitions. The repository may enable auxiliary services of the learning analytics stack. KYPO CRP then builds all sandboxes in the allocation pool in the cloud using sandbox and provisioning definitions.

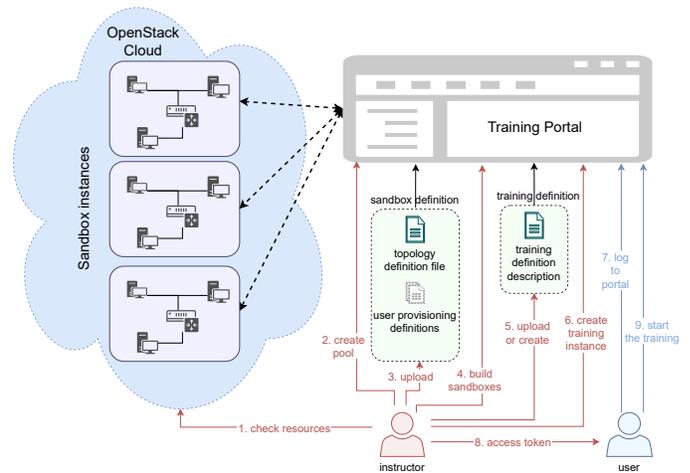


Fig. 4: The principle of how the learning environment KYPO CRP works.

Next, the instructor creates or imports a training definition describing the task assignments. As the last step, the instructor

TABLE I: The comparison of features and capabilities of KYPO Cyber Range Platform and Cyber Sandbox Creator

Feature	KYPO CRP	CSC
Required hardware	Cloud of tens of servers	Desktop/laptop
Virtualization technology	OpenStack	VirtualBox
Deployment	Advanced user	General user
Max. num. of hosts in a sandbox for each student	Up to free cloud resources	≈ 5
Number of parallel classes	Up to free cloud resources	1
Preparation effort at the students' side	None	Low
User access	Remote (web browser, SSH, RDP)	Local
Sandbox definition visible to students	No	Yes
Task presentation	Integrated	Separate

creates a training instance for a particular class. The instance defines the start and end times of the training session and a respective pool of sandboxes. KYPO CRP thus enables instructors to run multiple (different) classes in parallel.

The student logs into the web interface of KYPO CRP (training portal) and starts the training by entering an *access token* provided by the instructor. If the entered access token is correct, an available sandbox instance from the pool is assigned to the particular student. The student then starts solving the tasks defined in the training definition by interacting with the sandbox instance hosted in the cloud. KYPO CRP provides access to the sandbox host using a web browser, presents tasks, provides predefined on-demand hints, checks submitted answers, and collects events from its web interface and allocated sandboxes. The instructor can monitor the progress of all students in the class during the training at a dashboard as shown in Figure 6. When the training is over, analyses of student progress are available both to the student and instructor.

### B. Lightweight Learning Environment

The design of the learning environment CSC is shown in Figure 5. First, a superuser (instructor) writes or reuses a sandbox definition, from which an intermediate definition is generated. This can be further extended by provisioning definition, which specifies the software and configuration of the machines in the sandbox.

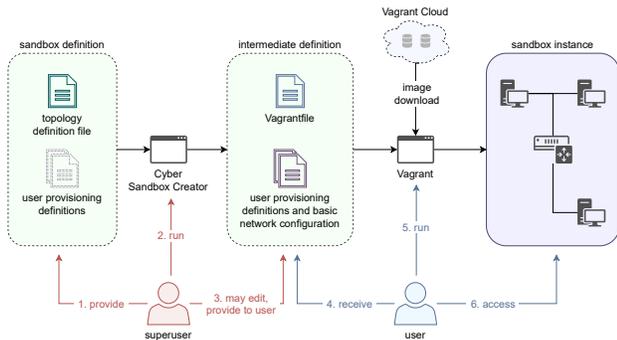


Fig. 5: The principle of how the learning environment CSC works.

The user (student) receives this intermediate definition and builds the sandbox locally using a single command. This is achieved by using Vagrant, which downloads images of operating systems for the VMs (base boxes), and provisions

them according to the superuser's configuration. As a result, the sandbox instance is ready on the student's computer.

The student performs the hands-on tasks in the sandbox instance. The task assignments and scaffolding are delivered via separate software, such as the training portal from the learning environment KYPO CRP or CTFd [29]. Moreover, the superuser can enable auxiliary services in the sandbox definition, such as command logging from the learning analytics stack. As a result, the commands submitted by the student in the sandbox are forwarded to a central storage, where they can be processed further or viewed by an instructor.

### C. Use Cases

Since both environments are based on common components, they enable the following use cases:

1) *Running the same training in either environment:* The common formats of training and sandbox definitions enable running training sessions in either KYPO CRP or CSC without additional instructor's effort. The instructor can consider which features of the particular environment better fulfill the needs of a particular class (see Table I) and choose a more suitable environment. An example training that demonstrates this feature is publicly available at [30].

2) *Sharing of created training between organizations:* The human-readable formats of training and sandbox definitions enable their use in other organizations in the KYPO CRP and CSC learning environments. Alternatively, it simplifies the adoption in different environments if they are based on the same or similar building blocks. For instance, training definitions could be easily transformed to a learning environment that is based on different technology for sandbox provisioning. Or, a sandbox definition of a typical small enterprise network could be used for different training sessions.

## VI. TEACHING EXPERIENCE

To demonstrate the usability of the presented learning environments for teaching cybersecurity, we now describe recent teaching experience. Since the year 2019, about 380 students used our learning environments in 31 training sessions. The students were undergraduates and graduates, professional learners, and selected high school students attending a cybersecurity competition. While most students received the training, some students were involved in creating the training and cybersecurity games as described in [31]. They were able to create sandbox and training definitions using the learning



Fig. 6: A part of the instructor’s dashboard showing progress of two students working in KYPO CRP. Colored bars represent different completed training phases. A dot means a hint displayed by a student, and a tick solution displayed. At first glance, the instructor can see the first player performs better than the second one who may need additional help. Here, we present a view with students’ avatars instead of their real names.

environment CSC and run the training for their peers in either the learning environment KYPO CRP or CSC.

Some sessions were held by other instructors at other institutions in several European countries. Both students and instructors were able to run the training in the environment CSC without our in-class assistance.

The learning environment KYPO CRP was successfully deployed in a private cloud at our university, Brno University of Technology, Czech Republic, and Swedish Defence Research Agency (FOI).

#### A. Classes Aimed at Formative Assessment

When teaching practical lab sessions at university courses, our goal is that students gain hands-on experience with using various cybersecurity tools. In this low-stakes context, we do not care if the students see the sandbox definition, so we usually decide to use CSC. Regardless of the number of students, everyone deploys the sandbox locally on their own computer. This way, we have taught classes from 10 up to 200 students.

However, the teachers need to allocate time for preparing detailed setup instructions, as well as be ready to troubleshoot the setup. The students use a wide range of host operating systems (e.g., Windows, MacOS, and different Linux distributions), and the hardware configuration of their devices varies widely. As a result, some students may experience early difficulties before the environment is ready to run on their machines. In particular, a few students encountered issues with the installation of VirtualBox.

Alternatively, we may opt in for the KYPO CRP. The setup is simpler for students, as the instructors prepare sandboxes in the cloud, and students access them remotely using a web browser or SSH. However, the number of sandboxes we can host is limited by the cloud’s resources. We usually teach classes from 10 to 30 students with this setup. An additional risk is that if the cloud space is shared, external users running other computational tasks in the cloud may jeopardize the stability of the environment and user experience. This risk is not associated with the CSC, because if the students experience low responsiveness, they only need to close other unnecessary applications running on their computers.

Regardless of the environment used, we can collect command histories of students solving the tasks [32] and provide them with formative feedback. This includes explaining what they did well and what they can improve, for example, how to address frequently occurring mistakes.

#### B. Classes Aimed at Summative Assessment

When we need to perform summative assessment, such as during a final exam or a competition, we need to hide the sandbox definitions from students. Therefore, we only use the KYPO CRP environment for this use case. We can control the visibility of hosts in the sandbox topology so that students are initially aware only of a limited number of machines. They also cannot directly see what applications are running there, and so the setup mimics more realistic situations.

### VII. CONCLUSIONS

Cybersecurity hands-on classes take place in a laboratory or an interactive learning environment featuring real-world tools, systems, and applications. The preparation of the classes requires not only pedagogical skills of choosing suitable instructional methods for the class but also proficiency in software development, deployment, and IT operations.

In this paper, we demonstrated a technical innovation for enhancing cybersecurity classes. We support cybersecurity educators by providing scalable and reusable building blocks of the classes. These range from open formats for the description of technical environment, through content, to ready-made artifacts applicable in various contexts.

We exploit these blocks while developing two learning environments, KYPO CRP and CSC. Each environment has its benefits and limitations, which determine its suitability for a particular use case. While KYPO CRP is more appropriate for summative assessment and classes requiring an extensive or complex networked environment, it relies on a dedicated cloud and personnel who can manage it. In contrast, CSC is more suitable for big classes featuring a small network environment that can be deployed at students’ hosts. Regardless, classes deployed in one environment can be easily deployed in the other if needed. In addition, open definitions of formats enable educators to enhance and edit the existing lab environments without much additional effort. Both environments KYPO CRP [2] and CSC [3] have been released as open-source software so educators can freely use them in their teaching practice.

### ACKNOWLEDGMENT

This research was supported by the Security Research Programme of the Czech Republic 2015–2022 (BV III/1–VS) granted by the Ministry of the Interior of the Czech Republic under No. VI20202022158 – Research of New Technologies to Increase the Capabilities of Cybersecurity Experts.

## REFERENCES

- [1] Joint Task Force on Cybersecurity Education, “Cybersecurity Curricular Guideline,” 2017. [Online]. Available: <http://cybered.acm.org>
- [2] Masaryk University, “KYPO Cyber Range Platform,” 2021. [Online]. Available: <https://crp.kypo.muni.cz>
- [3] —, “Cyber Sandbox Creator,” 2021. [Online]. Available: <https://gitlab.ics.muni.cz/muni-kypo-csc/cyber-sandbox-creator>
- [4] National Initiative for Cybersecurity Education. (2020) The Cyber Range: A Guide. National Initiative for Cybersecurity Education. [Online]. Available: <https://www.nist.gov>
- [5] M. Knüpfer, T. Bierwirth, L. Stiemert, M. Schopp, S. Seeber, D. Pöhn, and P. Hillmann, “Cyber Taxi: A Taxonomy of Interactive Cyber Training and Education Systems,” in *Model-driven Simulation and Training Environments for Cybersecurity*, G. Hatzivasilis and S. Ioannidis, Eds. Cham: Springer International Publishing, 2020, pp. 3–21.
- [6] Circadence. (2021) Circadence: Cyber Learning Platforms. Circadence. [Online]. Available: <https://www.circadence.com/>
- [7] IAI ELTA. (2021) Cyber Academy Onsite Training Solution. IAI ELTA. [Online]. Available: <https://www.iai.co.il/p/cyber-academy>
- [8] Airbus. (2021) CyberRange. Airbus. [Online]. Available: <https://airbus-cyber-security.com/products-and-services/prevent/cyberrange/>
- [9] Silensec. (2021) Cyber space, engaged. Silensec. [Online]. Available: <https://cyberranges.com/>
- [10] E. Ukwandu, M. A. B. Farah, H. Hindy, D. Brosset, D. Kavallieros, R. Atkinson, C. Tachtatzis, M. Bures, I. Andonovic, and X. Bellekens, “A Review of Cyber-Ranges and Test-Beds: Current and Future Trends,” *Sensors*, vol. 20, no. 24, 2020. [Online]. Available: <https://www.mdpi.com/1424-8220/20/24/7148>
- [11] N. Chouliaras, G. Kittes, I. Kantzavelou, L. Maglaras, G. Pantziou, and M. A. Ferrag, “Cyber ranges and testbeds for education, training, and research,” *Applied Sciences*, vol. 11, no. 4, 2021. [Online]. Available: <https://www.mdpi.com/2076-3417/11/4/1809>
- [12] M. M. Yamin, B. Katt, and V. Gkioulos, “Cyber ranges and security testbeds: Scenarios, functions, tools and architecture,” *Computers and Security*, vol. 88, p. 101636, 1 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0167404819301804>
- [13] C. Pham, D. Tang, K.-i. Chinen, and R. Beuran, “CyRIS: A Cyber Range Instantiation System for Facilitating Security Training,” in *Proceedings of the Seventh Symposium on Information and Communication Technology*, ser. SoICT ’16. New York, NY, USA: Association for Computing Machinery, 2016, p. 251–258. [Online]. Available: <https://doi.org/10.1145/3011077.3011087>
- [14] V. E. Urias, W. M. Stout, B. Van Leeuwen, and H. Lin, “Cyber Range Infrastructure Limitations and Needs of Tomorrow: A Position Paper,” in *2018 International Carnahan Conference on Security Technology (ICCST)*, 2018, pp. 1–5.
- [15] W. Du and R. Wang, “SEED: A Suite of Instructional Laboratories for Computer Security Education,” *J. Educ. Resour. Comput.*, vol. 8, no. 1, Mar. 2008. [Online]. Available: <https://doi.org/10.1145/1348713.1348716>
- [16] C. E. Irvine, M. F. Thompson, M. McCarrin, and J. Khosalim, “Live Lesson: Labtainers: A Docker-based Framework for Cybersecurity Labs,” in *2017 USENIX Workshop on Advances in Security Education (ASE 17)*. Vancouver, BC: USENIX Association, Aug. 2017. [Online]. Available: <https://www.usenix.org/conference/ase17/workshop-program/presentation/irvine>
- [17] M. F. Thompson and C. E. Irvine, “Individualizing Cybersecurity Lab Exercises with Labtainers,” *IEEE Security Privacy*, vol. 16, no. 2, pp. 91–95, 2018.
- [18] M. Thompson and C. Irvine, “Labtainers Cyber Exercises: Building and Deploying Fully Provisioned Cyber Labs That Run on a Laptop,” in *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education*, ser. SIGCSE ’21. New York, NY, USA: Association for Computing Machinery, 2021, p. 1353. [Online]. Available: <https://doi.org/10.1145/3408877.3432490>
- [19] K. Sanders, J. Boustedt, A. Eckerdal, R. McCartney, and C. Zander, “Folk pedagogy: Nobody doesn’t like active learning,” in *Proceedings of the 2017 ACM Conference on International Computing Education Research*, ser. ICER ’17. New York, NY, USA: Association for Computing Machinery, 2017, p. 145–154. [Online]. Available: <https://doi.org/10.1145/3105726.3106192>
- [20] C. Lang, G. Siemens, A. Wise, and D. Gašević, Eds., *Handbook of Learning Analytics*, 1st ed. Society for Learning Analytics Research (SoLAR), 2017. [Online]. Available: <https://doi.org/10.18608/hla17>
- [21] C. Romero, S. Ventura, M. Pechenizkiy, and R. S. Baker, Eds., *Handbook of educational data mining*. Boca Raton, FL, USA: CRC Press, 2010. [Online]. Available: <https://doi.org/10.1201/b10274>
- [22] C. Hundhausen, D. Olivares, and A. Carter, “IDE-Based Learning Analytics for Computing Education: A Process Model, Critical Review, and Research Agenda,” *ACM Transactions on Computing Education*, vol. 17, no. 3, pp. 11:1–11:26, Aug. 2017. [Online]. Available: <https://doi.org/10.1145/3105759>
- [23] Masaryk University, “KYPO Cyber Range Platform Topology Definition,” 12 2020. [Online]. Available: <https://zenodo.org/record/4314736>
- [24] RedHat. (2020) Ansible User Guide. [Online]. Available: [https://docs.ansible.com/ansible/latest/user\\_guide/basic\\_concepts.html](https://docs.ansible.com/ansible/latest/user_guide/basic_concepts.html)
- [25] HashiCorp. (2021) Website. [Online]. Available: <https://www.packer.io>
- [26] A. Davis, T. Leek, M. Zhivich, K. Gwinnup, and W. Leonard, “The fun and future of CTF,” in *2014 USENIX Summit on Gaming, Games, and Gamification in Security Education (3GSE 14)*. San Diego, CA: USENIX Association, Aug. 2014, pp. 1–9.
- [27] R. Gerhards, “The Syslog Protocol,” RFC 5424, Mar. 2009. [Online]. Available: <https://rfc-editor.org/rfc/rfc5424.txt>
- [28] Elasticsearch. (2020) What is the ELK Stack? Elasticsearch B.V. [Online]. Available: <https://www.elastic.co/what-is/elk-stack>
- [29] K. Chung, “Live Lesson: Lowering the Barriers to Capture The Flag Administration and Participation,” in *2017 USENIX Workshop on Advances in Security Education (ASE 17)*. Vancouver, BC: USENIX Association, 2017. [Online]. Available: <https://www.usenix.org/conference/ase17/workshop-program/presentation/chung>
- [30] Masaryk University. (2021) Junior hacker training. [Online]. Available: <https://gitlab.ics.muni.cz/muni-kypo-trainings/games/junior-hacker>
- [31] V. Švábenský, J. Vykopal, M. Cermak, and M. Laštovička, “Enhancing Cybersecurity Skills by Creating Serious Games,” in *Proceedings of the 23rd Annual ACM Conference on Innovation and Technology in Computer Science Education*, ser. ITiCSE 2018. New York, NY, USA: ACM, 2018, pp. 194–199. [Online]. Available: <http://doi.acm.org/10.1145/3197091.3197123>
- [32] V. Švábenský, J. Vykopal, D. Tovarňák, and P. Čeleda, “Toolset for Collecting Shell Commands and Its Application in Hands-on Cybersecurity Training [in press],” in *Proceedings of the 51st IEEE Frontiers in Education Conference*, ser. FIE ’21. New York, NY, USA: IEEE, 10 2020, pp. 1–9.