

Toolset for Collecting Shell Commands and Its Application in Hands-on Cybersecurity Training

Valdemar Švábenský
Masaryk University
Brno, Czech Republic
svabensky@ics.muni.cz

Jan Vykopal
Masaryk University
Brno, Czech Republic
vykopal@ics.muni.cz

Daniel Tovarňák
Masaryk University
Brno, Czech Republic
tovarnak@ics.muni.cz

Pavel Čeleda
Masaryk University
Brno, Czech Republic
celeda@ics.muni.cz

Abstract—This Full Paper in the Innovative Practice category presents and evaluates a technical innovation for hands-on classes. When learning cybersecurity, operating systems, or networking, students perform practical tasks using a broad range of command-line tools. Collecting and analyzing data about the command usage can reveal valuable insights into how students progress and where they make mistakes. However, few learning environments support recording and inspecting command-line inputs, and setting up an efficient infrastructure for this purpose is challenging. To aid engineering and computing educators, we share the design and implementation of an open-source toolset for logging commands that students execute on Linux machines. Compared to basic solutions, such as shell history files, the toolset’s novelty and added value are threefold. First, its configuration is automated so that it can be easily used in classes on different topics. Second, it collects metadata about the command execution, such as a timestamp, hostname, and IP address. Third, all data are instantly forwarded to central storage in a unified, semi-structured format. This enables automated processing of the data, both in real-time and post hoc, to enhance the instructors’ understanding of student actions. The toolset works independently of the teaching content, the training network’s topology, or the number of students working in parallel. We demonstrated the toolset’s value in two learning environments at four training sessions. Over two semesters, 50 students played educational cybersecurity games using a Linux command-line interface. Each training session lasted approximately two hours, during which we recorded 4439 shell commands. The semi-automated data analysis revealed different solution patterns, used tools, and misconceptions of students. Our insights from creating the toolset and applying it in teaching practice are relevant for instructors, researchers, and developers of learning environments. We provide the software and data resulting from this work so that others can use them in their hands-on classes.

Index Terms—cybersecurity education, host-based monitoring, command-line history, Syslog, virtual machines, sandbox, educational data mining, learning analytics, learning technology

I. INTRODUCTION

Hands-on training is vital for gaining expertise in computing disciplines. Topics such as cybersecurity, operating systems, and networking must be practiced in a computer environment so that students can try various tools and techniques. Such a learning environment is called a *sandbox*. It contains networked hosts that may be intentionally vulnerable to allow practicing cyber attacks and defense. These skills are grounded in the current cybersecurity curricular guidelines [1] to address the increasing shortage of cybersecurity workforce [2].

For the training, each student receives an isolated sandbox hosted locally or in a cloud. To solve the training tasks, students work with many tools, both in a graphical user interface (GUI) and a command-line interface (CLI). This paper focuses on the Linux CLI, which is common in higher education of computing, as well as software development in the industry practice.

Analyzing CLI interactions opens opportunities for educational research and classroom innovation. In traditional face-to-face classes, instructors must look at the students’ computer screens to observe the learning process. However, this approach does not scale for large classes, and it becomes difficult for distance education. Instead, if the students’ executed commands are logged, instructors and researchers may leverage them to support learning. By employing the methods of educational data mining [3] and learning analytics [4], the CLI data can help achieve important educational goals, such as to:

- better understand students’ approaches to learning, both in face-to-face and remote classes,
- objectively assess learning, and
- provide targeted instruction and feedback.

This paper examines the following research question relevant for instructors: *What can we infer from students’ command histories that is indicative of their learning processes?* Specifically, our goal is to understand how students solve cybersecurity assignments by analyzing their CLI usage. To address this question, we propose a generic method for collecting CLI logs from hands-on training. Then, we evaluate this method by gathering the logs from 50 students at four training sessions and investigating three sub-questions / use cases of the data:

- 1) *What does the command distribution indicate about the students’ approach to solving the tasks?* Our motivation is to analyze which tools are commonly used and how effective they are with respect to the training tasks.
- 2) *Which commands are used immediately after the student accesses the learning environment?* We can observe if the students started solving the initial task, familiarized themselves with the environment, or displayed off-task behavior. This allows for providing suitable scaffolding.
- 3) *How much time do students spend on the tasks, and how often do they attempt an action?* Observing the time differences between successive commands can indicate the students’ skill level and support assessment.

Although there are many cybersecurity learning environments, which we review in Section II, their logging support is often limited or non-existent. The current solutions do not allow instructors to uniformly collect CLI data and metadata with minimal setup and then correlate the logs from multiple sandboxes for advanced analyses.

This paper addresses this gap by presenting and evaluating a technical innovation for hands-on classes that employ CLI tools. We created a toolset that collects Linux shell commands in physical or virtual learning environments and stores them in a unified format. Compared to the previous practice, where observing students' learning was difficult or even impossible, the proposed innovation enables understanding student approaches at scale. It also allows employing educational data mining and learning analytics techniques to gain further insights.

The toolset design is explained in Section III. In Section IV, we introduce a study that deploys the toolset in practice and evaluates it in authentic educational contexts. Section V presents the results of the study and addresses the questions above. Section VI discusses the study and proposes multiple research ideas that further leverage the collected data. Finally, Section VII summarizes our contributions. We also publish the toolset as open-source software. Instructors, researchers, and developers can use it to enhance computing classes, such as teaching cybersecurity, operating systems, and networking.

II. RELATED WORK

This section contextualizes the work within the pedagogical literature. Then, it reviews learning environments and research in collecting and analyzing student data. We also point out the addressed gaps and explain how we differ from existing work.

A. Theoretical and Pedagogical Background

Active learning is “any instructional method that engages students in the learning process” [5]. Sanders et al. [6] reviewed 38 active instructional techniques, including hands-on labs, educational (serious) games, and automated tutoring systems that provide feedback to students. Unlike traditional lectures, active learning enables students to understand the topic more deeply [7]. Although the context matters, research strongly supports the notion that incorporating active methods improves learning [5]. That is why practice with CLI tools is important in engineering and computing education.

Zone of proximal development [8] is a learning theory about an imaginary “space” slightly beyond the student’s comfort zone, but where the student can learn if given some *scaffolding*. Scaffolding is instructional support such as feedback and advice from the teacher, assignment of tasks suitable for the learner’s skill level, and provision of hints. Applying this cognitive science theory improves the effectiveness of hands-on training. Although students often face ill-defined challenges that require complex problem-solving skills, given suitable scaffolding, they learn to solve practical problems in an authentic setting. Observing students’ learning processes by employing CLI logging is the first step toward providing scaffolding and timely feedback, which we demonstrate in Section VI-B.

B. Learning Environments for Cybersecurity Training

A *security testbed* is a physical or virtual environment that emulates computer systems for research, development, or education. It allows reliable, accurate, repeatable, and safe execution of experiments [9]. A related term is a *cyber range*, which is a platform for interactive simulation environments [10] for security training and experimentation [11]. Examples of testbeds and cyber ranges used in cybersecurity education are DETERLab [12], [13], EDURange [14], GENI [15], [16], and KYPO CRP [17]. For a thorough overview, see [18].

Testbeds and cyber ranges may be complex and costly to operate. However, packaged virtual machines (VMs) have similar benefits and are simpler to use. Thus, they are widely employed in cybersecurity courses [19]. Authors of [20] teach security with locally hosted VMs. Projects SecKnitKit [21] and SEED [22] provide VMs with labs for guided teaching of security topics. Some SEED labs were ported into Labtainers, a training framework based on Docker containers to allow easier use [23]. The containers, similarly to VMs, ensure that the lab software will always run the same. Finally, Vulnhub [24] and Metasploitable [25] share virtual targets for hacking.

Although many learning environments exist, few of them leverage the potential of collecting and analyzing training data. Only EDURange, DETERLab, and Labtainers include host-based logging of commands (see Section II-C for comparison with our approach). The VM solutions do not provide an easy way to monitor student actions [14], [26]. What is more, creating and configuring VMs is time-consuming, and they remain static once created [27].

C. Collection and Analysis of Cybersecurity Education Data

In Capture the Flag (CTF) challenges, the participants solve cybersecurity tasks to obtain text strings called “flags” as proof of solution. Flags include, for example, passwords from a breached service. Previous works collected quantitative data from CTF, such as the flag content or timing [28]. Although such data are relevant for assessing success or failure, they are insufficient for understanding *how* students learn. To illustrate, an incorrect flag means that the student did not complete the challenge. However, it does not reveal the approach to obtaining the flag, making it harder to address the student’s misconceptions. We need detailed information about how the student progressed while solving the tasks. Within the scope of our work, this means the entered shell commands.

Keyloggers such as `ttylog` [29] or recording software such as `asciinema` [30], [31] appear to be a universal solution. A keystroke logging approach has also been used in programming education research [32]. However, keyloggers capture individual characters, which is too fine-grained. Reconstructing the commands then requires parsing large amounts of data. Video recording software, on the other hand, does not allow easy correlation of data from multiple sandboxes. Native Linux solutions, such as commands stored in the `.bash_history` file at a local host, have the same limitation. Therefore, we need a compact machine-readable format suitable for logs from cybersecurity training.

Irvine et al. [23] collected files containing the `stdin` and `stdout` streams from Docker containers of students working in Labainers. Labainers software then used the files to check if a student achieved a certain goal (such as completing a learning task). However, the data were not aggregated in real-time from all containers to gain an overall understanding. Students had to send the log files manually after the training.

Weiss et al. [14], [26] logged shell commands in EDURange to visualize the students' steps. In follow-up work, Mirkovic et al. [33] presented the system for logging and analyzing commands in DETERLab. The system monitored terminal input and output and compared them with task milestones to automatically assess student progress. Similarly [34], [35] also implemented host-based logging to reconstruct the students' progress. Although the goal of the above-mentioned works is similar to ours, their approach is complementary. We extend the state of the art in the following aspects:

- 1) Focus on understanding of students' learning processes, not only the outcomes (succeeding or failing a task) when analyzing data from cybersecurity training (see Section V).
- 2) Real-time, automatic, and secure forwarding of logs.
- 3) Support of more Linux-based operating systems.
- 4) Enabling the logging of commands from Metasploit shell [25], an essential cybersecurity tool.
- 5) Easier deployment to other existing learning environments. The deployment is automated using Ansible [36], a well-established software for configuration management.
- 6) Generic output format that can be processed by custom software or integrated with an existing one, such as Weka [37] or ELK [38].
- 7) Demonstration of innovative analyses of command histories (see proof of concept in Section VI-B).

D. Tools for Teaching Linux Shell Commands

To provide a complementary viewpoint, we briefly examine systems for teaching Linux CLI in general, without an explicit focus on cybersecurity. TuxLab [39] and uAssign [40] are two such systems. Both leverage virtualization, since they are based on Docker containers. Students can use them for practicing Linux commands, which are immediately evaluated in the container and autograded. The difference from our work is that neither of them collects the commands for further analysis. TermAdventure [41], on the other hand, stores the command logs, but is not a standalone logging toolset as we present here.

III. TOOLSET FOR COMMAND LOGGING

Figure 1 shows the generic topology of a sandboxed network for cybersecurity training. It contains an attacker machine controlled by the trainee and several vulnerable hosts (black-colored elements). Since we assume one trainee per sandbox,¹ each trainee receives its identical copy. On top of that, we configure the (blue-colored) infrastructure elements responsible for collecting CLI logs from the hosts.

¹If teamwork is needed, a group can access the same sandbox. However, the toolset cannot distinguish individuals within the group. The mapping of logs to students is done only based on the sandbox ID.

A. How Does the Toolset Work?

Each host in the sandbox has two network interfaces: *training* (black) and *management* (blue). The former enables the trainee to interact with hosts in the sandbox and complete the training tasks. The latter exists in the network to initially configure the sandbox. We utilize it for host-based logging and forwarding of the logs to central storage.

The logs are generated on all sandboxed machines that a trainee accesses. Every time (s)he enters a command on a Linux host, it is stored on that host using the Syslog protocol [42]. The log is then instantly forwarded to the Syslog server of the sandbox. Finally, the logs from Syslog servers of all sandboxes are securely forwarded to a central storage (see Figure 1). The intermediate step with the Syslog server is optional and can be bypassed to log directly into the central storage. The data can be accessed continuously in real-time or after the training.

To verify that all submitted commands arrive correctly, we performed multiple manual tests. We deployed the data collection in two learning environments [17]: a complex cloud-based platform KYPO CRP and a lightweight, locally virtualized platform Cyber Sandbox Creator. Section IV and Section V present the study that analyzes the collected data.

The configuration files that supplement this paper (see Section VII-A) enable others to deploy this toolset in their physical or virtual environment. Instructors who do not have access to a cloud can prepare VMs provisioned with the logging. Students can download them and use them on their computers, and instructors can observe the commands at the central storage deployed at their institution or in a public cloud.

B. Implementation Details

Our approach works for the Bash shell [43] since it employs a Bash variable `$PS0`. Its value is expanded "after reading a command and before the command is executed" [43] to a call of the logger [44] command, which can enter arbitrary messages into the host's system log socket (typically `/dev/log`). This `AF_LOCAL` socket is widely used by standard logging daemons and services. Regardless of the used Linux distribution, modern logging daemons can transform, store, and forward log entries flexibly. We use `rsyslogd` daemon [45] to forward the logs to a remote location via the Syslog protocol, as described above. Since the logging is based on Syslog, it can be deployed generically. Any application can log into Syslog, or at least in a plain text file that can be forwarded to Syslog.

C. What Data Are Collected?

Figure 2 shows a command executed on the Attacker machine and the subsequently generated log record. We collect each line of the terminal input: the entire command, including all arguments. All inputs are timestamped (up to microsecond precision) and identify the host and sandbox in which the command was executed. The hosts must have a synchronized time to allow correlating the commands. We also collect essential metadata about the execution, such as the working directory and type of the command to distinguish between the Bash and Metasploit shell.

- *Machine-readable output*: the logs are stored in a universal, semi-structured format (JSON), which simplifies their further processing.
- *Installability*: the toolset is open-source and uses only standard Linux components. Thus, it is widely applicable and compatible.

E. Limitations of the Logging Toolset

The practical application of the toolset in two learning environments revealed some of its shortcomings:

- The data collection is constrained to the commands entered in a shell. Any edits of commands before the submission (pressing the Enter key) are not recorded. We also do not capture any other data indicating progress, such as mouse clicks, network traffic, or filesystem changes.
- We do not collect the command's return code or output, since logger is called before the command is executed. Neither we collect the submission of control sequences such as Ctrl + C or responses to prompts.
- The logging currently works only on Linux hosts. However, it could be analogously extended to Windows hosts to log PowerShell [47] instead of Linux shell commands.
- An advanced trainee with administrator privileges on a machine could tamper with or turn off the logging from that host. However, the instructor would notice the absence of logs. A future challenge is to hide the logging to prevent both intentional and accidental changes.

IV. STUDY SETUP AND METHODS

We now describe the study that involved deploying our toolset in authentic educational settings on multiple occasions. The study addresses our goal of understanding trainees' learning processes by examining their command histories. Moreover, we evaluate the applicability of the toolset and demonstrate the value of the collected data.

A. Teaching Context and Training Content

During the Fall 2019 and Spring 2020 semesters, the students of an undergraduate course Seminar on the Simulation of Cyber Attacks [48] at Masaryk University created eight educational cybersecurity games. The games involve practicing security skills using Linux Bash shell in a virtual network similar to the one in Figure 1. The game tasks focus on gradually compromising the security of a simulated IT system, following the stages of Mandiant's Attack Lifecycle [49, p. 27]. All the games exploit the same vulnerabilities, and they differ only in the fictitious background story and the task presentation.

The games' learning objectives and tasks are centered around using CLI tools in Linux distribution called Kali [50], [51]. Each game starts by scanning a target network using nmap [52]. Nmap can determine what services are running on the network hosts and their port numbers. This corresponds to the first stage *Initial Reconnaissance* in the Mandiant's Attack Lifecycle. The next task is to identify a vulnerable service on the target and exploit it using Metasploit [25], a penetration testing software that contains exploits for various vulnerabilities. After

exploiting the service, the attacker accesses the machine, copies a private SSH key, cracks its passphrase, and uses the key to connect to another host via SSH. For the passphrase cracking, John the Ripper tool was recommended. Given a text file with passphrases, the john command performs a dictionary attack.

When the students of the course finished creating their games, we added the logging module into the game configuration files. Although the games differed in technical details, the toolset was applicable to all of them. Afterward, all games were deployed in the KYPO CRP [17].

B. Participants

From December 2019 to July 2020, we hosted four events (game sessions), during which a total of 50 participants (*trainees*) played one of the eight games mentioned above. Each event lasted approximately two hours and had a diverse participant population:

- undergraduate or graduate students of computer science at our university, who have a basic cybersecurity background,
- security professionals, and
- senior high school students and bachelor students of other universities, the finalists of the Czech national cybersecurity competition.

The trainees attended voluntarily because of their interest in security and were not incentivized in any way. None of the trainees played any of the eight games before.

C. Privacy and Ethical Measures

Before the training sessions, we discussed the research with the institutional review board of our university. We obtained a waiver from the ethical committee since we intentionally do not collect any personal information; the data are anonymous and cannot be linked to individuals. The trainees agreed to the anonymized data collection via informed consent. We ensured they would not be negatively affected by the research, and they had the right to stop participating at any time without any restrictions. After the data collection, we manually checked whether they typed any personal information in the CLI and anonymized it. Only one student typed his real university login.

D. Methods of Data Processing and Analysis

Since all games had identical learning objectives and vulnerabilities, we merged all the logs. Then, we wrote Python scripts that parsed and analyzed the data. In addition to quantitative statistics, two security instructors (authors of this paper) exploratively analyzed the output of the automated data processing. Their goal was to manually uncover deeper insights. Since the scripts automatically presented aggregated results, their interpretation was straightforward.

The analysis focused on cybersecurity tools. When preprocessing the data, we removed extraneous whitespace and, when applicable, disregarded the order of the arguments, which was sometimes irrelevant. Apart from this sanitation, we did not modify the trainee data.

V. RESULTS

This section presents the analysis results to answer the three questions posed in Section I. Each question is addressed in a separate subsection. The results demonstrate the value of the collected command-line data. We also report our teaching experience, educational insights, and implications for practice.

A. Use Case 1: Executed Commands

Table I provides an overview of the collected commands. The most frequently used cybersecurity tool was `nmap` with 240 occurrences out of the 4439 commands in the whole dataset. This was expected, since the initial task in all games was to scan the target network. When looking not only at cybersecurity-specific commands, `ls` was the most frequent (568 times), followed by `cd` (320 times), which corresponds to the nature of the training that required traversing the filesystem.

TABLE I
OVERALL AND PER-TRAINEE STATISTICS OF COMMANDS ENTERED BY THE 50 TRAINEES (PLAYERS OF THE CYBERSECURITY GAMES).

Command type	Total	Min	Max	Median	Avg	Stdev
Linux Bash	3121	4	266	45	62.4	57.3
Metasploit	1318	0	213	17	26.4	36.3
Both types	4439	5	358	68	88.8	74.9

We now analyze how the trainees used three key command-line tools in the games: `nmap`, `ssh`, and `john`. We also examine the logs from the Metasploit console.

1) *Terminal Commands*: `Nmap` requires a target specification, such as a hostname or an IP address. However, in 14 cases, the trainees attempted an `nmap` scan but did not specify a target, which indicates misconceptions. Another interesting observation was that although the task was to scan ports, the option `-sn` that *disables* the port scan was used 13 times. An IP address that belonged to a router was scanned 7 times. In 8 cases, the trainee made a typo in the target IP address since it was off by one digit, but 9 other times they wanted to scan an address that was not even in the training sandbox.

Using these findings, instructors can provide targeted feedback to students, for example, explain the causes of the mistakes to the class. They can also use these insights to devise a scoring rubric for assessment (for example, scanning a correct IP address: 1 point).

Later in the game, the trainees obtained a private SSH key. They had to crack its passphrase by executing a dictionary attack using `john`. The most common mistakes among 201 `john` commands in the dataset were not specifying the dictionary or setting the option `--wordlist` to a folder instead of a file. After cracking the passphrase, the trainees usually had no problem connecting to a target host using `ssh`.

2) *Metasploit Commands*: Another task was to exploit a host using Metasploit. The correct course of action was:

- 1) find suitable exploit modules using `search`,
- 2) select a correct module from the list using `use`,
- 3) configure module options using `set`, and
- 4) run the module using `exploit` or `run`.

The logs revealed the following misconceptions and unusual activities of the trainees:

- instantly selecting the correct module without searching for it in Metasploit, which suggests previous knowledge of the attack or outside help,
- selecting a wrong exploit and attempting to run it,
- attempting to configure the module parameters with `set` before selecting the module with `use`,
- incorrectly setting the machine initiating the exploit (parameter `LHOST`) to the target host or other values,
- attempting to run the exploit before setting its parameters.

B. Use Case 2: First Commands Entered

We reconstructed the first high-level action of each trainee from the first few commands. To illustrate, Table II is an example of the first few commands of one trainee who started scanning the target network. The trainee's intention was apparent despite some errors and irrelevant commands.

TABLE II
TIMELINE OF THE TRAINING'S START FOR A SELECTED TRAINEE (SHORTENED), ALONG WITH A TIME DELAY BETWEEN TWO COMMANDS.

Time [m:ss]	Command delay [s]	Command
0:00	N/A	<code>nmap --help</code>
0:55	55	<code>nmap 172.18.1.5</code>
3:11	136	<code>pwd</code>
3:14	3	<code>ls</code>
3:26	12	<code>nmap --help</code>
4:10	44	<code>nmap -sV --p 10000 172.18.1.5</code>
4:17	7	<code>nmap -sV -p 10000 172.18.1.5</code>

This work pattern was also the most commonly observed. For 30 out of the 50 trainees, their first command was `nmap`, in which half of them correctly set a target. In five other cases, trainees started by familiarizing themselves with the network configuration using `ifconfig` or `ping`.

At the start of the training, the trainees did not deviate from the task and did not execute unusual commands. Nevertheless, analyzing the first actions can still be beneficial. If an instructor sees that a trainee started correctly, (s)he may devote more time to aid others. If an instructor observes off-task behavior at the beginning, (s)he can intervene early and help the trainee by providing targeted hints.

C. Use Case 3: Time-based Statistics

Table III shows the statistics for two variables. The first is the time spent playing, calculated as the time difference between the first and the last command. The trainees usually spent about an hour on the training, which is enough to record interactions in depth. A slight limitation is that we cannot determine the exact end time of the training from the command-line logs only. The last submitted command does not necessarily mean the end of the work on the task. The student could have continued by using a GUI tool, consulting their progress with another person, or searching for online help, for example.

TABLE III

TOTAL TIME SPENT ON THE TRAINING AND THE TIME DIFFERENCES (GAPS) BETWEEN TWO SUCCESSIVE COMMANDS FOR THE 50 TRAINEES.

	Min	Max	Median	Avg	Stdev
Game time [m:ss]	16:10	1169:20	77:12	108:12	159:05
Avg gap [m:ss]	0:00	631:57	0:15	1:13	10:17

As expected, the time spent in the game positively correlates with the number of submitted commands (Spearman's $\rho = 0.79$). This correlation is statistically significant (two-sided p-value $< 10^{-11}$, calculated by the Python `scipy.stats` module [53]).

The second variable in Table III is the time difference between two successive commands. The median, average, and standard deviation were first computed individually for each trainee. The table reports the averages of these individual quantities. The trainees usually submitted a command every minute or so, which seems appropriate since they also read the documentation and contemplated their approach.

The near-zero time differences occurred when a trainee copy-and-pasted a command or used shortcuts such as “arrow up” to repeat the previous command. Minor time differences were observed when a trainee made an error, immediately recognized it, and fixed it (see the last two lines in Table II). We also observed differences smaller than the 15-second median when the trainee was unsure how to use a command and applied a trial-and-error method. Finally, a burst of commands in a short period may indicate scripting or attempts to tamper with the system, although we did not observe this in our data. On the other hand, long idle periods can mean disengagement, need for help, or technical issues with the learning environment.

VI. STUDY DISCUSSION

This section discusses the study limitations and proposes future work that will use the command-line data.

A. Limitations of the Study

The results have two main limitations. First, the logs can only indicate a problem, not explain why it occurs. We proposed possible explanations, but they are yet to be confirmed in an in-depth qualitative study with students. Nevertheless, the logs authentically record the trainees' progress, creating a basis for targeted educational interventions.

Second, we focused only on the teacher's perspective and enhancing the instructor's awareness of the classroom. It would be interesting to close the feedback loop and evaluate how this educational intervention impacts the students.

B. Educational Implications and Future Work

A CLI is difficult for beginners since it lacks a graphical interface and is strict about using the commands properly. Our toolset for collecting and analyzing the commands allows instructors to better understand students' behaviors, which is a prerequisite for helping them learn.

This paper demonstrated three initial use cases. For the future, we see more advanced research opportunities that require high-quality command-line data. The possibilities are to:

- Qualitatively analyze the causes of student mistakes and misconceptions (for example, using interviews and think-aloud protocols) and address them in the classroom.
- Perform pattern mining to find interesting association rules and sequences in the student data and visualize them.
- Provide automated, personalized feedback to many students at scale, both in real-time and after the training.
- Assess students, including partial credit for unfinished solutions. Students can be graded for using (or not using) a specific set of commands and their arguments.
- Measure learning in a pre-test/post-test study framework to see whether the interventions achieved by analyzing command-line data helped improve student achievement.

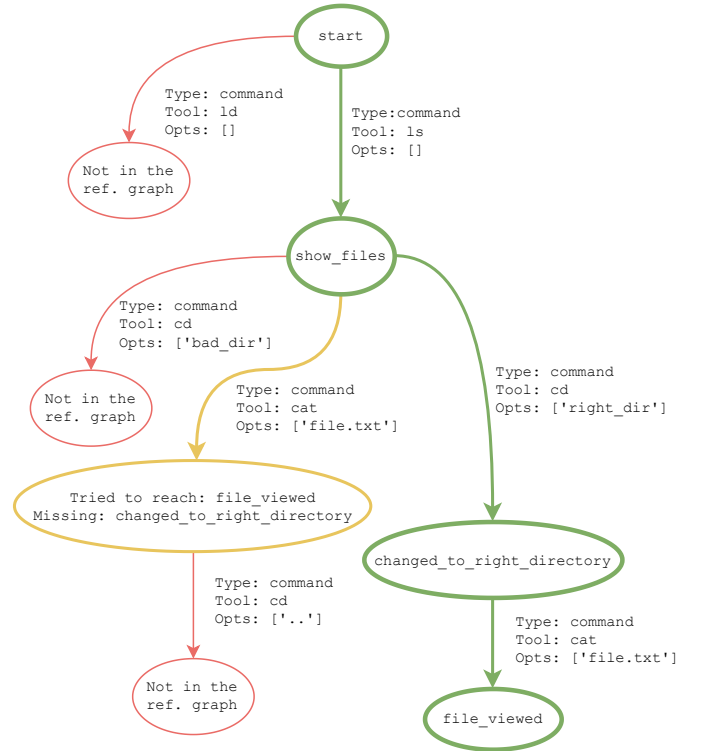


Fig. 4. Graph visualization of student's progress through an exercise that employs command-line tools.

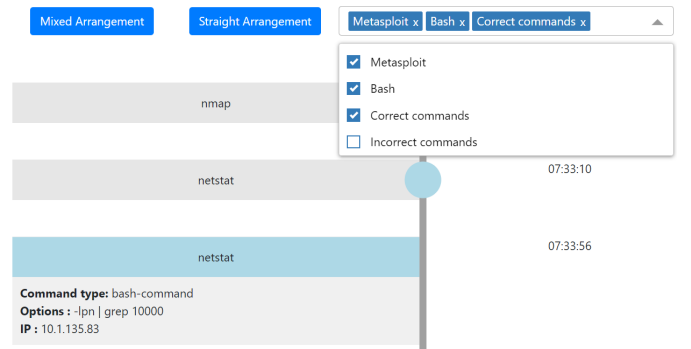


Fig. 5. Timeline visualization of student's progress through an exercise that employs command-line tools.

In our further ongoing research, we investigate methods for graphical modeling of student's steps (see Figure 4). The proof-of-concept application automatically processes students' command histories and maps them to a reference solution of the training. It can display correctly achieved steps, errors, and potential omissions. Another proof-of-concept web application in Figure 5 displays an interactive timeline of the submitted commands, allowing to filter correct or erroneous attempts.

VII. CONCLUSIONS

Command-line tools are used in almost all areas of computing. In cybersecurity, they are essential both to learn about cyber attacks and to set up defense mechanisms. Even today, when security experts may also use sophisticated graphical interfaces, tools controlled via a shell still form an essence of cybersecurity operations. Therefore, learning and understanding these tools is crucial to mitigate the current cyber threats.

We presented a modular toolset for logging commands from Bash and Metasploit shell. It enables a data-driven understanding of students' approaches to practicing cybersecurity, system administration, and networking. We deployed the toolset in two learning environments and showed the value of the gained data by analyzing student approaches. Since the toolset can be applied in many other contexts, it can foster further research and development, as we showed in Section VI-B.

Observing student actions in the training session is invaluable for instructors to provide timely help to students during the session or detailed feedback afterward. This is a substantial innovation of typical practice when no information or only the (in)correctness of the task solutions are recorded. Therefore, this paper contributes to addressing a challenge that computing instructors face. Although tinkering with a CLI is common for students who are learning, if errors occur repeatedly, instructors should know about that to intervene appropriately.

The logging toolset helps not only instructors and students of cybersecurity, but also in other courses that feature command-line tools, such as operating systems or networking. Observing student actions is especially relevant in distance education when instructors have limited awareness of what students do on their computers. This proved invaluable in our online teaching practice during the COVID-19 pandemic in 2020–2021. Most of the analyzed data were collected from remotely held training sessions, in which we would otherwise not have any awareness about the students' progress.

A. Contributed Artifacts

This paper's supplementary materials contain four core contributions for instructors, educational researchers, and developers of interactive learning environments.

- 1) The toolset that can be integrated with existing learning environments. We share configuration scripts (Ansible [36] roles) that automatically deploy host-based logging to the training sandbox similar to the one in Figure 1.
- 2) Educators can familiarize themselves with the toolset. They can automatically set up a learning environment from scratch on a personal computer using a single command.

We enable this by providing a Vagrant [54] definition file that builds a locally virtualized learning environment for a cybersecurity game. Moreover, the game includes exercise assignments that can be completed in the learning environment to practice cybersecurity skills.

- 3) We publish the collected data from the training sessions studied in this paper, as well as others we conducted.
- 4) We provide analytical software that was used to process and analyze the data.

All artifacts are linked from a single Zenodo repository [55] and are available under permissive open-source licenses. Each artifact contains detailed instructions to allow educators, researchers, and developers to use our results in their contexts.

ACKNOWLEDGMENT

This research was supported by ERDF “CyberSecurity, CyberCrime and Critical Information Infrastructures Center of Excellence” (No. CZ.02.1.01/0.0/0.0/16_019/0000822).

REFERENCES

- [1] Joint Task Force on Cybersecurity Education. (2017) Cybersecurity Curricular Guideline. ACM/IEEE/AIS SIGSEC/IFIP. [Online]. Available: <http://cybered.acm.org>
- [2] (ISC)². (2020) Cybersecurity Workforce Study. [Online]. Available: <https://www.isc2.org/Research/Workforce-Study>
- [3] C. Romero, S. Ventura, M. Pechenizkiy, and R. S. Baker, Eds., *Handbook of educational data mining*. Boca Raton, FL, USA: CRC Press, 2010.
- [4] C. Lang, G. Siemens, A. Wise, and D. Gašević, Eds., *Handbook of Learning Analytics*, 1st ed. Society for Learning Analytics Research (SoLAR), 2017.
- [5] M. Prince, “Does Active Learning Work? A Review of the Research,” *Journal of Engineering Education*, vol. 93, no. 3, pp. 223–231, 2004. [Online]. Available: <https://doi.org/10.1002/j.2168-9830.2004.tb00809.x>
- [6] 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>
- [7] G. Petty, *Teaching Today: A Practical Guide*. Nelson Thornes, 2009.
- [8] A. V. Robins, L. E. Margulieux, and B. B. Morrison, “Novice programmers and introductory programming,” in *The Cambridge Handbook of Computing Education Research*, S. A. Fincher and A. V. Robins, Eds. Cambridge, United Kingdom: Cambridge University Press, 2019, ch. 9, pp. 231–275.
- [9] C. Siaterlis and B. Genge, “Cyber-physical testbeds,” *Commun. ACM*, vol. 57, no. 6, pp. 64–73, Jun. 2014. [Online]. Available: <http://doi.acm.org/10.1145/2602575>
- [10] European Cyber Security Organisation, “Understanding Cyber Ranges: From Hype to Reality,” (ECSO), Tech. Rep., 2020. [Online]. Available: <https://ecs-org.eu/documents/publications/5ea47a760f24f.pdf>
- [11] J. Davis and S. Magrath, “A survey of cyber ranges and testbeds,” Defence Science and Technology Organisation (DSTO), Tech. Rep. DSTO-GD-0771, 2013. [Online]. Available: <https://apps.dtic.mil/sti/citations/ADA594524>
- [12] T. Benzel, “The Science of Cyber Security Experimentation: The DETER Project,” in *Proceedings of the 27th Annual Computer Security Applications Conference*, ser. ACSAC '11. New York, NY, USA: ACM, 2011, pp. 137–148. [Online]. Available: <https://doi.org/10.1145/2076732.2076752>
- [13] J. Mirkovic and P. A. H. Peterson, “Class Capture-the-Flag Exercises,” in *2014 USENIX Summit on Gaming, Games, and Gamification in Security Education (3GSE 14)*. San Diego, CA: USENIX Association, 2014. [Online]. Available: <https://www.usenix.org/conference/3gse14/summit-program/presentation/mirkovic>
- [14] R. Weiss, F. Turbak, J. Mache, and M. E. Locasto, “Cybersecurity Education and Assessment in EDURange,” *IEEE Security & Privacy*, vol. 15, no. 3, pp. 90–95, 2017.

- [15] M. Berman, J. S. Chase, L. Landweber, A. Nakao, M. Ott, D. Raychaudhuri, R. Ricci, and I. Seskar, "GENI: A federated testbed for innovative network experiments," *Computer Networks*, vol. 61, pp. 5–23, 2014. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1389128613004507>
- [16] X. Mountrouidou, X. Li, and Q. Burke, "Cybersecurity in liberal arts general education curriculum," 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. 182–187. [Online]. Available: <http://doi.acm.org/10.1145/3197091.3197110>
- [17] J. Vykopal, P. Čeleda, P. Seda, V. Švábenský, and D. Továřík, "Scalable Learning Environments for Teaching Cybersecurity Hands-on [in press]," in *Proceedings of the 51st IEEE Frontiers in Education Conference*, ser. FIE '21. New York, NY, USA: IEEE, 10 2021, pp. 1–9.
- [18] M. M. Yamin, B. Katt, and V. Gkioulos, "Cyber ranges and security testbeds: Scenarios, functions, tools and architecture," *Computers & Security*, vol. 88, p. 101636, 2020. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0167404819301804>
- [19] M. O'Leary, "Innovative pedagogical approaches to a capstone laboratory course in cyber operations," in *Proceedings of the 2017 ACM SIGCSE Technical Symposium on Computer Science Education*, ser. SIGCSE '17. New York, NY, USA: ACM, 2017, pp. 429–434. [Online]. Available: <http://doi.acm.org/10.1145/3017680.3017720>
- [20] M. Timchenko and D. Starobinski, "A simple laboratory environment for real-world offensive security education," in *Proceedings of the 46th ACM Technical Symposium on Computer Science Education*, ser. SIGCSE '15. New York, NY, USA: ACM, 2015, pp. 657–662. [Online]. Available: <http://doi.acm.org/10.1145/2676723.2677225>
- [21] A. Siraj, B. Taylor, S. Kaza, and S. Ghafoor, "Integrating security in the computer science curriculum," *ACM Inroads*, vol. 6, no. 2, pp. 77–81, May 2015. [Online]. Available: <http://doi.acm.org/10.1145/2766457>
- [22] W. Du, "SEED: Hands-On Lab Exercises for Computer Security Education," *IEEE Security & Privacy*, vol. 9, no. 5, pp. 70–73, 2011.
- [23] C. E. Irvine, M. F. Thompson, M. McCarrin, and J. Khoslim, "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, 2017. [Online]. Available: <https://www.usenix.org/conference/ase17/workshop-program/presentation/irvine>
- [24] VulnHub. (2021) Vulnerable By Design. [Online]. Available: <https://www.vulnhub.com>
- [25] Offensive Security. (2021) Metasploit Unleashed. OffSec Services Limited. [Online]. Available: <https://www.offensive-security.com/metasploit-unleashed/>
- [26] R. Weiss, M. E. Locasto, and J. Mache, "A Reflective Approach to Assessing Student Performance in Cybersecurity Exercises," in *Proceedings of the 47th ACM Technical Symposium on Computing Science Education*, ser. SIGCSE '16. New York, NY, USA: ACM, 2016, pp. 597–602. [Online]. Available: <http://doi.acm.org/10.1145/2839509.2844646>
- [27] Z. C. Schreuders, T. Shaw, M. Shan-A-Khuda, G. Ravichandran, J. Keighley, and M. Ordean, "Security Scenario Generator (SecGen): A Framework for Generating Randomly Vulnerable Rich-scenario VMs for Learning Computer Security and Hosting CTF Events," 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/schreuders>
- [28] R. Ošlejšek, V. Rusňák, K. Burská, V. Švábenský, and J. Vykopal, "Visual Feedback for Players of Multi-Level Capture the Flag Games: Field Usability Study," in *Proceedings of the IEEE Workshop on Visualization for Computer Security (VizSEC)*. New York, NY, USA: IEEE, 2019.
- [29] Rocasa Code. (2018) Ttylog. GitHub. [Online]. Available: <https://github.com/rocasat/ttylog>
- [30] M. Kulik and contributors. (2021) asciinema. [Online]. Available: <https://asciinema.org/>
- [31] I. Hassan, "Teaching Cybersecurity to Computer Science Students Utilizing Terminal Sessions Recording Software as a Pedagogical Tool," in *2020 IEEE Frontiers in Education Conference (FIE)*. New York, NY, USA: IEEE, 2020, pp. 1–8.
- [32] P. Ihtola, A. Vihavainen, A. Ahadi, M. Butler, J. Börstler, S. H. Edwards, E. Isohanni, A. Korhonen, A. Petersen, K. Rivers, M. A. Rubio, J. Sheard, B. Skupas, J. Spacco, C. Szabo, and D. Toll, "Educational Data Mining and Learning Analytics in Programming: Literature Review and Case Studies," in *Proceedings of the 2015 ITICSE on Working Group Reports*, ser. ITICSE-WGR '15. New York, NY, USA: ACM, 2015, pp. 41–63. [Online]. Available: <https://doi.org/10.1145/2858796.2858798>
- [33] J. Mirkovic, A. Aggarwal, D. Weinman, P. Lepe, J. Mache, and R. Weiss, "Using Terminal Histories to Monitor Student Progress on Hands-on Exercises," in *Proceedings of the 51st ACM Technical Symposium on Computer Science Education*, ser. SIGCSE '20. New York, NY, USA: ACM, 2020, p. 866–872. [Online]. Available: <https://doi.org/10.1145/3328778.3366935>
- [34] M. Andreolini, V. Colacino, M. Colajanni, and M. Marchetti, "A Framework for the Evaluation of Trainee Performance in Cyber Range Exercises," *Mobile Networks and Applications*, vol. 25, 2019.
- [35] Z. Tian, Y. Cui, L. An, S. Su, X. Yin, L. Yin, and X. Cui, "A Real-Time Correlation of Host-Level Events in Cyber Range Service for Smart Campus," *IEEE Access*, vol. 6, pp. 35 355–35 364, 2018.
- [36] RedHat. (2021) Ansible. [Online]. Available: <https://www.ansible.com/>
- [37] E. Frank, M. A. Hall, and I. H. Witten, "The WEKA Workbench. Online Appendix for Data Mining: Practical Machine Learning Tools and Techniques," 2016. [Online]. Available: <https://waikato.github.io/weka-wiki/>
- [38] Elasticsearch. (2021) What is the ELK Stack? Elasticsearch B.V. [Online]. Available: <https://www.elastic.co/what-is/elk-stack>
- [39] D. Brown. (2017) TuxLab. Carnegie Mellon University. [Online]. Available: <http://tuxlab.org/>
- [40] J. Bailey and C. Zilles, "UAssign: Scalable Interactive Activities for Teaching the Unix Terminal," in *Proceedings of the 50th ACM Technical Symposium on Computer Science Education*, ser. SIGCSE '19. New York, NY, USA: ACM, 2019, p. 70–76. [Online]. Available: <https://doi.org/10.1145/3287324.3287458>
- [41] M. Šuppa, O. Jariabka, A. Matejov, and M. Nagy, "TermAdventure: Interactively Teaching UNIX Command Line, Text Adventure Style," in *Proceedings of the 26th ACM Conference on Innovation and Technology in Computer Science Education*, ser. ITICSE '21. New York, NY, USA: Association for Computing Machinery, 2021, pp. 108–114. [Online]. Available: <https://doi.org/10.1145/3430665.3456387>
- [42] R. Gerhards, "The Syslog Protocol," RFC 5424, Mar. 2009. [Online]. Available: <https://rfc-editor.org/rfc/rfc5424.txt>
- [43] C. Ramey and B. Fox. (2020, 12) The GNU Bash Reference Manual, for Bash, Version 5.0. Free Software Foundation. [Online]. Available: <https://www.gnu.org/savannah-checkouts/gnu/bash/manual>
- [44] Man7. logger(1) – Linux manual page. [Online]. Available: <https://man7.org/linux/man-pages/man1/logger.1.html>
- [45] Adiscon GmbH. (2020) The Rocket-Fast Syslog Server – Rsyslog. [Online]. Available: <https://www.rsyslog.com/>
- [46] International Organization for Standardization, "Systems and software engineering – Systems and software Quality Requirements and Evaluation (SQuaRE)," ISO/IEC, Tech. Rep. 25010:2011(en), 2011. [Online]. Available: <https://www.iso.org/standard/35733.html>
- [47] Microsoft. (2021) Get-History. [Online]. Available: <https://docs.microsoft.com/en-us/powershell/module/microsoft.powershell.core/get-history>
- [48] 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>
- [49] Mandiant, "APT1: Exposing One of China's Cyber Espionage Units," FireEye, Inc, Tech. Rep., 2013. [Online]. Available: <https://www.fireeye.com/content/dam/fireeye-www/services/pdfs/mandiant-apt1-report.pdf>
- [50] OffSec Services Limited. (2021) Kali Linux Tools Listing. [Online]. Available: <https://tools.kali.org/tools-listing>
- [51] A. Ghafarian, "Using Kali Linux Security Tools to Create Laboratory Projects for Cybersecurity Education," in *Proceedings of the Future Technologies Conference (FTC) 2018*. Cham: Springer International Publishing, 2019, pp. 358–367.
- [52] G. Lyon. (2021) Nmap Network Scanning. Nmap. [Online]. Available: <https://nmap.org/book/man.html>
- [53] SciPy. (2021) scipy.stats.pearsonr. [Online]. Available: <https://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.pearsonr.html>
- [54] HashiCorp. (2021) Vagrant. [Online]. Available: <https://www.vagrantup.com/>
- [55] V. Švábenský, J. Vykopal, D. Továřík, and P. Čeleda, "Dataset: Toolset for Collecting Shell Commands and Its Application in Hands-on Cybersecurity Training," 2021. [Online]. Available: <https://doi.org/10.5281/zenodo.5126693>