

# Use of Cyber Attack and defense agents in Cyber Ranges: A Case Study

Muhammad Mudassar Yamin and Basel Katt

Norwegian University of Science and Technology, Gjøvik 2815, Norway  
{muhammad.m.yamin,basel.katt}@ntnu.no

**Abstract.** With the ever-changing cybersecurity landscape, the need for continuous training for a new cybersecurity skill set is a requirement. This training can be delivered on platforms like cyber ranges. Cyber ranges support training by providing a simulated or emulated representation of computer network infrastructure besides additional training and testing services. Cyber attack and defense skills can be gained by attacking and defending the infrastructure; however, to provide more realistic training, there is a need for necessary friction in the environment, which can be related to both the attacker’s and defender’s actions. The actions of human teams—both attackers and defenders—provide this friction. Involving human teams in large-scale cybersecurity exercises is relatively inefficient and not feasible for standardizing training because different teams apply different tactics. Currently, the proposed solutions for cyber range training platforms focus on automating the deployment of the cybersecurity exercise infrastructure but not on the execution part. This leaves room for improving exercise execution by adding realism and efficiency. This research presents an agent-based system that emulates cyber attack and defense actions during cybersecurity exercise execution; this helps provide realistic and efficient cybersecurity training. To specify agents’ behavior and decision making, a new model, called the execution plan (EP), was developed and utilized in this work.

**Keywords:** Cyber attack agent · Cyber defense agent · Cyber Range · Security Exercise.

## 1 Introduction

Conducting operational cybersecurity exercises is a difficult and inefficient task [YK18a]. We have found that automating the different roles involved in cybersecurity exercises can reduce these inefficiencies [YKT<sup>+</sup>18]. These roles primarily involve a human team required to set up the exercise technical network infrastructure. Additionally, there is a team that attacks the deployed infrastructure as an attacker and a team that defends it as a defender [YKG20]. There can be multiple ways with which a cybersecurity exercise can be executed that may or may not involve both attackers and defenders at the same time. However, in a realistic environment, to train attackers, the systems being attacked are expected to be defended by somebody. Because of shortage in the cybersecurity skills, it is very difficult to

find people with the relevant skill set [YK19a] to conduct continuous cybersecurity exercises. Moreover, different people have different tactics and techniques in cybersecurity operations, making a standardized assessment of cybersecurity exercises difficult [HWD<sup>+</sup>17]. Therefore, there is a need for automating attack and defense roles in cybersecurity exercises. Despite its importance, there is a lack of research dealing with realism and efficiency in cybersecurity exercise execution in cyber ranges. Most of the related work deals with automating the creation and deployment of the exercise infrastructure. This leaves room for researchers to improve the realism and efficiency of cybersecurity exercise execution. We tackle this issue by proposing an agent-based system, one in which we model the attacker and defender roles and automate their execution as required. In particular, we developed a new modeling technique: execution plan (EP). EP is a multi-level model for specifying behavior and decision-making process for attacking and defending agents. We argue that such agents will add the necessary friction in the cybersecurity exercise environment to make them realistic and reduce the human input of attackers and defenders to make exercise execution more efficient. Therefore, in the current research, we aim to answer the following research question (RQ):

***RQ:** How can cybersecurity attack and defense scenario models be executed autonomously in a cybersecurity exercise to make cybersecurity exercise execution realistic and efficient?*

We present our experience in developing and using cyber-attack and defense agents during cybersecurity exercises against human adversaries. The current paper focuses on the conceptual design, agent decision modeling, practical implementation and user experience with cyber-attack and defense agents. The system is evaluated using a case study against defined benchmarks. The case study involved an operational cybersecurity exercise in which the attack and defense agents were deployed along with human participants. The attack agents were used to create forensic traces for blue team members, which were verified in their forensic reports. At the same time, the defense agents were used to add friction or realism in the exercise environment and were evaluated based on the compromised status of the systems on which they were deployed. The paper is structured as follows: first, we share the research background and related work. After that, we state our research methodology. We then present the conceptual design and technical implementation of cyber attack and defense agents. Finally, we present the experimental results and conclude the article.

## 2 Research Background and Related Work

Cybersecurity is important both against individual inexperienced hackers and against coordinated teams of hackers that might or might not have governmental support. The conventional methods of teaching cybersecurity include lectures, assignments, seminars, and hands-on labs. Hands-on methods include competitions, challenges, and exercises such as the following:

1. Capture the flag (CTF), which focus on attacking, defending, or attack and defense at the same time.
2. Cyber defense exercises (CDX), which focus primarily on defending.

These competitions, challenges, and exercises are conducted in isolated safe environments, which are called *cyber ranges*. Cyber ranges can host single standalone challenges for CTF competitions or represent a sector/organization’s complex computer network infrastructure for CDX exercises [YKG20].

An important element missing from virtual environments is an element of active opposition. A training simulation environment has static defenses (access control, firewalls, fixed set of intrusion methods, etc.) but lacks active opposition (e.g., monitor logs, blocked connections, exploit switching, or information gathering). This results in the cyber operators behaving as though opponents do not have a tangible existence and higher-level goals. Second, it ignores an opportunity to tailor the student’s learning experience through adjustable adversary behavior [JON<sup>+</sup>15].

Cyberwarfare is an imminent threat to military and civilian systems; it could damage the economy and national security. Cyber aggressors are guided by cognitive behavior (script-kiddies, ideological activists, investigators, financial criminals, intelligence agents, or cyber warfighters). Building an effective training system for cyberwarfare currently faces many barriers. Current training environments are unable to capture the purpose, creativity, and adaptability of cyber warfighters, and cyber warfighters need to be effectively trained against a cunning and adaptive opponents.

We conducted a detailed study on cyber ranges [YKG20] and identified that after 2014, different operations in cybersecurity exercises have become automated

## 2.1 Related Work

Multiple researchers have worked in the development of cyber attack and defense agents. Here, we discuss some of the work relevant to our research. For emulating attacker steps, a lot of research work has been carried out, resulting in open source, free, and commercially available solutions. Some of them are the following:

Splunk attack range [spl] is a limited cyber range deployment tool in which a small infrastructure can be deployed on cloud and local machines. The infrastructure is monitored by various Splunk attack monitoring and detection engines. Different attacks of the infrastructures are simulated using ART (*Atomic Red Team*) [red]. ART follows the MITRE attack chain model [MIT] and can simulate an attack on Windows, Linux, and Mac OS systems. It uses YAML-based inputs to execute atomic tests for adversary actions on the target systems.

APTSimulator [APT] is an open source advanced persistence threat simulation tool. The tool uses batch scripts with various hacking utilities to create system compromise traces like command and control, defense evasion, lateral movement, and so forth. It is used for endpoint detection agent assessment, testing security monitoring and detection capabilities, and preparing the digital forensic class environment. It roughly follows the MITRE attack chain model and is also limited to emulating attacks on Windows-based host machines.

Metta [ube] is an open source information security preparedness tool. The tool uses Virtualbox, with different development tools like Redis/Celery, python, and vagrant, to simulate adversarial actions. Input is given to the tool in the form of a YAML file, which is parsed to execute the attacker’s action on the host- and network-based systems. Metta follows the MITRE attack chain model and is limited to emulating attacks on Windows-based systems.

In the case of cybersecurity exercises, there is a need for the repeatable [HS16] execution of attacker steps for standardized training. Moreover, there are legal and ethical concerns in developing autonomous cyber-attack agents [YB18], so for an attack agent, the attack execution steps need to be planned while keeping a human in the loop before executing them in a cybersecurity exercise environment.

For the defense agent, most related work has focused on network- and host-based detection systems, while some have looked into introducing active attack-repellent features in defense agents. Randolph et al. [JON<sup>+</sup>15] conducted research about modeling and integrating cognitive agents in the cyber domain. The purpose was to develop agents that can produce the necessary friction during cybersecurity exercises to create realism. They developed a novel model for cyber offense and defense; they used the model to create software adapters that translate from task-level actions to network-level events to support agent-network interoperability for cybersecurity operations. They presented a high-level defender goal hierarchy in which the defender has to (1) establish a baseline, (2) detect an ongoing attack, (3) stop an ongoing attack, and (4) prevent future attacks.

Kott et al. [KTD<sup>+</sup>18b] put forward the idea of the development of a “*Hello world*” program for the intelligent autonomous defense agent. The researchers stated that the autonomous agent, should have the following six capabilities to be considered intelligent:

1. The agent should be strictly associated with its environment and should be useless outside its designed environment.
2. The agent should interact with its environment using appropriate sensors.
3. The agent should act to achieve its stated goals.
4. The agent activities should be sustained overtime to be autonomous.
5. The agent should have an internal world model that can express its states with performance measures.
6. The agent should learn new knowledge and modify its model and goals over time based on new knowledge.

Paul et al. [TKD<sup>+</sup>18] proposed the reference architecture for autonomous intelligent cyber defense agents (AICA). In their proposed architecture, the researchers stated that AICA has five main high-level functions:

1. Sensing world state information
2. Planning and action selection
3. Collaboration and negotiation
4. Action execution
5. Learning and knowledge improvement

The researchers also suggested a functional architecture for AICA and stated that according to their assessment, the development of such agents is not beyond the current technical capabilities and can be developed within ten years.

Although other researchers are also investigating autonomous intelligent cyber defense agents, the work is still in the design stage [KTD<sup>+</sup>18b,TKD<sup>+</sup>20]. One implementation of such an agent was proposed by Randolph et al. [JON<sup>+</sup>15], in which the researchers suggested the idea of adding friction in cybersecurity exercises. Their approach utilized cognitive modeling of human cybersecurity experts to model the behavior of the agent based on human expertise. We consider such an approach not suitable for agent development because human experts have different subjective experiences, which can result in unintentional bias in their behavior, as observed in data-driven AI algorithms [YUUK21]. Kott et al. [KTD<sup>+</sup>18b] and Paul et al. [TKD<sup>+</sup>18] provided the baseline requirement and functional needs of the cyber defense agent, which we considered suitable to fulfill our requirements. For the attack agent, we find the ART, APTSimulator, and META approaches suitable for usage in cybersecurity to conduct cybersecurity exercises. We integrated multiple ideas and approaches suggested by various researchers for the development of our attack and defense agents. Our attack agent follows a systemic step-by-step execution of attacks similar to ART, APTSimulator, and META, but it is integrated with a cybersecurity exercise orchestrator, making it suitable for computationally repeatable cybersecurity exercises and experiments. Our defense agent is also integrated with our cybersecurity exercise orchestrator and provides example implementation of the design presented by Kott et al. [KTD<sup>+</sup>18b].

### 3 Methodology

We employed numerous research methodologies during this research activity. We used the **DSR** *design science research methodology* [HC10] for the overall development of the necessary artifacts for this research. DSR is a very well-known research methodology that has five phases 1) *awareness of the problem*, 2) *suggestion*, 3) *development*, 4) *evaluation*, and 5) *conclusion*. These phases are iterative in nature, and the results of these phases are used to improve the overall design to produce a research artifact that addresses the research problem [KV08].

In awareness of the problem, the research problem can be identified, so we leaned on certain studies for this phase [YK18a,YKT<sup>+</sup>18,YKG20]. In the suggestion phase, the solution's designs are proposed to address the research problem, and we conducted this step using certain studies [YKG20,YK19b]. We are currently developing and evaluating the artifacts, and the present paper is presenting the results of this phase. For the development of such an artifact, model-driven approaches are widely employed [BCD<sup>+</sup>19]. In such an approach, a complex problem is abstractly defined in the form of a model, and automation techniques are used to generate low-level artifacts from the abstract model. In our previous work, we developed a **DSL** *domain specific language* to formally specify the cybersecurity exercise environment [YK22]. The environment contains the exercise infrastructure and agents

running within the environment. In this work, we used the DSL developed for automating the creation and deployment of the exercise environment. Additionally, we augmented the DSL with a new modeling technique based on attack/defense trees that we call execution plans (EPs). EPs enable a designer to model the behavior and decision making of attack and defense agents. Finally, in this work, we use **applied experimentation** in operational cybersecurity exercises against defined benchmarks to gather the analytical data for evaluating the performance of developed cybersecurity defense agents [EM17].

## 4 Conceptual Design

This research work is a part of a larger initiative in which the whole process of the cybersecurity exercise life cycle is automated. To achieve this, a DSL is developed to transform abstract concepts related to the cybersecurity exercise life cycle [YK18a] into concrete artifacts. These abstract concepts include defining the network topology, defining the vulnerabilities in the nodes connected to the network, defining benign network traffic, and defining attacker and defender behavior. In this work, our scope is limited to attacker and defender behavior, so we only focus on the concepts involved in attack and defense agent development. Figure 1 illustrates the DSL meta-model for defining attacker and defender properties used to generate agent artifacts. Later in section 5, a concrete syntax will be presented, which will give an example of an instantiation of this DSL meta-model.

The attack agent comprises of a total six concepts. The first one is the *attack-action-id*, which is an identifier of a potential attack action. In case the attack action causes a particular tool to be executed, the tool name can be used as an identifier. The second, third, and fourth concepts are *Agent IP*, *Agent User ID*, and *Agent User Password*. These concepts provide the information about the agent from which the attacker action is going to be performed. The fifth concept is *Argument*, which represents one of the properties that are needed for the attacker’s action to be executed. In terms of an a tool that represents an attacker’s action, this concept represents the tool’s specific arguments. The sixth and final concept for the attack agent is *Target*, which is the IP address of a machine on which the attack agent performs its actions. Similarly, the defense agent has five properties. The first three are *Agent IP*, *Agent User ID*, and *Agent User Password*, which provide the information necessary to install the defense agent on a system. The fourth property is the *Operating System*, which provides the information of the operating system the defense agent is working on. The final property is the *Parameter*, which provides the defender with a specific knowledge base. The parameter contains a list of pairs, each of which consists of an attacker action and the defender reaction to it.

The DSL uses an orchestrator that implements the abstract concepts defined in it to concretely create operational artifacts for establishing the necessary exercise infrastructures, generating network traffic, emulating benign users, launching cyber attacks, and defending against these attacks. The orchestrator has a master control unit connected to the attack agent and used to control them in a semi-autonomous

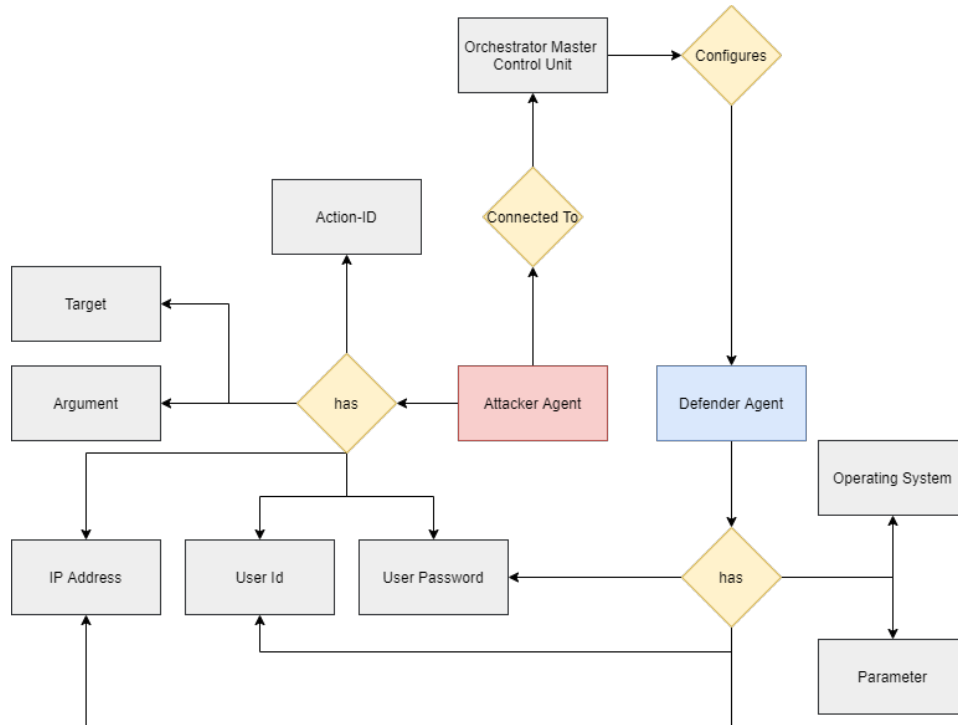


Fig. 1: A meta-model of the attack and defense agents

manner, while the orchestrator configures the defense agents before deployment so that they can work in an autonomous mode. The operational artifacts contain network topology templates for cloud deployment, specification of the vulnerabilities in the form of software, service and configuration, and the specification of the benign, attacker, and defense agents' behaviors present in the deployed network. The specification is given to the orchestrator in JSON, and it starts generating the necessary artifacts in five steps, as shown in Figure 2. First, the exercise network infrastructure is generated through a HEAT template for infrastructure deployment in Opstack-based cloud. In the second step, software, service, and configuration are invoked in the deployed infrastructure using a custom SSH-based installation and configuration module to make the infrastructure vulnerable. In the third step, part of the deployed infrastructure is used to generate benign traffic using various automated tools like TCP relay and VNCD tool. In the fourth step, an attack agent is used to test and verify the vulnerabilities present in the exercise infrastructure. In the fifth and final step, a defense agent is deployed in part of the infrastructure to add the necessary friction in the cybersecurity exercise. The DSL implementation related to exercise network infrastructure generation and generating benign user

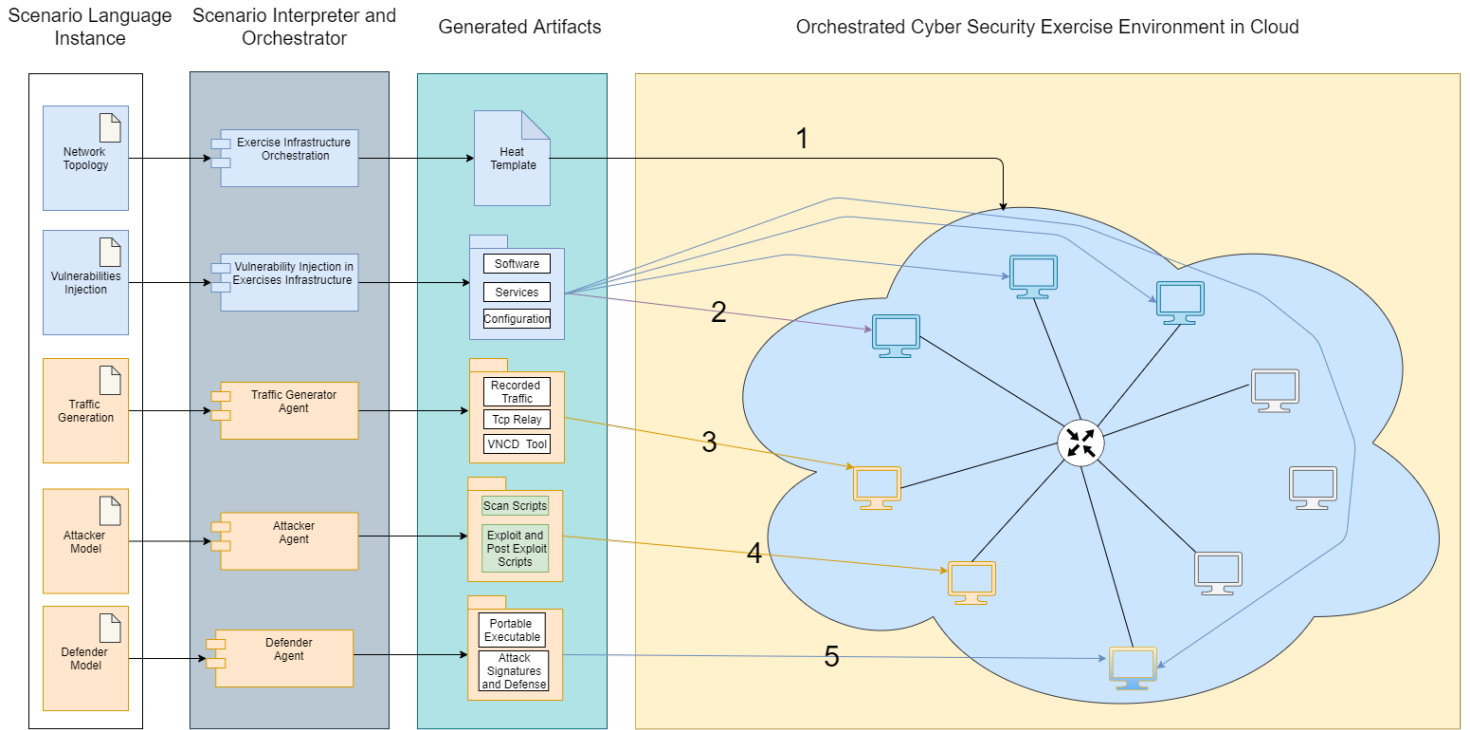


Fig. 2: Cyber security exercise operation orchestrator

behavior and traffic is part of another research work. In the current work, we are only focusing on the DSL instance of attackers and defenders.

There can be multiple ways attack and defense agents can be designed and deployed. This depends on the eventual goal of the agent, that is, what is expected from the agent. We can model the goals of the attack and defense agents based on the responsibilities of the red and blue teams. Lockheed Martin's *cyber kill chain* course of action matrix [HCA<sup>+</sup>11] provides a simplified way to model the attack and defense phases. For the attacker, there are seven phases *reconnaissance, weaponization, delivery, exploitation, installation, command and control, and actions on objective*. These attack phases utilize a set of tools and techniques to achieve the attacker's eventual objectives and goals, which could be the disruption of services or extraction of data. On the defense side, there are six phases to stop the attacker: *detect, deny, disrupt, degrade, deceive, and destroy*. The defender uses different network/host intrusion detection and prevention systems, firewalls, antivirus software, and honeypots to achieve the objective of stopping the attacker.

There are other models like MITRE ATT&CK that can be applied for modeling attacker actions and relevant defender reactions. However, MITRE ATT&CK strictly focuses on concrete actions, tactics, and techniques that are specific to an



operating system. The cyber kill chain is very general and can easily model attack and defense in a different operating environment. We consider the generality of the cyber kill chain suitable for modeling attack and defense scenarios for cybersecurity exercises.

These attack and defense phases are executed by utilizing different tools and techniques. These tools and techniques provide relevant information to the agents so that they can make intelligent decisions. However, it should be noted that this intelligent decision making is strictly dependent on the amount of information being shared, which is related to the agent's goals.

If we look carefully at the information sources, there are primarily two types: external and internal. For an attack agent, external information can be gathered from scanning networks and identifying software, services, and configurations, while its internal information can be the knowledge about the exploits on the identified software, services, and configurations. A defense agent's external information sources contain information from the environment such as network activity, while internal information sources contain information about the system's internal activities such as event logs, which are widely used to detect system exploitation and lateral movement [Cen17]. These information sources are combined to provide *security information and event management* capabilities for defending against the attacks by correlating information from multiple sources. However, information correlation requires manually defining security events to look for and manually take actions against them to stop the attacker in its tracks. A defense agent can assess the type of traffic to identify whether it is benign or malicious. An attacker can overwhelm the defender by launching multiple attacks at the same time, which could make intelligent decision making very difficult. It will also create a new threat vector for the defender because decision making depends on external sources that can be manipulated.

Let us analyze the course of action matrix [HCA<sup>+</sup>11] for attackers and defenders. Here, the attacker's reconnaissance and weaponization goals can be detected by external information sources like web analytics and NIDS (network intrusion detection systems). In contrast, exploitation and installation can be detected by HIDS (host intrusion detection system). Although detecting an attack is desirable at an early stage, a host-based system can be better suited to respond directly to the attack; it can detect a security event and automatically respond to it by making operational changes such as applying local firewall rules and installing patches through its knowledge base without relying on an external input to deny the attacker actions. The knowledge base can contain information about the expected attacker's actions and the appropriate defender response. This knowledge can be useful for known attack tactics and techniques; however, it needs to be updated for new attack detection and responses, which require some intelligence. This intelligent behavior can be learned by analyzing the attack vectors and implementing security actions against them, manually first and automatically later. The attack vectors can be learned by constantly monitoring the system state and detecting changes. When a change is detected, the events that lead to that change can be fetched for identifying the malicious actions. A set of predefined reactions can be

specified for implementation against a particular set of actions to deny the attacker from using them for further exploitation.

All components and parts of a cybersecurity exercise environment are considered a system, and each system is running software and services with system-specific configurations. The system for the attack agent is a Kali Linux machine controlled by another system running our developed orchestrator software and using SSH as a service for communication with the Kali machine. The orchestrator can control multiple Kali machines to launch multiple attacks at the same time. Similarly, for the defender part, the orchestrator can inject a defense agent with its knowledge base as software that can independently run on the injected system to protect its software, service, and configuration. Additionally, there are traffic generators that are present in the cybersecurity exercise environment, which are basically attack agents performing benign activities such as replying PCAP files and automating GUI user behavior using VNCDtool. The agents and their interactions are presented in 3, which is mapped with the third, fourth, and fifth steps of the orchestrator, as presented in Figure 2.

The developed agents operate in a cloud-based cybersecurity exercise environment. The environment has attack and vulnerable machines on which the attack and defense agents are operating. The vulnerable machines have vulnerabilities related to software, services, and configurations that an attack agent can remotely sense. The behavior of the agents in the environment is governed by the world state. The *world state* is the software's, services', and configurations' specific information provided to the agents. New information about the world state is gathered from the environment in which the agents are operating by using their sensing capability. The sensing capability indicates which type of systems the agent is interacting with and triggers an action when it finds some specific information about the world state [KTD<sup>+</sup>18b, KTD<sup>+</sup>18a]. For the attack agent, active and passive reconnaissance tools like arp-scan and Nmap are used to gather information about the services running in the network. This information is used to create the world model for the attack agent, and changes in these services will update the world state for the attacker. When the attack agent senses a vulnerable service, it triggers a change in its world state, upon which an attack action is selected and executed on the vulnerable machine. Similarly for the defense agent, Windows security logs provides an active sensing capability to create a model of the world state, which includes the type of software and services running on the system and any changes in their configuration. When an attacker interacts with the system defended by the defense agent, it creates logs that are then used to update the world state and trigger a reaction from the defense agent. The decision-making process of the proposed agents are discussed in section 4.1.

When the attacker's action is executed, an event is created in the vulnerable machine environment. The defense agent has a list of malicious events, the current world state, and responses to those events. The defense agent can sense the known events generated by attacker actions, which can trigger a change in its world state. The defense agent then selects an appropriate response and executes it to counter the attack agent's actions. The interactions between the attack and defense agents

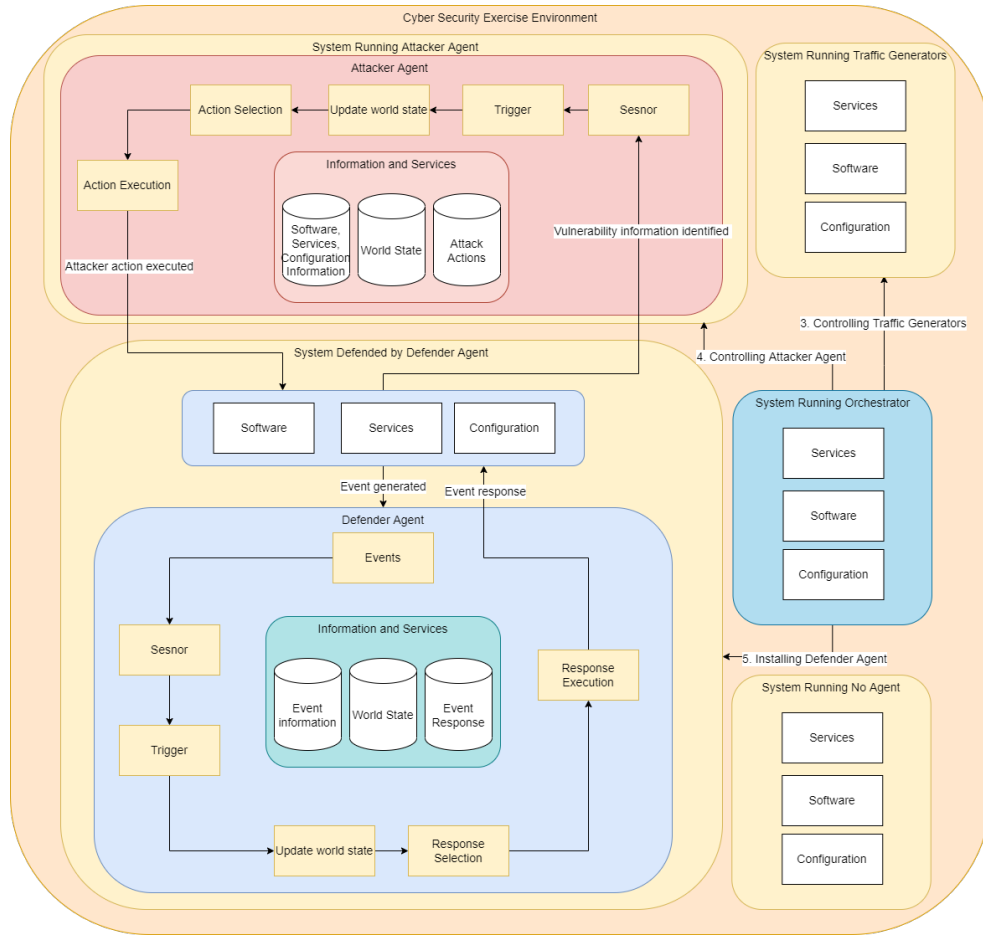


Fig. 3: Attack and defense agent environment and interactions

in the cybersecurity exercise environment is presented in Figure 3, and details of the attack and defense agent work flows are presented in Figure 6 and Figure 7 respectively:

In terms of deploying the agents in the cybersecurity exercise environment, there are certain considerations. Some researchers have implemented script execution techniques on the machine [Gen] to generate attack logs; however, we deployed the attack agents on a remote machine to emulate realistic adversary behavior. On the other hand, the defense agents were deployed on the machines because a central command and control unit could have been compromised to disable the defenders' functionality.

## 4.1 Agent Decision Modeling

On an abstract level, our agents have five properties— *Knowledge*, *Goals*, *Conditions*, *Actions*, and *Commands*, where the knowledge is provided through the DSL. The *Knowledge* of an agent contains information about the world state, like the software, services, and configuration running in the environment. An agent’s *Condition* is used to perform condition-specific actions on the software, services, and configuration. These conditions are triggered based on events that change the world state. An *Action* is executed using a set of *Commands*. A set of successful action executions result in the achievement of a *Goal*, which is modeled based on the *CKC* course of action matrix. For the attacker, these goals are *Reconnaissance*, *Delivery*, *Exploitation*, *Installation*, *Command and Control* and *Actions on Objective*. Similarly, for the defender, these goals are *Detect* and *Deny*.

The DSL instance is translated into EPs (Execution Plans) for the achievement of specific *Goals*. We adapted the concepts from the attack and defense trees [KMRS14], as well as the hierarchy of action plans [Kot05], to develop the EP models. The EP models consist of three levels of decisions *high*, *medium*, and *low*.

**4.1.1 EP Model** EPs are tree-structured models that represent the agent’s decision-making process. An EP describes the goals, conditions, and commands of an agent, as well as showing the path that needs to be taken to reach the final conditions and fulfill the goals. These conditions result in one of the following EP outputs: plan fulfilled, plan not fulfilled, or plan maybe impractical.

The root of an EP tree is the goal of an agent, and an end-leaf of an EP tree represents the commands that will fulfill an agent’s goal. An EP consists of three decision levels—Level 1, Level 2, and Level 3—and each level is represented as a sub-tree of the EP tree. The root of Level 1 of an EP tree is the root of the EP tree. The leaves of one sub-tree are the roots of the next level sub-tree. A parent node is connected with its children nodes using two possible operators, *AND* and *OR*, which are represented by  $\wedge$  and  $\vee$ , respectively. The semantics of the nodes and operators in an EP tree depend on the level where the node exists.

- Level 1** The Level 1 sub-tree of an EP tree is the first high level sub-tree of the EP tree. The root node of a Level 1 sub-tree of an EP tree represents the main goal of the EP tree, and the leaves represent a set of sub-goals. The operator  $\wedge$  is used if all the sub-goals need to be fulfilled for the parent goal to be achieved. On the other hand,  $\vee$  can be used if the fulfillment of one sub-goal will result in the fulfillment of the parent goal.
- Level 2** The Level 2 sub-tree of an EP tree is the second medium-level sub-tree of the EP tree. The root of the Level 2 sub-tree is a leaf in the Level 1 tree or the root of the EP tree where the Level 1 sub-tree consists of one node. Level 2 represents a sequence of conditions that need to be checked to decide which actions should be executed. The nodes of a Level 2 sub-tree are conditions with two possible outputs *Yes/No*, meaning that only  $\vee$  operators are allowed in a Level 2 sub-tree. Each parent node is connected to, at most, two children

nodes, which represents the next conditions to be checked. A special type of condition is called "Not Fulfilled." It is a final condition with no children, and it is denoted by the symbol  $\text{—}$ ; reaching this condition means that the plan is not or cannot be fulfilled. A leaf in a Level 2 sub-tree can be either an NF "Not Fulfilled" condition or an action.

**Level 3** The Level 3 sub-tree of an EP tree is the third low-level sub-tree of the EP tree. Level 3 roots are actions represented in the leaf nodes of the Level 2 sub-tree, and the nodes represent concrete commands. Both  $\wedge$  and  $\vee$  are allowed in a Level 3 sub-tree.  $\wedge$  means all the children's commands need to be executed to achieve the action, while  $\vee$  means that the execution of any of the commands can achieve the action.

**Output** Plan output "Fulfilled" is reached when all the commands in the EP tree leaves are executed successfully and the goal is achieved. Plan output "Impractical" is reached when the result of the execution of one command leaf is not successful. Plan output "Not Fulfilled" is reached when the agent cannot reach an action leaf because of knowledge or resource limitations. The EP plan in this case will stop at the Level 2 sub-tree.

**4.1.2 EP Formal Model** We use Datalog [Dat] for formal modeling of the agents decisions and to verify the different decision properties like: *is the goal fulfilled or not?*. Datalog is a programming language based on a declarative logic [Llo12]. It is employed by researchers for large-scale software analyses [Nai20], automatic evaluations of cybersecurity matrices [ZN20], and the verification of cybersecurity exercise scenarios [RCA20], making it suitable as a formal model for cybersecurity exercise scenarios. It consists of two parts: facts and clauses. A fact conforms to the parts of the elements of the predicated phenomenon. A clause refers to information deriving from other subsets of information. Clauses rely on terms, which can contain variables; however, facts cannot. It adjudicates whether the specific term is adherent to the specified facts and clauses. If it happens to be so, the specific query is validated via a query engine, providing the prerequisite facts and clauses.

When running a Datalog operation, the specified conditions include a combination of two facts along with a singular clause. We assign a condition that if the query is valid, a specific response is to be expected at the end. The conclusion of the said experiment is that the specific response is received and that the query is satisfied. By utilizing the clauses via their variables, the engine can pinpoint and find the result. For a concrete example [CGT<sup>+</sup>89], consider the facts "John is the father of Harry" and "Harry is the father of Larry". A clause will allow us to deduce facts from other facts. In this example, consider we want to know "Who is the grandfather of Larry?". We can use three variables  $X, Y$  and  $Z$  and make a deductive clause: If  $X$  is the parent of  $Y$  and  $Y$  is the father of  $Z$ , then  $X$  will be the grandfather of  $Z$ . To represent facts and clauses, Datalog uses *horn clauses* in a general shape:

$$L_0 : -L_1 \dots, L_n$$

Each instance of  $L$  represents a *literal* in the form of a *predicate* symbol that contains one or multiple *terms*. A *term* can have a constant or variable value. A

Datalog clause has two parts: the left hand side part is called the head, while the right hand side part is called the body. The body of the clause can be empty, which makes the clause a fact. A body that contains at least literal represents the rules in the clause. Lets us represent the above mentioned facts that "John is the father of Harry" and "Harry is the father of Larry" as follows:

$$Father(John, Harry)$$

$$Father(Harry, Larry)$$

The clause if  $X$  is the father of  $Y$  and  $Y$  is the father of  $Z$ , then  $X$  will be the grandfather of  $Z$  can be represented as follows:

$$GrandFather(Z, X) : -Father(Y, X), Father(Z, Y)$$

For our agents we define 4 basic predicates for decision modeling which are 1) Goal, 2) Condition, 3) Action and 4) Fulfilled. The facts for the decision model with their definitions are as follows:

**Definition 1.** The Goals predicate is logically presented as  $Goals(Goal, SubGoal)$ , and it has two variables Goal and SubGoal. The Goal is a string value which indicates attack and defense goals like 'Exploit System' for attack and 'Prevent Attacks' for defense. The SubGoal is also a string values which contains sub goals for achieving the Goal like 'Reconnaissance' and 'Exploitation' for attack and 'Detect' and 'Deny' for defense. A concrete example of Goals predicate for attack is presented as follows:

$$Goals('ExploitSystem', 'Reconnaissance')$$

$$Goals('ExploitSystem', 'Exploitation')$$

Similarly for defense Goals can be represented as fallow:

$$Goals('PreventAttacks', 'Detect')$$

$$Goals('PreventAttacks', 'Deny')$$

**Definition 2.** The Condition predicate is logically presented as  $Condition(Goal, Parameter, Command)$  and it has three variables Goal, Parameter and Command. Goal is a string value which is the leaf of Level 1 tree consisting goals and sub goals like 'Reconnaissance', Parameter is a string value that is condition specif to achieve the goal like 'Active' for performing active reconnaissance, Command contain the parameter for performing a low level action like service and version scan using Nmap '-sS -sV'. A concrete example of Condition predicate is presented as follows:

$$Condition('Reconnaissance', 'Active', '-sS -sV')$$

Similarly a defense Condition can be represented as fallow:

$$Condition('Deny', 'shell.exe', 'taskkill/IM" shell.exe"/F')$$

**Definition 3.** The *Action* predicate is logically presented as  $Action(ActionName, ActionTarget)$ , and it has two variables *ActionName* and *ActionTarget*. *ActionName* is a string value which contain low level goal action execution like *Nmap* for attacker and *Deny* for defender. *ActionTarget* is string value which contain the information of machine address on which the action is to be executed for the attacker and the pattern on which the action is triggered by the defender. A concrete example of *Action* predicate is presented as fallows:

For attacker:

$$Action('Nmap', '172.168.2, 2')$$

For defense:

$$Action('Deny', 'shell.exe')$$

**Definition 4.** The *Fullfilled* predicate is logically presented as  $Fullfilled(Goal, ActionName)$  and it has two variables *Goal* and *ActionName*. *Goal* is a string value that indicates high level goal like '*Reconnaissance*' and *ActionName* is a string value that indicates concrete tool or action to achieve the high level goal like '*Nmap*'. A concrete example of *Fullfilled* predicate is presented as fallows:

$$FullFilled('Reconnaissance', 'Nmap')$$

For defense:  $FullFilled('Deny', 'shell.exe')$

**4.1.3 EP Model Verification** The decision model presented in section 4.1 help us to verify various agent properties before their execution, like:

- How high level goal can be translated into to low level actions
- Can the agent fulfill the given goal?
- What information is missing to achieve the goal?

To verify the decision model we define a new predicate *CheckGoals* which takes two variables *Goal* and *SubGoal* and is logically presented as  $CheckGoals(Goal, SubGoal)$ . A logical relationship is defined between the *Goal* and *SubGoal* so it can be established whether the *Goal* is the leaf for *Level 2* conditions or the *SubGoal*.

$$CheckGoals(Goal, SubGoal) \leq CheckGoals(SubGoal, Goal)$$

*FullFilled* predicate is used to link the *Goal* and *SubGoals* which is presented as fallows:

$$FullFilled(Goal, SubGoal) \leq CheckGoals(Goal, SubGoal)$$

Furthermore, it is defined whether a *SubGoal* needs to be completed in order to achieve the *Goal*. A relationship is established between *Goal*, *SubGoal* and *Action* using transitive property which is presented as fallows:

$$FullFilled(Goal, SubGoal) \leq CheckGoals(Goal, Action) \\ \& FullFilled(SubGoal, Action) \& (Goal \neq SubGoal)$$

To verify the attack decision for achieving high level goal using low level action based upon certain conditions following clause can be defined:

$$\text{FullFilled('Reconnaissance', SubGoal) \& Condition(Goal, 'Active', Commands)} \\ \& \text{Action(SubGoal, ActionTarget)}$$

The clause will return the high level goal, and the low level specific action to be executed based upon the specific command. Similarly for the defense decision following clause can be defined:

$$\text{FullFilled('Deny', SubGoal) \& Condition('Deny', SubGoal, Commands)}$$

The clause will return the low level pattern through which the action is triggered and the low level command to deny that action.


**4.1.4 EP Model Representation** A schematic representation of the attack and defense agent EP models all three levels is presented in Figure 4 and Figure 5, respectively.

In Figure 4, a high-level *Goal* is given to an agent that has the aim of performing system exploitation. The sub-goals are *Reconnaissance* and *Exploitation*. The agent will check information in its knowledge base to make medium-level decisions, for example, whether information about the target is provided or not. If the target information is provided, it will check whether it is accessible or not; if it is accessible, then the agent will check if some specific argument is present to perform a specific kind of reconnaissance like *nbtscan*, which is a low-level decision. Otherwise, it will use a default reconnaissance technique like *nmap* or *netcat*. Similarly, if the target information is not provided, the agent will check whether there is a network interface and if on that network interface it can perform *arp-scan* a medium-level decision. If the condition returns true, then the agent can perform an arp-scan using default commands, which is a low-level decision.

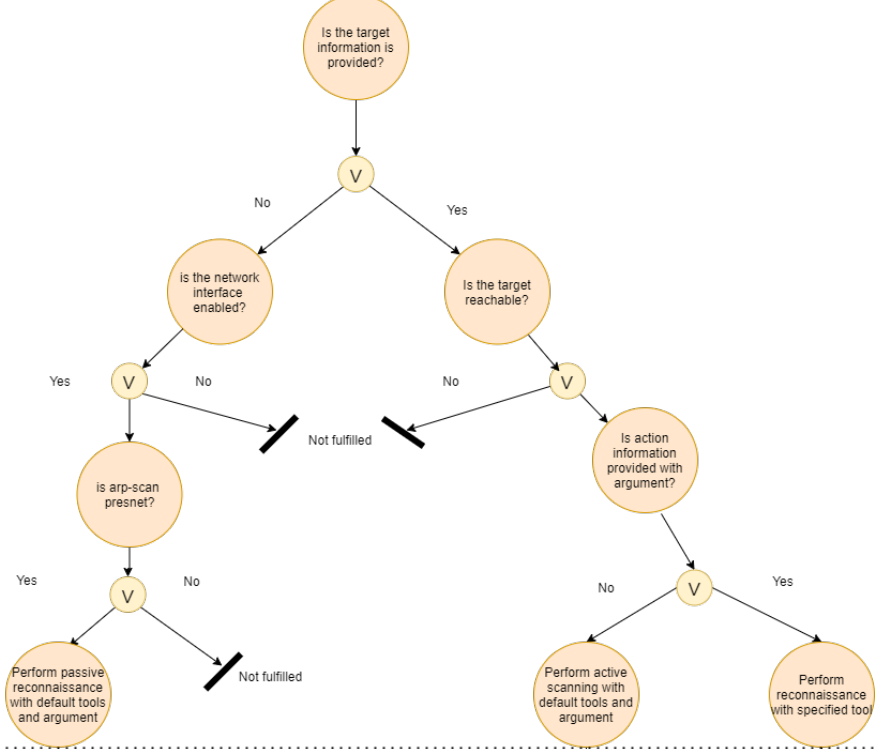
In Figure 5, a high-level *Goal* is given to an agent to prevent the attackers from performing any actions. The sub-goals are to *Detect* and *Deny* attacker actions. Whenever an attacker performs an action, it creates an event. If the event is detected, then it is checked in the knowledge base of the defender; this is considered a medium-level decision. If the knowledge base contains information about the reaction to this specific action, then it will execute a specific command to the counter-attacker's action; this is considered a low-level decision, for example, killing a specific malicious process using *taskkill*. In another case, if there is no specific command to react to the attacker's action, then the agent will execute a general defense command to counter the action, such as closing the ports using *netstat* or *npx-kill-port*. If the action of the attacker is not detected, then the agent will fail to defend the system.



Level 1



Level 2



Level 3




Fig. 4: Attack agent EP

Level 1

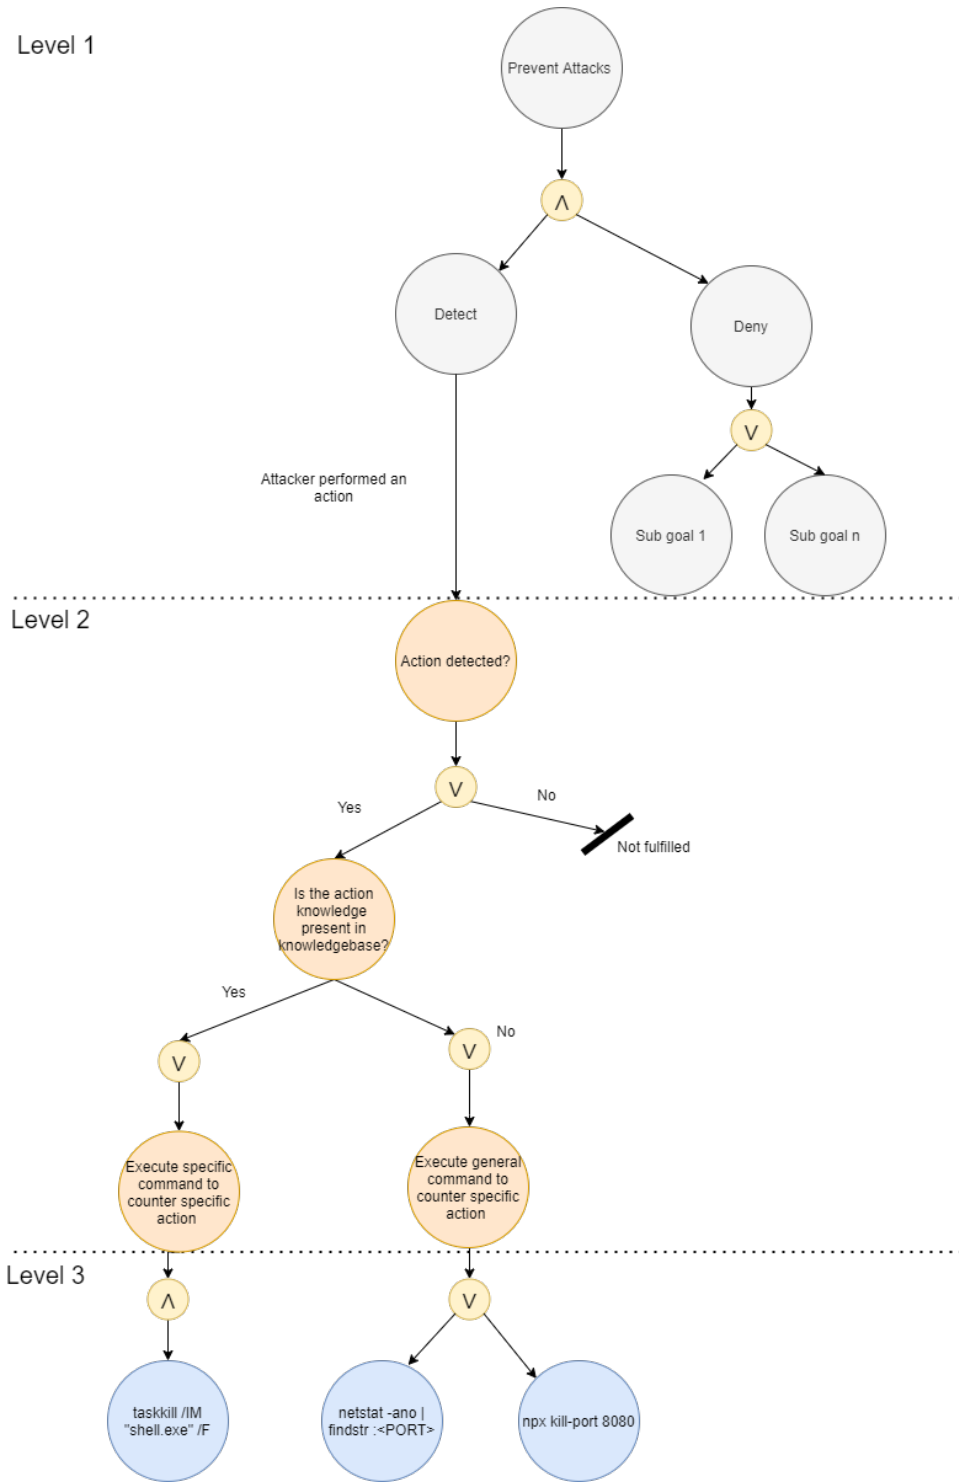


Fig. 5: Defense agent EP

**4.1.5 Attack Agent** Conceptually, the attack agent’s goal is to perform the steps involved in exploiting vulnerable systems during a cybersecurity exercise. The steps involve performing scanning, identifying vulnerable services, and launching an attack on them. Multiple adversary emulation tools already exist in academia and industry, here using various techniques ranging from logical programming [Yue15] to AI [Sto18] for achieving this goal.

A model-driven approach for executing the attacks during a cybersecurity exercise can provide repeatable and standardized training. The model needs to follow penetration testing execution standards [PTE] to leave realistic attack and exploitation traces for the defender or blue team members to identify. This can be achieved by specifying the attacker’s actions in a DSL, hence enabling the precise execution of attack steps and helping in the evaluation by the defenders in incident response and forensic analysis.

We combined complex attacker operations into six components of a DSL, whose concrete syntax instance is presented in Listing 1. These components specify the attack techniques that are going to be used by the attacker and on which target it needs to be performed. The DSL instance components are used to provide the necessary information to the EP model to specify the behavior of the attacker based on the *cyber kill chain*. Once the model has been created, it is verified by executing it in different operational cybersecurity exercises with the same network topology to check whether its execution is repeatable for standardized training.

Listing 1: Concrete syntax for attacker DSL instance

```
[
  {
    "nbtscan": {
      "AgentIP": "192.168.81.128",
      "AgentUserID": "root",
      "AgentUserPassword": "toor",
      "Argument": "",
      "Target": "172.168.2.17"
    }
  }
]
```

The attack agent DSL model has the following properties:

1. It can perform actions from the following list: *Reconnaissance, Delivery, Exploitation, Installation, Command and Control*, and *Actions on Objectives*.
2. It has six attributes: agent action name, agent to use, agent credential user ID and password, action-specific parameters and the target on which action is needed to performed.
3. It runs in a separate attack machine in the exercise network environment, where it has network-level access to vulnerable machines present in the exercise network.

4. It interacts with exercise networks using the specific actions, which can collect information about the software, services, and configurations present in exercise machines and then perform other actions to exploit those machines. Deciding on which actions and commands to execute is specified in the EP model.

Utilizing the above-mentioned properties, the attack agent launches attacks in a semi-autonomous manner, as defined in the EP models. These EPs consist of executing the attack phases presented in *CKC* by utilizing its attack agent properties. The attack agent's overall workflow is represented in Figure 6. It represents how the

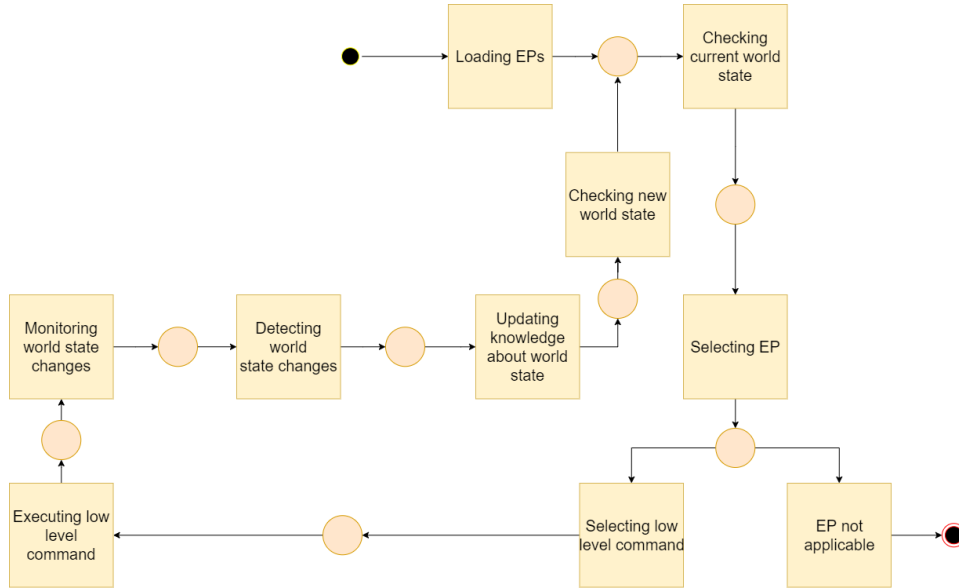


Fig. 6: Attack agent's work flow

agent functions. The agent then does the following: (1) First, the agents loads the goals and related EPs. (2) The agent checks the world state. Initially, the world state is empty for a newly deployed agent. (3) An EP model for the loaded goal is utilized by the agent to decide which actions to perform and fulfill the goal. (4) The agent decides which low-level commands to execute based on the EP model (4.1), and it will fail in case the goal was not fulfilled or was impractical (4.2). (5) The selected command(s) will be executed. (6) The agent (continuously) monitors the environment to detect any changes in the world state and to get the output and results of executing its commands. (7) The sensors detect new knowledge. (8) New knowledge updates the world state. (9) The agent checks the new world state.

**4.1.6 Defense Agent** The defense agent's primary goal is to defend its system from external and internal attackers. This primary goal has additional sub-goals

in which the defense agent has to *Detect, Deny, Disrupt, Degrade, Deceive, and Destroy* an attacker. The focus of this work is on *detect* and *deny*. These sub-goals are achieved by the usage of different tools and techniques at different stages of an attack. These tools and techniques include but are not limited to network and host intrusion detection and the prevention system, web analytics, security configuration, and system user training.

Our DSL instance is used to specify the defense agent's properties. Based on these properties, the EPs are executed by our orchestrator. The orchestrator inserts the agent in the machine present in the exercise network with the knowledge base of the events generated by the attacker's actions. Conceptually, the defense agent has the following properties:

1. It can detect and deny the actions performed by the attacker.
2. It has a knowledge base that contains information about the attacker's actions and their countermeasures.
3. It runs on the exercise machines being attacked.
4. It interacts with the events generated by the attacks and implements specified countermeasures on the machine it is running.

Utilizing the above-mentioned properties, the defense agent can perform defense measures against launched attacks in a semi-autonomous manner, as defined in the EP models. These EPs consist of executing the defense phases presented in *CKC* by utilizing the defense agent properties. One key difference between the attack and defense agent is that the defense agent is not controlled by a Master and is independent in its execution. The Master only configures the knowledge base of the defense agent one time and uploads it on the machine that needs to be defended. The defense agent's overall workflow uploaded to a machine is represented in Figure 7. In the work flow, the agent does the following: (1) First, the agent loads the goals and related EPs. (2) The agent checks the world state based on the EP model of detecting the attacker. Initially, the world state is empty for a newly deployed agent. (3) The agent will use its sensors to detect an event generated by attacker actions based on its knowledge base. (4) In the detection phase, if the attacker's action event is detected, then the agent will check its knowledge base to counter the attacker's action (4.1). If the attacker's actions were not detected by the agent, then it will fail to deny the attacker (4.2). (5) The agent will update its world state and act on new EP models to deny the attacker. (6) To deny the attacker, the agent will check its knowledge base. (7) For denying the attacker, the agent will execute a low-level command that changes the world state. If the agent's knowledge base has information about the specific action, then it will execute a specific command (7.1). If the agent does not have specific countering information, then it will execute a general command (7.2). (8) New knowledge is updated in the world state of the agent. (9) The agent checks the new world state.

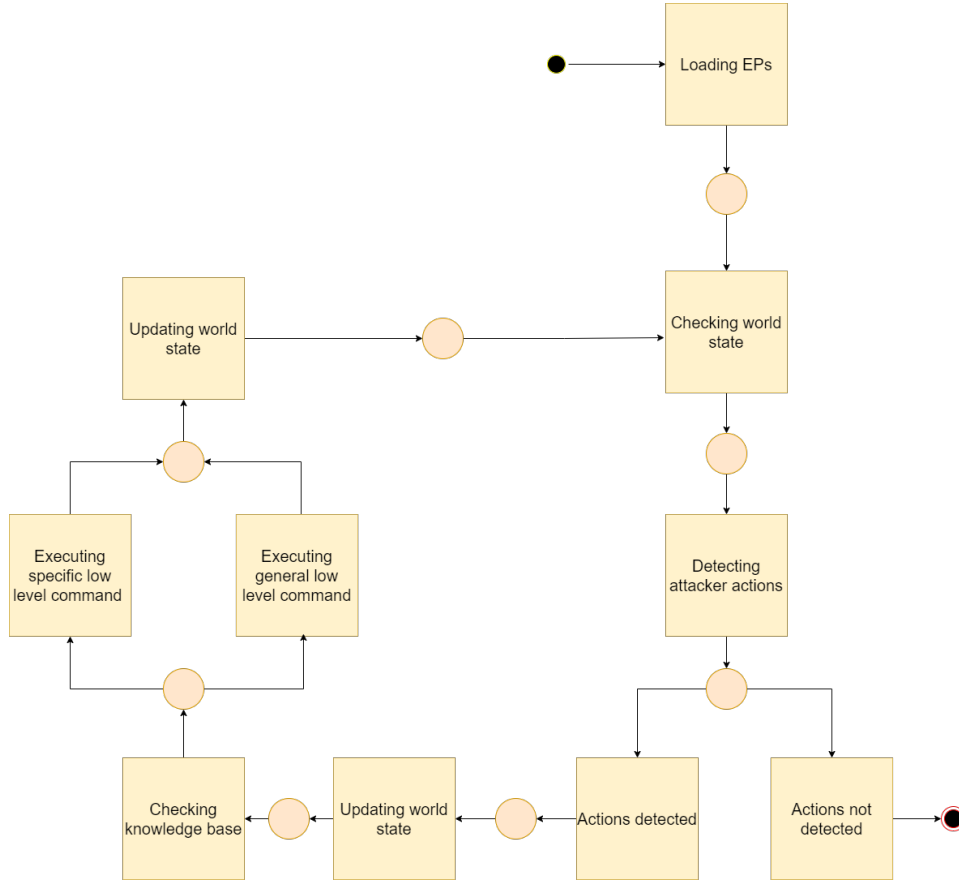


Fig. 7: Defense agent's work flow

## 5 Technical Implementation

### 5.1 Attack Agent

The attack agent is a Kali Linux automation utility that can automate most of the Kali Linux environment tools. These tools can perform red team operations such as reconnaissance, weaponization, delivery exploitation, installation, command and control, and actions on objectives. In a cybersecurity exercise environment, one or more Kali Linux machines can be deployed to perform the attacker's actions. The orchestrator has a remote Master control unit that controls these machines using a dedicated SSH connection. The Kali Linux agents and the Master control unit work in a *Botnet Command and Control* [ZM09,FSR09] manner. The attacker's actions are modeled in the DSL, which the Master control unit interprets and forwards to the Kali Linux machines as *EP*.

The Mastercontrol unit contains the *EP* of various attack stages, such as Nmap scripts for scanning and Metasploit scripts for exploitations. The resource script [Res] of Metasploit is used to perform post-exploitation on the exploited machines. An extensive logging mechanism is integrated into the Master control unit, which can collect logs from the Kali machines to confirm whether the launched attack steps were successful or not. A schematic diagram presenting the main components existing in an attack agent is shown in Figure 8. Each component represents a category of tools that can be used by the attack agent. For example, these components include scanners to collect information about the exercise environment and vulnerability executors to perform a particular attack or exploit. The results of the attacks are logged to update the world state and to inform the Master control unit about successful attacks.

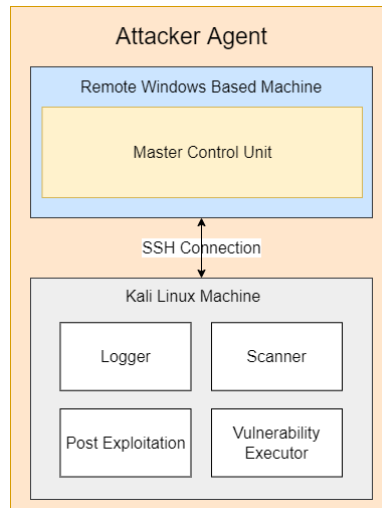


Fig. 8: Attack agent

**5.1.1 Scanner** The scanner can work both passively and actively. In passive mode, the scanner uses ARP resonance techniques for the identification of targets using Netdiscover [net]. When the targets are identified, it can switch into an active scanning mode and use Nmap for the identification of vulnerable services running on the system. The information of vulnerable services and their exploitation methods are provided in the DSL parameters.

**5.1.2 Vulnerability Executor** The vulnerability executor can take the information from the scanner to launch an attack based on predefined conditions, or it can follow the concrete action steps provided in DSL and the EP. DSL contains the

general static information for the attack agent, and the EP contains the execution plan to fulfill the attacker's goals. The conditions include finding a specific service or application signature and launching the relevant, very well-known approach. On the other hand, the concrete action steps from EP provides a repeatable execution of vulnerability exploits. The DSL constructs include (1) the tool or action name needed to be executed, (2) the agent's IP and credentials from which the action is to be executed, and (3) the specific vulnerability parameters and the target address on which the attack is executed.

**5.1.3 Post Exploitation** When a vulnerability is exploited, the post-exploitation module performs different tasks like credentials and memory dumps, backdoor injection, pivoting and lateral movement, and so forth. The post-exploitation steps are predefined, and because the cyber kill chain does not incorporate post-exploitation steps, concepts from MITRE Attack [MIT] are incorporated in it. For a Windows-based environment, most of the post-exploitation is performed through predefined Mimikatz commands with standard Meterpreter modules [Mim] and Powershell scripts [Pow]. For a Linux-based environment, a set of bash scripts [Lin] is used in an automated manner for post-exploitation.

**5.1.4 Logger** The logger logs all the different agents' activities with respect to time, the commands used, and their results in textual format. The logs are used to verify different attack agent success parameters in scanning, exploiting, and post-exploitation of the vulnerabilities in the cybersecurity exercise environment.

## 5.2 Defense Agent

The defense agent is a portable executable that can be deployed in a Windows-based machine. The defense agent's *EP* is generated based on the DSL instance, that is, on which system it is needed to be deployed on and what kind of action it needs to take, as presented in the Listing 2.

Listing 2: Concrete syntax for defender behavior emulation

```
[
  {
    "Defender 1": {
      "MachineIP": "192.168.81.132",
      "MachineUserID": "root",
      "MachineUserPassword": "toor",
      "OS": "Windows",
      "Parameter": "Actions1.csv"
    }
  }
]
```



The defense agent has multiple components, such as knowledge base, monitoring, analysis engine, event collector, and event responder. The knowledge base of the defender can be configured to adjust the agent behavior based on the scenario requirement, the details of which are given below and presented in Figure 9. The defense agent is deployed in a *Windows*-based environment. It uses a custom monitor and analysis engine that collects security events from Windows event logs and that act as sensors to collect information about the environment. When an event is detected whose information is present in the defense agent knowledge base, a trigger will change its world state, resulting in the selection and execution of different responses against the attacks using the information present in the knowledge base. This defense step execution is mapped with different *CKC* phases and is conceptually represented in Figure 3.

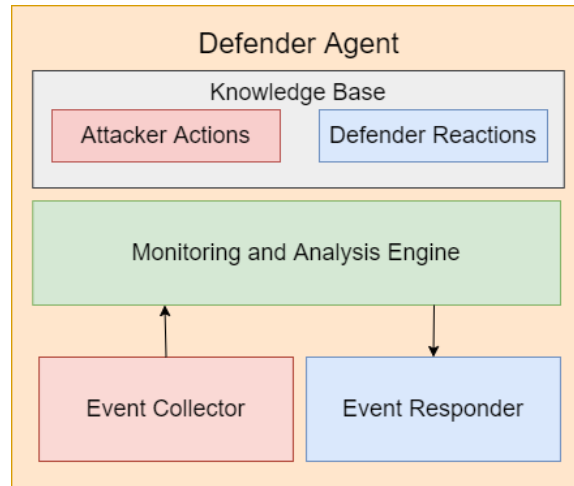


Fig. 9: Defense agent

**5.2.1 Knowledge Base** The knowledge base for the defense agent is a simple CSV file that contains the list of attacker actions and defender reactions. This approach of segregating the knowledge of the defense agent from the program provides the flexibility to use different levels of knowledge against the different skill set levels of attacker. For example, in a cybersecurity exercise for novice and expert hackers, the knowledge base can be adjusted to create a balanced environment [MTWP15]. We analyze some example attacker actions and defender reactions below.

**5.2.2 Attacker Actions** We can consider the example of *Pass the hash attack* on a remote Windows-based system. The attack will generate a 4688 Windows security event log that contains the following command line information:

Listing 3: Process command line information

```
C:\Windows\System32\wbem\WMIC.exe
```

This event and command line information can be inserted into the knowledge base to give the capability to the defender to detect such an attack signature. If the attacker is skilled enough, then the attacker can use various payload obfuscation techniques to bypass the defender's detection.

**5.2.3 Defender Reaction** The defender's reaction can be variable based on the scenario requirement, the knowledge for preventing the above attack can be added in the CSV file.

Listing 4: Defender reaction to the detected event

```
Reg add
↪ \HKEY_LOCAL_MACHINE\SYSTEM\CurentControlSet\Control\Terminal
↪ Server" /v fDenyTSConnections /t REG_DWORD /d 1 /f
```

The defense agent can disable the service being exploited to prevent the known attack; however, it will not be able to prevent the attacks that it has no information about. To address this, the defense agent has a monitoring and analysis engine to detect and respond to new attack patterns. Most of the attacks result in events that have similar patterns as the defined attack action; for example, opening a port from different exploits will have similar signature. The defense agent can utilize such information to prevent the attack.

**5.2.4 Monitoring and Analysis Engine** The monitoring and analysis engine provides the defense agent the capability to acquire and apply new knowledge. We can consider a case where the attacker was able to bypass the detection mechanism of the defender; then, the attacker will try to achieve its goals and objectives. For instance, an attacker can try to tamper with the content of the file, which was specified in the knowledge base to be monitored. This action will also generate an event log that can be traced back to the original process using our exploit chain detection algorithm [YKG19]. Using this information, the knowledge base of the defense agent can be updated automatically to kill the exploited process or close the vulnerable port using standard system commands.

Listing 5: Sample commands used for detecting file manipulation

```
type <filename.txt>
more <filename.txt>
```

Listing 6: Sample commands used for detecting user manipulation

```
whoami /all
net user
net user <user> <pass> /add
net localgroup administrators <user> /add
net user <pass> /del
```

Listing 7: Sample commands used for detecting host information gathering

```
systeminfo
driverquery
tasklist
fsutil fsinfo drives
net time
net file
net session
net use
```

**5.2.5 Event Collector** The event collector collects *Windows* security event logs. The event logs contain a lot of information that can be used in event correlation and for a process analysis. The event collectors parse the event logs, remove irrelevant information, and forward them to the monitoring and analysis engine for further processing. The event collector can be configured to collect the security logs from the active directory to perform network-centric cyber defense. However, currently, it is only working on Windows-based host systems.

Figure 10 represents sample attacker actions for the event collector in the defense agent in which attacker is trying to retrieve clear text WLAN credentials. One attacker action is to open a CMD shell on a compromised system; then, the attacker can fetch the WLAN profiles present on the system through a CMD command. The WLAN profiles contain information about the WIFI networks with which the system is or was connected. The attacker then fetches the clear text credentials of a WIFI network SSID through another command. The attack scenario involves the total execution of three commands. In Figure 10, it can be seen that the process ID for initiating the CMD console is 0x34f8, which then executed two child processes with the process ID 0x1a44 and 0x82c.

**5.2.6 Event Responder** The event responder is running on the host system with system-level privileges. It merely takes input from the monitoring and analysis engine and performs relevant operating system security tasks. These tasks involve executing *Windows command line* and *Powershell* for managing and implementing security setting changes on the system.

All these processes in Figure 10 create events in the Windows security logs and can be monitored by the monitoring and analysis engine. If the command entered



Fig. 10: Sample event collector

in the CMD console is being monitored and detected, then the parent process ID 0x34f8 will be used to kill the process. If the action of the attacker is not detected and a secure resource is retrieved, here being the credentials of the WLAN SSID, then the process is traced back, and the command parameters used for leading to the information retrieval are saved in the knowledge base to prevent future exploitation.

An attacker can bypass the defender command and event monitoring capability using different command-line obfuscation techniques [asi]. The techniques use special characters and encoding schemes to evade any pattern matching algorithms; an example of such technique is presented in Figure 11, where the command *whoami* is executed in a CMD shell of a Windows 10 machine using various obfuscation techniques. Although machine learning-based techniques are developed to detect such obfuscated commands [YK18b,HKR18], we did not integrate such a solution in the defense agent yet.



```
Command Prompt
C:\Users\mudda>whoami
desktop-13b7nks\mudda

C:\Users\mudda>w^h^o^a^m^i
desktop-13b7nks\mudda

C:\Users\mudda>"w""ho""am""i
desktop-13b7nks\mudda

C:\Users\mudda> w""ho^am""i
desktop-13b7nks\mudda

C:\Users\mudda>
```

Fig. 11: Sample techniques for defense agent detection bypass

## 6 Experimentation

### 6.1 Experimental Setup of the Cyber Range

The experimental environment setup was created using our cybersecurity exercises scenario modeling language [YK19b]. The experimental setup was used to conduct three cybersecurity exercises, in which one was against the attack agent and two were against the defense agent. A total of 101 people who were from 20–25 years old and from Norway participated in the exercises; quantitative methods were used to evaluate the agents' performance.

The attack agents were used in a digital forensic and incident response cybersecurity exercise at the Norwegian University of Science and Technology [Cou], in which 84 people participated in 17 groups on a multi-subnet exercise network environment of 408 machines. Each group was provided with individual networks compromised by human attackers and attack agents. Each network contained 11 Windows- and Linux-based machines, and 2 out of 17 networks were compromised by the attack agent. The participants did not have any knowledge about the attacker and how they exploited the machine.

The defense agents were used in two cybersecurity exercises, which were conducted at the Norwegian Cyber Range [OmN]. The first exercise was a 48-hour long qualification round for the Norwegian national team for the European Cyber Security Challenge [Nor,Eur], in which 17 people participated in 5 groups on a multi-subnet exercise network environment of 75 machines. The second exercise was a 2-week-long exercise conducted during the Ethical Hacking course taught at the Norwegian University of Science and Technology, in which 84 people participated in 17 groups on a multi-subnet exercise network environment of 408 machines. Both exercises were focused on a penetration testscenario of a small organization.

The organization infrastructure has a multi-subnet network containing 11 Windows- and Linux-based machines. The participants had access to the public network through 5 dedicated Kali machines. For the attack agents, two Kali machines were used for launching attacks and creating forensic traces. In the organization's infrastructure, 2 out of those 11 machines had the same vulnerabilities, but a single

machine was running the proposed defense agent. Each group was assigned a segregated replica of the organization network and tasked with doing a pen-test. The scenario's ultimate goal was to tamper with the content of a file in the scenario machine running the defense agent, and the participants were incentivized with extra points to achieve the goal. However, they did not know the presence of the defense agent. A schematic diagram representing the experimental infrastructure is presented in the Figure 12.

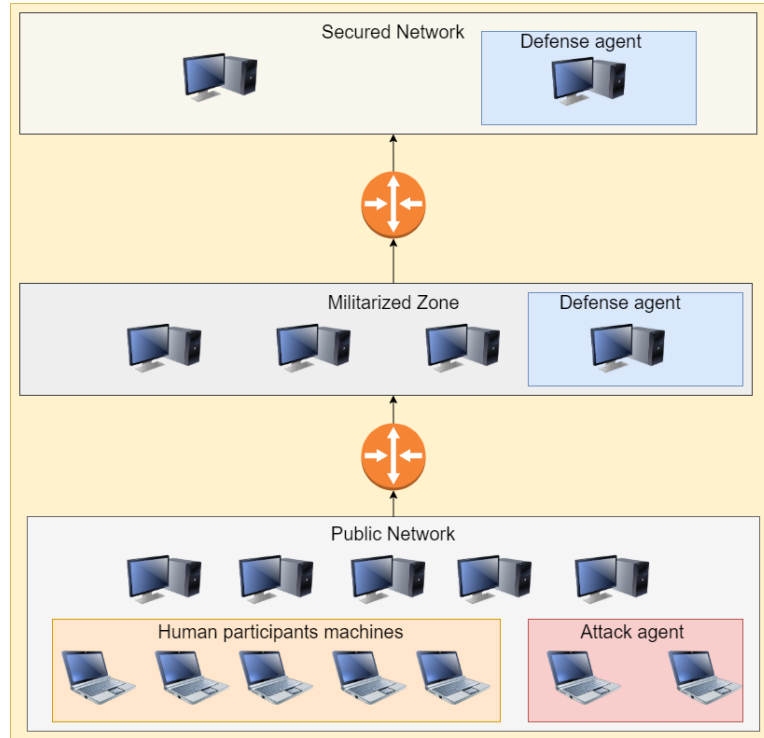


Fig. 12: Experimental setup of the cyber range

## 6.2 Test Cases

**6.2.1 Attack Agent** We defined four test cases to evaluate the attack agent performance. These test cases were selected based on the information we gathered from cybersecurity exercises [YKG20]. These include scanning the network, performing exploitation, post-exploitation, and launching attacks that were not successful. Details of the test cases are as follows:

1. Perform network scan on all the machines present in the exercise network.

2. Exploit  $n$  of the machines present in the exercise network.
3. Perform post-exploitation on  $n$  of the machines present in the exercise network.
4. Launch unsuccessful attacks on  $n$  of the machine present in the exercise network.

In Listing 8, a snippet of the test cases execution is provided. The agent's goal according to *EP* was to exploit system. The agent with the IP address of 10.10.4.96 first performed a full network scan using Nmap on subnet 10.10.1.1/24 to emulate a realistic adversary. Then, the second action was to launch a successful FTP exploit on 10.10.1.4; the FTP exploit was designed as Metasploit resource script, so it performed the post-exploitation steps automatically. After that, the agent launched an unsuccessful attack on 10.10.1.5 using the same exploit and a successful attack on 10.10.1.6 using a different exploit.

Listing 8: Concrete syntax for attacker behavior emulation

```
[
{
  "ActiveScan": {
    "AgentIP": "10.10.4.96",
    "AgentUserID": "root",
    "AgentUserPassword": "toor",
    "Argument": "SV",
    "Target": "10.10.1.1/24"
  },
  "MetaSploit": {
    "AgentIP": "10.10.4.96",
    "AgentUserID": "root",
    "AgentUserPassword": "toor",
    "Argument": "FTPexploit",
    "Target": "10.10.1.4"
  },
  "MetaSploit": {
    "AgentIP": "10.10.4.96",
    "AgentUserID": "root",
    "AgentUserPassword": "toor",
    "Argument": "FTPexploit",
    "Target": "10.10.1.5"
  },
  "MetaSploit": {
    "AgentIP": "10.10.4.96",
    "AgentUserID": "root",
    "AgentUserPassword": "toor",
    "Argument": "Vulnserver",
    "Target": "10.10.1.6"
  }
}
```

```
}
]
```

**6.2.2 Defense Agent** In case of the defense agent, we used three cases to evaluate their performance indicators:

1. Number of machines exploited not running the defense agent.
2. Number of machines exploited running the defense agent.
3. Files that are tampered with and that were monitored by the defense agent.

We used the knowledge base similar to Listing 4, 5, 7, and 6 in the case study for test case execution. The defender’s goal according to *EP* was to detect and deny the attacker using its knowledge base. In the knowledge base, different attacker actions such as information gathering and user and file manipulation were presented, and the defender’s actions against those activities were given.

### 6.3 Evaluation

We employed both quantitative and qualitative evaluation metrics to evaluate the agents. The quantitative metrics were used formally analyses the agent properties and to measure the efficiency of the developed artifacts in which the performance of humans was compared with the proposed agent in similar task with respect to time and resources; this is discussed in sections 6.3.2 and 6.3.3. The qualitative metrics used to measure the realism through a survey conducted on the participants who took part in the experimental scenario. The questions we asked are given in Appendix A, and its details are presented in section 6.3.4.

**6.3.1 Agent Decision Modeling and Verification** The EP model presented in section 4.1 helps us to analyze different test cases before their actual execution by the agents. This analysis helps us to verify different model properties like

- How high level goal can be translated into to low level actions
- Can the agent fulfill the given goal?
- What information is missing to achieve the goal?

This enables us to fine-tune agent decisions based upon model analysis for their precise execution. Listing 9 and 10 presents the implementation and logical verification conditions of the EP model for attack and defense decisions in a PyDatalog [pyD]. While listing 11 presents a sample analysis of the presented models. PyDatalog is an implementation of Datalog and in python and is used for its easy interoperability with rest of the technology stack used in this research for artifact development.

Listing 9: EP Decision Model implementation for attack



```
#Defining necessary term for the model
pyDatalog.create_terms('Goal','SubGoal','Condition','Action',
'CheckGoals','FullFilled','ActionName','ActionTarget','Commands')

#Defining root goal with sub goals of attack agent with an AND
↪ relation
+Goal('Exploit System','Reconnaissance')
+Goal('Exploit System','Exploitation')

#Defining sub goal of attack agent with an OR relation
+Goal('Exploitation','Service' or 'Configuration')

#Defining attack actions
+Action('Default', '172.168.2.1')
+Action('ping', '172.168.2.1')
+Action('Nmap', '172.168.2.1')

#Defining fulfillment requirements for attack agent
+FullFilled('Reconnaissance','ping')
+FullFilled('Reconnaissance','Nmap')
+FullFilled('Reconnaissance','WirShark')
+FullFilled('Reconnaissance','NetCat')
+FullFilled('Reconnaissance','Default')

#Defining conditions for attack agent

##Checking action target information is provided
FullFilled('Reconnaissance', SubGoal) &
↪ Action(SubGoal,ActionTarget) or print ('NotFullFilled')

##Condition to check is target accessible or not
+Condition('Reconnaissance', 'TargetAccess', '-i 4' )

##Validating target access
Condition(Goal, 'TargetAccess', Commands) &
↪ FullFilled('Reconnaissance', SubGoal) &
↪ Action(SubGoal,ActionTarget) or print ('NotFullFilled')

##Condition to check is network interface enabled
+Condition('Reconnaissance', 'NetworkInterface', 'netstat -i' )

##Validating network interface is enabled
```

```

Condition(Goal, 'NetworkInterface', Commands) &
↳ FullFilled('Reconnaissance', SubGoal) &
↳ Action(SubGoal,ActionTarget) or print ('NotFullFilled')

##Condition to check is arp-scan is present
+Condition('Reconnaissance', 'arp-scan-check', 'man arp-scan' )

##Validating network arp-scan is present
Condition(Goal, 'arp-scan-check', Commands) &
↳ FullFilled('Reconnaissance', SubGoal) &
↳ Action(SubGoal,ActionTarget) or print ('NotFullFilled')

##Checking action name information is provided
FullFilled('Reconnaissance', SubGoal) &
↳ Action(SubGoal,ActionTarget) or Action('Default',ActionTarget)

##Defining the attacker action to execute
+Condition('Reconnaissance', 'Active', '-sS -sV' )
+Condition('Reconnaissance', 'Passive', 'arp-scan -interface=eth0
↳ -localnet' )
+Condition('Reconnaissance', 'Default', 'nc -zv ' )

```

Listing 10: EP Decision Model implementation for defense

```

##Defining goals and sub goals of defense agent
+Goal('Prevent Attacks','Detect')
+Goal('Prevent Attacks','Deny')

##Defining defense actions
+Action('Deny', 'shell.exe')
+Action('Deny', 'port 8080')

##Defining fulfillment requirements for defense agent
+FullFilled('Detect','shell.exe')
+FullFilled('Detect','rootkit.exe')
+FullFilled('Detect','chroot.exe')
+FullFilled('Detect','port 8080')

##Check the attack action is detected or not
FullFilled('Detect', SubGoal) & Condition('Deny', SubGoal,
↳ Commands) or print ('NotFullFilled')

```

```
#Defining conditions for defense agent to prevent attack action
+Condition('Deny', 'shell.exe', 'taskkill /IM "shell.exe" /F' )
+Condition('Deny', 'Default', 'npx kill-port 8080' )
```

Listing 11: EP Decision Model analysis

```
#Establishing links between goals, sub goals
CheckGoals(Goal,SubGoal) <= CheckGoals(SubGoal,Goal)
FullFilled(Goal,SubGoal) <= CheckGoals(Goal,SubGoal)

#Establishing links between sub goals and conditions
CheckGoals(SubGoal,Condition) <= CheckGoals(SubGoal,Condition)
FullFilled(SubGoal,Condition) <= CheckGoals(SubGoal,Condition)

#Establishing links between conditions and action
CheckGoals(Condition, Action) <= CheckGoals(Action,Condition)
FullFilled(Condition, Action) <= CheckGoals(Action,Condition)

#Check a goal can be full filled to perform specific action
FullFilled(Goal,SubGoal) <= CheckGoals(Goal,Action) &
↳ FullFilled(SubGoal,Action) & (Goal != SubGoal)

#Sample analysis condition for attack agent EP decision
Condition(Goal, 'Actvie', Commands) & FullFilled('Reconnaissance',
↳ SubGoal) & Action(SubGoal,ActionTarget)

#Sample analysis condition for defense agent EP decision
FullFilled('Detect', SubGoal) & Condition('Deny', SubGoal,
↳ Commands)
```

While translating high-level goals to low-level action and task is a complex and challenging task, our model can perform it based upon the given conditions. Furthermore, the agents' different decisions were verified, highlighting the decisions that can result in goals not fulfilled or impractical. This allowed us to plan agent decisions based upon the scenario requirement. Like in some scenarios, the agents need to make wrong decisions against human adversaries to maintain realism. Similarly, in some scenarios, the agents were required to execute the actions as fast as possible, like performing dry runs on exercise infrastructure, so their decision model can be tuned to avoid unfulfilled and impractical decisions to save time. A sample decision model verification for attack and defense agent decisions highlighted in Figure 4 and Figure 5 is presented in Figure 13.

### 6.3.2 Attack agent Results

```

C:\Python\python.exe
#####
#   Attack Agent EP Model Decision   #
#####
INFO:pyDatalog.pyEngine:New fact : Condition('Reconnaissance','Active','-sS -sV')
INFO:pyDatalog.pyEngine:New fact : Fullfilled('Reconnaissance','NetCat')
INFO:pyDatalog.pyEngine:New fact : Fullfilled('Reconnaissance','WirShark')
INFO:pyDatalog.pyEngine:New fact : Fullfilled('Reconnaissance','Default')
INFO:pyDatalog.pyEngine:New fact : Fullfilled('Reconnaissance','ping')
INFO:pyDatalog.pyEngine:New fact : Fullfilled('Reconnaissance','Nmap')
INFO:pyDatalog.pyEngine:New fact : Action('Nmap','172.168.2.1')
INFO:pyDatalog.pyEngine:New fact : pyD_query1('Reconnaissance','-sS -sV','Nmap','172.168.2.1')
Goal      | Commands | Action | ActionTarget
-----|-----|-----|-----
Reconnaissance | -sS -sV | Nmap   | 172.168.2.1
#####
#   Defense Agent EP Model Decision   #
#####
INFO:pyDatalog.pyEngine:New fact : Fullfilled('Detect','rootkit.exe')
INFO:pyDatalog.pyEngine:New fact : Fullfilled('Detect','chroot.exe')
INFO:pyDatalog.pyEngine:New fact : Fullfilled('Detect','shell.exe')
INFO:pyDatalog.pyEngine:New fact : Condition('Deny','shell.exe','taskkill /IM "shell.exe" /F')
INFO:pyDatalog.pyEngine:New fact : _pyD_query2('shell.exe','taskkill /IM "shell.exe" /F')
Action    | Commands
-----|-----
shell.exe | taskkill /IM "shell.exe" /F
#####
Press any key to continue . . .

```

Fig. 13: Agent decision model verification

**6.3.2.1 Task performed by humans:** Human teams of attackers were given the task to perform penetration testing on the segregated exercise networks presented in Figure 12. They had to discover, exploit, analyze, and report the identified vulnerabilities in a penetration testing report. The penetration testing report was used for their evaluation and comparison with the attack agent’s performance.

**6.3.2.2 Task performed by agent:** The attack agent was tasked with performing penetration testing on the similar segregated exercise network presented in Figure 12. First, the attack agent performed a full network scan of the network to emulate an adversary’s scanning. Then, the attack agent created forensic evidence by launching attacks and performing post-exploitation. Out of the 11 machines, the attack agent was tasked to compromise 4 machines, perform post-exploitation, launch failed attacks on 3 machines, and launch no attack on 4 machines. The attack agent was programmed to emulate a human adversary, so it created successful and unsuccessful attack traces for forensic investigators.

**6.3.2.3 Comparison of human and attack agent performance:** The humans and attack agents were given same task, but the humans participated in teams of five, and most teams compromised a minimum of four machines, while one team compromised eight machines in the exercise network. The human teams took around 50 hours to complete the assigned task, while the attack agent were able to emulate the human performance in approximately 10 minutes.

**6.3.2.4 Verification of performance:** Human and attack agent performance was measured in the same way. The human participants in the digital forensic and incident response exercise were tasked with performing the forensic analysis of the compromised network by the human teams and the attack agent and to present

their findings in a digital forensic and incident response report. The report was used to assess the performance of the attack agent in the cybersecurity exercise, the summary of the findings are presented in Table 1:

Exercise 1			
Group task	Compromised machines identified	Post-exploitation identified	Attack attempts identified
Forensic analysis of machine compromised by humans	3	3	3
Forensic analysis of machine compromised attack agent	4	4	3

Table 1: Results of the cybersecurity exercise against the attack agent

**6.3.2.5 Summary of the results:** The human participants detected most of the successful attacks, post-exploitation, and unsuccessful attack attempts. The attack traces were identical to the attack traces generated by an human attacker, and the participants were not able to identify that they were generated by an attack agent. This indicates that the attack agent was providing the necessary realism and removing the need for a red team member for launching attacks, thus increasing the efficiency by automating different CKC phases like *Reconnaissance*, *Delivery*, *Exploitation*, *Installation*, *Command and Control*, and *Actions on Objectives*. It was identified that the proposed agent was working as expected and suitable in a cybersecurity exercise for creating digital forensic traces. We didn't include the factor of time in this work, which can be used by an experienced forensics analysts for distinction between human and attack agent generated forensic traces. However, over the passage of time we are incorporating the concepts of technical injects that are executed according to a scenario timeline to make the generation of forensic traces as realistic as possible.

### 6.3.3 Defense Agent Results

**6.3.3.1 Tasks performed by humans:** Teams of humans were tasked to compromise a vulnerable machine that was or was not running the defense agent. In the scenario, the human teams were incentivized with additional points to exploit a particular machine known as the *CEO Machine*, which had the same vulnerabilities present in another machine *Machine9* but was running the defense agent.

**6.3.3.2 Tasks performed by agents:** The defense agent was tasked to block or prevent the attacks launched by the human attackers. The defense agent had

a knowledge base to prevent particular attacks from a specific vulnerability also present in *Machine9*. Additionally, the defense agent was also monitoring a local file to prevent access to the attacker in case the attacker exploited a vulnerability not present in the defense agent’s knowledge base.

**6.3.3.3 Comparison of human and attack agent performance:** In the exercise environment, there were no machines actively defended by human adversaries, so comparing the performance regarding efficiency of the human and defense agents is a bit difficult. However, the machines that were defended by the proposed agents were difficult to exploit compared with the undefended machines. According to IBM, the average detection time for a data breach is 206 days [Cos19]; we are not arguing that our proposed solution will drastically improve this, but we are arguing that having an active agent running in a system will restrict the attacker’s actions and improve system security.

**6.3.3.4 Verification of performance:** In the second exercise, three out of five groups compromised the vulnerable machine not running the defense agent. In comparison, one group compromised the machine running the defense agent using a vulnerability that was not in the defense agent’s knowledge base but could not tamper with the file because of the defense agent’s actions. Similarly, in the third exercise, 8 out of 17 groups compromised the vulnerable machine, and 2 groups were able to compromise the machine running the defense agent using a vulnerability not present in the defense agent’s knowledge base but were not able to tamper with the content of the file. The results of the experiments are presented in Table 2.

Exercise 2						
Number of Groups	Groups Vulnerable	Exploited Machine	Groups Vulnerable Running Agent	Exploited Machine Defense	Groups Tampered with the File	
5	3		1		0	
Exercise 3						
17	8		2		0	

Table 2: Results of the cybersecurity exercise against the defense agent

**6.3.3.5 Summary of the results:** The results indicate that the defense agent created the necessary friction and added realism by preventing attacks present in its knowledge base during an operational cybersecurity exercise. The defense agent removed the need for an active blue team member during the exercise, thus increasing the efficiency by automating certain tasks in CKC like *Detect* and *Deny* for the defender. The inclusion of the defense agent made the exercise environment more dynamic and challenging because some groups were not able to compromise the machines running the defense agents.

**6.3.4 Qualitative feedback from the exercise participants** The qualitative feedback consists of survey responses from the 3 teams, which represent 15 students. The survey was conducted in a semi-informal way. The participants were asked to answer a list of pre-defined open-ended questions using digital communication apps, hence following COVID-19 restrictions. Follow-up questions were asked if the answers needed further explanation. The survey was conducted in a relaxed and friendly environment, and the participants were given sufficient time to reflect on their experience and properly answer the question to avoid any biases. No personally indefinable information was collected during the survey to avoid any GDPR (*General Data Protection Regulation*)-related issues. During the survey, when we asked the participants if they found the scenario realistic, one survey participant stated the following:

*Scenarios were pretty realistic for the hacking phase*

This sentiment was shared by the majority of the survey participants. The exercise was conducted in two phases *Ethical Hacking* and *Incident response & Forensics*. The participants were very impressed by the complexity and dynamism of the scenario in the *Ethical Hacking* phase; this can be attributed to the presence of a dynamic defense agent in the exercise environment. However, for the *Incident response & Forensics*, they were not that impressed because the environment was static. The participants expected continuing attacks during this phase to make it more dynamic. The feedback of the participants was noted for implementing active attack execution in the *Incident response & Forensics* phase in the future. When we asked about the difficulty of the scenario, one participant stated the following:

*I think it is good that the scenario is large and consist of both easy machines and more difficult ones. This allows weaker students to be able to get points and provides a challenge for stronger students with much experience. In my opinion, the project is good from a grading perspective*

This sentiment was also shared by most of the survey respondents. The participants indicated that they found the machines to exploit having a variety of difficulty levels *easy*, *medium*, and *hard*, which allowed the participants with ranging skill sets to practice their skills. This indicates that the presence of the developed agents in the scenario provided balance in the lab, which made some machines difficult to exploit because of them having similar vulnerabilities to other easily exploitable machines. When we asked about the number of machines exploited by their teams, two teams stated that they exploited four machines, while one team exploited nine machine. Continuing from this, we asked a specific question about the machine that was not running the agent *Machine9* and the machine that was running the defense agent *CEO machine*. Two teams were not able to exploit both machines, while one team was able to exploit it *Machine9*. When we asked why they were not able to exploit, it they responded with the following:

*No exactly each planned attack went through except for one where we were trying to do an smb exploitation but we couldn't figure out and came to the conclusion that it was rabbit hole and moved on*

The rest of the questions were asked to improve the quality of the exercise scenario and are not relevant to this study. From the qualitative feedback, it can be concluded that the exercise scenario that incorporated our agents are quite realistic and offer the opportunity to exercise participants to practice their skills against realistic computational adversaries.

## 7 Discussion and Conclusion

In this work, we investigated the cyber-attack and defense agents' usage of cyber ranges for improving the realism and efficiency of cybersecurity exercise execution. We identified that such agents could provide the necessary level of friction during an exercise. For example, in a red team exercise, a defense agent will try to prevent attackers from achieving their objectives. On the other hand, in a blue team exercise, an attack agent will conduct various attacks and create forensics traces, such that the need for a human red team is reduced. This makes cybersecurity exercise execution more realistic and efficient.

We proposed EP models for specifying the agents' decision making. An EP model contains three levels of decisions: *high*, *medium*, and *low*. These decisions were translated using a DSL into goals, actions, and commands. We presented the workflows of the attack and defense agents to showcase how they made their decisions during the execution of a cybersecurity exercise. We employed the proposed agent-based system in cybersecurity exercises and presented their performance results in the form of a case study.

For the attack agent, we consider its performance satisfactory when applied in a semi-autonomous manner. The attack agents create realistic forensic traces during a cybersecurity exercise, which is verified by the human participants. On the other hand, our developed cyber defense agents currently have the six capabilities highlighted by Kott [KTD<sup>+</sup>18b]. However, we do not consider these agents suitable for deployment in the actual production environment for active cyber defense because they were tested in a controlled environment of cyber ranges. Yet, they can be considered a first step to achieve autonomous cyber defense. One of the limitation of the proposed defense agents is in the way they respond to attacks. If an agent detects a new attack on a specific port, it will update the knowledge base with some information about the attack's signature. Then, it will close the port or service for the next attack with a similar signature without analyzing the impact of its action, which is not suitable for real-world systems. We will address this issue in our future research, and we are planning to develop state machines that can help the defense agent take the optimal action to deal with the attacker.

In the future, we are also planning to integrate specific test cases with the attack agent to automatically perform the security assurance of systems and provide a quantitative score of the state of system security [KP18,KP19]. Finally, we plan to use our developed cyber-attack and defense agents in additional cybersecurity exercises with different network topology positions. This will help us analyze the impact of such agents at different stages of attacks and defenses, for example, an



exercise designer can use such agents in an outer network on dispensable systems to obtain the attacker's tactics and techniques for preventing attacks in the internal network.

## References

- [APT] *APTSimulator/test-sets at master · NextronSystems/APTSimulator*. <https://github.com/NextronSystems/APTSimulator/tree/master/test-sets>. – (Accessed on 01/27/2021)
- [asi] *Dosfuscation*. <https://i.blackhat.com/briefings/asia/2018/asia-18-bohannon-invoke-dosfuscation-techniques-for-fin-style-dos-levelcmd-obfuscation-wp.pdf> - (Accessed on 01/28/2021)
- [BCD<sup>+</sup>19] BRAGHIN, Chiara ; CIMATO, Stelvio ; DAMIANI, Ernesto ; FRATI, Fulvio ; MAURI, Lara ; RICCOBENE, Elvinia: A model driven approach for cyber security scenarios deployment. In: *Computer Security*. Springer, 2019, S. 107–122
- [Cen17] CENTER, JPCERT C.: *Detecting lateral movement through tracking event logs*. 2017
- [CGT<sup>+</sup>89] CERI, Stefano ; GOTTLOB, Georg ; TANCA, Letizia u. a.: What you always wanted to know about Datalog (and never dared to ask). In: *IEEE transactions on knowledge and data engineering* 1 (1989), Nr. 1, S. 146–166
- [Cos19] *Cost of a Data Breach Study — IBM*. <https://www.ibm.com/security/data-breach>. Version: 2019. – (Accessed on 06/16/2021)
- [Cou] *Course - Incident Response, Ethical Hacking and Forensics - IMT3004 - NTNU*. <https://www.ntnu.edu/studies/courses/IMT3004#tab=omEmnet>
- [Dat] *Datalog: Deductive Database Programming*. <https://docs.racket-lang.org/datalog/index.html>. – (Accessed on 09/30/2020)
- [EM17] EDGAR, Thomas W. ; MANZ, David O.: *Research methods for cyber security*. Syngress, 2017. – 271–297 S.
- [Eur] *European Cyber Security Challenge — ECSC*. <https://europeancybersecuritychallenge.eu/> - (Accessed on 01/28/2021)
- [FSR09] FEILY, Maryam ; SHAHRESTANI, Alireza ; RAMADASS, Sureswaran: A survey of botnet and botnet detection. In: *2009 Third International Conference on Emerging Security Information, Systems and Technologies IEEE*, 2009, S. 268–273
- [Gen] *Generating Realistic Non-Player Characters for Training Cyberteams*. <https://insights.sei.cmu.edu/blog/generating-realistic-non-player-characters-for-training-cyberteams/>. – (Accessed on 04/19/2021)
- [HC10] HEVNER, Alan ; CHATTERJEE, Samir: Design science research in information systems. In: *Design research in information systems*. Springer, 2010, S. 9–22
- [HCA<sup>+</sup>11] HUTCHINS, Eric M. ; CLOPPERT, Michael J. ; AMIN, Rohan M. u. a.: Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains. In: *Leading Issues in Information Warfare & Security Research* 1 (2011), Nr. 1, S. 80
- [HKR18] HENDLER, Danny ; KELS, Shay ; RUBIN, Amir: Detecting malicious PowerShell commands using deep neural networks. In: *Proceedings of the 2018 on Asia Conference on Computer and Communications Security*, 2018, S. 187–197
- [HS16] HOLM, Hannes ; SOMMESTAD, Teodor: Sved: Scanning, vulnerabilities, exploits and detection. In: *MILCOM 2016-2016 IEEE Military Communications Conference IEEE*, 2016, S. 976–981

- [HWD<sup>+</sup>17] HEROLD, Nadine ; WACHS, Matthias ; DORFHUBER, Marko ; RUDOLF, Christoph ; LIEBALD, Stefan ; CARLE, Georg: Achieving reproducible network environments with INSALATA. In: *IFIP International Conference on Autonomous Infrastructure, Management and Security* Springer, Cham, 2017, S. 30–44
- [JON<sup>+</sup>15] JONES, Randolph M. ; O’GRADY, Ryan ; NICHOLSON, Denise ; HOFFMAN, Robert ; BUNCH, Larry ; BRADSHAW, Jeffrey ; BOLTON, Ami: Modeling and integrating cognitive agents within the emerging cyber domain. In: *Proceedings of the Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC)* Bd. 20 Citeseer, 2015
- [KMRS14] KORDY, Barbara ; MAUW, Sjouke ; RADOMIROVIĆ, Saša ; SCHWEITZER, Patrick: Attack–defense trees. In: *Journal of Logic and Computation* 24 (2014), Nr. 1, S. 55–87
- [Kot05] KOTENKO, Igor: Agent-based modeling and simulation of cyber-warfare between malefactors and security agents in Internet. In: *19th European Simulation Multiconference “Simulation in wider Europe, 2005*
- [KP18] KATT, Basel ; PRASHER, Nishu: Quantitative security assurance metrics: REST API case studies. In: *Proceedings of the 12th European Conference on Software Architecture: Companion Proceedings*, 2018, S. 1–7
- [KP19] KATT, Basel ; PRASHER, Nishu: Quantitative Security Assurance. In: *Exploring Security in Software Architecture and Design*. IGI Global, 2019, S. 15–46
- [KTD<sup>+</sup>18a] KOTT, Alexander ; THÉRON, Paul ; DRAŠAR, Martin ; DUSHKU, Edlira ; LEBLANC, Benoît ; LOSIEWICZ, Paul ; GUARINO, Alessandro ; MANCINI, Luigi ; PANICO, Agostino ; PIHEL GAS, Mauno u. a.: Autonomous Intelligent Cyber-defense Agent (AICA) Reference Architecture. Release 2.0. In: *arXiv preprint arXiv:1803.10664* (2018)
- [KTD<sup>+</sup>18b] KOTT, Alexander ; THOMAS, Ryan ; DRAŠAR, Martin ; KONT, Markus ; POYLISHER, Alex ; BLAKELY, Benjamin ; THERON, Paul ; EVANS, Nathaniel ; LESLIE, Nandi ; SINGH, Rajdeep u. a.: Toward Intelligent Autonomous Agents for Cyber Defense: Report of the 2017 Workshop by the North Atlantic Treaty Organization (NATO) Research Group IST-152-RTG. In: *arXiv preprint arXiv:1804.07646* (2018)
- [KV08] KUECHLER, Bill ; VAISHNAVI, Vijay: On theory development in design science research: anatomy of a research project. In: *European Journal of Information Systems* 17 (2008), Nr. 5, S. 489–504
- [Lin] *Linux Post Exploitation Command List · mubix/post-exploitation Wiki*. <https://github.com/mubix/post-exploitation/wiki/Linux-Post-Exploitation-Command-List>. – (Accessed on 01/25/2021)
- [Llo12] LLOYD, John W.: *Foundations of logic programming*. Springer Science & Business Media, 2012. – 1–31 S.
- [Mim] *Mimikatz - Metasploit Unleashed*. <https://www.offensive-security.com/metasploit-unleashed/mimikatz/>. – (Accessed on 01/25/2021)
- [MIT] *MITRE ATT&CK®*. <https://attack.mitre.org/>. – (Accessed on 01/25/2021)
- [MTWP15] MIRKOVIC, Jelena ; TABOR, Aimee ; WOO, Simon ; PUSEY, Portia: Engaging Novices in Cybersecurity Competitions: A Vision and Lessons Learned at {ACM} Tapia 2015. In: *2015 {USENIX} Summit on Gaming, Games, and Gamification in Security Education (3GSE 15)*, 2015

- [Nai20] NAIK, Mayur: Petablock: Large-Scale Software Analysis and Analytics Using Datalog / GEORGIA TECH RESEARCH INST ATLANTA ATLANTA United States. 2020. – Forschungsbericht
- [net] *netdiscover*. <https://manpages.debian.org/unstable/netdiscover/netdiscover.8.en.html>. – (Accessed on 01/23/2021)
- [Nor] *Norwegian Cyber Security Challenge - NCSC - NTNU*. <https://www.ntnu.no/ncsc>. – (Accessed on 01/28/2021)
- [OmN] *Om Norwegian Cyber Range - NTNU*. <https://www.ntnu.no/ncr>. – (Accessed on 01/28/2021)
- [Pow] *PowerShellMafia/PowerSploit: PowerSploit - A PowerShell Post-Exploitation Framework*. <https://github.com/PowerShellMafia/PowerSploit>. – (Accessed on 01/25/2021)
- [PTE] *PTEs Technical Guidelines - The Penetration Testing Execution Standard*. <https://tinyurl.com/6cgn3cu>. – (Accessed on 01/20/2021)
- [pyD] *pyDatalog*. <https://sites.google.com/site/pydatalog/home> - (Accessed on 09/03/2021)
- [RCA20] RUSSO, Enrico ; COSTA, Gabriele ; ARMANDO, Alessandro: Building Next Generation Cyber Ranges with CRACK. In: *Computers & Security* (2020), S. 101837
- [red] *Atomic-red-team: Small and highly portable detection tests based on MITRE's ATT&CK*. <https://github.com/redcanaryco/atomic-red-team>. – (Accessed on 01/27/2021)
- [Res] *Resource Scripts — Metasploit Documentation*. <https://docs.rapid7.com/metasploit/resource-scripts/>. – (Accessed on 01/19/2021)
- [spl] *splunk/attack\_range: A tool that allows you to create vulnerable instrumented local or cloud environments to simulate attacks against and collect the data into Splunk*. [https://github.com/splunk/attack\\_range](https://github.com/splunk/attack_range). – (Accessed on 01/27/2021)
- [Sto18] STOECKLIN, Marc P.: Deeplocker: How AI can power a stealthy new breed of malware. In: *Security Intelligence, August 8* (2018)
- [TKD<sup>+</sup>18] THERON, Paul ; KOTT, Alexander ; DRAŠAR, Martin ; RZADCA, Krzysztof ; LEBLANC, Benoît ; PIHEL GAS, Mauno ; MANCINI, Luigi ; PANICO, Agostino: Towards an active, autonomous and intelligent cyber defense of military systems: The NATO AICA reference architecture. In: *2018 International conference on military communications and information systems (ICMCIS)* IEEE, 2018, S. 1–9
- [TKD<sup>+</sup>20] THERON, Paul ; KOTT, Alexander ; DRAŠAR, Martin ; RZADCA, Krzysztof ; LEBLANC, Benoît ; PIHEL GAS, Mauno ; MANCINI, Luigi ; DE GASPARI, Fabio: Reference architecture of an autonomous agent for cyber defense of complex military systems. In: *Adaptive Autonomous Secure Cyber Systems*. Springer, 2020, S. 1–21
- [ube] *Uber-Common: An information security preparedness tool to do adversarial simulation*. <https://github.com/uber-common/ubemta>. – (Accessed on 01/27/2021)
- [YB18] YAMIN, Muhammad M. ; BASEL, KATT: Ethical Problems and Legal Issues in Development and Usage Autonomous Adversaries in Cyber Domain. (2018)
- [YK18a] YAMIN, Muhammad M. ; KATT, Basel: Inefficiencies in Cyber-Security Exercises Life-Cycle: A Position Paper. In: *AAAI Fall Symposium: ALEC*, 2018, S. 41–43

- [YK18b] YAMIN, Muhammd M. ; KATT, Basel: Detecting malicious windows commands using natural language processing techniques. In: *International Conference on Security for Information Technology and Communications* Springer, 2018, S. 157–169
- [YK19a] YAMIN, Muhammad M. ; KATT, Basel: Cyber Security Skill Set Analysis for Common Curricula Development. In: *Proceedings of the 14th International Conference on Availability, Reliability and Security*, 2019, S. 1–8
- [YK19b] YAMIN, Muhammad M. ; KATT, Basel: Modeling attack and defense scenarios for cyber security exercises. In: *5th interdisciPlinary cyber research conference*, 2019, S. 7
- [YK22] YAMIN, Muhammad M. ; KATT, Basel: Modeling and executing cyber security exercise scenarios in cyber ranges. In: *Computers & Security* 116 (2022), S. 102635
- [YKG19] YAMIN, Muhammad M. ; KATT, Basel ; GKIOULOS, Vasileios: Detecting Windows Based Exploit Chains by Means of Event Correlation and Process Monitoring. In: *Future of Information and Communication Conference* Springer, 2019, S. 1079–1094
- [YKG20] YAMIN, Muhammad M. ; KATT, Basel ; GKIOULOS, Vasileios: Cyber ranges and security testbeds: Scenarios, functions, tools and architecture. In: *Computers & Security* 88 (2020), S. 101636
- [YKT<sup>+</sup>18] YAMIN, Muhammad M. ; KATT, Basel ; TORSETH, Espen ; GKIOULOS, Vasileios ; KOWALSKI, Stewart J.: Make it and break it: An IoT smart home testbed case study. In: *Proceedings of the 2nd International Symposium on Computer Science and Intelligent Control*, 2018, S. 1–6
- [Yue15] YUEN, Joseph: Automated cyber red teaming / DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION EDINBURGH (AUSTRALIA) CYBER AND . . . . 2015. – Forschungsbericht
- [YUUK21] YAMIN, Muhammad M. ; ULLAH, Mohib ; ULLAH, Habib ; KATT, Basel: Weaponized AI for cyber attacks. In: *Journal of Information Security and Applications* 57 (2021), S. 102722
- [ZM09] ZEIDANLOO, Hossein R. ; MANAF, Azizah A.: Botnet command and control mechanisms. In: *2009 Second International Conference on Computer and Electrical Engineering* Bd. 1 IEEE, 2009, S. 564–568
- [ZN20] ZABER, Matthew ; NAIR, Suku: A framework for automated evaluation of security metrics. In: *Proceedings of the 15th International Conference on Availability, Reliability and Security*, 2020, S. 1–11

## A Survey Question

1. Did you find the scenario realistic?
2. Did you find the scenario difficulty hard, medium, or easy?
3. How many machines did you exploited?
4. Did you find similarities between Machine9 and CEO machines?
5. Did you identify any of your attacks get blocked?
6. If yes, did you exploited both or only one and why?
7. What can be improved in the scenario?