

Automating the Generation of Cyber Range Virtual Scenarios with VSDL

Gabriele Costa¹, Enrico Russo^{2*}, and Alessandro Armando²

¹IMT School for Advanced Studies, Lucca 55100 Italy
gabriele.costa@imtlucca.it

²University of Genoa, DIBRIS, Genoa 16145 Italy
enrico.russo@unige.it, alessandro.armando@unige.it

Received: October 17, 2022; Accepted: December 4, 2022; Published: December 31, 2022

Abstract

A *cyber range* (CR) is an environment used for training security experts and testing attack and defense tools and procedures. Usually, a cyber range simulates one or more critical infrastructures that attacking (red) and defending (blue) teams must compromise and protect, respectively. The infrastructure can be physically assembled, but much more convenient is to rely on the Infrastructure as a Service (IaaS) paradigm. Although some modern technologies support the IaaS, the design and deployment of scenarios of interest are mostly manual. As a consequence, it is a common practice to have a cyber range hosting few (sometimes only one), consolidated scenarios. However, reusing the same scenario may significantly reduce the effectiveness of the training and testing sessions. In this paper, we propose a framework for automating the definition and deployment of arbitrarily complex cyber range scenarios. The framework relies on the *virtual scenario description language* (VSDL), i.e., a domain-specific language for defining high-level features of the desired infrastructure while hiding low-level details. The semantics of VSDL is given in terms of constraints that must be satisfied by the virtual infrastructure. These constraints are then submitted to an SMT solver to check the satisfiability of the specification. If satisfiable, the specification gives rise to a model that is automatically converted to a set of deployment scripts to be submitted to the IaaS provider.

Keywords: Cyber range, cybersecurity training, scenario description language

1 Introduction

Cyber defense (as well as offense) relies on security experts that must be properly trained through hands-on activities [1] and equipped with adequate security tools. This simple fact is boosting the interest of the international community toward creating *cyber ranges*. In short, a cyber range is an environment where trainees compete or cooperate to achieve specific security goals. Needless to say, they should interact with a realistic environment that accurately mimics real-world settings. A common practice is to have a defending, aka *blue*, team and an attacking, aka *red*, team. For instance, the blue team can be asked to enhance the security of an Information Technology (IT) infrastructure in a limited amount of time. Afterward, the red team must violate the security of the infrastructure by accessing target data or compromising a certain resource. Specific security vulnerabilities of interest can also be injected into

Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications, 13(4):61-80, Dec. 2022
DOI: 10.58346/JOWUA.2022.14.004

*Corresponding author: Department of Informatics, Bioengineering, Robotics, and Systems Engineering (DIBRIS), Viale Causa, 13, 16145 Genoa, Italy

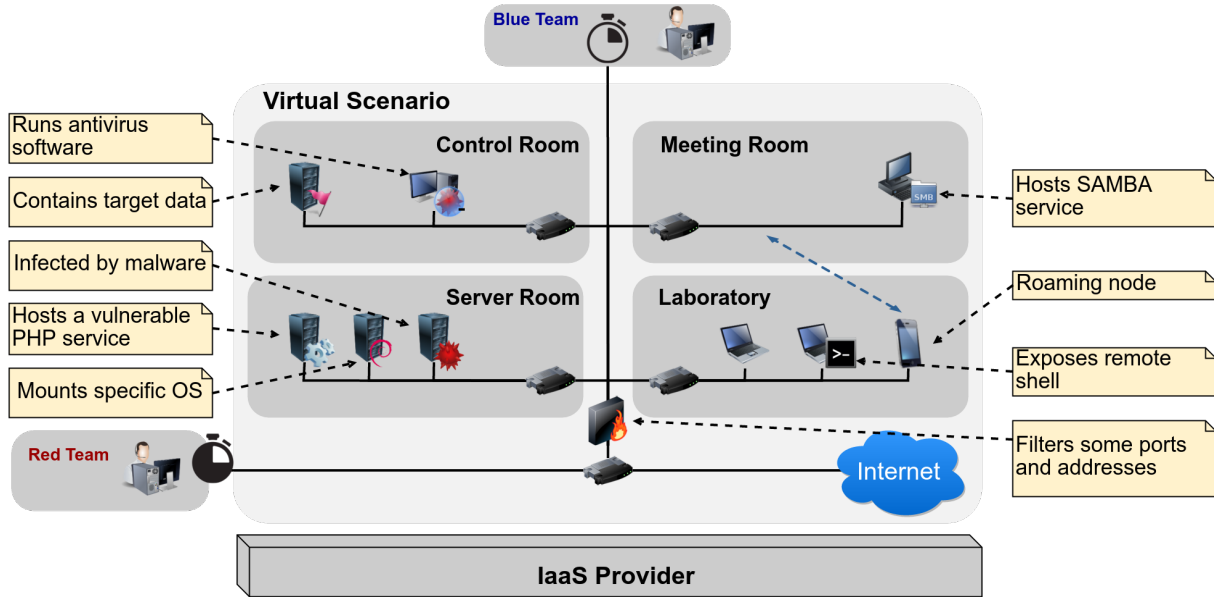


Figure 1: A virtual scenario with annotated requirements.

the original infrastructure to organize aimed sessions. The operational environment, including networks, hardware, software, and how they behave during the session, is called a *scenario*.

Infrastructure as a Service (IaaS) is a convenient paradigm for defining and deploying the elements of a scenario. Virtualization technologies can emulate both networks and computational nodes. For instance, OpenStack [2] can be used to emulate large-scale, heterogeneous networks of virtual machines. However, defining virtual scenarios for a cyber range poses several specific issues. The main limitation is the relatively short lifetime of a scenario. In principle, a cyber range should permit the definition of a scenario that is used for a session lasting a few hours. Ideally, the scenarios should not be reused as they might become repetitive and rapidly lose interest. Thus, the cyber range should provide a mechanism for rapidly generating new scenarios while guaranteeing that they include the desired features. Unfortunately, the process of defining a scenario using the existing IaaS solutions is nontrivial. Indeed, IaaS platforms are usually designed for defining and hosting long-term infrastructures. As a consequence, existing cyber ranges tend to reuse few scenarios. For instance, the Michigan cyber range [3] has a single virtual infrastructure of four buildings of a fictional town called Alphaville. Similarly, the Open Cyber Challenge Platform (OCCP) [4] includes a limited number of scenarios.¹ Instead, the annual NATO Locked Shields initiative [5] relies on a virtual scenario renewed yearly and only used for two days. As mentioned above, reusing the same scenario drastically reduces the effectiveness of the training activity and, thus, the usefulness of the cyber range.

We propose an example to highlight better the structure of a virtual scenario and its features. We will use it as a working example along the paper to constructively show the steps of our approach.

Example 1. Consider the scenario graphically depicted in Figure 1. It consists of a network composed of four sub-networks, i.e., Server Room, Laboratory, Meeting Room, and Control Room. The blue team is deployed within the network perimeter and goes offline after a certain time. Instead, the red team access is delayed and takes place from outside the perimeter, i.e., from the public network.

The features of the elements appearing in the scenario (written inside note labels) are heterogeneous

¹For the time being, only one scenario is documented, and few others have been announced.

and partial. Most of them are straightforward. For instance, the Laboratory network must consist of two laptops and a mobile phone. One of the laptops is accessible through a remote shell, and the mobile phone, at a certain time, must migrate to the network Meeting Room. Notice that no other aspects of these nodes are relevant to the scenario, e.g., the structure of the filesystem of the three devices.

Reasonably, the features that one wants to specify when defining a scenario belong to the following categories.

Networking. What are the existing channels and connections? Are there active firewalls? What rules do they apply?

Hardware. What are the hardware capabilities of the nodes (CPU speed, disk size)? What is their role in the scenario (e.g., servers, mobile phones, or laptops)?

Software. What OS runs on a node? What applications and libraries are installed? Is there a software monoculture [6]?

Data. What information is stored in a node? Does the file system contain any relevant files?

Users and privileges. Who can access a certain node? What are the privileges of the users over the file system of a node?

Time. How does the infrastructure change over time? Are there nomadic nodes? Are there failures of nodes or networks?

We claim that domain-specific languages [7] (DSL) are in order. As a matter of fact, their tailored syntax can precisely describe the desired features of any infrastructure. Also, the formal semantics of a DSL supports and drives automatic validation, refinement, and implementation processes.

In this paper, we present a framework for the automatic validation and implementation of virtual scenarios for cyber ranges. The framework relies on a *virtual scenario description language* (VSDL) for the high-level specification of the scenario properties. The semantics of a VSDL specification is given in terms of (quantifier-free) linear integer arithmetic (QFLIA, see [8]) assertions. Assertions are then processed through a *satisfiability modulo theories* (SMT) solver, which checks whether they admit a model. The model assigns values to constants and functions, and it is automatically translated into a corresponding virtual scenario.

Our approach provides several advantages over the manual modeling of a scenario. Among them, the most important ones are the following.

Verifiability. Ensuring that a scenario exposes some relevant features, e.g., the presence of a vulnerability, is usually nontrivial. Satisfiability checking returns a model with the requested properties. The model is automatically translated into a scenario script, so avoiding errors that might derive from a manual implementation.

Expressiveness. VSDL permits the description of a scenario using a rich syntax. Expressible statements cover many aspects of the scenario, from the high, abstract level to low, concrete details. Temporal conditions can guard statements so that scenarios evolving in time can also be modeled.

Compositionality. Existing scenarios can be modified and extended by adding new statements, elements, or even entire infrastructures. For instance, a scenario can be created by combining the infrastructures used in other previously defined scenarios. The verification process guarantees that the composition does not invalidate the required features.

Integration. The result of the instantiation process is a set of scripts for the deployment engine of an IaaS provider. As far as they refer to distinct entities, the scripts can be combined with those produced through other channels (e.g., manually written).

A further contribution of this work is the implementation of a working prototype that has been integrated with state-of-the-art technologies like OpenStack, Terraform, and Packer.

This paper is structured as follows. Section 2 describes some related works on cyber ranges and virtual infrastructures. In Section 3 we provide an overview of the architecture of our framework, while in Section 4, we present the virtual scenario description language and its interpretation. Finally, Section 5 presents the virtual scenario generation procedure, and Section 7 concludes the paper.

2 Related work

A cyber range hosts one or more virtual infrastructures used for training and testing purposes (see [9] for a survey). The number of such facilities is rapidly growing and many active projects exist, e.g., see [10, 11, 3, 4, 12].

The *NATO Locked Shield* [5] and the *National Cyber Range* [12] (NCR) are among the most prominent initiatives. We already mentioned in the introduction the first one and the related yearly events. Instead, the NCR relies on a physical infrastructure partially documented in [13]. However, at the best of our knowledge, more recent proposals tend to avoid bare metal implementation as it is more costly and less flexible.

In general, virtual infrastructures are attracting major interest by both academia and industry. The main reason is that they decouple the computational elements from the physical infrastructure hosting them. This favors their re-usability, maintainability, adaptiveness and resilience.

The technologies supporting Infrastructure as a Service (IaaS) are the main candidates for the implementation of a cyber range. For instance, IBM's *Softlayer* [14] and VMWare's *vCloud* [15] can be used to deploy a private cloud. Some environments, e.g., Cisco's *Fog Computing* [16], even support mixed cloud-IoT infrastructures. Nevertheless, these solutions target long-term infrastructures, while cyber range scenarios can have a very short life.

Some authors proposed specification languages for describing virtual infrastructures. For instance, in [17], the *Infrastructure and Network Description Language* (INDL) is presented. There the authors show that INDL is expressive enough to model two virtual infrastructures of interest studied by two EU projects. Another proposal is the description language VXDL [18]. With VXDL one can define the requirements that the infrastructure must satisfy to achieve its goal, e.g., in terms of latency. These languages can precisely describe a virtual infrastructure. However, because of their different purpose, they are not adequate for the definition of virtual scenarios for the cyber range whereby structural aspects may not be strictly defined, e.g., disk size of a server or connection bandwidth, while some requirements must be satisfied, e.g., the presence of a piece of vulnerable software on some node.

Network virtualization can be carried out through *Software Defined Networking* [19, 20] (SDN). SDN allows for the definition and the centralized management of virtual networks abstracted from the physical layer but does not support the description of computational nodes connected to the virtual networks.

A number of languages for the definition and orchestration of services have been put forward, see, e.g., [21, 22, 23, 24, 25]. Some of these frameworks can automatically generate service compositions that satisfy functional or security goals. Although web services can be relevant or even central in a scenario, these languages do not provide adequate support to model the infrastructural elements.

The possibility of injecting vulnerabilities is crucial for the scenarios of the cyber range. Several applications and systems for testing vulnerability scanners and training security analysts have been released in the last few years. For instance, *Damn Vulnerable Web Application* (DVWA) [26], *WebGoat* [27] and

Gruyere [28] deliberately include vulnerabilities and challenges. Similar projects target other environments of interest, e.g., OSes and mobile apps. Even though they are relevant, the training with these applications tends to be artificial and repetitive and, thus, partially incompatible with the requirements of a cyber range.

A further crucial aspect is the number of available resources. As a matter of fact, a virtual infrastructure is executed by a physical platform. In principle, given a scenario, there is no guarantee that the physical environment has enough computational resources to execute it effectively. Some existing cyber ranges [3, 4] avoid this check by using a few scenarios that have been extensively tested. However, it is important to notice that, due to their strategic role, national authorities tend not to divulge the internals of their cyber ranges.

3 Architecture and Workflow

In this section, we describe our approach in terms of the used technologies and how we compose them into a unified workflow.

3.1 Involved technologies

We integrate our proposal with state-of-the-art technologies supporting the creation and management of virtual infrastructures. Such integration guarantees that our approach can be readily applied to the real world.

OpenStack. OpenStack [29] is a platform for the execution of private and public clouds. Many providers, e.g. IBM, VMWare, Cisco, Citrix, etc.², joined the initiative by integrating the OpenStack API into their products. The OpenStack framework consists of a collection of core services dedicated to all the aspects of a virtual infrastructure. Moreover, it provides APIs and a dashboard application that an administrator can use to create, modify and monitor the existing infrastructures.

Terraform. The OpenStack dashboard is designed for manually defining a virtual infrastructure. Terraform [30] provides a convenient way for creating and managing them. In particular, Terraform relies on a scripting language that one can use to describe the virtual elements, e.g., nodes and networks. Then, Terraform translates the script content into a sequence of OpenStack API invocations to create the defined objects. Moreover, a script can be used to update existing elements. Indeed, Terraform automatically checks whether differences exist between the running infrastructure and the new script and only submits the needed modifications. Clearly, a Terraform script must precisely describe the elements to be created.

Packer. Another task that one might want to automate is the creation and configuration of node images. This operation requires generating an OS image that must be installed on a node and configuring it with the required software. Similarly to Terraform, Packer [31] offers a convenient scripting language for defining and customizing OS images. A terraform script can exploit one or more images created with Packer for initializing a computational node.

3.2 Workflow

Figure 2 depicts the abstract workflow of our approach. We start from a running instance of OpenStack. Such instances can host one or more virtual infrastructures, created by different administrators, called

²see <https://www.openstack.org/marketplace/drivers/> for a complete list.

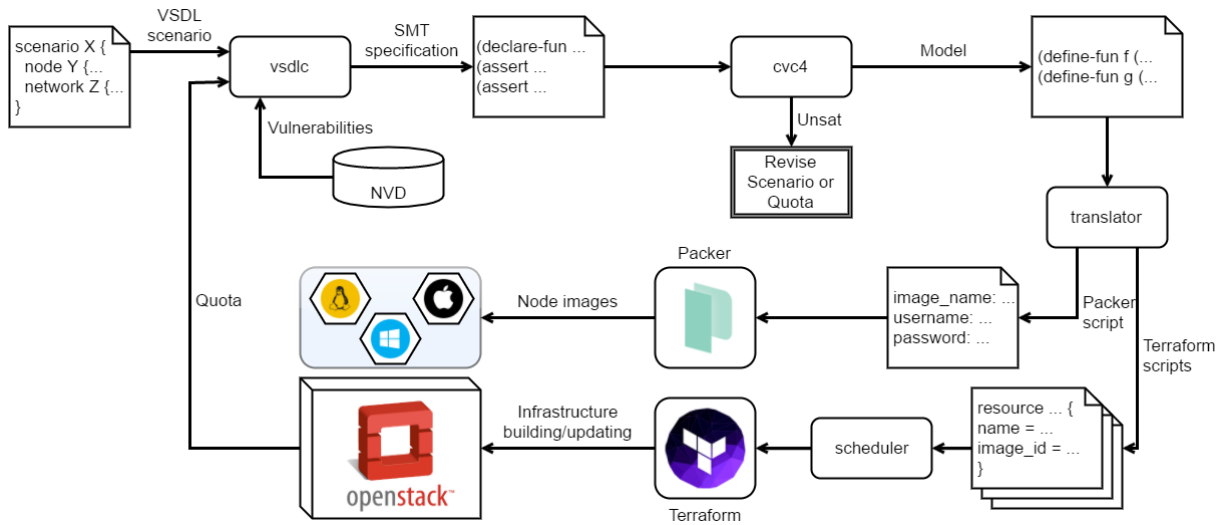


Figure 2: Schematic representation of the VSDL-based workflow.

tenants. Each tenant has a *quota* of assigned, virtual resources (e.g., virtual CPUs and virtual disk size). A tenant wanting to create a new scenario writes a VSDL specification as described in Section 4.1. Then, the specification is processed by the VSDL compiler (i.e., *vsdlc*) which also retrieves the quota information from OpenStack. Moreover, the compiler collects the definition of the vulnerabilities mentioned in the specification file from a local database. Briefly, the repository is a copy of the online NVD repository (see Section 4.3 for more details) where each vulnerability identifier is associated with a statement. Statements are replaced through a pre-processing step. The output is an SMT specification (see Section 4.2) that we process with the SMT solver CVC4 [32].³ CVC4 returns either *unsat* or a model satisfying the specification. In the first case, the specification cannot be instantiated. Unsatisfiability means that the scenario is either contradictory or exceeds the quota of the tenant.⁴ In both cases, the tenant has useful feedback for refining her specification.

Instead, if a model is generated it is used to feed the infrastructure instantiation process (translator). The translator converts the definitions of the model into corresponding entries of the Packer and Terraform scripts (see Section 5). These scripts are the input for infrastructure initialization and updating components. The packer script is only executed once before the initialization to create the OS images needed for the scenario. Instead, the Terraform scripts are passed to a scheduler that executes them at the right time. Every script (except for the first one which initializes the infrastructure) causes a modification of the infrastructure, e.g., by disconnecting or adding a node. The scheduler launches the scripts following a precise timeline (see Section 5).

4 Virtual Scenario Description Language

In this section, we introduce VSDL. Due to its rich syntax, we only present part of it and we provide the basic intuition through the application to our working example.

³Since SMT specification language is standard, CVC4 can be replaced or put in parallel with any other solver.

⁴The two cases can be easily disambiguated by checking the satisfiability of the scenario without quota constraints.

Table 1: Excerpt of the VSDL syntax.

S ::=	scenario Id TTU { E* }
TTU ::=	ϵ duration TI
TI ::=	Nat m Nat h
E ::=	NODE NET
NODE ::=	node Id { GNO* }
GNO ::=	[UN] -> NO NO
NO ::=	NOA not NO NO and NO NO or NO (NO)
NOA ::=	type is (compute ... same as Id) flavour is (mobile ... same as Id) cpu is (equal to Freq ... same as Id) disk is (equal to Size ... same as Id) OS is (Id same as Id ...) mounts software Id exists user Id user Id can (read read exec) Path contains (file directory) Path suffers from Vuln ...
NET ::=	network Id { GNE* }
GNE ::=	[UN] -> NE NE
NE ::=	NEA not NE NE and NE NE or NE (NE)
NEA ::=	bandwidth is (equal to BW ... same as Id) gateway has direct access to the Internet addresses range from Addr to Addr firewall blocks (port Nat IP Addr) firewall forwards (port Nat to Nat ↔ IP Addr to Addr) node Id (is connected has IP Addr) ...
UN ::=	UNA not UN UN and UN UN or UN (UN)
UNA ::=	switch (on off) at Id . TExp

4.1 VSDL syntax

A VSDL specification describes a scenario in terms of its core elements, i.e., *nodes* and *networks*. Both of them are defined through a group of statements. Statements cover a plethora of aspects, e.g., connectivity, firewall rules, and hardware profiles, and they can be composed through standard logic connectives. Also, a statement can lay under a temporal guard establishing when, during the scenario, it must hold.

The syntax of VSDL is given in Table 1, where—for the sake of simplicity—we omit some statements and we focus on the most illustrative ones. A scenario S has a name (Id), a duration (TTU) and a sequence of elements E .⁵ The duration can be unspecified (ϵ) or equal to a given interval (TI) in hours or minutes. An element E can be either a node (NODE) or a network (NET). Both nodes and networks have a unique identifier and a sequence of guarded node statements (GNO and GNE). The GNO and GNE statements can either have a guard UN or not (unguarded statements NO and NE). Guards can be atomic (UNA) or obtained by applying the standard boolean connectives, i.e., negation, conjunction

⁵We use a semicolon to separate the terms of a sequence.

and disjunction. An atomic guard `switch on at $t.P(t)$` says that the guarded statement becomes true at time t . Also, t must satisfy a predicate P defined through a TExp expression, i.e., a boolean expression on time intervals possibly including other (previously declared) time variables t', t'', \dots . The guard `switch off at $t.P(t)$` works symmetrically by stating that the guarded statement becomes false at t .

Unguarded statements NO and NE can be atomic statements (NOA and NEA, respectively) or composed ones, i.e., by means of the logical connectives. Atomic statements for nodes and networks are different in order to capture the respective peculiarities. Node statements are mostly self-explanatory and include type (compute vs. storage), hardware (CPU speed, disk size, etc.), OS and installed software, users and privileges (read, write, execute), and content of the file system (files and directories). The only statement that requires more attention is `suffers from` and we will discuss it in Section 4.3. Network statements include bandwidth, access to the public network, the range of addresses that can be assigned to connected nodes, firewall rules (e.g., port forwarding and address filtering), and network participants. The following examples illustrate the use of the VSLD.

Example 2. Consider the following node blocks.

```
node Phone {
  flavour is mobile;
  not (disk is larger than 8 GB);
  not (cpu is faster than 2 GHz);
  (OS is Android-21) or (OS is Android-19);
}

node ApacheS {
  flavour is server;
  disk is larger than 200 GB;
  cpu is faster than 8 GHz;
  OS is Debian-8;
  mounts software apache2;
  mounts software php5;
  mounts software dvwa-setup.sh;
}
```

Briefly, it contains the statements for nodes “Phone” and “ApacheS”. The first node represents a smartphone in the scenario and must have an adequate hardware profile, i.e., mobile flavor. Also, specific hardware constraints can be specified. For instance, here we force Phone to have 8 GB of disk space at most. Moreover, we want Phone to have a CPU slower than 2 GHz. For what concerns the software running on Phone, the requirement is that it mounts Android version 5.0 (API level 21) or 4.4 (API level 19).

ApacheS represents a server hosting an Apache/PHP web application. Thus, we require the node to have a server flavor; with more than 200 GB of disk and a CPU faster than 8 GHz. Also, the OS of the server is Debian Linux version 8. Moreover, we force the server to install three pieces of software: Apache 2.x HTTP server, PHP5, and `dvwa-setup.sh`. The last one is a script setting up the Apache server with DVWA [26].

Example 3. Consider the following VSDL fragment.

```
network Laboratory {
  addresses range from 8.8.8.1 to 8.8.8.64;
  node RSLaptop has IP 8.8.8.3;
  [switch off at  $t.t < 40$  m] -> node Phone is connected;
}
```



```

network Main {
  gateway has direct access to the Internet;
  node Laboratory is connected;
  firewall blocks port 22;
  firewall forwards port 80 to 8080;
  firewall blocks IP 8.8.8.1;
}

```

The fragment contains two network elements, i.e., *Laboratory* and *Main*. Nodes connected to *Laboratory* will have IP addresses ranging from 8.8.8.1 to 8.8.8.64. Both *Phone* and *RSLaptop* are connected to *Laboratory*. *RSLaptop* must have the specific address 8.8.8.3. The last statement states that *Phone* will leave the sub-network after 40 minutes.

The *Main* network is connected to the Internet. Also, *Laboratory* is a sub-network of *Main*. The last three statements define the firewall rules: the firewall must block any incoming connection on port 22, redirect messages using port 80 to port 8080 and prevent any communication with address 8.8.8.1.

4.2 VSDL semantics

The semantics of VSDL is given through a translation into a SMT [33] specification. An SMT specification is a sequence of assertions over the values assumed by functions and constants. The domains of functions define the *theory* under which the formula must be satisfiable. VSDL statements refer to several complex data types, e.g., IP addresses, time, node, and network identifiers. We reduce all of them to the domain of positive numbers.⁶ All the relevant aspects of the scenario are encoded through one or more dedicated functions, called *description functions*. Below we list some of the most interesting ones along with their meaning in natural language.

- $\text{node.disk}(t : \mathbb{N}_0, n : \mathbb{N}_0) = s : \mathbb{N}_0$. At time t node (identified by) n has a disk size of s MB.
- $\text{node.cpu}(t : \mathbb{N}_0, n : \mathbb{N}_0) = s : \mathbb{N}_0$. At time t node n has a CPU speed of s MHz.
- $\text{node.app}(t : \mathbb{N}_0, n : \mathbb{N}_0, s : \mathbb{N}_0) = b : \mathbb{B}$. At time t node n mount software s if and only if $b = \text{true}$.
- $\text{node.user.canr}(t : \mathbb{N}_0, n : \mathbb{N}_0, u : \mathbb{N}_0, r : \mathbb{N}_0) = b : \mathbb{B}$. At time t , on node n , user u can read resource r if and only if $b = \text{true}$.
- $\text{network.gateway.internet}(t : \mathbb{N}_0, n : \mathbb{N}_0) = b : \mathbb{B}$. At time t network n is directly connected to the Internet if and only if $b = \text{true}$.
- $\text{network.node.address}(t : \mathbb{N}_0, n : \mathbb{N}_0, m : \mathbb{N}_0) = a : \mathbb{N}_0$. At time t node m is connected to network n with address a (if $a = 0$ the node m is not connected to n).
- $\text{network.firewall.address.forward}(t : \mathbb{N}_0, n : \mathbb{N}_0, a : \mathbb{N}_0) = b : \mathbb{N}_0$. At time t the firewall of network n forwards incoming packets with destination a to b (if $b = 0$ the packets is blocked).
- $\text{network.firewall.port.forward}(t : \mathbb{N}_0, n : \mathbb{N}_0, p : \mathbb{N}_0) = q : \mathbb{N}_0$. At time t the firewall of network n forwards incoming packets on port p to q (if $q = 0$ the packet is blocked).

Assertions belong to three groups that we describe below.

⁶For time intervals one might prefer to consider real numbers rather than positive ones (with the time unit set to 1 minute). This would require a change of theory, but this does not affect to rest of the dissertation.

Scenario. These constraints consist of a direct translation of the VSDL statements into SMT assertions.

Most of the statements have a straightforward interpretation, e.g., if the specification of node n says that, at time t , the CPU speed is faster than 800 MHz, it will result in `(assert (> (node.cpu t n) 800))`.

Resources. Some constraints are posed by the computational resources available to the physical infrastructure, i.e., the quota of the tenant. Such constraints are obtained by translating the OpenStack quota information (see Section 3) into corresponding assertions. For instance, if the OpenStack quota limits the total CPU frequency to K THz, we must include an assertion stating that the summation of the speed of all the CPUs appearing in the scenario cannot exceed K .

Invariants. This category includes the constraints that must be constantly satisfied by a scenario and, possibly, are implicitly entailed by the specification. For instance, we often want to ensure that all the nodes defined in a scenario are distinct (notice that this might not be always required). Another assumption might be that nodes belonging to the same network do not share their IP address.

Example 4. Consider again the specifications given in Examples 3 and 2. The statements for *Lababory*, *Phone* and *ApacheS* are translated into a specification resembling that given in Table 2.

The meaning of the specification is the following. Each scenario element, i.e., nodes and networks, declared in the VSDL specification is translated into a corresponding constant (lines 2-4). Description functions are constrained to assume certain values depending on the VSDL statements. For instance, the assertions at lines 7-10 encode the statements defining the hardware profile of *Phone* (see Example 2).⁷ Lines 13-28 encode the properties of the *Laboratory* network. In particular, any node n must have either an IP address taken from the defined interval (lines 16-17)⁸ or the constant 0 for “disconnected” (line 19). Intuitively, line 21 states that node *RSLaptop* is connected to a specific address. Lines 22-28, encoding the conditional statement of Example 2, require more attention. First of all, a constant t is declared (line 22). Then t is constrained by the `switch off` guard expression, i.e., $t < 40$ (line 23). Similarly to the case for *RSLaptop*, we translate the statement *node Phone is connected* by asserting that `network.node.address(u, Phone, Laboratory)` must be greater than 0. This assertion is put in the scope of a double implication (lines 26 and 27) such that the assertion, initially true, must be false after t . Lines 31-35 encode the statements for network *Main*. Briefly, *Main* is connected to the Internet (line 30) and *Laboratory* is connected to *Main* (line 32). Also, the firewall of *Main* blocks (i.e., forwards to 0) packets on port 22 (line 33), forwards packets on port 80 to port 8080 (line 34) and blocks packets directed to address 8.8.8.1. Finally, the invariants block states that all the nodes in the specification must be distinct.

4.3 Vulnerability injection

Enabling a vulnerability in existing infrastructures while keeping them realistic is a rather complex task. As discussed in Section 2, nowadays a common approach is to run on a certain node a piece of software, e.g., a web application, where many vulnerabilities can be enabled through a proper configuration. This substantially simplifies the work of the attackers and defenders who only have to discover them.

In the last years, several actors put considerable effort into compiling vulnerability reports and keeping repositories up to date. Among them, the National Vulnerability Database⁹ (NVD) represents a major proposal for the standardization of vulnerability reports. NVD records include a unique id, a textual description, various scores and, more interestingly, a *list of known, vulnerable configurations*. Each

⁷Disk size and CPU speed are given in MB and MHz, respectively.

⁸The IP address $a.b.c.d$ is encoded through the formula $d + 2^8c + 2^{16}b + 2^{24}a$.

⁹<https://nvd.nist.gov/>

Table 2: Fragment of the SMT specification obtained from Examples 2 and 3.

```

1 ; Scenario elements
2 (declare-fun Phone () Int)
3 (declare-fun ApacheS () Int)
4 (declare-fun Laboratory () Int)
5 ; ...
6 ; Hardware constraints: Phone
7 (assert (forall ((u Int)) (and (< (node.cpu u Phone) 16192) (< (node.disk u Phone) 32768))))
8 (assert (forall ((u Int)) (and (>= (node.cpu u Phone) 512) (>= (node.disk u Phone) 2048))))
9 (assert (forall ((u Int)) (< (node.disk u Phone) 8192)))
10 (assert (forall ((u Int)) (< (node.cpu u Phone) 2048)))
11 ; ...
12 ; Network constraints: Laboratory
13 (assert (forall ((u Int) (n Int))
14   (or
15     (and
16       (<= (network.node.address u n Laboratory) 134744065)
17       (>= (network.node.address u n Laboratory) 134744128)
18     )
19     (= (network.node.address u n Laboratory) 0)
20   )))
21 (assert (forall ((u Int)) (= (network.node.address u RSLaptop Laboratory) 134744067)))
22 (declare-fun t () Int)
23 (assert (< t 40))
24 (assert (forall ((u int))
25   (and
26     (=> (<= u t) (> (network.node.address u Phone Laboratory) 0))
27     (=> (> u t) (not (> (network.node.address u Phone Laboratory) 0)))
28   )))
29 ; ...
30 ; Network constraints: Main
31 (assert (forall ((u Int)) (network.gateway.internet u Main)))
32 (assert (forall ((u Int)) (> (network.node.address u Laboratory Main) 0)))
33 (assert (forall ((u Int)) (= (network.firewall.port.forward u Main 22) 0)))
34 (assert (forall ((u Int)) (= (network.firewall.port.forward u Main 80) 8080)))
35 (assert (forall ((u Int)) (= (network.firewall.address.forward u Main 134744065) 0)))
36 ; ...
37 ; Invariants
38 (assert (not (= Phone ApacheS)))
39 (assert (not (= Phone Laboratory)))
40 (assert (not (= ApacheS Laboratory)))
41 ; ...

```

configuration consists of a (often) simple formula, i.e., a disjunction of affected components in CPE (Common Platform Enumeration) format.

Intuitively, a CPE is a unique identifier for hardware/software configuration. The basic scheme of a CPE is

$$\text{cpe}://\{\text{prt}\}:\{\text{vnd}\}:\{\text{prd}\}:\{\text{ver}\}:\{\text{upd}\}:\{\text{edt}\}:\{\text{lan}\}$$

where

- {prt} is a single character indicating whether the CPE refers to a class of applications (a), operating systems (o) or hardware (h);
- {vnd} is the vendor of the product;
- {prd} is the product name;
- {ver} is the version number of the product;
- {upd} is the product update identifier;
- {edt} is the product edition name, and;
- {lan} is the product language tag.

Also, notice that some fields are optional (meaning “any value” when omitted) and few wildcards are allowed.

Example 5. *The NVD record CVE-2015-0235¹⁰ reports a heap-based buffer overflow vulnerabilities that permits the remote execution of arbitrary, unauthorized instructions. The record indicates the following two vulnerable configurations.*

```
<vuln:vulnerable-configuration>
  <cpe-lang:logical-test operator="OR" negate="false">
    <cpe-lang:fact-ref name="cpe:/a:oracle:communications:13.1"/>
    <cpe-lang:fact-ref name="cpe:/a:oracle:pillar_axiom:6.1"/>
    <cpe-lang:fact-ref name="cpe:/a:oracle:pillar_axiom:6.2"/>
    <cpe-lang:fact-ref name="cpe:/a:oracle:pillar_axiom:6.3"/>
  </cpe-lang:logical-test>
</vuln:vulnerable-configuration>
<vuln:vulnerable-configuration>
  <cpe-lang:logical-test operator="OR" negate="false">
    <cpe-lang:fact-ref name="cpe:/a:gnu:glibc:2.0"/>
    ...
    <cpe-lang:fact-ref name="cpe:/a:gnu:glibc:2.17"/>
  </cpe-lang:logical-test>
</vuln:vulnerable-configuration>
```

The first configuration happens when some specific Oracle software is installed. In particular, three versions of Pillar Axiom [34], i.e., from 3.1 to 3.3, suffer from the vulnerability. Moreover, a second family of software modules enables CVE-2015-0235. Such a family includes several versions of the GNU C library `glibc`.

¹⁰<https://web.nvd.nist.gov/view/vuln/detail?vulnId=CVE-2015-0235>

The translation from the NVD vulnerable configuration to VSDL is quite straightforward and it is obtained through the statement `suffers from`. The statement contains a reference to a vulnerability identifier taken from the NVD and its meaning is that the corresponding node must host one of the vulnerable configurations described above. In practice, it is a shorthand for a composition of `OS is` and `mounts software` statements corresponding to the structure of the NVD vulnerable configuration.

Example 6. Consider again the vulnerability of Example 5 and the following VSDL node specification.

```
node N {
  suffers from "CVE-2015-0235";
}
```

It is equivalent to the following one.

```
node N {
  (
    mounts software communications-13.1
    or mounts software pillar_axiom-6.1
    or mounts software pillar_axiom-6.2
    or mounts software pillar_axiom-6.3
  )
  or
  (
    mounts software glibc-2.0
    ...
    or mounts software glibc-2.1
  );
}
```

5 Automatic Generation of Scenarios

The output of the SMT solvers follows a standard syntax [8]. In particular, a model consists of a finite sequence of function definitions given through `define-fun` statements. Under our assumptions, models contain (i) a list of node identifiers, i.e., a constant for each node and network, (ii) a list of time switches, i.e., constants identifying the instants at which the scenario changes its state, and (iii) a definition for each of the description functions introduced in Section 4.2.

It is important to notice that SMT solvers tend to generate minimal models, i.e., those that satisfy the input specification by assigning smaller values to the variables. In principle, this is a desirable property as it guarantees that our approach generates compact scenarios using a minimal amount of computational resources to satisfy the given constraints. Pragmatically, scenario designers might find the output model simplistic. In this respect, the model provides useful information that can be used to refine the original specification. If the model obtained from scenario S is satisfactory, terraform and packer scripts are generated as follows.

1. For each time switch t_i (including $t_0 = 0$) a terraform script called " $S.t_i.tf$ " is created. All the scripts are initialized with OpenStack access details, including user and tenant name, password, and authorization URL (necessary for establishing a valid session).
2. For each node identifier n (i) a corresponding terraform node resource is added to each script " $S.t_i.tf$ " and (ii) a packer JSON script $n.json$ is created. The latter defines the OS image to be used for the initialization of node n . The former contains the `image_name` of the packer and the `name` field.

Table 3: Excerpt of the model generated from the specification of Example 4.

```

1 (model
2 (define-fun Phone () Int 1)
3 (define-fun ApacheS () Int 2)
4 (define-fun RSLaptop () Int 3)
5 (define-fun Laboratory () Int 4)
6 (define-fun Main () Int 5)
7 ;...
8 (define-fun t () Int 1)
9 ;...
10 (define-fun node.cpu ((p1 Int) (p2 Int)) Int (ite (= p2 1) 512 (ite (= p2 2) 8193 (ite ...)))
11 (define-fun node.disk ((p1 Int) (p2 Int)) Int (ite (= p2 1) 2048 (ite (= p2 2) 204801 (ite ...)))
12 (define-fun network.gateway.internet ((p1 Int) (p2 Int)) Bool (ite (= p2 4) false (ite (= p2 5) true (ite ...)))
13 ;...
14 )

```

3. For each network identifier m a corresponding terraform router resource, together with a suitable interface and port, is added to each script " $S_{t_i} . \tau f$ ".
4. For each description function f a group of corresponding terraform commands is added to resource k in script " $S_{t_i} . \tau f$ ". Such commands depend on the value that f assumes on t_i and k .

Below we show the outcome of the procedure previously described when applied to our working example.

Example 7. Consider the SMT specification of Example 4. The output generated by an SMT solver invoked over it is similar to that given in Table 3. It consists of a list of functions and constants definitions. Constants for nodes and networks are assigned to distinct identifiers, i.e., positive numbers (lines 2-6), while time constants are mapped to specific, possibly overlapping minutes of the scenario duration (line 8).¹¹

Functions (lines 10-12) are slightly more complex. They consist of a finite composition of conditional statements (if-then-else, ite) testing the value of (some of) the formal parameters of a function to decide the result. For instance, *node.cpu* (line 10) is assigned to the partial function $f : \mathbb{N}_0 \times \mathbb{N}_0 \rightarrow \mathbb{N}_0$ defined as follows.

$$f = \lambda p1, p2. \begin{cases} 512 & \text{if } p2 = 1 \\ 8193 & \text{if } p2 = 2 \\ \vdots & \\ \perp & \text{otherwise} \end{cases}$$

Assuming no other time constants exist, the model given above results in two terraform scripts " $S_{t_0} . \tau f$ " and " $S_{t_1} . \tau f$ ". The content of the script " $S_{t_0} . \tau f$ " includes the fragment shown in Table 4.

Briefly, it contains the declarations of OpenStack resources for the initialization of the infrastructure. Among them, routers (lines 2-9) are labeled with "openstack_networking_router_v2". Simply, they include a name attribute identifying them. Moreover, the router *Main* has a reference to a (predefined) gateway for accessing the Internet (line 4). The script binds routers with sub-networks (lines 11-15) through a specific interface (lines 17-20). Sub-networks also include an address mask in CIDR notation.

¹¹Notice that t can also be legally assigned to 0, which implies (see Examples 3 and 1) that the Phone node is never connected to (or immediately disconnected from) the Laboratory network. If this is not the intention of the designer, it means that the scenario is underspecified.

Table 4: Excerpt of the terraform script "S_0.tf".

```

1 # ...
2 resource "openstack_networking_router_v2" "main" {
3   name = "Main"
4   external_gateway="b998c866-f909-48a3-a5d6-7837fe91354d"
5 }
6
7 resource "openstack_networking_router_v2" "laboratory" {
8   name = "Laboratory"
9 }
10 # ...
11 resource "openstack_networking_subnet_v2" "laboratory" {
12   name = "Laboratory"
13   network_id = "${openstack_networking_network_v2.laboratory.id}"
14   cidr = "8.8.8.1/26"
15 }
16 # ...
17 resource "openstack_networking_router_interface_v2" "laboratory_router" {
18   router_id = "${openstack_networking_router_v2.main.id}"
19   subnet_id = "${openstack_networking_subnet_v2.laboratory.id}"
20 }
21 # ...
22 resource "openstack_networking_port_v2" "phone_laboratory" {
23   network_id = "${openstack_networking_network_v2.laboratory.id}"
24   fixed_ip {
25     subnet_id = "${openstack_networking_subnet_v2.laboratory.id}"
26   }
27 }
28
29 resource "openstack_compute_instance_v2" "phone" {
30   name = "Phone"
31   image_name = "android-4.4-x86_64"
32   flavour_name = "mobile.phone"
33   network {
34     port = "${openstack_networking_port_v2.phone_laboratory.id}"
35   }
36 }
37 # ...

```

Finally, computational nodes are labeled with “`openstack_compute_instance_v2`” (lines 29-36) and connected to a network through a port (lines 22-27). The attributes of a node define its name, the OS image to be installed on it, and its hardware profile.¹² Once submitted, the script results in the virtual infrastructure appearing in Figure 3. Also, as shown in Figure 4, DVWA is actually running on the appointed node. Exploiting a command injection vulnerability, for instance, an attacker can run `nmmap` to check which IP address in the local sub-network correspond to active hosts.

¹²Currently, Terraform does not support detailed hardware description, e.g., CPU speed. We avoid this issue by dynamically customizing the OpenStack flavours, e.g., “mobile.phone”.

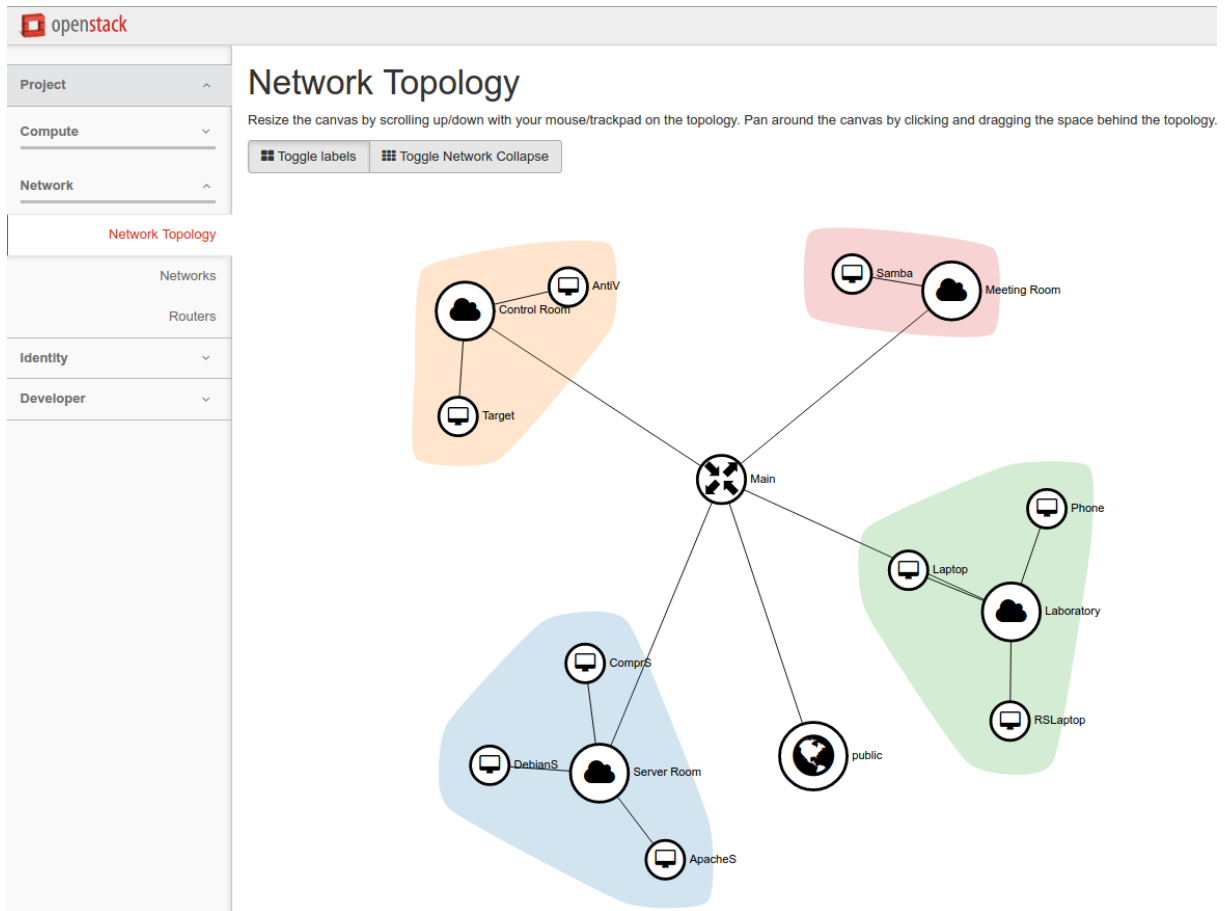


Figure 3: The virtual infrastructure as displayed by the OpenStack dashboard.

6 Discussion

In this section, we critically discuss our proposal’s main novelties and limitations.

Verifiability. As we highlighted in the previous sections, SMT-based verification allows one to formally check the correctness of the VSDL scenarios. Nevertheless, other approaches are viable and one might want to consider other languages. As a matter of fact, existing SMT theories mostly deal with domains that are not of great interest for developing cyber range scenarios, e.g., think of real numbers.

Expressiveness. As for the previous case, the high expressiveness of SMT may not always be desirable. Indeed, less expressive languages could provide better performances and more scalability in large scenarios. For instance, in [35] Datalog is used in place of SMT.

Compositionality. Our high-level DSL supports and simplifies the creation of reusable components and configurations. Moreover, our verification process allows designers to quickly and automatically validate scenarios so avoiding composition errors. This substantially reduces the effort for the creation of training environments.

Integration. By relying on a mainstream IaaS provider, our scenarios can be integrated into many con-

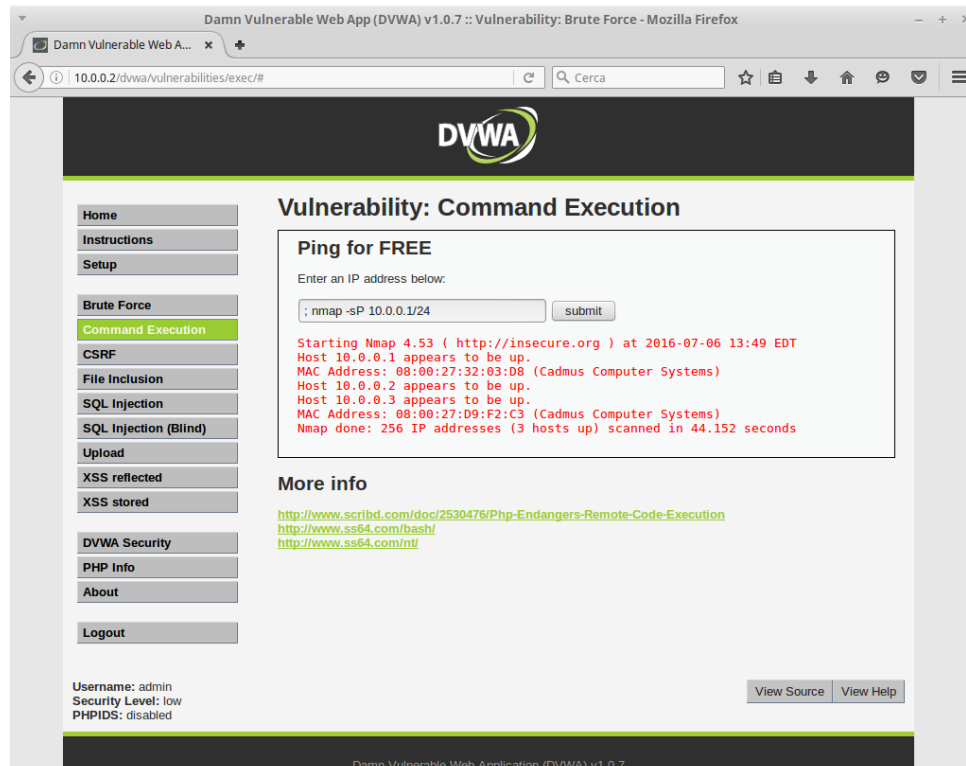


Figure 4: The working instance of DVWA exposing a command injection vulnerability.

texts and they can take advantage of existing infrastructures. In general, one might want to achieve integration with other types of infrastructures, e.g., made of physical devices. For the time being this is not supported by our approach, but technologies such as Software-Defined Networking [36] could be considered for this aim.

7 Conclusion

We presented a framework for the definition, validation, and generation of virtual scenarios, being at the very core of every cyber range. To the best of our knowledge, this is the first proposal for such a framework. Our approach offers several desirable features in terms of verifiability and maintainability. Moreover, we developed and integrated it with state-of-the-art technologies. Last but not least, we plan to apply our framework to the forthcoming Italian national cyber range.

This is only the first step toward a fully automated system and many aspects still need to be considered and investigated. We schematically report those that are the more challenging and relevant. We account for all of them as future work.

Traffic simulation. Although an infrastructure can be detailed from an architectural point of view, the scenarios based on it might lack of realism in terms of network traffic. Traffic generators exist, e.g., *Ostinato*¹³, but they need to be configured to correctly simulate the real behavior of the infrastructure. This will also require extending the syntax of VSDL with statements for describing the network activity of a node.

¹³<http://ostinato.org/>

Infrastructure inference. For the time being, virtual scenarios are designed by experts to resemble a real infrastructure. Obtaining a VSDL model from monitoring/observing an actual infrastructure would lead to an easier and more realistic modeling process. For instance, the output of tools like *EtherApe*¹⁴ and *Xplico*¹⁵ could be (partially) translated to VSDL.

Attack trees. Vulnerabilities play a central role, in particular, in the training sessions. Although we can effectively inject vulnerabilities process, forcing their exploitation through predefined steps following a didactic purpose is not trivial. We plan to extend our framework with predefined sets of attack trees that a scenario designer can include in her VSDL specification. Since attack trees include sub-goals, we must make sure that the scenario permits the exploration of the tree and includes at least one attack pattern.

Infrastructure fuzzing. As discussed in Section 5, our framework generates a minimal infrastructure satisfying the specification. Nevertheless, it is a common practice to have rich scenarios with myriads of nodes having no specific/active roles in the attack/defense process. To this aim, we aim at including fuzzing methodologies that can add complexity to a scenario without compromising its key features.

References

- [1] L. Demetrio, G. Lagorio, M. Ribauda, E. Russo, and A. Valenza. ZenHackAdemy: Ethical hacking @ DIBRIS. In *Proc. of the 11th International Conference on Computer Supported Education (CSEDU'19), Heraklion, Crete, Greece*, pages 405–413. SciTePress, January 2019.
- [2] S. Omar, A. Mohammed, and E. Mohsine. Openstack: toward an open-source solution for cloud computing. *International Journal of Computer Applications*, 55(3):38–42, October 2012.
- [3] Michigan Educational Research Information Triad. Cyber Education Training, 2022. <https://www.merit.edu/security/training/> [Online; Accessed on December 15, 2022].
- [4] R.H. Wagner. *Designing a network defense scenario using the open cyber challenge platform*. University of Rhode Island, 2013.
- [5] F.G.J. Absil, A.D. Dijk, and J.M.G. Meulendijks. Lessons learned from NATO's cyber defence exercise locked shields 2015. *Militaire Spectator*, 185(2):65–74, 2016.
- [6] D. Geer, R. Bace, P. Gutmann, P. Metzger, C.P. Pflieger, J.S. Quarterman, and B. Schneier. *CyberInsecurity: The Cost of Monopoly*. Technical report, Computer & Communications Industry Association, September 2003.
- [7] A. Van Deursen, P. Klint, and J. Visser. Domain-specific languages: An annotated bibliography. *ACM Sigplan Notices*, 35(6):26–36, June 2000.
- [8] C. Barrett, P. Fontaine, and C. Tinelli. The SMT-LIB Standard: Version 2.5. Technical report, Department of Computer Science, The University of Iowa, 2015.
- [9] J. Davis and S. Magrath. A Survey of Cyber Ranges and Testbeds. Technical Report DSTO-GD-0771, Defence Science and Technology Organization (DSTO), Australian Department of Defence, October 2013.
- [10] AZCWR. Arizona Cyber Warfare Range, 2016. <http://azcwr.org/> [Online; Accessed on December 15, 2022].
- [11] T. Benzel. The Science of Cyber Security Experimentation: The DETER Project. In *Proc. of the 27th Annual Computer Security Applications Conference (ACSAC'11), Orlando, Florida, USA*, pages 137–148. ACM, December 2011.
- [12] B. Ferguson, A. Tall, and D. Olsen. National Cyber Range Overview. In *Proc. of the 2014 IEEE Military Communications Conference (MILCOM'14), Baltimore, MD, USA*, pages 123–128. IEEE, October 2014.

¹⁴<http://etherape.sourceforge.net/>

¹⁵<http://www.xplico.org/>

- [13] M. Rosenstein and F. Corvese. A Secure Architecture for the Range-Level Command and Control System of a National Cyber Range Testbed. In *Proc. of the 5th Workshop on Cyber Security Experimentation and Test (CSET'12)*, pages 1–9. USENIX, August 2012.
- [14] IBM. Softlayer Web Page, 2022. <http://www.softlayer.com/> [Online; Accessed on December 15, 2022].
- [15] VMWare. vCloud Web Page, 2022. <http://vcloud.vmware.com/> [Online; Accessed on December 15, 2022].
- [16] M.H. Syed, E.B. Fernandez, and M. Ilyas. A pattern for fog computing. In *Proc. of the 10th Travelling Conference on Pattern Languages of Programs (VikingPLoP'16)*, Leerdam, AA, Netherlands, pages 1–10. ACM, April 2016.
- [17] M. Ghijsen, J. van der Ham, P. Grosso, and C. de Laat. Towards an Infrastructure Description Language for Modeling Computing Infrastructures. In *Proc. of the 10th IEEE International Symposium on Parallel and Distributed Processing with Applications (PDPA'12)*, Leganes, Spain, pages 207–214. IEEE, July 2012.
- [18] G. Piegas G.P. Koslovski, P.V. Primet, and A.S. Charao. VXDL: Virtual resources and interconnection networks description language. In *Proc. of the 2nd International Conference on Networks for Grid Applications (GridNets'08)*, Beijing, China, volume 2 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, pages 138–154. Springer, October 2008.
- [19] N. McKeownand, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner. OpenFlow: Enabling Innovation in Campus Networks. *SIGCOMM Computer Communication Review*, 38(2):69–74, March 2008.
- [20] D. Kreutz, F. M. V. Ramos, P. E. Veríssimo, C. E. Rothenberg, S. Azodolmolky, and S. Uhlig. Software-Defined Networking: A Comprehensive Survey. *Proceedings of the IEEE*, 103(1):14–76, January 2015.
- [21] M. Bartoletti, P. Degano, and G.L. Ferrari. Types and Effects for Secure Service Orchestration. In *Proc. of the 19th IEEE Computer Security Foundations Workshop (CSFW'06)*, Venice, Italy, pages 57–69. IEEE, July 2006.
- [22] G. Costa, P. Degano, and F. Martinelli. Modular plans for secure service composition. *Journal of Computer Security*, 20(1):81–117, January 2012.
- [23] F. Montesi, C. Guidi, R. Lucchi, and G. Zavattaro. JOLIE: a Java Orchestration Language Interpreter Engine. *Electronic Notes in Theoretical Computer Science*, 181:19–33, 2007.
- [24] G. Castagna, N. Gesbert, and L. Padovani. A Theory of Contracts for Web Services. In *Proc. of the 35th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages (POPL'08)*, San Francisco, California, USA, pages 261–272. ACM, January 2008.
- [25] M.A. Barbosa and L.S. Barbosa. A Perspective on Service Orchestration. *Science of Computer Programming*, 74(9):671–687, 2009.
- [26] DVWA team. Damn Vulnerable Web Application (DVWA), 2022. <https://github.com/digininja/DVWA> [Online; Accessed on December 15, 2022].
- [27] WebGoat Team. WebGoat Project. Technical report, The Open Web Application Security Project, 2022.
- [28] B. Leban, M. Bendre, and P. Tabriz. Gruyere: Web Application Exploits and Defenses. Technical report, Google, 2022.
- [29] A. Corradi, M. Fanelli, and L. Foschini. Vm consolidation: A real case based on openstack cloud. *Future Generation Computer Systems*, 32:118–127, March 2014.
- [30] HashiCorp. Terraform Online Documentation, 2022. <https://www.terraform.io/docs/index.html> [Online; Accessed on December 15, 2022].
- [31] HashiCorp. Packer Online Documentation, 2022. <https://www.packer.io/docs/> [Online; Accessed on December 15, 2022].
- [32] C. Barrett, C.L. Conway, M. Deters, L. Hadarean, D. Jovanović, T. King, A. Reynolds, and C. Tinelli. CVC4. In *Proc. of the 23rd International Conference on Computer Aided Verification (CAV'11)*, Snowbird, UT, USA, volume 6806 of *Lecture Notes in Computer Science*, pages 171–177. Springer, July 2011.
- [33] C. Barrett, R. Sebastiani, S. Seshia, and C. Tinelli. *Satisfiability Modulo Theories*. IOS Press, 2009.
- [34] Oracle Corporation. Pillar Axiom Path Manager 3.1. Technical report, Oracle Corporation, March 2012.

- [35] E. Russo, G. Costa, and A. Armando. Building next generation Cyber Ranges with CRACK. *Computers & Security*, 95:101837, August 2020.
- [36] D. Kreutz, F.M. V.Ramos, P.E. Verissimo, C.E. Rothenberg, S. Azodolmolky, and S. Uhlig. Software-defined networking: A comprehensive survey. *Proceedings of the IEEE*, 103(1):14–76, January 2015.
-

Author Biography



Gabriele Costa is Associate Professor at the SySMA Group of the IMT School for Advanced Studies. He received his M.Sc. in Computer Science in 2007 and his Ph.D. in Computer Science in 2011. He was a member of the cybersecurity group of the Istituto di Informatica e Telematica (IIT) of the CNR. His appointments include a period as visiting researcher at ETH Zurich in 2016-2017. He has been co-founder of the Computer Security Laboratory (CSec) at DIBRIS. He is co-founder and CRO of Talos, a spin-off of DIBRIS focused on Cybersecurity. His focus is on studying and applying formal methods for the automatic verification and testing of mobile and modular systems.



Enrico Russo is Assistant Professor in Computer Engineering and received his M.Sc. in Computer Science in 2001 and his Ph.D. in Computer Science in 2021. His work is focused on Cyber Range systems with particular emphasis on investigating techniques for simplifying the generation of the training environments and for automatizing the tasks of personnel involved in exercises. His research interests also include maritime cyber security. In 2019, he co-founded ZenHack, the Capture the Flag team of the University of Genova, and is personally involved in different live-fire cyber exercises as a Green/Blue team member.



Alessandro Armando is Full Professor in Computer Engineering and received his Ph.D. in Computer Engineering at the University of Genoa. His appointments include a position as research fellow at the University of Edinburgh and one at INRIA-Lorraine (France). He is professor at the University of Genoa where he teaches Computer Security. He has contributed to founding and directing the Master in Cybersecurity and Data Protection of the University of Genoa and the Master in Digital Forensic and Cyber Technologies at the School of Telecommunications of the Defence. He is currently serving as director of the Ph.D. Program in Security, Risk and Vulnerability of the University of Genoa. He founded and led the Security & Trust Research Unit of the Bruno Kessler Foundation in Trento. He has been coordinator and/or team leader in several national and EU research projects, including a European Industrial Doctorate in partnership with SAP. He contributed to the discovery of authentication flaws in the Single Sign-On standards and implementations, including a serious man-in-the-middle attack on the SAML-based SSO for Google Apps. He is currently serving as deputy director of the Cybersecurity National Laboratory of the National Inter-university Consortium on Informatics (CINI).