

A reproducible methodology for investigating network attacks on corporate infrastructure using an open-source cyber range

Maksym Delembovskyi^{1,*}

¹Kyiv National University of Construction and Architecture, Air Force Ave., 31, Kyiv, 03037, Ukraine

Abstract

The growing number and sophistication of network attacks on corporate infrastructure, together with a global shortage of security professionals, make realistic hands-on training and reproducible research a strategic priority. Cyber ranges provide isolated, instrumented environments for this purpose, yet state-grade and commercial platforms are financially inaccessible to most universities, while open alternatives often lack an integrated monitoring stack or depend on heavyweight cloud infrastructure. This paper presents the design, implementation and evaluation of a fully open-source cyber range dedicated to the study of network attacks on corporate infrastructure. From an analysis of the contemporary threat landscape and a comparative study of existing platforms, we derive explicit functional and non-functional requirements and propose a six-layer architecture with a four-zone segmented network that reproduces a realistic small-to-medium enterprise. The environment is provisioned entirely through infrastructure-as-code using Proxmox VE, Terraform and Ansible, and is equipped with pfSense, Suricata and the Elastic Stack and a catalogue of six attack scenarios mapped to MITRE ATT&CK. Experimental evaluation shows that the range provisions in five to ten minutes, operates at 62% CPU and 84% RAM utilisation with fifteen concurrent virtual machines, sustains 99.86% availability over a 72-hour test, and detects attacks with an aggregate F1-score of 97.3%; an integral efficiency assessment yields a score of 0.87. The results confirm that an accessible, reproducible and effective cyber range can be built entirely from open-source components.

Keywords

cyber range, network attacks, corporate infrastructure, intrusion detection, virtualization, cybersecurity, penetration testing, MITRE ATT&CK

1. Introduction

Over the past decade, corporate computer networks have become the principal target of an increasingly professionalised cyber-threat ecosystem. According to Check Point Research, organisations worldwide faced an average of more than 1,150 cyberattacks per week during 2023, and this figure continued to grow by double digits quarter over quarter into 2024 [1]. The financial consequences are equally severe: the IBM Cost of a Data Breach Report 2024 places the global average cost of a single breach at USD 4.88 million, the highest value ever recorded [2]. At the same time, the ISC2 Cybersecurity Workforce Study estimates a global shortage of roughly four million security professionals [3], which makes the practical, hands-on preparation of specialists a strategic priority rather than a purely academic concern.

The Ukrainian context illustrates the stakes with particular clarity. In December 2015, a coordinated intrusion against three regional electricity distribution companies left approximately 230,000 customers without power and became the first publicly confirmed cyberattack to cause a blackout [4]. One year later, the Industroyer/CrashOverride malware targeted a transmission substation near Kyiv, demonstrating purpose-built capabilities for manipulating industrial protocols [5]. In 2017, the NotPetya wiper – initially distributed through a compromised Ukrainian accounting software update – propagated globally and caused damages estimated in the billions of dollars [6]. These events, together with

CMiGIN'26: Fifth International Workshop on Conflict Management in Global Information Networks, June 26–27, 2026, Krakow, Poland

*Corresponding author.

✉ maksdel2@gmail.com (M. Delembovskyi)

ORCID 0000-0002-6543-0701 (M. Delembovskyi)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

the 2021 Colonial Pipeline ransomware incident that disrupted fuel supply across the United States east coast [7] and the 2017 Equifax breach that exposed the records of 147 million people through an unpatched Apache Struts vulnerability [8], show that network attacks on corporate and critical infrastructure are neither hypothetical nor rare. The State Service of Special Communications and Information Protection of Ukraine reported thousands of registered incidents in 2023 alone, a substantial share of which targeted organisational network infrastructure [9].

Defending against such threats requires practitioners who can recognise multi-stage intrusions and operate detection and response tooling under realistic conditions. Adversaries increasingly follow the tactics and techniques catalogued in the MITRE ATT&CK framework [10], chaining reconnaissance, initial access, privilege escalation, lateral movement and impact into coherent campaigns. Reproducing this behaviour for training and research cannot be done safely on production systems. Cyber ranges – isolated, instrumented environments that emulate realistic infrastructure – have therefore emerged as the standard instrument for cybersecurity education, exercises and applied research [11, 12].

However, the cyber-range landscape presents a practical dilemma for academic institutions. State-grade and commercial platforms offer high fidelity but are financially inaccessible to most universities, whereas open solutions frequently lack an integrated monitoring stack, depend on heavyweight cloud infrastructure, or do not cover a contemporary corporate technology stack and attack catalogue. This paper addresses that gap by presenting the design, implementation and evaluation of a fully open-source cyber range dedicated to the study of network attacks on corporate infrastructure.

The contributions of this work are fourfold. First, we derive an explicit set of functional and non-functional requirements for an academic cyber range from an analysis of the threat landscape and of existing platforms. Second, we propose a layered architecture with a four-zone segmented network that reproduces a realistic small-to-medium enterprise and is provisioned entirely through infrastructure-as-code. Third, we implement a reproducible environment based on Proxmox VE, Terraform, Ansible, pfSense, Suricata and the Elastic Stack, together with a catalogue of six attack scenarios mapped to MITRE ATT&CK. Fourth, we evaluate the resulting system experimentally with respect to deployment time, resource efficiency, availability, and intrusion-detection effectiveness, and we summarise the outcome through an integral efficiency metric.

2. Related work

Research on cyber ranges and security testbeds has grown rapidly. Ukwandu et al. provide a comprehensive review of cyber ranges and testbeds, segmenting documented platforms by type, underlying technology, threat scenarios and application domain, and observe a diminishing boundary between information-technology and operational-technology ranges [11]. Yamin et al. propose a systematic taxonomy describing the scenarios, functions, tools and architecture of cyber ranges, which has become a common reference point for subsequent designs [12]. The European Cyber Security Organisation summarises the architectural components of a range – user interface, service provider and orchestration – and frames the transition of cyber ranges from hype to operational reality [13]. Earlier surveys by Davis and Magrath catalogued the first generation of military and academic testbeds [14].

Among national and military platforms, the United States National Cyber Range, originally developed under a DARPA programme, supports large-scale exercises and the testing of security technologies with high fidelity [15]. Commercial offerings such as Cyberbit and the IBM Cyber Range provide turnkey security-operations simulations but at significant licensing cost. These platforms set the benchmark for realism yet remain closed and expensive, motivating the search for open alternatives.

The KYPO Cyber Range Platform, developed at Masaryk University, is one of the most mature open-source platforms; it builds on OpenStack and describes scenarios declaratively, supporting the full life cycle from definition through deployment to assessment [16, 17]. Vykopal et al. report lessons learned from running complex hands-on defence exercises on KYPO, highlighting the importance of automation and reproducibility [18]. The AIT Cyber Range emphasises a comprehensive environment for exercises, education and research, integrating realistic traffic generation and scenario orchestration [19]. CyRIS,




Figure 1: Comparative analysis of representative cyber-range platforms across key characteristics (scalability, scenario realism, open-source availability, ease of deployment, ICS/SCADA support, documentation).

the cyber-range instantiation system, automates the creation of training environments from high-level descriptions and directly influenced the infrastructure-as-code approach adopted in later work [20]. Chouliaras et al. survey ranges and testbeds specifically for education, training and research, and identify monitoring integration and scenario realism as recurring weaknesses of lightweight academic deployments [21].

More recent contributions extend the concept in several directions. Karjalainen and Kokkonen describe a comprehensive cyber arena as the next generation of range [22]; Russo et al. enable next-generation ranges with mobile security components [23]; and Smyrlis et al. propose a model-driven assurance platform that couples ranges with security certification [24]. Urias et al. critically examine the limitations and future needs of cyber-range infrastructure, arguing for reproducible experimentation protocols [25]. Importantly, Noponen et al. study the cybersecurity of cyber ranges themselves, cataloguing threats and mitigations and underscoring the necessity of strict isolation [26], while Ziubina et al. demonstrate a hybrid range for testing security [27]. Finally, Yaroviy et al. propose quantitative models for evaluating the efficiency and resilience of an educational cyber range, combining throughput, resource-utilisation, success-probability and resilience indicators into a single integral objective function; we adopt their model for the summative assessment in Section 7 [28].

Across this body of work a consistent pattern emerges: commercial platforms maximise fidelity at the expense of cost and openness, whereas open platforms either require heavyweight cloud infrastructure (KYPO/OpenStack) or provide flexible network emulation without an integrated detection and SIEM stack. The design presented in this paper deliberately targets the middle ground – an entirely open-source, on-premise range with first-class monitoring – suited to the constraints of higher-education institutions.

Table 1

Comparison of Representative Cyber-Range Platforms

Platform	Scalab.	Realism	Open src.	ICS/SCADA
NCR (US DoD)	High	Very high	No	Yes
Cyberbit	High	High	No	Yes
IBM Cyber Range	High	Very high	No	Partial
KYPO CRP	High	High	Yes	Partial
GNS3 / EVE-NG	Medium	High	Yes	Yes
This work	High	High	Yes	Planned

3. Network attacks on corporate infrastructure

3.1. Threat landscape and notable incidents

The corporate network is the connective tissue of modern digital business, hosting directory services, databases, enterprise resource-planning systems and communication platforms. Its centrality also makes it the primary attack surface. Analyses of recent incidents reveal a consistent multi-stage structure aligned with MITRE ATT&CK [10]: reconnaissance and initial access (often through exposed web services or phishing), establishment of persistence, privilege escalation, defence evasion, lateral movement across the internal network, and finally an impact stage such as data exfiltration or ransomware deployment.

The 2017 Equifax breach exemplifies this chain: attackers exploited an unpatched Apache Struts deserialisation vulnerability in a public web application to obtain initial access, then moved laterally to reach databases containing the personal data of 147 million individuals [8]. The 2015–2016 attacks on the Ukrainian power grid combined spear-phishing and the BlackEnergy toolkit with later, protocol-aware malware to manipulate operational systems [4, 5]. NotPetya weaponised a software-update channel and credential-harvesting plus the EternalBlue exploit to achieve rapid internal propagation [6]. These events confirm that effective training environments must reproduce not isolated exploits but complete attack chains across segmented network zones.

3.2. Classification of network attacks

To inform the scenario catalogue, network attacks relevant to corporate infrastructure are systematised along four dimensions: the layer of the OSI model, the execution technique, the adversary objective and the degree of automation. Table 2 summarises seven principal categories that the cyber range is required to support. Volumetric and protocol-level denial-of-service attacks aim to exhaust computational resources or link capacity; man-in-the-middle techniques intercept and manipulate traffic; application-layer attacks such as SQL injection target data and code execution; authentication attacks such as brute force and pass-the-hash seek unauthorised access; and advanced persistent threats combine lateral-movement techniques to spread within the network. Malware, and ransomware in particular, represents the impact stage of many contemporary campaigns.

4. Requirements and design methodology

The requirements for the cyber range were derived from two sources: the analysis of the threat landscape (Section 3) and the comparative study of existing platforms (Section 2). They are expressed as five functional requirements (F1–F5) and five non-functional requirements (N1–N5), which subsequently serve as evaluation criteria in Section 7.

Table 2
Classification of Network Attacks on Corporate Infrastructure

Category	Example techniques	OSI	Objective
DoS / DDoS	SYN/UDP/HTTP flood, amplification	L3–L7	Service unavailability
Network-layer	ARP/IP spoofing, ICMP redirect	L2–L3	Traffic interception
MITM	SSL stripping, DNS spoofing	L3–L7	Data interception
Application	SQL injection, XSS, CSRF	L7	Data / code execution
Authentication	Brute force, pass-the-hash, Kerberoasting	L4–L7	Unauthorised access
APT / Lateral mov.	SMB relay, WMI exec, PsExec	L4–L7	Network propagation
Malware	Ransomware, trojan, rootkit, C2	L3–L7	Encryption / control

4.1. Functional requirements

- F1 reproduce a three-tier network topology (external / DMZ / internal) with realistic services (AD, DNS, HTTP, SMTP, database, SMB).
- F2 support at least five categories of network attack, including DDoS, MITM, SQL injection, brute force and lateral movement / ransomware.
- F3 integrate an IDS/SIEM subsystem with real-time event collection, storage and correlation.
- F4 provide automated, reproducible deployment and reset through infrastructure-as-code.
- F5 document each scenario for repeatability, capturing attack parameters, expected effects and detection results.

4.2. Non-functional requirements

- N1 resource efficiency: deployable on a mid-range server (32 GB RAM, 8-core CPU) or equivalent cloud instance.
- N2 scalability: the architecture must allow additional nodes and scenarios without restructuring.
- N3 openness: rely exclusively on open-source software.
- N4 isolation: operate fully isolated from production networks; no traffic may cross the boundary.
- N5 documentation: every component and scenario must be supported by sufficient technical documentation.

5. Cyber range architecture

5.1. Layered architecture

The architecture of the cyber range is organised into six interdependent layers, shown in Figure 2, following the component model articulated by the European Cyber Security Organisation [13] and the NICE workforce framework [29]. At the base, a physical or cloud host provides compute, storage and networking. Above it, a virtualisation and isolation layer abstracts the hardware and isolates concurrent scenarios. The target-infrastructure layer is the core of the range: it reproduces a realistic corporate network divided into an external (attacker) zone, a demilitarised zone and an internal corporate zone. An orchestration layer automates provisioning and configuration; a monitoring and SIEM layer collects and correlates telemetry; and, at the top, a management interface supports scenario registration, scheduling and reporting.

This layering serves two purposes. It enforces the isolation requirement (N4) by ensuring that no traffic crosses between the range and external production networks, addressing the self-protection concerns raised by Noponen et al. [26]; and it enables reproducibility by separating immutable infrastructure definitions from per-scenario configuration.

Layered architecture of the cyber range for researching network attacks




Figure 2: Six-layer architecture of the proposed cyber range, from the physical/cloud host through virtualization, target infrastructure, orchestration and monitoring to the management interface.

Table 3

IP Addressing Scheme of the Cyber-Range Zones

Zone	Subnet	Gateway	Key hosts
External / Attacker	10.0.1.0/24	10.0.1.1	attacker-01 (Kali Linux)
DMZ	10.0.2.0/24	10.0.2.1	web-srv-01, dns-smtp-01
Internal (LAN)	10.0.3.0/24	10.0.3.1	ad-dc-01, db-srv-01, workstations
Management / SIEM	10.0.0.0/24	10.0.0.1	proxmox, siem-elastic, suricata, ansible

5.2. Network segmentation and addressing

The target network is partitioned into four addressed zones, each governed by an explicit set of pfSense firewall rules. Figure 3 depicts the detailed addressing and segmentation, and Table 3 records the corresponding subnets. The attacker zone (10.0.1.0/24) hosts the offensive workstation; the demilitarised zone (10.0.2.0/24) contains the public-facing web, DNS and mail services; the internal zone (10.0.3.0/24) hosts the Active Directory domain controller, the database server and the end-user workstations; and a dedicated management zone (10.0.0.0/24) carries the hypervisor, the SIEM, the IDS sensor and the automation controller.

The management zone is fully isolated from the target zones and is unreachable from scenario participants. The Suricata sensor receives a mirrored copy of inter-zone traffic through a SPAN port on the firewall, observing all traffic without participating in routing. Inter-zone policy permits traffic from the attacker zone to the DMZ on standard service ports, restricts DMZ-to-internal communication to the minimum required (for example, the web server reaching the database on TCP/3306), and by default blocks attacker-to-internal traffic except where a specific scenario requires it.




Figure 3: Detailed IP addressing and network segmentation into four isolated zones, with the Suricata sensor receiving mirrored inter-zone traffic via a SPAN port on the pfSense firewall.

6. Implementation

6.1. Virtualization platform and orchestration

The range is deployed on a single mid-range host (AMD Ryzen 7 5800X, 64 GB ECC RAM, NVMe storage) running Proxmox Virtual Environment 8.x [30]. Proxmox was selected after a comparative analysis of virtualisation platforms because it combines a fully open-source licence, simultaneous support for KVM virtual machines (required for Windows and complex Linux hosts) and LXC containers, snapshotting, and an official Terraform provider. Provisioning follows an infrastructure-as-code pipeline: Terraform [31] declaratively creates virtual machines and containers through the Proxmox API, and Ansible [32] configures each host according to its role through idempotent playbooks. The complete environment can be provisioned, and reset to a clean state between sessions, in five to ten minutes, satisfying requirement F4 and echoing the automation philosophy of CyRIS [20].

6.2. Target services and monitoring stack

Target services reproduce a representative small-to-medium enterprise. The DMZ hosts an Apache web server running the Damn Vulnerable Web Application [33], a BIND9 DNS server and a Postfix/Dovecot mail gateway, all on Ubuntu Server 22.04 LTS. The internal zone hosts an Active Directory domain controller (Windows Server 2022 or Samba AD), a MySQL 8.0 database with a deliberately vulnerable schema, a Samba file server and Windows workstations. The offensive workstation runs Kali Linux with Metasploit, Nmap, Hydra, sqlmap and hping3 [34, 35].

The monitoring stack combines a network IDS with a SIEM. Suricata 7.x [36] operates in passive mode on the management sensor, applying both the Emerging Threats ruleset and a set of custom rules written specifically for the range’s scenarios; its EVE JSON output is shipped to the Elastic Stack [37]. Elasticsearch stores and indexes security events, Logstash parses and normalises them to the Elastic Common Schema, and Kibana provides dashboards and detection rules. The pfSense

Attack scenario matrix of the cyber range (MITRE ATT&CK alignment)

Attack scenario	MITRE tactic	Target host	Attacker tool	IDS detection	Difficulty
DDoS (SYN/UDP/HTTP)	TA0040 Impact (T1498)	web-srv-01 DMZ	hping3 Scapy	Yes (Suricata)	●●○○○
ARP spoofing / MITM	TA0006 Credential Access (T1557)	web-srv-01 ws-01	arp spoof Bettercap	Yes (Suricata)	●●●○○
SQL injection (DWA)	TA0009 Collection (T1190)	web-srv-01 MySQL	sqlmap Burp Suite	Yes (ModSec)	●●○○○
Brute force / credential	TA0006 Cred.Access (T1110)	ad-dc-01 SSH servers	Hydra Medusa	Yes (Suricata)	●●●○○
Lateral movement (PtH)	TA0008 Lateral Mov. (T1550)	ad-dc-01 db-srv-01	Impacket PsExec	Partial (ELK)	●●●●○
Ransomware simulation	TA0040 Impact (T1486)	ws-01, ws-02 db-srv-01	Custom script + OpenSSL	Yes (Wazuh)	●●●●●

● = difficulty level (1-5)

Figure 4: Matrix of the six implemented attack scenarios with their MITRE ATT&CK alignment, target host, offensive tooling, detection mechanism and relative difficulty.

Table 4
Attack Scenarios and Their Mapping to MITRE ATT&CK Tactics

Scenario	ATT&CK tactic	Tooling	Detection
DDoS (SYN/UDP/HTTP)	Impact (T1498)	hping3	Suricata
ARP spoofing / MITM	Credential Access (T1557)	Bettercap	Suricata
SQL injection	Initial Access (T1190)	sqlmap	Suricata / WAF
Brute force	Credential Access (T1110)	Hydra	Suricata
Lateral movement (PtH)	Lateral Movement (T1550)	Impacket	SIEM (ELK)
Ransomware	Impact (T1486)	controlled script	SIEM / HIDS

firewall [38] implements segmentation, NAT and traffic shaping, and Docker Compose [39] hosts the SIEM micro-services. The entire toolchain is open-source, satisfying requirement N3.

6.3. Attack scenarios

Six attack scenarios constitute the methodological core of the range. Each scenario is documented with a target host, technical prerequisites, a step-by-step procedure, the expected indicators of compromise and the corresponding Suricata detection rule, and each is mapped to one or more MITRE ATT&CK tactics [10]. Figure 4 presents the scenario matrix; Table 4 lists the scenarios with their ATT&CK alignment and detection mechanism.

The scenarios span the principal attack categories. The DDoS scenario generates SYN, UDP and HTTP floods with hping3 to study link saturation and mitigation. The man-in-the-middle scenario performs ARP spoofing to intercept internal traffic. The SQL-injection scenario uses sqlmap against the vulnerable web application to demonstrate data exfiltration, recalling the initial-access stage of the Equifax breach [8]. The brute-force scenario targets SSH and directory services. The lateral-movement scenario reproduces a pass-the-hash campaign across the Active Directory domain – the most realistic emulation of advanced persistent threats and the most challenging to detect – while the ransomware scenario simulates the impact stage through controlled file encryption.




Figure 5: CPU load (left) and WAN throughput (right) of the DMZ web server during a 120-second SYN flood, showing saturation during the attack and recovery within 30–40 s after mitigation.

7. Experimental evaluation

7.1. Testbed and methodology

The range was evaluated on the host described in Section 6 with respect to the requirements of Section 4. Four experiments were conducted: behaviour under a denial-of-service attack, scalability with respect to the number of concurrent virtual machines, availability under sustained operation, and the effectiveness of the intrusion-detection subsystem. Performance telemetry was collected at five-second intervals; detection metrics were computed by comparing Suricata/SIEM alerts against the ground-truth attack schedule.

7.2. Performance under a Denial-of-Service Attack

The first experiment subjected the DMZ web server to a 120-second SYN flood. Figure 5 shows the CPU load and the WAN interface throughput. Before the attack the system operated nominally, with CPU utilisation near 15% and traffic around 50 Mbit/s. When the flood began, CPU utilisation rose to approximately 85% and inbound traffic peaked near 920 Mbit/s, effectively saturating the link. After kernel-level SYN-cookie protection and pfSense traffic shaping were activated, the indicators returned to baseline within 30–40 seconds. The experiment confirmed both that the range can detect and visualise volumetric attacks and that standard mitigations are effective.

7.3. Scalability and availability

The second and third experiments examined scalability and reliability. Figure 6 (left) plots host CPU and RAM utilisation against the number of concurrent virtual machines. At the maximum design configuration of fifteen concurrent machines, CPU utilisation reached 62% and RAM utilisation 84%, both within the 80% load threshold for the processor and leaving headroom for snapshots. This result is consistent with the findings of Yaroviyy et al., who report CPU utilisation below 60% even for configurations with many more virtual machines [28]. Figure 6 (right) reports a 72-hour soak test in which a probe queried all key services every 60 seconds; overall availability was 99.86%, exceeding the 99% service-level target. The two brief dips at hours 24 and 48 correspond to scheduled environment resets and did not exceed 40 seconds; no unplanned failures or memory leaks were observed.

7.4. Detection effectiveness

The fourth experiment assessed detection effectiveness. Figure 7 shows the detection latency and the precision/recall achieved for each attack type. Attacks with clear network signatures were detected




Figure 6: Host resource utilisation versus the number of concurrent virtual machines (left) and system availability over a 72-hour soak test (right).




Figure 7: Detection latency (left) and precision/recall (right) of the Suricata + Elastic SIEM subsystem for each attack type.

fastest – SQL injection in 0.9 s and the SYN flood in 1.2 s – whereas multi-stage attacks required event correlation and therefore longer latency: 8.7 s for ransomware and 12.4 s for lateral movement. Precision and recall exceeded 94% for the simpler attacks; for lateral movement they were 82% and 78% respectively, consistent with the known difficulty of detecting advanced persistent-threat techniques such as pass-the-hash. Figure 8 (left) presents the aggregated confusion matrix across all scenarios, yielding an accuracy of 97.7%, a precision of 98.1%, a recall of 96.5% and an F1-score of 97.3%, with only eight false positives among 985 analysed events.

7.5. Integral efficiency assessment

For a summative assessment, the integral efficiency model of Yaroviy et al. was applied [28]. The model combines normalised indicators into a single objective function,

$$F_1 = \alpha \cdot E + \beta \cdot S + \gamma \cdot R, \quad (1)$$

where E denotes efficiency (performance and resource utilisation), S denotes success (reliability and absence of failures), R denotes resilience (tolerance to overload), and the weights satisfy $\alpha + \beta + \gamma = 1$; equal weights of 0.33 were adopted. The normalised component values obtained from the experiments were throughput $T = 0.86$, resource-utilisation efficiency $\eta = 0.78$, session-success probability $P_s = 0.96$ and resilience coefficient $C = 0.88$, shown in Figure 8 (right). Substituting the aggregated




Figure 8: Aggregated confusion matrix of the detection subsystem (left) and normalised components of the integral efficiency function with the resulting score $F_1 = 0.87$ (right).

Table 5
Evaluation Outcomes Against the Original Requirements

Req.	Description	Result	Status
F1	Three-tier topology with realistic services	4 zones, 9 services	Met
F2	At least five attack categories	6 scenarios	Met
F3	IDS/SIEM with event correlation	Suricata + ELK	Met
F4	Automated deployment (IaC)	5–10 min	Met
N1	Resource efficiency	CPU 62%, RAM 84%	Met
N3	Open-source only	100% FOSS	Met
N4	Network isolation	4 isolated zones	Met
—	Detection effectiveness	F1 = 97.3%	—
—	Integral efficiency	0.87	High

indicators yields an integral efficiency of 0.87, which corresponds to a high level and confirms the suitability of the range for educational and research use. Table 5 summarises the evaluation outcomes against the original requirements; all functional (F1–F5) and non-functional (N1–N5) requirements were satisfied.

8. Discussion and limitations

The results demonstrate that a fully open-source, on-premise cyber range can reproduce realistic corporate-network attacks, detect them with high accuracy, and operate reliably on commodity hardware. Three findings merit emphasis. First, the infrastructure-as-code pipeline reduced the time to a clean, reproducible environment to a few minutes, which is decisive for both teaching throughput and repeatable experimentation – a property repeatedly identified as essential in prior work [18, 20, 25]. Second, the integration of a network IDS with a SIEM, rather than network emulation alone, addresses the monitoring gap reported for lightweight academic ranges [21]. Third, the strict four-zone isolation directly mitigates the self-protection risks that Noponen et al. attribute to cyber ranges themselves [26].

Several limitations remain. The detection latency for multi-stage attacks (8–12 s) is dominated by the need to correlate events over time; behavioural analytics based on machine learning could reduce it. The current catalogue focuses on information-technology services and does not yet include operational-technology or IoT targets, which Ziubina et al. show to be valuable for critical-infrastructure training [27]. Finally, the single-host deployment, while cost-effective, bounds the maximum scenario size; a Proxmox cluster would raise this ceiling and add fault tolerance.

9. Conclusions and future work

This paper presented the design, implementation and evaluation of an open-source cyber range for researching network attacks on corporate infrastructure. From an analysis of the contemporary threat landscape – including the Ukrainian power-grid attacks, NotPetya, Colonial Pipeline and the Equifax breach – and a comparative study of existing platforms, we derived an explicit set of requirements and a six-layer architecture with a four-zone segmented network. The range was implemented entirely with open-source technologies – Proxmox VE, Terraform, Ansible, pfSense, Suricata and the Elastic Stack – and equipped with six attack scenarios mapped to MITRE ATT&CK.

Experimental evaluation confirmed that the range provisions in five to ten minutes, operates at 62% CPU and 84% RAM utilisation with fifteen concurrent virtual machines, sustains 99.86% availability over a 72-hour test, and detects attacks with an aggregate F1-score of 97.3%. An integral efficiency assessment yielded a score of 0.87, indicating a high level of efficiency and resilience. The range thus offers an accessible, reproducible and effective platform for cybersecurity education and applied research.

Future work will integrate machine-learning-based behavioural analysis to accelerate the detection of multi-stage attacks, add host-based detection agents to shorten the response time to internal threats, extend the scenario catalogue to ICS/SCADA and IoT targets, and migrate to a Proxmox cluster to improve scalability and fault tolerance.

Declaration on Generative AI

The author have not employed any Generative AI tools.

References

- [1] Check Point Research, Cyber Security Report 2024, Technical Report, Check Point Software Technologies, 2024. URL: <https://www.checkpoint.com/security-report/>.
- [2] IBM Security, Cost of a Data Breach Report 2024, Technical Report, IBM Corporation, 2024. URL: <https://www.ibm.com/reports/data-breach>.
- [3] ISC2, ISC2 Cybersecurity Workforce Study 2023, Technical Report, International Information System Security Certification Consortium, 2023.
- [4] Electricity Information Sharing and Analysis Center (E-ISAC), Analysis of the Cyber Attack on the Ukrainian Power Grid, Defense Use Case, SANS ICS, 2016.
- [5] A. Cherepanov, Win32/Industroyer: A New Threat for Industrial Control Systems, White Paper, ESET, 2017.
- [6] US-CERT, Alert TA17-181A: Petya Ransomware, Technical Report, Cybersecurity and Infrastructure Security Agency, 2017. URL: <https://www.cisa.gov/>.
- [7] Cybersecurity and Infrastructure Security Agency, FBI, Joint Cybersecurity Advisory AA21-131A: DarkSide Ransomware, Technical Report, CISA, 2021. URL: <https://www.cisa.gov/>.
- [8] U.S. Government Accountability Office, Data Protection: Actions Taken by Equifax and Federal Agencies in Response to the 2017 Breach, Technical Report GAO-18-559, U.S. Government Accountability Office, 2018.
- [9] State Service of Special Communications and Information Protection of Ukraine, Russia's Cyber Tactics: Annual Report 2023, Technical Report, SSSCIP / CERT-UA, 2024. URL: <https://cert.gov.ua/>.
- [10] MITRE Corporation, MITRE ATT&CK: Adversarial tactics, techniques, and common knowledge, <https://attack.mitre.org/>, 2024. URL: <https://attack.mitre.org/>.
- [11] E. Ukwandu, M. A. Ben Farah, H. Hindy, D. Brosset, D. Kavallieros, R. Atkinson, C. Tachtatzis, M. Bures, I. Andonovic, X. Bellekens, A review of cyber-ranges and test-beds: Current and future trends, *Sensors* 20 (2020) 7148. doi:10.3390/s20247148.
- [12] M. M. Yamin, B. Katt, V. Gkioulos, Cyber ranges and security testbeds: Scenarios, functions, tools and architecture, *Computers & Security* 88 (2020) 101636. doi:10.1016/j.cose.2019.101636.

- [13] European Cyber Security Organisation, Understanding Cyber Ranges: From Hype to Reality, WG5 Paper, ECSO, 2020.
- [14] J. Davis, S. Magrath, A Survey of Cyber Ranges and Testbeds, Technical Report, DSTO Defence Science and Technology Organisation, 2013.
- [15] B. Ferguson, A. Tall, D. Olsen, National cyber range overview, in: Proceedings of the IEEE Military Communications Conference (MILCOM), IEEE, 2014, pp. 123–128.
- [16] J. Vykopal, R. Oslejsek, P. Celeda, M. Vizvary, D. Tovarnak, KYPO cyber range: Design and use cases, in: Proceedings of the 12th International Conference on Software Technologies (ICSOFT), SciTePress, 2017, pp. 310–321.
- [17] P. Čeleda, J. Čegan, J. Vykopal, D. Tovarnák, KYPO — a platform for cyber defence exercises, in: STO-MP-MSG-133, NATO Science and Technology Organization, 2015, pp. 1–12.
- [18] J. Vykopal, M. Vizvary, R. Oslejsek, P. Celeda, D. Tovarnak, Lessons learned from complex hands-on defence exercises in a cyber range, in: Proceedings of the IEEE Frontiers in Education Conference (FIE), IEEE, 2017, pp. 1–8. doi:10.1109/FIE.2017.8190713.
- [19] M. Leitner, M. Frank, W. Hotwagner, G. Langner, O. Maurhart, T. Pahi, L. Reuter, F. Skopik, P. Smith, M. Warum, Enabling exercises, education and research with a comprehensive cyber range, Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications 12 (2021) 37–61. doi:10.22667/JOWUA.2021.12.31.037.
- [20] C. Pham, D. Tang, K.-i. Chinen, R. Beuran, CyRIS: A cyber range instantiation system for facilitating security training, in: Proceedings of the 7th International Symposium on Information and Communication Technology (SoICT), ACM, 2016, pp. 251–258. doi:10.1145/3011077.3011087.
- [21] N. Chouliaras, G. Kittes, I. Kantzavelou, L. Maglaras, G. Pantziou, M. A. Ferrag, Cyber ranges and testbeds for education, training, and research, Applied Sciences 11 (2021) 1809. doi:10.3390/app11041809.
- [22] M. Karjalainen, T. Kokkonen, Comprehensive cyber arena: The next generation cyber range, in: Proceedings of the IEEE European Symposium on Security and Privacy Workshops (EuroS&PW), IEEE, 2020, pp. 11–16.
- [23] E. Russo, L. Verderame, A. Merlo, Enabling next-generation cyber ranges with mobile security components, in: Testing Software and Systems (ICTSS), Springer, 2020, pp. 150–165.
- [24] M. Smyrlis, I. Somarakis, G. Spanoudakis, G. Hatzivasilis, S. Ioannidis, CYRA: A model-driven cyber range assurance platform, Applied Sciences 11 (2021) 5165. doi:10.3390/app11115165.
- [25] V. E. Urias, W. M. S. Stout, B. Van Leeuwen, H. Lin, Cyber range infrastructure limitations and needs of tomorrow: A position paper, in: Proceedings of the International Carnahan Conference on Security Technology (ICCST), IEEE, 2018, pp. 1–5. doi:10.1109/CCST.2018.8585460.
- [26] S. Noponen, J. Pärssinen, J. Salonen, Cybersecurity of cyber ranges: Threats and mitigations, International Journal of Information Security Research 12 (2022) 1032–1040.
- [27] R. Ziubina, G. Litawa, R. Szklarczyk, P. Biel, O. Veselska, J. Nikodem, Detection of viruses using machine learning method, in: PACIS 2020 Proceedings, 2020.
- [28] R. Yaroviy, O. Skliarenko, A. Kozynets, V. Lakhno, Quantitative models for evaluating the efficiency and resilience of a cyber range for educational and research purposes, CEUR Workshop Proceedings 4133 (2025).
- [29] M. Aleksander, G. Litawa, V. Karpinskyi, Distributed computing system which solves an elliptic curve discrete logarithm problem, in: Proceedings of the 10th International Conference “The Experience of Designing and Application of CAD Systems in Microelectronics”, 2009, pp. 378–380.
- [30] Proxmox Server Solutions GmbH, Proxmox VE Administration Guide, 2024. URL: <https://pve.proxmox.com/pve-docs/>.
- [31] HashiCorp, Terraform Documentation, 2024. URL: <https://developer.hashicorp.com/terraform/docs>.
- [32] Red Hat, Ansible Documentation, 2024. URL: <https://docs.ansible.com/>.
- [33] R. Wood, Damn vulnerable web application (DVWA), <https://github.com/digininja/DVWA>, 2024. URL: <https://github.com/digininja/DVWA>.
- [34] OffSec Services Limited, Kali Linux Documentation, 2024. URL: <https://www.kali.org/docs/>.

- [35] Rapid7, Metasploit Framework Documentation, 2024. URL: <https://docs.metasploit.com/>.
- [36] Open Information Security Foundation, Suricata User Guide, 2024. URL: <https://docs.suricata.io/>.
- [37] Elastic N.V., Elastic Stack Documentation, 2024. URL: <https://www.elastic.co/guide/>.
- [38] Netgate, pfSense Documentation, 2024. URL: <https://docs.netgate.com/pfsense/>.
- [39] Docker Inc., Docker Documentation, 2024. URL: <https://docs.docker.com/>.