Hybrid Capture the Flag platform for cybersecurity education

Anuar Askarov[1,*,†], Samat Mukhanov[1,†], Saule Amanzholova[1,†], Dauren Sagidullauly[1,*,†] and Kymbat Seilkhanova[1,*,†]

[1] *Astana IT University, Prospekt Mangilik Yel., Astana 010000, Kazakhstan*

Abstract

Cybersecurity education favors scenario-driven, experiential models that aim to close the gap between theory and practice. Yet many CTFs skew toward either rigid guided tasks or highly realistic attack-defense play, and often separate Red and Blue skills or misjudge difficulty. The article (a) reviews CTF-based learning, platforms, and learning evidence; (b) integrates existing design heuristics into a case-informed methodology that operationalizes a hybrid CTF format that combines Red and Blue learning goals; and (c) presents an implementation delivered on virtualized infrastructure and a mixed audience of 36 participants at a seven-hour, 12-team event. Results show strong engagement, appropriate challenge, and coverage of offensive and defensive techniques. We discuss realism–manageability trade-offs, scaffolding for novices, and scaling, and offer design guidance plus future directions on outcomes, fair scoring, and retention.

Keywords

cybersecurity education, capture-the-flag, gamification, red team, blue team, cyber range, CTFd, VMware ESXi, network forensics, social engineering, pedagogy

1. Introduction

The growing frequency and sophistication of cyber incidents have placed sustained pressure on the education pipeline to produce graduates who are ready to perform under realistic conditions. Traditional lecture-centric formats, while valuable for foundational theory, fall short of capturing the messy, multi-stage, adversarial nature of real-world cybersecurity operations. This gap has fueled widespread adoption of hands-on learning approaches, exemplified by CTF competitions and cyber ranges that immerse students in concrete tasks with high fidelity to operational tools and scenarios[1]. These activities demand not just procedural knowledge but also creative problem solving and time-boxed decision-making [1, 2, 3].

CTF competitions now span a broad spectrum from Jeopardy-style freeform puzzle solving to team-based attack–defense battles. Jeopardy-style CTFs present a battery of independent challenges across categories (web exploitation, cryptography, reverse engineering, forensics, etc.), typically emphasizing guided problem-solving in a controlled environment. In contrast, attack–defense CTFs increase realism by requiring participants to alternate between offensive and defensive roles, approximating the duel of tactics and counter-tactics found in real cyber conflicts [4, 8, 9, 10]. However, most existing CTF platforms and challenge libraries exhibit an offensive bias, focusing on hacking and exploitation, while omitting critical Blue Team behaviors like log monitoring, network triage, and incident response. Mapping studies have shown that human and organizational aspects such as social engineering and security culture receive comparatively less attention in standard CTF

[1] *CISN 2025: Workshop on Cybersecurity, Infocommunication Systems and Networks, November 19-20, 2025, Almaty, Kazakhstan*

[*] Corresponding author.

[†] These authors contributed equally.

255160@astanait.edu.kz (A. Askarov); s.mukhanov@astanait.edu.kz (S. Mukhanov); s.amanzholova@astanait.edu.kz (S. Amanzholova); d.sagidullauly@astanait.edu.kz (D. Sagidullauly); k.seilkhanova@astanait.edu.kz (K. Seilkhanova)

0009-0009-2025-8838 (A. Askarov); 0000-0001-8761-4272 (S. Mukhanov); 0000-0002-6779-9393 (S. Amanzholova); 0009-0004-0996-0275 (D. Sagidullauly); 0009-0004-0996-0275 (K. Seilkhanova)

challenge corpora[2]. These gaps suggest an opportunity for hybrid CTF designs that stitch together offensive and defensive elements within a controlled, classroom-friendly format.

Research Questions: To address these gaps, we explore a hybrid CTF approach guided by the following research questions:

- RQ1: How can a hybrid CTF event be designed to balance offensive and defensive learning objectives effectively, providing both realism and structured guidance within practical classroom constraints?
- RQ2: What evidence of learning or skill development (e.g., engagement metrics, self-reported gains) emerges from participation in a hybrid CTF event, and how does this relate to the event's learning objectives?

Contributions: This work makes two contributions. First, we review and synthesize the literature on CTF-based cybersecurity education, competition platforms, and prior empirical evaluations, extrapolating design principles for hybrid CTFs that foreground both Red and Blue team learning outcomes while remaining tractable for instructors. Second, we present a case-informed methodology and deployment of a hybrid CTF event, along with its evaluation in a mixed-ability cohort. We discuss the results in terms of learner engagement, skills coverage, and practical lessons, and we relate these findings directly to educational objectives (e.g., which competencies were exercised and what knowledge gains were observed). The article is organized as follows: Section Related Work surveys existing CTF platforms, hybrid formats, and educational studies. Section Methodology details the pedagogical framing, infrastructure, platform configuration, and challenge design of our hybrid CTF. Section Results and Discussion presents the evaluation outcomes, including participant engagement, perceived difficulty, and how these results map to learning objectives. Section Conclusion offers design guidance for educators and outlines future work, including a placeholder for incorporating pre- and post-assessment data on student skills.

2. Related work

CTF Platforms and Hybrid Formats: Over the past decade, numerous platforms and frameworks have been developed to support cybersecurity competitions. Open-source and commercial platforms such as Facebook's open-source FBCTF and the popular CTFd framework provide scalable infrastructure for Jeopardy-style challenges, with studies highlighting their usability and assessment features [4, 5, 18]. Other platforms like TryHackMe (THM) and Hack The Box (HTB) offer extensive libraries of guided challenges and realistic virtual machines, respectively, illustrating different points on the guided-vs-realism spectrum. Table 1 (in Section Infrastructure) compares THM and HTB, highlighting how guided content can lower entry barriers for novices, whereas realistic full-system challenges better mirror real-world scenarios. These differences carry educational implications: highly scaffolded environments can efficiently build fundamental skills, while more open-ended environments foster independent problem-solving and realism. Designing a hybrid CTF requires capturing the strengths of both – providing guidance and scaffolding without sacrificing operational authenticity.

Academic competitions like the International CTF (iCTF) have long demonstrated the feasibility of large-scale attack-defense exercises, albeit with significant organizational overhead [8]. Vigna et al. (2014) chronicle ten years of iCTF, noting the good, bad, and ugly lessons of scaling such events [8]. More recent efforts have explored hybrid or red vs. blue formats in controlled settings. For example, Chindrus and Caruntu (2023) describe a red-and-blue team competition scenario and emphasize the "virtuous cycle" of skill growth when offensive and defensive teams learn from each other's tactics[3][4]. Their case study underscores the value of combining adversarial (Red) and defensive (Blue) training methodologies to improve overall cybersecurity proficiency. Canitano (2022) similarly developed a framework for attack/defense CTF competitions, demonstrating how hybrid formats can be implemented at the university level (e.g., by structuring exercises where teams swap between attacking others' systems and defending their own) [10]. These studies highlight the primary trade-off in hybrid CTF design: realism vs. manageability. Fully interactive red/blue exercises provide high realism but can overwhelm instructors due to complex infrastructure support and the need to prevent inter-team chaos [8, 10]. Simplified hybrids, like the one we implement, capture slices of realism (segmented networks, realistic service endpoints, adversarial elements) without cross-team attacks, yielding a compromise that is logistically tractable for small teams of organizers.

Gamified Learning and Educational Impact: CTFs are a form of gamified learning, applying game elements to an educational context, and they have been shown to significantly boost student motivation and engagement. Prior work has found that integrating CTF-style exercises into cybersecurity courses increases student interactivity and perceived usefulness of the material, without harming objective learning outcomes [2]. Cole (2022) reported that students in an introductory security class found CTF labs more engaging and equally effective as traditional exercises in terms of exam performance[5][6]. Beyond motivation, researchers are now examining the knowledge gains from CTF participation. A recent empirical study by Schafeitel-Tähtinen and Lazarov (2025) measured the effectiveness of a gamified CTF scenario using pre- and post-surveys in two universities. They found that participating in CTFs led to significant increases in students' cybersecurity knowledge, practical skills, self-efficacy, and interest levels[7]. Notably, these positive effects were observed across different student groups regardless of prior skill, indicating that well-designed CTF activities can benefit both novices and more experienced students[8]. This evidentiary support for learning gains addresses a known concern in gamification research: it is not enough for an activity to be "fun" or engaging; it should also produce measurable improvements in competence [16, 17].

Cyber Ranges and Platform Innovations: In parallel with CTFs, cyber ranges (controlled, virtual environments for cybersecurity training) have become prominent. Recent studies focus on building flexible cyber ranges that can host complex scenarios for training and testing. Park et al. (2022) propose a multi-cyber range architecture that can isolate multiple concurrent training missions, reflecting military training needs [1]. Shin et al. (2024) describe design principles for cyber training ranges that effectively respond to evolving cyber threats, emphasizing the integration of realistic threat vectors and defensive monitoring in training scenarios [3]. These works contribute to the infrastructure aspect of cybersecurity education, demonstrating how advanced virtualization and orchestration can provide safe yet realistic learning environments. Our approach leverages similar ideas on a smaller scale: using virtualized infrastructure (ESXi) and network segmentation to create an environment where both red-team style exploits and blue-team style investigations can occur safely.

In summary, existing research and practice underscore the importance of balanced design in cybersecurity competitions. A hybrid CTF that explicitly links challenges to learning objectives can address gaps identified in prior work (e.g., inclusion of defensive and human-centric content [14]) while maintaining the engagement benefits of gamification. The methodology we present next draws on these lessons: we aim to achieve an effective balance between offensive and defensive challenges, guided scaffolding and open problem-solving, and realistic complexity versus manageable scope.

3. Methodology

3.1. Pedagogical framing and design goals

The hybrid CTF design was guided by three pedagogical goals distilled from prior work. First, balance: we deliberately coupled offensive tasks that cultivate adversarial (Red Team) thinking with defensive tasks that require evidence-driven reasoning and operational hygiene [4, 9, 14, 19]. By ensuring participants face both attack and defense challenges, we aimed to broaden skill coverage beyond what a purely offensive or purely defensive event would achieve. Second, realism under constraints: we sought to simulate operational cues and realistic scenarios (e.g. server logs, network traces, role-play elements) without the heavy orchestration overhead or unpredictability of a fully interactive attack–defense tournament [8, 10]. In practice, this meant incorporating real-world artifacts (traffic captures, vulnerable VMs, etc.) and adversarial storylines, but isolating teams from directly attacking each other. Third, accessibility with rigor: we integrated scaffolded hints and tiered challenge difficulty to support novices while still providing meaningful challenges for advanced learners [7, 13]. Hints and guidance were designed to teach rather than simply give away answers, in line with evidence that immediate feedback can sustain motivation if calibrated properly to avoid short-circuiting learning [13, 16, 17].

We adopted a problem-based learning stance: each challenge was conceived as a self-contained micro-case with a narrative arc that surfaces a core cybersecurity competency and encourages reflection on generalizable tactics, techniques, and procedures. Hints were deployed just-in-time (on a timed unlock or upon request) and were crafted to reveal the thought process or method rather than the solution outright. This scaffolding approach aligns with pedagogical recommendations to maintain a balance between difficulty and approachability for the target audience[9][10]. Ultimately, our design goals serve the overarching learning objectives of the event: to engage participants in applying both offensive and defensive security skills, and to do so in an environment where they can learn from the experience (through hints, post-challenge write-ups, and debriefings) rather than simply be tested.

3.2. Infrastructure and isolation

To ensure a safe and controlled environment, all challenge artifacts were deployed on a virtualized cyber range using VMware ESXi. Each team operated within a segregated network segment to prevent any unintended interference between teams. The infrastructure plan (Figure 1) consisted of a demilitarized zone (DMZ) network where the CTFd platform and challenge VMs resided, isolated from external networks and from each other. Participants connected via a VPN gateway to the DMZ, which granted access only to the CTFd web interface and the specific service endpoints of each challenge VM. This network segmentation minimizes the "blast radius" of any misconfiguration or abuse, mirroring best practices in production environments for containing threats [11]. We selected ESXi for its support of fast snapshotting and reset of VMs between sessions, as well as flexible resource provisioning – features important for scaling the event to different cohort sizes and quickly recovering from any issues during the competition.

Table 1
The comparative analysis of two main cybersecurity training platforms, TryHackMe and Hack the Box

Characteristic	TryHackMe (THM)	Hack The Box (HTB)
Primary format & focus	Guided "rooms" and learning paths; also Jeopardy-style categories; includes competitive "King of the Hill"	Vulnerable VMs plus Jeopardy-style challenges (web, crypto, RE, forensics)
Access / hosting	Fully browser-based via AttackBox or OpenVPN	VPN or browser PwnBox (cloud Parrot Linux)
Pricing / license	Freemium; subscription unlocks full catalog and longer AttackBox sessions	Freemium; VIP/VIP+ expand access to retired machines/challenges and extend PwnBox time (up to unlimited with VIP+)
Guidance level	High: step-by-step lessons, embedded tutorials and hints; beginner-friendly curricula	Moderate: realistic, self-directed practice on full-system VMs; less guided
Content scope / categories	Jeopardy-style categories (web exploitation, forensics, etc.) within structured learning paths	Full-system exploitation labs + Jeopardy-style tasks across web, crypto, RE, forensics
Educational strengths	Excellent for foundational skills; classroom-suitable due to guidance and hints; lowers entry barrier	High educational relevance; mirrors real-world skills; strong engagement for universities/self-learners
Notes / limitations	Geared to beginner–intermediate depth in the excerpt; not focused on advanced attack-defense tournaments	Framed as a training portal rather than a self-hosted event framework in the excerpt

Within this infrastructure, we set up three dedicated Linux VMs (one per challenge, see Section "Challenge Craftsmanship") and the CTFd server. Each challenge VM was pre-configured with all necessary services and isolated so that teams could only interact with it via specific ports/protocols relevant to the challenge. This approach provided a slice of realism (teams had to use actual tools and network connections to solve challenges) while maintaining control: teams could not, for example, attack the CTFd server or each other's VMs because the network rules did not permit it. All traffic was confined to the virtual range.

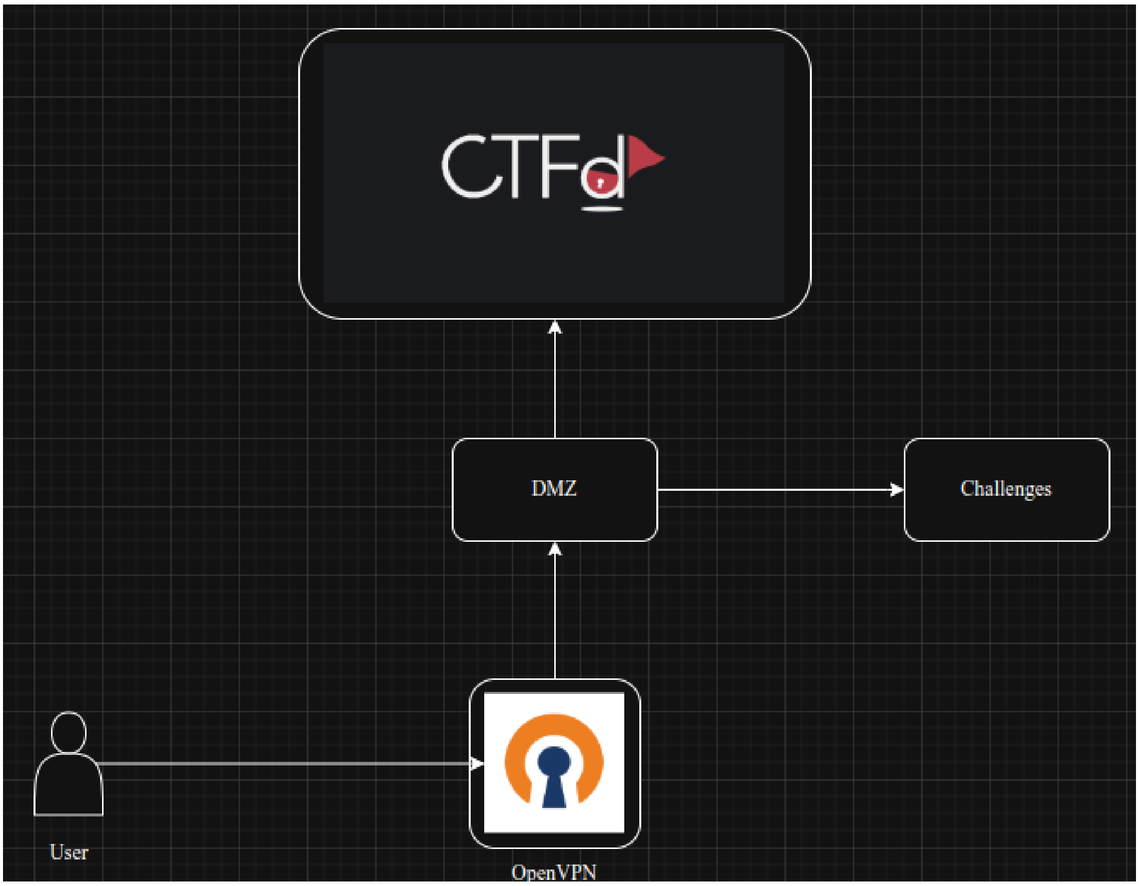

Figure 1: The plan of infrastructure.

3.3. Platform selection and configuration

CTFd was selected for administration convenience, intuitive team management, and scope for extensibility and customization, all advantages cited in the literature [4,5,18]. The instance was configured with fixed-value challenges in Jeopardy mode and a live scoreboard. Hints were staged on time-based unlocks to reduce brute forcing and nudge metacognition ("what class of vulnerability is this?") over memorized steps.

Although our format was not a full attack–defense competition, we incorporated Blue Team elements and role-play aspects into the Jeopardy-style framework. For instance, one challenge simulated an incident response scenario by providing log files and requiring participants to interpret them, and another involved a social engineering interaction with a mock "AI guard". By leveraging such Blue-leaning tasks and creative prompts within a Jeopardy structure, we heeded calls in the literature to include a broader range of competencies (beyond pure hacking) in cybersecurity exercises [9, 14]. This hybrid use of CTFd allowed us to blend offensive and defensive content smoothly: teams solved challenges through the CTFd interface as usual, but the nature of some challenges was defensive or human-focused, which is atypical for many CTFs.

3.4 Challenge craftsmanship

Three challenges were created to span scenarios and map back to both red and blue behaviors.

- "Unserialize Me" (Web Exploitation): A medium-hard web challenge anchored in PHP object injection via unsafe deserialization. The challenge plants a serialized PHP object in a cookie and uses a magic method side effect to read server files. Learners must infer the vulnerability by inspecting HTTP responses, decoding the cookie data, and crafting a malicious payload to exploit the deserialization flaw and retrieve a flag from the server. This reinforces secure coding principles and input validation practices, as participants see how a seemingly innocuous serialized object can lead to arbitrary file reads [14]. Learning objective: Identify and exploit an insecure deserialization vulnerability, and understand its implications for web application security (Bloom's Apply/Analyze level).
- "Network Analysis" (Packet Forensics/OSINT): A Blue Team flavored task built around a PCAP-NG network traffic capture containing game server discovery packets. Learners have to filter the capture for a specific UDP port related to the Steam gaming protocol, then pivot from the discovered data to an online profile, and finally decode an embedded Base64 artifact in that profile to recover the flag. This challenge rehearses practical skills in packet filtering, protocol recognition, external intelligence gathering, and artifact decoding—behaviors analogous to what a Security Operations Center (SOC) analyst might do when investigating network telemetry [14]. Learning objective: Practice network forensic analysis and intelligence-gathering techniques to extract hidden information from real traffic (Bloom's Apply level progressing to Analyze).
- "Soldier Swofford" (Adversarial Dialogue/Social Engineering): A low-tech, interactive role-play challenge designed to surface human-centric security risks and policy constraints. An AI-driven "gate guard" character only responds to a precise, rank-authenticated command sequence. Learners must deduce the correct phrases by reasoning about military protocol, chain-of-command, and social engineering cues provided in the scenario. Successfully eliciting the secret from the guard requires understanding the human context and carefully crafting the request. This scenario broadens the scope beyond technical exploits, echoing calls to include human factors in cybersecurity training [14]. Learning objective: Develop skills in social engineering and policy-compliant reasoning, highlighting the human element of security (Bloom's Evaluate level, as participants must assess the scenario and formulate a correct strategy).

Each challenge is authored with a consistent narrative arc, explicit learning objective, environmental notes, and a post-hoc write-up template which teams are encouraged to complete. Authoring artifacts (code snippets, Dockerfiles, and setup notes) are versioned for reuse.

3.5. Participants, procedure, and instrumentation

The hybrid CTF event, nicknamed "Digital Fortress," was a one-day competition lasting approximately seven hours. A total of 36 participants took part, organized into 12 teams of three members each. The participant cohort was diverse in background: teams included advanced high school students, undergraduate students (primarily cybersecurity or IT majors), and industry professionals from local companies. All participants had at least a basic familiarity with CTF-style challenges prior to the event (many had played in online CTFs or training platforms like THM/HTB before), ensuring that no team was completely new to the CTF format. This mix of experience levels meant that some teams had strong technical skills while others were more novice, a factor we monitored in the competition dynamics. We did not collect detailed demographic data (such as age or gender) due to the informal nature of the event, but anecdotally the majority of student participants were in their third or fourth year of university, and the professionals had 1-5 years of work experience in IT or cybersecurity.

Procedure. Teams received an orientation pack (VPN setup, CTFd accounts, rules). At 09:00 CTFd opened with a live scoreboard; teams worked at their own pace. Organizers monitored service status, the CTFd admin panel, and a Q&A chat. When many teams stalled on a task, brief just-in-time hints were broadcast. The event ended at 16:00, followed by an anonymous survey and a 30-minute discussion to gather feedback on clarity, realism, and perceived learning.

Instrumentation & data. CTFd logs captured logins, solves, time-to-first-solve, and hint usage; VPN logs indicated connectivity issues. The survey used Likert items on enjoyment and difficulty plus multiple-choice/free-text on challenge preferences and learning. As a voluntary event, no formal pre/post tests were administered; we note this limitation and plan to include such assessments in future iterations.

3.6. Outcome measures and analysis approach

Evaluation approach. Our analysis is deliberately descriptive and formative, appropriate for a single extracurricular deployment. We organize outcomes into Engagement, Progress, and Perception. Engagement covers participation (attendance, team formation), time on task across the 7-hour window, and hint-usage patterns. Progress uses objective CTFd metrics—solves per challenge/category, time-to-first-solve, and full time-to-solve distributions—to gauge relative difficulty and steady advancement. Perception draws on the post-event survey and discussion, capturing enjoyment, perceived difficulty/relevance, and qualitative comments.

Analysis & interpretation. With 12 teams (36 participants) and no pre/post tests, we avoid causal claims and treat findings as exploratory. We report simple descriptive statistics (e.g., share rating difficulty "just right," average solves per team) and conduct a thematic analysis of comments. Where useful, we compare patterns to prior CTF-education studies [2, 11–14] to note alignments or divergences (e.g., a preference for web challenges).

Limitations. We acknowledge selection bias in voluntary participation, the lack of longitudinal follow-up for retention, and any event-specific technical issues. These limitations are presented alongside results to temper conclusions and inform design improvements.

4. Results and discussion

Participants rated overall enjoyment highly and the live scoreboard visibly sustained effort across the event. In the survey, web exploitation was the most preferred category, with networking and forensics in the middle tier. This distribution is consistent with prior classroom experience where the immediate payoff and narrative clarity of web tasks tend to attract early attention [11,12,13]. Qualitative comments praised the "freshness" of the AI role-play and the realism of the network trace, indicating appetite for diversification beyond standard web-heavy sets.

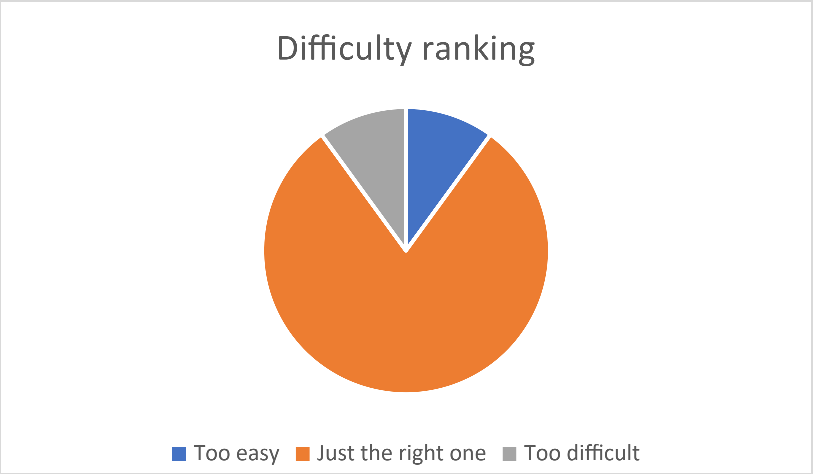

Figure 2: The overall evaluation of difficulty.

Eighty percent of respondents rated the overall difficulty “Just the right one,” with a minority at each tail labeling it too easy or too difficult. Observationally, teams pursued different approaches: some maximized early reward by front-loading easier tasks while others stayed focused on the harder web exploit. Hint unlocks triggered visible second winds, especially on the deserialization chain. These dynamics support the design choice to scaffold within challenges and keep inter-challenge dependency minimal so that teams can self-pace and avoid deadlock.

The hybrid set elicits distinct cognitive moves. In “Unserialize Me,” teams are observed inspecting HTTP responses, base64 decoding, class inference, payload crafting, and verification; these are classic offensive steps but have defensibility implications (input validation, safe serialization practices). In “Network Analysis,” teams practice filtering, protocol recognition, external look-ups, and artifact decoding; these are behaviors linked to SOC investigation flows. In “Soldier Swofford,” teams must engage in protocol reasoning and constraint navigation, useful proxies for policy-aware operations and adversarial thinking that involves human rules. This triangulation of measures reflects prior literature mapping of CTFs to technical skills but also speaks to the often-noted lack of human-centric content [14].

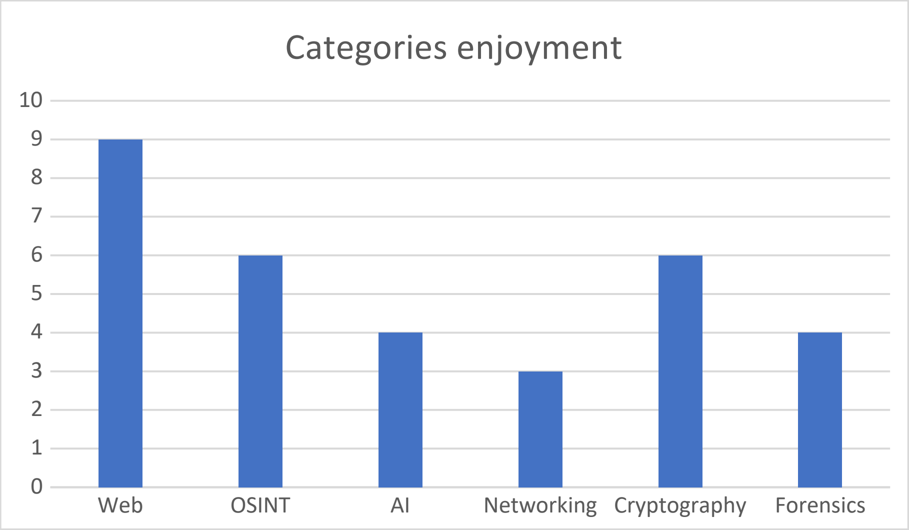

Figure 3: The enjoyment of the categories of challenges.

While interactive attack-defense is a dynamic learning setting, the orchestration costs (service health monitoring, scoring uptime, abuse containment) can overwhelm small instructional teams [8,10]. Our hybrid design captures slices of operational realism (segmented networks, service endpoints, adversarial cues) while eschewing cross-team assaults. The trade-off is lower exposure to adversarial disruption and defense-under-fire. For many courses, especially mixed-ability cohorts, this seems a productive compromise: it delivers operational signals and affords reflective problem solving, yet remains logistically tractable.

CTFd’s administrative surface and hinting mechanics supported just-in-time scaffolding and formative assessment. As in the platform literature, ease of use and assessment features are not niceties but drivers of educational value when used deliberately [4,5,18]. Scoring in this event remained simple (fixed points) which was adequate for a formative event. Future iterations can experiment with decaying-difficulty scoring or partial-credit designs, mindful of existing evidence on fairness and strategic behavior in scoring algorithms [19].

This evaluation is small-N, non-experimental, and subject to selection biases common to voluntary competitions. The survey instrument was minimal by design. We did not collect longitudinal data on retention or transfer. The attack surface of each challenge was intentionally limited, and a broader “fog of war” (alert fatigue, competing signals) was not simulated. These limits constrain inference but are appropriate to this article’s goal: actionable design guidance grounded in a realistic classroom deployment.

Clarity matters: narrative prompts must be short and unambiguous. Diversity is rewarding: mixing modalities (web, network, role-play) combats fatigue and spreads wins. Hints should teach: hint tiers should model thought process not just steps. Instrument early: capture solve times and hint clicks to identify friction points. Right-size the ops: a VM/ESXi + VPN + CTFd stack is a sweet spot for many programs - enough realism without constant firefighting.

4.1. Pre- and post-event skill assessment

In a future version of this study, this section will present a comparison of participants' cybersecurity knowledge and skills before and after the CTF event. For example, we plan to administer a short pre-test quiz covering key concepts addressed by the challenges (web vulnerabilities, network analysis, social engineering) and an equivalent post-test after the event. The expectation is that participants' post-test scores will improve relative to pre-test, indicating knowledge gains attributable to the CTF experience. Any improvements (or lack thereof) will be analyzed in relation to the learning objectives.

5. Conclusion

CTFs remain one of the most effective delivery modes for practice-ready cybersecurity skills but their educational impact hinges on design choices. A hybrid approach that deliberately couples offensive puzzle-solving with defensive and human-centered reasoning can broaden skill coverage while preserving manageability. The case demonstrated here shows that a modest, well-engineered infrastructure and a targeted set of challenges can deliver high engagement and well-balanced exposure for mixed-ability cohorts.

The recommendation to educators is pragmatic: start with a small hybrid set anchored in clear learning goals, bias toward reproducible infrastructure and strong write-ups, instrument for formative insight, iterate on hints and narratives, and when capacity grows, add attack-defense elements selectively. For researchers, priorities include measuring learning gains across Red and Blue competencies, evaluation of scoring schemes for fairness and motivation, and studying retention and transfer to authentic tasks.

Beyond single events, maturing hybrid CTFs into course infrastructure requires operating playbooks and reproducible artifacts. We suggest a public artifact bundle (Dockerfiles or VM snapshots, sanitized PCAPs, and seed datasets) with version tags; an outcomes-aligned rubric that separately scores Red and Blue competencies (e.g., reconnaissance, exploit design, log triage, and hypothesis testing) and integrates write-ups as authentic assessment; accessibility and ethics guardrails (codes of conduct, non-harm principles, and accommodations for differently-abled students as well as graded "safe mode" tasks); and operational metrics (service availability SLOs, hint latency, and resource utilization) so that instructors can anticipate cost and staffing. Coupled with pre/post quizzes and delayed retention probes, these practices would let institutions compare cohorts and iterate with evidence while keeping risk, effort, and cost manageable.

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT-5 and Grammarly in order to: Grammar and spelling check. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] M. Park, H. Lee, Y. Kim, K. Kim, and D. Shin. "Design and implementation of multi-cyber range for cyber training and testing." Applied Sciences 12(24):12546, 2022. DOI: 10.3390/app122412546.

[2] S. V. Cole. "Impact of CTF-style vs. traditional exercises in an introductory computer security class." In Proc. 27th ACM Conf. on Innovation and Technology in Computer Science Education (ITiCSE), pp. 470–476, 2022. DOI: 10.1145/3502718.3524806.

[3] Y. Shin, H. Kwon, J. Jeong, and D. Shin. "A study on designing cyber training and cyber range to effectively respond to cyber threats." Electronics 13(19):3867, 2024. DOI: 10.3390/electronics13193867.

[4] M. Swann, J. Rose, G. Bendiab, S. Shiaeles, and F. Li. "Open source and commercial Capture The Flag cybersecurity learning platforms: A case study." In Proc. 2021 IEEE Intl. Conf. on Cyber Security and Resilience (CSR), 2021, pp. 168–175. DOI: 10.1109/CSR51186.2021.9527915.

[5] R. G. Chicone and S. Ferebee. "A comparison study of two cybersecurity learning systems: Facebook's open-source CTF and CTFd." Issues in Information Systems 21(1): 202–212, 2020.

[6] D. Michael and S. Chen. Serious Games: Games that Educate, Train, and Inform. Thomson Course Technology PTR, 2006.

[7] T. Balon and I. Baggili. "Cybercompetitions: A survey of competitions, tools, and systems to support cybersecurity education." Education and Information Technologies 28(9): 11759–11791, 2023. DOI: 10.1007/s10639-022-11451-4.

[8] G. Vigna, K. Borgolte, J. Corbetta, A. Doupé, Y. Fratantonio, L. Invernizzi, D. Kirat, and Y. Shoshitaishvili. "Ten years of iCTF: The good, the bad, and the ugly." In Proc. USENIX Summit on Gaming, Games, and Gamification in Security Education (3GSE), 2014.

[9] C. Chindrus and C.-F. Caruntu. "Securing the network: A red and blue cybersecurity competition case study." Information 14(11):587, 2023. DOI: 10.3390/info14110587.

[10] G. Canitano. Development of a Framework for Attack/Defense Capture the Flag Competition. Master's thesis, Politecnico di Torino, 2022.

[11] H. Gonzalez and R. Llamas. "Cybersecurity teaching through gamification: Aligning training resources to our syllabus." Research in Computing Science 146: 35–43, 2017.

[12] M. Adams and M. Makramalla. "Cybersecurity skills training: An attacker-centric gamified approach." Technology Innovation Management Review 5(1): 5–14, 2015.

[13] M. Malone, Y. Wang, K. James, M. Anderegg, J. Werner, and F. Monrose. "To gamify or not? On leaderboard effects, student engagement and learning outcomes in a cybersecurity intervention." In Proc. 52nd ACM SIGCSE, pp. 1135–1141, 2021.

[14] V. Švábenský, P. Čeleda, J. Vykopal, and S. Brišáková. "Cybersecurity knowledge and skills taught in capture the flag challenges." Computers & Security 102: 102154, 2021. DOI: 10.1016/j.cose.2020.102154.

[15] A. D. Cybulski. Sim-Cyberpunk: Serious Play, Hackers and Capture the Flag Competitions. PhD thesis, University of Toronto, 2023.

[16] J. J. Lee and J. Hammer. "Gamification in education: What, how, and why bother?" Academic Exchange Quarterly 15(2): 1-5, 2011.

[17] R. N. Landers. "Developing a theory of gamified learning: Linking serious games and gamification of learning." Simulation & Gaming 45(6): 752–768, 2015.

[18] R. G. Chicone, T. Burton, and J. Huston. "Using Facebook's Capture the Flag platform as a hands-on learning and assessment tool for cybersecurity education." International Journal of Conceptual Structures and Smart Applications 6(1): 18–32, 2018.

[19] P.-V. Besson, R. Brisse, H. Orsini, N. Talon, J.-F. Lalande, F. Majorczyk, A. Sanchez, and V. Viet Triem Tong. "CERBERE: Cybersecurity exercise for red and blue team entertainment, reproducibility." In Proc. IEEE Intl. Conf. on Big Data (BigData), pp. 2980–2988, 2023. DOI: 10.1109/BigData59044.2023.10386953.

[20] T. Schafeitel-Tähtinen and W. Lazarov. "Teaching and learning cybersecurity using capture the flag: Effectiveness comparison between university students in Finland and Czechia." Computer Applications in Engineering Education 30(5): e70082, 2025. DOI: 10.1002/cae.70082.