TRIDENT: HOW TO BREAK DEEP REINFORCEMENT LEARNING CYBER DEFENSES (AGENTIC)

Anonymous authors
Paper under double-blind review

ABSTRACT

Autonomous cyber defense systems based on Deep Reinforcement Learning (DRL) have attracted significant research attention, yet remain evaluated almost exclusively against static, heuristic red agents, leaving their robustness against adaptive threats critically understudied. Meanwhile, recent advances in Reinforcement Learning with Verifiable Rewards (RLVR) have improved LLM reasoning, but their integration into cybersecurity remains elusive due to the absence of suitable benchmark environments and interaction datasets. To bridge this gap, we introduce **Trident**, an agentic LLM red teaming framework comprising three components: a dynamic benchmark with isolated sandbox servers spanning CybORG CAGE 4 and CyberWheel, a dataset comprising over 13,000 high-fidelity red-blue interaction trajectories for RLVR, and a "Code-as-Policy" RLVR agentic architecture (**Trident Agentic**). The latter reformulates red agent training as a contextual bandit via a tripartite Log Summarizer–Planner–Coder design, where a trainable Planner generates complete attack strategies from compressed execution logs, which a frozen Coder translates into executable Python policies deployed against live DRL defenders. Empirical evaluations reveal a consistent brittleness in existing defenses: a single trainable 7B planner degrades blue-agent performance by an average of **628%** across four state-of-the-art CAGE 4 defenders and eliminates over **99%** of the defender's episodic reward in CyberWheel, while autonomously discovering novel behaviors such as decoy avoidance and adaptive state prioritization that static heuristics entirely fail to uncover. Anonymized code is available here[1].

1 INTRODUCTION

Modern computer networks face escalating cyber threats that increasingly outpace manual defense capabilities (DARPA, 2026). This has driven a long line of research in autonomous cyber defense (ACD): from early rule-based expert systems (Neumann & Porras, 1999; Stakhanova et al., 2007) and anomaly detection (Lin et al., 2022; Zhao et al., 2023), to deep reinforcement learning (DRL) agents (Singh et al., 2025; King et al., 2025; Lyu et al., 2021; Castro et al., 2025) that actively monitor network states and recover compromised hosts in real-time. DRL-based defenders, in particular, remain critically understudied against adaptive threats (Kiely et al., 2025; Oesch et al., 2024): they are evaluated almost exclusively against static, heuristic red agents with fixed behavioral signatures, leaving their robustness against adaptive attackers unknown. A recent systematization of DRL for cybersecurity raises the same concern, identifying the fixed, non-adaptive attacker as one of the field's most prevalent evaluation pitfalls (McFadden et al., 2026).

The ideal red agent for rigorous evaluation should be adaptive: capable of observing defensive behavior, reasoning about weaknesses, and synthesizing novel attack strategies. Recent progress in agentic Large Language Models (LLMs) makes this vision increasingly plausible and urgent: Claude Mythos (Anthropic, 2026) has demonstrated that large-scale frontier agentic LLMs can autonomously identify and exploit zero-day vulnerabilities across major operating systems and web browsers. Research on agentic cybersecurity frameworks has made significant advancements across Capture The Flag (CTF) challenges (Abramovich et al., 2025; Shao et al., 2024a; Zhang et al., 2025a), penetration testing (Deng et al., 2024), and vulnerability reproduction (Zhu et al., 2025; Wang et al., 2026c; Zhuo et al., 2026; Suryanto et al., 2026). However, all of these works deploy

[1] https://anonymous.4open.science/r/Trident-A934

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

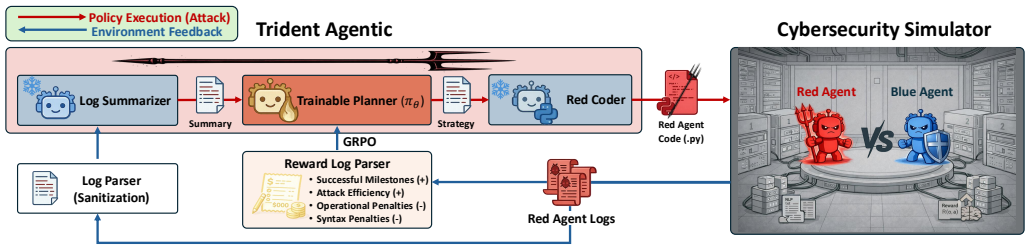

Figure 1: **Overview of the Trident Agentic framework.** A Log Summarizer observes and summarizes sanitized logs, enabling a trainable Planner (π_θ) to formulate abstract tactics. A frozen Coder translates the tactic into an executable Python policy deployed against a live DRL defender. Execution logs provide a verifiable reward signal to update π_θ via GRPO.

models via prompting or supervised fine-tuning on static datasets, and none train an agent through direct interaction with an active defender using Reinforcement Learning (RL), nor do they provide the infrastructure to do so.

DeepSeek-R1 (Guo et al., 2025) has popularized Reinforcement Learning with Verifiable Rewards (RLVR), where an LLM policy learns to maximize rewards derived from objective, environment-provided rules rather than relying solely on human-labeled examples. RLVR has since shown significant advances across mathematical reasoning (Cobbe et al., 2021; He et al., 2026; Shao et al., 2024b), software engineering (Wei et al., 2025; Zeng et al., 2025), information retrieval (Jin et al., 2025), and computer vision (Liu et al., 2025). Yet its integration into cybersecurity has remained elusive because execution environments are often ephemeral or restricted (Zhuo et al., 2026), even though the latest large-scale studies identify RLVR as a promising direction yet leave its realization entirely unexplored (Yu et al., 2025). A key reason is the absence of a suitable benchmark: one that provides stable, repeatable adversarial environments, verifiable reward signals, and active defenders against which red agents can be trained and evaluated.

High-fidelity simulators, such as CybORG CAGE 4 (Kiely et al., 2025) and CyberWheel (Oesch et al., 2024), offer a concrete path forward. They provide stable, repeatable adversarial environments designed to closely replicate real-world network operations. CAGE 4 models a realistic segmented enterprise network with cooperative multi-agent defenders operating under partial observability, while CyberWheel grounds its attack actions in real-world MITRE ATT&CK (MITRE, 2026) techniques via Atomic Red Team (Red Canary, 2026), supporting higher attack realism. Together, they offer realistic network attack and defense dynamics with deterministic, verifiable reward signals from system events, the properties needed to support RLVR training against active defenders. This motivates a critical question: *Can agentic LLMs be trained via RLVR to compromise state-of-the-art DRL defenders?*

To answer this question, we introduce **Trident** (Figure 1), a dynamic red teaming benchmark spanning CAGE 4 and CyberWheel. Trident provides isolated sandbox servers that execute generated attack policies against live DRL defenders to produce verifiable reward signals, alongside a dataset of over 13,000 high-fidelity red-blue interaction trajectories for RL-based post training, to the best of our knowledge, the first resource to support RLVR training against active DRL defenders.

To demonstrate the utility of this infrastructure, we further introduce **Trident Agentic**, an agentic LLM framework trained via RLVR under a "Code-as-Policies" (Liang et al., 2023; Wei et al., 2025) paradigm, where an agent's policy is instantiated as executable, LLM-generated code. Rather than selecting actions step-by-step under a Decentralized Partially Observable Markov Decision Process (Dec-POMDP), Trident Agentic reformulates red agent training as a contextual bandit. A trainable *Planner* receives a compressed summary of past execution logs from a *Log Summarizer* and generates an attack strategy in a single forward pass. A frozen *Coder* then translates this strategy into an executable Python policy.

Using Trident Agentic, we reveal a critical gap in how autonomous cyber defenders are currently assessed. Trident Agentic systematically compromises state-of-the-art DRL blue agents by discovering attack sequences that expose weaknesses in their learned policies, weaknesses that traditional heuristic red agents entirely fail to uncover. On average, it degrades blue-agent returns by **628%** across the four CAGE 4 defenders and reduces the CyberWheel defender's reward by over **99%** relative to the default heuristic baseline. These results suggest that current autonomous defenders

may be more brittle against adaptive attackers than prior evaluations have indicated. Our primary contributions are as follows:

- **Trident evaluation benchmark and dataset**: We introduce a cross-environment red teaming benchmark spanning CAGE 4 and CyberWheel. This framework utilizes isolated sandbox servers to produce verifiable reward signals against live DRL defenders. We release over 13,000 red-blue interaction trajectories, supporting RLVR training on LLMs.
- **Trident Agentic**: We introduce a tripartite agentic LLM architecture that leverages the "Code-as-Policies" paradigm to reformulate red agent training from a Dec-POMDP into a contextual bandit via RLVR, enabling the direct generation of executable attack strategies.
- **Empirical evidence of a consistent DRL brittleness**: Trident Agentic causes an average **628%** performance loss on four state-of-the-art CAGE 4 defenders and a **99%** reward collapse in CyberWheel, with a trained 7B SLM surpassing frontier-scale models, revealing that current static-adversary evaluation practices critically underestimate the brittleness of autonomous cyber defenses.

Beyond its immediate contributions, Trident's RESTful sandbox architecture, where generated policies are submitted and rewards are returned, enables seamless integration of future, higher-fidelity environments as drop-in servers. This plug-and-play design not only supports the continued development of adaptive red agents, but equally opens the door to training blue agents robust against generative attackers, establishing Trident as an extensible platform for the long-term co-evolution of autonomous cyber offense and defense.

2 PRELIMINARIES

2.1 ADVERSARIAL CYBER ENVIRONMENT AS DEC-POMDP

Previous studies model the autonomous cyber defense task from the defender's perspective as a Dec-POMDP (Singh et al., 2025; King et al., 2025), defined by the tuple $\mathcal{M} = (\mathcal{I}, \mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \Omega, \mathcal{O}, \gamma)$. Here, $\mathcal{I}$ represents the team of agents including $M \geq 1$ defending agents. $\mathcal{S}$ is the true network state space, $\mathcal{A}$ is the joint defense action space, and $\mathcal{R}(s, a)$ is the shared reward for the defending team. Each defending agent $i \in \mathcal{I}$ perceives the environment only through its local observation via the observation function $\mathcal{O}(o|s', a)$, reflecting the partial observability inherent to realistic network defense. Crucially, in these formulations, red agents and optionally, benign users are not treated as independent learning agents within $\mathcal{I}$, but rather as fixed logic embedded within the environment. The state transition dynamics are conditioned on the red agent policy code $c \in \mathcal{C}$, written as $\mathcal{T}(s'|s, a, c)$, reflecting that the network state trajectory is jointly determined by the environment dynamics and the static logic embedded in c. In existing benchmarks such as CAGE 4 ($M > 1$) and CyberWheel ($M = 1$), c is instantiated as a heuristic agent such as a finite state machine, meaning blue agents are evaluated against a non-adaptive adversary[2]. This fixity leaves critical vulnerabilities in learned blue policies undetected. Since c is executable code that modern LLMs can naturally generate, we propose training an LLM meta-policy π_θ to directly synthesize c conditioned on execution feedback from prior trajectories.

2.2 STRATEGIC RED TEAMING AS A CONTEXTUAL BANDIT

Directly applying π_θ within the step-by-step POMDP loop is intractable due to prohibitive computational cost (Castro et al., 2025). To circumvent this, Trident adopts a **Code-as-Policies** (Liang et al., 2023) paradigm: rather than sequentially emitting discrete primitive actions, π_θ expresses its entire strategy as executable source code that captures complex control flows and semantic reasoning. Specifically, π_θ receives a context C consisting of environment descriptions and historical feedback, and generates a complete, executable Python policy $c \in \mathcal{C}$ in a single forward pass $c \sim \pi_\theta(c \mid C)$. Since π_θ generates c in a single step and receives a reward only upon full execution of c within $\mathcal{M}$, this naturally reduces to a **Contextual Bandit** (Ahmadian et al., 2024; Wang et al., 2026a; Cobbe et al., 2021; Shao et al., 2024b) problem. The training objective is to maximize $J(\theta)$, the expected cumulative return when c is executed within $\mathcal{M}$:

$$J(\theta) = \mathbb{E}_{c \sim \pi_\theta(\cdot|C)} \left[\mathbb{E}_{\tau \sim \pi_c} \left[\sum_{t=0}^{T} \mathcal{R}(s_t, a_t) \right] \right] \tag{1}$$

[2]CyberWheel also supports DRL-based red agents; however, we fix c to a static heuristic in both environments for experimental consistency with CAGE 4.

The inner expectation $\mathbb{E}_{\tau\sim\pi_c}$ reflects multiple layers of stochasticity: even for a deterministic c, the resulting trajectory τ remains stochastic due to the non-deterministic transition dynamics $\mathcal{T}$ and the probabilistic interventions of the active DRL defender. Note that, for simplicity, in multi-agent red team settings, c is deployed uniformly across all red agents such that they share identical behavioral logic throughout each episode.

2.3 Optimization via GRPO

To optimize the meta-agent π_θ under this formulation, we utilize Group Relative Policy Optimization (GRPO) Shao et al. (2024b) as a surrogate objective for Equation 1. The GRPO objective and the associated relative advantage are defined as:

$$\mathcal{J}_{\text{GRPO}}(\theta) = \mathbb{E}\left[\frac{1}{G}\sum_{i=1}^{G}\min\left(\frac{\pi_\theta(c_i|C)}{\pi_{\theta_{\text{old}}}(c_i|C)}A_i, \text{clip}\left(\frac{\pi_\theta(c_i|C)}{\pi_{\theta_{\text{old}}}(c_i|C)}, 1-\epsilon, 1+\epsilon\right)A_i\right) - \beta\mathbb{D}_{\text{KL}}(\pi_\theta\|\pi_{\text{ref}})\right], A_i = \frac{r_i - \text{mean}(r_1,\ldots,r_G)}{\text{std}(r_1,\ldots,r_G)} \tag{2}$$

For each context C, GRPO samples a group of red policies $\{c_1,\ldots,c_G\}$ from the old policy $\pi_{\theta_{\text{old}}}$. The advantage A_i is computed by normalizing the verifiable rewards r_i relative to the group performance, where β and ϵ are hyperparameters for stability, and π_{ref} is a frozen reference model to ensure training stability. This objective allows the meta-agent to iteratively refine its strategic reasoning, discovering effective attack logics against adaptive DRL defenders through direct environmental feedback.

3 Trident

Optimizing $J(\theta)$ via GRPO requires a distribution of high-quality red-blue interaction trajectories to sample contexts C from. However, no existing resource provides such trajectories against active DRL defenders. Trident addresses this by providing both the dataset and the benchmark infrastructure needed to realize RLVR training in this setting.

3.1 Dataset preparation

As summarized in Table 1, existing LLM-driven cybersecurity benchmarks lack long-horizon adversarial dynamics. Large-scale datasets predominantly provide static security corpora (Yu et al., 2025) or QA pairs for supervised fine-tuning (Suryanto et al., 2026; Yu et al., 2025; Alam et al., 2024; Wan et al., 2024), while interactive benchmarks focus primarily on isolated CTF (Zhuo et al., 2026; Yang et al., 2023) or single-target Common Vulnerabilities and Exposures (CVE) exploitation (Zhu et al., 2025). None of these capture the continuous, multi-step interactions required to defeat active network defenders. To fill this gap, we construct $\mathcal{D}$, a large-scale dataset of over 13,000 high-fidelity red-blue interaction trajectories spanning CAGE 4 and CyberWheel, whose construction is detailed below.

Table 1: Comparison of **Trident** Dataset against existing benchmarks. By focusing on core capabilities, Trident uniquely provides a massive-scale dataset combined with an interactive environment supporting Agentic evaluation, RLVR, and Active DRL defenses.

Benchmark	Task	Scale (Train / Total)	Agentic?	RLVR?	Active Defense?
NYU CTF Shao et al. (2024a)	CTF	- / 200	✓	✗	✗
InterCode-CTF Yang et al. (2023)	CTF	- / 100	✓	✗	✗
Cybench Zhang et al. (2025a)	CTF	- / 40	✓	✗	✗
Cyber-Zero Zhuo et al. (2026)	CTF	6,188 / 6,188	✓	✗	✗
CVE-bench Zhu et al. (2025)	CVE Exploitation	- / 509	✓	✗	✗
CyberSecEval 3 Wan et al. (2024)	Security Eval.	- / 6	✓	✗	✗
CTI-bench Alam et al. (2024)	Security QA	2,500 / 2,500	✗	✗	✗
RedSage Suryanto et al. (2026)	Security QA	266K / 266K	✗	✗	✗
PRIMUS Yu et al. (2025)	Security Corpus	674K / 674K	✗	✗	✗
CyberBattleSim Team. (2021)	Network Attack	N/A	✗	✗	✗
CybORG CAGE 4 Kiely et al. (2025)	Network Defense	N/A	✗	✗	✓
CyberWheel Oesch et al. (2024)	Network Attack/Defense	N/A	✗	✗	✓
Trident (Ours)	**Network Attack**	**13,056+ / 13,744+**	✓	✓	✓

Sandbox infrastructure : To enable high-concurrency parallel rollouts, we decouple policy execution from the simulation runtime via a client-server architecture. Trident operates on the client side, transmitting generated Python policies to isolated, containerized sandbox servers that execute them natively against live DRL defenders and return verifiable execution logs with deterministic reward

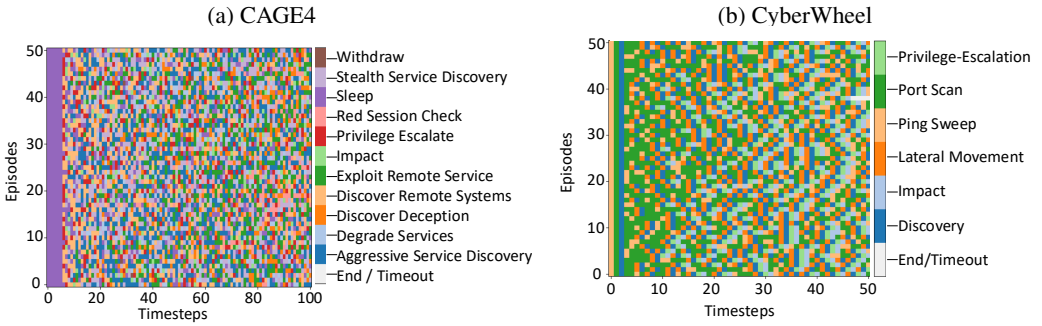

Figure 2: **Tactical branching in CybORG and CyberWheel.** Each row represents one of 50 sampled trajectories from a baseline heuristic against a fixed DRL defender. This seed variation ensures diverse data for RLVR training.

signals. Each sandbox exposes a standardized RESTful interface, allowing concurrent rollouts across independent simulator instances without shared state, guaranteeing reproducibility and integrity of $\mathcal{D}$.

Trajectory rollout: We leverage a diverse suite of N red agent against a fixed DRL blue defender (Details in Appendix C), iteratively sampling a heuristic $c_j \sim \text{Uniform}(1, N)$ and a randomized seed $s_i \sim \mathcal{S}$ to generate trajectories, collected as:

$$\mathcal{D} = \left\{ \tau_i' = \left(o_t^{(i)}, a_t^{(i)} \right)_{t=0}^{T} \sim P(\cdot \mid c_j, s_i) : 1 \leq i \leq n, \;\; j \sim \text{Uniform}(1, N) \right\} \tag{3}$$

This rollout process yields 3,556 training and 188 validation trajectories for CyberWheel, and 9,500 training and 500 validation trajectories for CAGE 4, totaling 13,744 samples. As the framework supports arbitrary seed randomization, additional trajectories can be generated on demand. As illustrated in Figure 2, even fixed heuristic logic produces significant tactical branching across seeds and defender interventions, capturing the complexity of the adversarial landscape.

Natural language log serialization: Each trajectory τ_i' is serialized into a structured natural language logs via a fixed, deterministic template. Each timestep encodes the agent's partial observation, differential host intelligence updates, and executed actions at every step t, followed by an episode-level impact summary, as illustrated in Figure 14. This prompt dataset bootstraps the meta-agent with a structural map of the adversarial landscape. A more detailed analysis of the Trident Dataset, including its coverage, diversity, outcome distribution, redundancy, and train–validation consistency, is provided in Appendix F.

3.2 Trident Agentic

Trident Agentic decomposes the attack agent into three specialized modules, each with a distinct role: compressing observations, reasoning about strategy, and generating executable code. Following prior multi-agent designs (Li et al., 2026; Feng et al., 2026; Wu et al., 2024; Chen et al., 2024), this separation of concerns isolates the RL optimization signal to a single trainable strategic component, while keeping the remaining modules frozen.

Log summarizer: A hybrid module that uses a rule-based parser combined with a frozen LLM to compress a prohibitive amount of serialized logs (Figure 14). To ensure the input remains within token limits without losing critical information, the parser identifies and preserves all

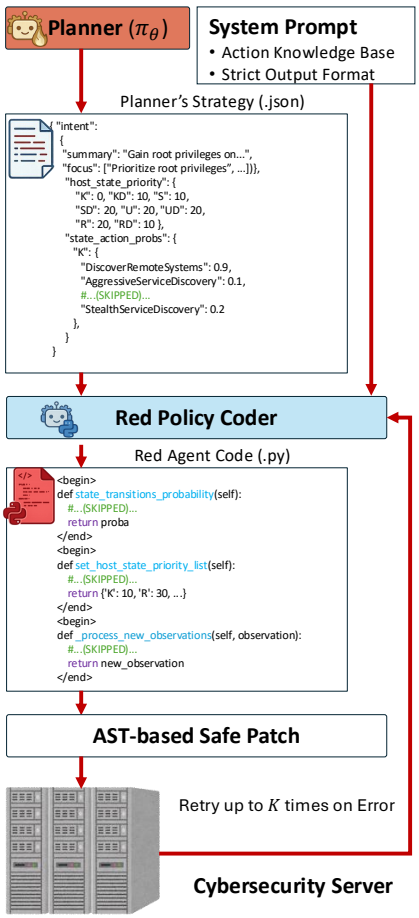

Figure 3: Planner-Coder Pipeline.

"milestone" steps—such as successful exploits, privilege escalations, or service impacts—while randomly sampling from failed attempts and reconnaissance noise. This heuristic filtering ensures that the most informative parts of the attack trajectory are always presented to the LLM. The model then synthesizes these filtered logs into a dense semantic summary for the Planner.

Planner (π_θ): The sole trainable component of our architecture. Conditioned on the condensed summary as its partially observable state, it formulates abstract tactical logic, outputting a structured JSON payload as its strategic action. By delegating exact programming syntax to the Coder, the Planner focuses exclusively on high-level decision-making, smoothing the optimization landscape. Its weights are updated directly via GRPO (Equation 2) driven by verifiable execution outcomes, ensuring the policy learns generalized attack strategies strictly from observable feedback without any privileged environment access.

Coder: As illustrated in Figure 3, a frozen model translates the Planner's strategic intent into executable Python code. Following recent advancements in robust LLM-driven code generation (Zhu et al., 2026a; Wang et al., 2026b), we incorporate an Abstract Syntax Tree (AST)-based patching mechanism to verify structural integrity and ensure execution stability before deployment.

Furthermore, an iterative self-repair loop is employed to mitigate minor syntax errors that could disrupt the learning. A sandbox performs a fast-pass sanity check; if an execution error occurs, the traceback is provided to the Coder for correction. Permitting up to K(=3) repair attempts absorbs trivial syntax errors, ensuring the Planner is not penalized for incidental implementation noise. However, if the *Coder* exhausts all retries, this indicates the Planner's strategy is fundamentally unimplementable, and the resulting execution penalty is intentionally applied.

3.3 Domain-grounded prompt design

A critical challenge in applying LLMs to cybersecurity is preventing hallucinated actions and ensuring the agent operates strictly within the observable state. To address this, each module in Trident is governed by heavily engineered, domain-grounded prompts detailed in Appendix G. First, we inject an action knowledge base directly into the system prompts, providing the LLMs with explicit prerequisites and expected outcomes for complex killchain actions. Second, we enforce Strict API Boundaries. The Coder and Planner are explicitly forbidden from accessing unobservable simulation internals (e.g., querying decoy flags). They are constrained to use only legal APIs, forcing them to rely on behavioral inference rather than omniscient state access. Finally, to guarantee deterministic pipeline execution, we utilize rigid output schemas—forcing the Planner to generate abstract strategy JSON, and constraining the Coder to output exactly delineated Python blocks.

3.4 Reward design

To accurately evaluate the generated algorithmic policy c, we design a dense, verifiable reward function derived directly from the execution logs of each trajectory τ. Let $\phi(e, o_{0:T})$ count occurrences of event e across the observation sequence of length T, and let $\mathcal{E}^+$, $\mathcal{E}^-$ denote positive milestone events (e.g., system impacts, privilege escalations) and negative events (e.g., failed exploits, invalid actions), respectively. We define the tactical reward components as:

$$R^+ = \sum_{e \in \mathcal{E}^+} q_e \phi(e, o_{0:T}) + \frac{r}{T} \sum_{e \in \mathcal{E}^+} \phi(e, o_{0:T}), \qquad R^- = \sum_{e \in \mathcal{E}^-} s_e \phi(e, o_{0:T}) \tag{4}$$

where q_e is a per-event weight reflecting the relative importance of each positive milestone, r is a scalar efficiency coefficient that normalizes total achievements by episode length T, and s_e is a per-event penalty weight controlling the severity of each negative event. Thus R^+ accumulates weighted killchain milestones augmented by an efficiency bonus, and R^- penalizes wasteful behavior.

Crucially, to penalize syntactically invalid or non-executable logic, the final episodic reward is determined via a conditional branch. If the generated policy c fails to execute due to unrecoverable syntax or runtime errors, the episode is aborted and a strict execution penalty p is applied. Otherwise, the policy receives a tactical reward clipped within $[R_{\min}, R_{\max}]$:

$$R(c, o_{0:T}) = \begin{cases} p & \text{if } c \text{ fails to execute or violates constraints} \\ \text{clip}(R^+ - R^-,\ R_{\min},\ R_{\max}) & \text{otherwise} \end{cases} \tag{5}$$

Full specifications of $\mathcal{E}^+$, $\mathcal{E}^-$, and p for each environment are provided in Appendix B.1.3, Appendix B.2.3, and Appendix B.3, respectively.

4 EXPERIMENTS

Table 2: Performance comparison across CAGE4 and CyberWheel settings. Lower returns indicate worse blue-agent performance. Values are reported as mean ± standard deviation. Parenthesized values denote the relative degradation (Δ) from the Default Agent; positive values indicate stronger attacks. The final column averages Δ over the four CAGE4 defenders only.

Method	CAGE4				CyberWheel	Average CAGE4 Degradation (Δ)
	GNN King et al. (2025)	HMARL3 Singh et al. (2025)	HMARL4 Singh et al. (2025)	MARL Lyu et al. (2021)	PPO (NN) Oesch et al. (2024)	
Default Agent	-202.34 ± 1.42	-132.63 ± 0.87	-940.87 ± 15.75	-171.09 ± 4.32	4574.59 ± 3.65	–
GPT-4o Hurst et al. (2024)	-351.43 ± 65.30 (73.68%)	-191.56 ± 42.17 (44.44%)	-844.99 ± 196.51 (-10.19%)	-323.91 ± 9.85 (89.33%)	1919.31 ± 149.59 (58.04%)	49.32%
Zero-shot	-498.74 ± 65.85 (146.48%)	-598.09 ± 26.15 (350.96%)	-1701.01 ± 113.20 (80.79%)	-435.84 ± 29.82 (154.75%)	2107.84 ± 59.64 (53.92%)	183.25%
CoT	-367.19 ± 19.58 (81.47%)	-200.09 ± 22.76 (50.87%)	-737.94 ± 171.87 (-21.57%)	-340.55 ± 42.79 (99.05%)	1729.16 ± 69.83 (62.20%)	52.46%
DRL Red	-311.11 ± 25.24 (53.76%)	-271.88 ± 9.47 (104.99%)	-314.52 ± 14.99 (-66.57%) †	-262.13 ± 13.27 (53.22%)	**-66.35 ± 27.52 (101.45%)**	36.35%
Trident Agentic (Ours)	**-2083.68 ± 110.89 (929.79%)**	**-849.85 ± 6.46 (540.78%)**	**-2241.72 ± 194.82 (138.26%)**	**-1716.73 ± 33.28 (903.43%)**	35.70 ± 48.89 (99.22%)	**628.07%**

To evaluate the effectiveness and robustness of Trident Agentic, we conduct our experiments within multiple environments with different victim models. We design our evaluation to answer two primary Research Questions (RQs): (**RQ1**) Can the algorithmic policy generated by Trident Agentic defeat state-of-the-art DRL blue agents where traditional heuristic attackers fail? (**RQ2**) Is RLVR training beneficial for generating effective attacks against DRL blue agents?

4.1 EXPERIMENTAL SETUP AND BASELINES

Implementation Details: We used Qwen2.5-7B-Instruct for our planner policy(π_θ), Qwen2.5-14B-Instruct-AWQ Yang et al. (2024); Team (2024) for summarizer and Qwen2.5-Coder-7B-Instruct Hui et al. (2024) for coder agent. We used a single NVIDIA-H200 GPU to host frozen LLM agents and 2 NVIDIA-H200 GPUs for training our planner policy. The group size of GRPO (G) is set to 8, learning rate of 5.0×10^{-7} and batch size of 2. The KL penalty term (β) is set to 0.05 for CyberWheel and 5.0×10^{-4} for CAGE4, respectively. The implementation was mainly built upon verl Sheng et al. (2025) for the RLVR training loop, utilizing vllm Kwon et al. (2023) for efficient policy generation and Ray Moritz et al. (2018) for distributed orchestration.

Environment Setup: Both CAGE 4 Kiely et al. (2025) and CyberWheel Oesch et al. (2024) model enterprise networks as adversarial POMDPs where a killchain-based red agent attacks hosts across subnets while a DRL blue agent defends in real-time — producing deterministic, verifiable reward signals from discrete-event logs. Despite this shared structure, they differ in network topology, defender architecture, and mission complexity. CAGE 4 features a multi-agent military network with dynamic mission phases and cooperative blue defenders, while CyberWheel grounds its red agent actions in real-world MITRE ATT&CK MITRE (2026) techniques, providing higher attack realism. We evaluate across both to test the generalizability of Trident. Full environment details are provided in Appendix B.

Victim Models: We evaluate the performance of Trident Agentic against three state-of-the-art DRL blue agents. "GNN" King et al. (2025); Kiely et al. (2025) represents the network as a temporal attributed graph processed by a graph convolutional network optimized by Proximal Policy Optimization (PPO) Schulman et al. (2017), with actions formulated as graph edits on nodes and edges, enabling zero-shot generalization to unseen topologies. "H-MARL" Singh et al. (2025) employs a hierarchical architecture optimized by PPO that decomposes cyber defense into sub-tasks, coordinated by either an expert-rule or a learned master policy across five blue agents. "MARL" Lyu et al. (2021) is a decentralized multi-agent PPO baseline where each blue agent maintains a separate actor-critic trained solely on its local observations. All three are evaluated in CAGE 4. For CyberWheel, we evaluate against "PPO(NN)" Oesch et al. (2024),

CAGE4: Dynamic State Priority Adaptation

[Default Code (CAGE4)]

```
def set_host_state_priority_list(self):
  return None # equal priority (default
agent)
```

[Planner's Thought]

Shift from discovery to exploitation. Prioritize high-value targets and increase focus on privilege escalation and root compromise.

[Coder's Implementation]

```
def set_host_state_priority_list(self):
  return {
    'K': 10, 'KD': 10, 'S': 10, 'SD': 10,
    'U': 20, 'UD': 20, 'R': 30, 'RD': 30
  }
```

Figure 4: CAGE4 policy generation example. Trident Agentic coder introduces a structured priority that shifts focus toward high-value targets.

a PPO-based blue agent from the official CyberWheel repository trained to minimize server downtime against a heuristic red agent.

Dataset and Training Split: For each target DRL blue agent, we trained a dedicated red policy (π_θ) via RLVR exclusively on the training split of the Trident Dataset. Specifically, policies evaluated in CAGE 4 are trained solely on the CAGE 4 training split, and those evaluated in CyberWheel are trained solely on the CyberWheel training split. The performance reported in our evaluation is then derived from the corresponding validation splits to ensure a rigorous assessment of generalization across unseen network states and defensive configurations.

Baselines: We compare Trident against a rule-based default agent, three alternative planner configurations that isolate the contribution of RLVR training, and a simulator-native DRL baseline that tests whether generic policy adaptation alone is sufficient. The "Default Agent" uses the default `FiniteStateRedAgent` from the CybORG CAGE 4 repository for CAGE 4 and the `BFSServerDowntime` strategy for CyberWheel. The "DRL Red (PPO)" baseline trains a separate PPO red agent against each frozen blue defender using the simulator-native observation and action interfaces. Unlike Trident, it directly learns a step-wise numerical policy without the Log Summarizer, Planner, Coder, or GRPO components. This baseline therefore distinguishes the benefit of Trident's agentic policy-generation framework from generic adaptation to a fixed defender. Its architecture, training budget, and evaluation details are provided in Appendix E.

The remaining LLM-based configurations vary only the planner. "Zero-shot" uses the same base model without task-specific training; "CoT" Kojima et al. (2022) uses an identical base model and prompt augmented with a chain-of-thought instruction; and "GPT-4o" uses GPT-4o Hurst et al. (2024) to evaluate the effect of frontier-model capacity without task-specific training. These LLM-based configurations use identical observations and evaluation protocols, and all prompts are identical except for the additional CoT instruction. The PPO baseline uses its native numerical interface but is evaluated using the same benchmark-specific protocol.

CyberWheel: Decoy Avoidance

[Default Red Agent]
Simple BFS prioritizing unimpacted servers.
[Planner's Thought]
"Based on the analysis, we can summarize the key issues with the current strategy: 1. Fixation on single hosts, leading to redundant actions. 2. Inefficient exploitation of the early window. 3. Lack of effective behavioral detection for decoys. We need to reformulate our strategy to ensure that the agent can efficiently reach and impact unimpacted servers within the early window and avoid fixating on already impacted or suspicious hosts. Additionally, we should develop a mechanism to detect and avoid decoys."
[Generated Logic (Python)]

```
class ImprovedRedStrategy:
  # Learned parameter avoiding decoys
  DECOY_THRESHOLD = 3
  def select_target(cls, agent):
    if history.svc_count == 0:
      cls.failures[host] += 1
      if cls.failures >= cls.THRESHOLD:
        return cls.select_next_target() # Abort
-------Skipped-------
```

Figure 5: Example generated logic for decoy avoidance in CyberWheel that abandons targets.

Evaluation Metrics: We adopt the blue agent's cumulative episodic reward as our primary metric, where a lower value indicates more effective red agent attacks, consistent with prior DRL defense benchmarks. This captures disruption to green agent activity across mission phases in CAGE 4, and defender success via decoys in CyberWheel. CAGE 4 results are reported at training step 150, averaged over 500 randomized episodes of 500 steps each (T=500 in Equation 1); CyberWheel results are reported at training step 1,000, averaged over 188 randomized episodes of 100 steps each. All baselines are reported as the mean and standard deviation across 3 independent runs. We note a deliberate asymmetry favoring baselines: GPT-4o, Zero-shot, and CoT metrics are computed *exclusively over successful executions*, filtering out fatal errors, while Trident's evaluation is strictly comprehensive, absorbing all execution failures and their associated penalties.

4.2 Results

Quantitative Analysis: Table 2 presents cumulative rewards against all defensive baselines. Heuristic red agents are weak evaluators: blue agents trained exclusively against them condition on their fixed behavioral signatures, leaving systematic blind spots invisible under current evaluation practices. Trident exploits these blind spots by dynamically synthesizing executable attack scripts tailored

to the evolving network state, degrading the four CAGE 4 defenders by an average of 628.07% and eliminating over 99% of the CyberWheel defender's reward, affirmatively answering **RQ1**. Notably, the Zero-shot SLM outperforms GPT-4o in several settings, which we attribute to prompt optimization bias: our scaffolding was iteratively developed around the 7B model, and correcting this disadvantage via extensive prompt-tuning was computationally infeasible over thousands of interactive DRL episodes. Despite this inadvertent advantage to the Zero-shot baseline, Trident Agentic still substantially surpasses it: static prompt engineering alone only partially compromises DRL defenders, while RLVR training provides a decisive advantage, answering **RQ2**.

Qualitative Analysis: Figure 4 and Figure 5 illustrate emergent attack behaviors autonomously discovered by Trident Agentic through RLVR. In CyberWheel, Trident Agentic identifies the baseline agent's naive BFS strategy as vulnerable to honeypots and autonomously synthesizes a decoy-avoidance mechanism that aborts attacks on suspicious hosts after repeated failures. In CAGE 4, rather than treating all network states equally, Trident Agentic learns to prioritize high-value states such as Root and User access over basic reconnaissance states. These behaviors confirm that the "Code-as-Policies" paradigm provides the expressivity necessary to discover novel, environment-specific attack vectors that static heuristics cannot represent.

Ablation Studies: Due to space constraints, detailed ablation studies regarding the reward formulation and the design choices of our agentic components are deferred to Appendix D.

5 Limitations

While Trident Agentic demonstrates exceptional efficacy against state-of-the-art DRL defenses, two engineering choices limit out-of-the-box generalization. First, the framework depends heavily on domain-specific prompt engineering: the tripartite architecture relies on hand-crafted templates and rule-based log parsers, meaning adaptation to new simulation engines requires manual recalibration of these scaffolds. Second, the reward design remains environment-dependent; although the underlying concepts are universal, their implementation is tied to the specific log formats and action spaces of CAGE 4 and CyberWheel. Future work will focus on automating log parsing and reward discovery to move toward a truly zero-shot, adaptive red agent.

6 Broader impacts and ethical considerations

This work demonstrates that sophisticated offensive capabilities against state-of-the-art DRL defenses can be realized with compact, open-weight SLMs, introducing a dual-use risk that will grow as simulators more closely approximate real-world networks. Assuming adversaries will not independently develop these methods is not a viable security posture. We argue that transparency is preferable to obscurity: by openly demonstrating Trident Agentic's capabilities, we expose a critical evaluation gap and call on the community to develop equally adaptive, RL-driven blue agents. This framework provides the automated adversaries and verifiable metrics needed to train more robust cyber defenses before these attack strategies are independently deployed.

7 Conclusions

In this paper, we introduced **Trident**, a dynamic red teaming benchmark comprising isolated sandbox servers, red-blue interaction datasets, and **Trident Agentic**, a tripartite RLVR agentic architecture to reframe autonomous red teaming as a contextual bandit problem. Our evaluations reveal a consistent brittleness in existing autonomous defenders: a single trainable 7B LLM planner achieves an average **628%** reduction in blue-agent performance in CAGE 4 and erases over **99%** of defender reward in CyberWheel, while autonomously discovering novel behaviors such as decoy avoidance and adaptive state prioritization. Beyond its immediate contributions, Trident opens several exciting future directions: its sandbox infrastructure naturally supports a self-improving curriculum where trained red agent policies generate richer adversarial trajectories to bootstrap progressively stronger attackers, whose pressure drives more robust blue agents in a game-theoretic co-evolution loop, while its RESTful design invites integration with higher-fidelity environments toward real-world autonomous cyber defense.

AI USE STATEMENT

Generative AI tools, including ChatGPT, Claude, and Gemini, were used to assist with software implementation and debugging, drafting and editing portions of the manuscript, and creating or refining graphical assets such as icons and pictograms. In particular, generative AI provided assistance with portions of the experimental implementation, while all generated or modified code was reviewed, tested, and integrated by the authors. The research questions, core methodological ideas, experimental design, interpretation of results, and final editorial decisions were developed and made by the authors. All AI-assisted text, code, and graphical content was reviewed by the authors, who take responsibility for the final content of this work.

REPRODUCIBILITY STATEMENT

To support reproducibility, we provide the source code and experimental artifacts in an anonymous repository at https://anonymous.4open.science/r/Trident-A934, together with the dataset used in this work at https://huggingface.co/datasets/ALG0603/Trident. The manuscript and appendix report the model architectures, training procedures, hyperparameters, environment settings, evaluation protocols, and computational resources used in our experiments. Random seeds and aggregation procedures are also specified where applicable.

ETHICS STATEMENT

This work involves dual-use considerations associated with automated offensive cybersecurity research. We discuss the associated risks, intended defensive applications, and broader ethical considerations in Section.6.

REFERENCES

Talor Abramovich, Meet Udeshi, Minghao Shao, Kilian Lieret, Haoran Xi, Kimberly Milner, Sofija Jancheska, John Yang, Carlos E Jimenez, Farshad Khorrami, Prashanth Krishnamurthy, Brendan Dolan-Gavitt, Muhammad Shafique, Karthik R Narasimhan, Ramesh Karri, and Ofir Press. EnIGMA: Interactive tools substantially assist LM agents in finding security vulnerabilities. In *Forty-second International Conference on Machine Learning*, 2025. URL https://openreview.net/forum?id=Of3wZhVv1R.

Arash Ahmadian, Chris Cremer, Matthias Gallé, Marzieh Fadaee, Julia Kreutzer, Olivier Pietquin, Ahmet Üstün, and Sara Hooker. Back to basics: Revisiting REINFORCE-style optimization for learning from human feedback in LLMs. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 12248–12267, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.662. URL https://aclanthology.org/2024.acl-long.662/.

Md Tanvirul Alam, Le Nguyen, Dipkamal Bhusal, and Nidhi Rastogi. Ctibench: a benchmark for evaluating llms in cyber threat intelligence. In *Proceedings of the 38th International Conference on Neural Information Processing Systems*, NIPS '24, Red Hook, NY, USA, 2024. Curran Associates Inc. ISBN 9798331314385.

Anthropic. Claude mythos preview. https://red.anthropic.com/2026/mythos-preview/, 2026. [Accessed 17-04-2026].

Manish Bhatt, Sahana Chennabasappa, Cyrus Nikolaidis, Shengye Wan, Ivan Evtimov, Dominik Gabi, Daniel Song, Faizan Ahmad, Cornelius Aschermann, Lorenzo Fontana, et al. Purple llama cyberseceval: A secure coding benchmark for language models. *arXiv preprint arXiv:2312.04724*, 2023.

Manish Bhatt, Sahana Chennabasappa, Yue Li, Cyrus Nikolaidis, Daniel Song, Shengye Wan, Faizan Ahmad, Cornelius Aschermann, Yaohui Chen, Dhaval Kapil, et al. Cyberseceval 2: A wide-ranging cybersecurity evaluation suite for large language models. *arXiv preprint arXiv:2404.13161*, 2024.

Blake Bullwinkel, Eugenia Kim, Amanda Minnich, and Mark Russinovich. Learning to attack and defend: Adaptive red teaming of language models via grpo. *arXiv preprint arXiv:2606.09701*, 2026.

Sebastián R Castro, Roberto Campbell, Nancy Lau, Octavio Villalobos, Jiaqi Duan, and Alvaro A Cardenas. Large language models are autonomous cyber defenders. In *2025 IEEE Conference on Artificial Intelligence (CAI)*, pp. 1125–1132. IEEE, 2025.

Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chi-Min Chan, Heyang Yu, Yaxi Lu, Yi-Hsin Hung, Chen Qian, Yujia Qin, Xin Cong, Ruobing Xie, Zhiyuan Liu, Maosong Sun, and Jie Zhou. Agentverse: Facilitating multi-agent collaboration and exploring emergent behaviors. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=EHg5GDnyq1.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.

DARPA. Darpa's artificial intelligence cyber challenge (aixcc). https://aicyberchallenge.com/overview/, 2026. [Accessed 17-04-2026].

Gelei Deng, Yi Liu, Víctor Mayoral-Vilches, Peng Liu, Yuekang Li, Yuan Xu, Tianwei Zhang, Yang Liu, Martin Pinzger, and Stefan Rass. PentestGPT: Evaluating and harnessing large language models for automated penetration testing. In *33rd USENIX Security Symposium (USENIX Security 24)*, pp. 847–864, Philadelphia, PA, August 2024. USENIX Association. ISBN 978-1-939133-44-1. URL https://www.usenix.org/conference/usenixsecurity24/presentation/deng.

Richard Fang, Rohan Bindu, Akul Gupta, and Daniel Kang. Llm agents can autonomously exploit one-day vulnerabilities. *arXiv preprint arXiv:2404.08144*, 2024a.

Richard Fang, Rohan Bindu, Akul Gupta, Qiusi Zhan, and Daniel Kang. Llm agents can autonomously hack websites. *arXiv preprint arXiv:2402.06664*, 2024b.

Tao Feng, Haozhen Zhang, Peixuan Han, Zijie Lei, and Jiaxuan You. Graphplanner: Graph-based agentic routing for LLMs. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=ZdGB7MNQDT.

Chengquan Guo, Xun Liu, Chulin Xie, Andy Zhou, Yi Zeng, Zinan Lin, Dawn Song, and Bo Li. Redcode: Risky code execution and generation benchmark for code agents. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024. URL https://openreview.net/forum?id=mAG68wdggA.

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

Andreas Happe and Jürgen Cito. Getting pwn'd by ai: Penetration testing with large language models. In *Proceedings of the 31st ACM joint european software engineering conference and symposium on the foundations of software engineering*, pp. 2082–2086, 2023.

Andreas Happe, Aaron Kaplan, and Juergen Cito. Llms as hackers: Autonomous linux privilege escalation attacks. *Empirical Software Engineering*, 31(3):70, 2026.

Zhiwei He, Tian Liang, Jiahao Xu, Qiuzhi Liu, Xingyu Chen, Yue Wang, Linfeng Song, Dian Yu, Zhenwen Liang, Wenxuan Wang, Zhuosheng Zhang, Rui Wang, Zhaopeng Tu, Haitao Mi, and Dong Yu. Deepmath-103k: A large-scale, challenging, decontaminated, and verifiable mathematical dataset for advancing reasoning. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=kHB5Te5IWm.

Binyuan Hui, Jian Yang, Zeyu Cui, Jiaxi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Kai Dang, et al. Qwen2. 5-coder technical report. *arXiv preprint arXiv:2409.12186*, 2024.

Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024.

Guangyu Jiang, Mahdi Imani, Nathaniel D Bastian, and Tian Lan. Agentic ai for cyber defense: Llm-guided hierarchical multi-agent reinforcement learning. In *MILCOM 2025-2025 IEEE Military Communications Conference (MILCOM)*, pp. 1518–1523. IEEE, 2025.

Bowen Jin, Hansi Zeng, Zhenrui Yue, Jinsung Yoon, Sercan O Arik, Dong Wang, Hamed Zamani, and Jiawei Han. Search-r1: Training LLMs to reason and leverage search engines with reinforcement learning. In *Second Conference on Language Modeling*, 2025. URL https://openreview.net/forum?id=Rwhi91ideu.

Mitchell Kiely, Metin Ahiskali, Etienne Borde, Benjamin Bowman, David Bowman, Dirk Van Bruggen, KC Cowan, Prithviraj Dasgupta, Erich Devendorf, Ben Edwards, et al. Exploring the efficacy of multi-agent reinforcement learning for autonomous cyber defence: A cage challenge 4 perspective. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 28907–28913, 2025.

Isaiah J King, Benjamin Bowman, and H Howie Huang. Automated cyber defense with generalizable graph-based reinforcement learning agents. *arXiv preprint arXiv:2509.16151*, 2025.

Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35: 22199–22213, 2022.

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*, 2023.

Zhuofeng Li, Haoxiang Zhang, Seungju Han, Sheng Liu, Jianwen Xie, Yu Zhang, Yejin Choi, James Zou, and Pan Lu. In-the-flow agentic system optimization for effective planning and tool use. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=Mf5AleTUVK.

Jacky Liang, Wenlong Huang, Fei Xia, Peng Xu, Karol Hausman, Brian Ichter, Pete Florence, and Andy Zeng. Code as policies: Language model programs for embodied control. In *2023 IEEE International conference on robotics and automation (ICRA)*, pp. 9493–9500. IEEE, 2023.

Xinjie Lin, Gang Xiong, Gaopeng Gou, Zhen Li, Junzheng Shi, and Jing Yu. Et-bert: A contextualized datagram representation with pre-training transformers for encrypted traffic classification. In *Proceedings of the ACM Web Conference 2022*, pp. 633–642, 2022.

Zefang Liu, Jialei Shi, and John F Buford. Cyberbench: A multi-task benchmark for evaluating large language models in cybersecurity. In *AAAI 2024 Workshop on Artificial Intelligence for Cyber Security*, 2024.

Ziyu Liu, Zeyi Sun, Yuhang Zang, Xiaoyi Dong, Yuhang Cao, Haodong Duan, Dahua Lin, and Jiaqi Wang. Visual-rft: Visual reinforcement fine-tuning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 2034–2044, 2025.

Xueguang Lyu, Yuchen Xiao, Brett Daley, and Christopher Amato. Contrasting centralized and decentralized critics in multi-agent reinforcement learning. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '21, pp. 844–852, Richland, SC, 2021. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9781450383073.

Yecheng Jason Ma, William Liang, Guanzhi Wang, De-An Huang, Osbert Bastani, Dinesh Jayaraman, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Eureka: Human-level reward design via coding large language models. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=IEduRUO55F.

Shae McFadden, Myles Foley, Elizabeth Bates, Ilias Tsingenopoulos, Sanyam Vyas, Vasilios Mavroudis, Chris Hicks, and Fabio Pierazzi. SoK: The pitfalls of deep reinforcement learning for cybersecurity. In *35th USENIX Security Symposium (USENIX Security 26)*, Baltimore, MD, August 2026. USENIX Association. URL https://www.usenix.org/conference/usenixsecurity26/presentation/mcfadden.

MITRE. Mitre att&ck. https://attack.mitre.org/, 2026. [Accessed 20-04-2026].

Philipp Moritz, Robert Nishihara, Stephanie Wang, Alexey Tumanov, Richard Liaw, Eric Liang, Melih Elibol, Zongheng Yang, William Paul, Michael I Jordan, et al. Ray: A distributed framework for emerging {AI} applications. In *13th USENIX symposium on operating systems design and implementation (OSDI 18)*, pp. 561–577, 2018.

Peter G Neumann and Phillip A Porras. Experience with emerald to date. In *Workshop on Intrusion Detection and Network Monitoring*, pp. 73–80, 1999.

Sean Oesch, Amul Chaulagain, Brian Weber, Matthew Dixson, Amir Sadovnik, Benjamin Roberson, Cory Watson, and Phillipe Austria. Towards a high fidelity training environment for autonomous cyber defense agents. In *Proceedings of the 17th Cyber Security Experimentation and Test Workshop*, CSET '24, pp. 91–99, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400709579. doi: 10.1145/3675741.3675752. URL https://doi.org/10.1145/3675741.3675752.

Red Canary. Atomic red team. https://www.atomicredteam.io/, 2026. [Accessed 20-04-2026].

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

Minghao Shao, Sofija Jancheska, Meet Udeshi, Brendan Dolan-Gavitt, haoran xi, Kimberly Milner, Boyuan Chen, Max Yin, Siddharth Garg, Prashanth Krishnamurthy, Farshad Khorrami, Ramesh Karri, and Muhammad Shafique. Nyu ctf bench: A scalable open-source benchmark dataset for evaluating llms in offensive security. In *Advances in Neural Information Processing Systems*, volume 37, pp. 57472–57498, 2024a. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/69d97a6493fbf016fff0a751f253ad18-Paper-Datasets_and_Benchmarks_Track.pdf.

Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024b.

Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. Hybridflow: A flexible and efficient rlhf framework. In *Proceedings of the Twentieth European Conference on Computer Systems*, EuroSys '25, pp. 1279–1297, New York, NY, USA, 2025. Association for Computing Machinery. ISBN 9798400711961. doi: 10.1145/3689031.3696075. URL https://doi.org/10.1145/3689031.3696075.

Aditya Vikram Singh, Ethan Rathbun, Emma Graham, Lisa Oakley, Simona Boboila, Peter Chin, and Alina Oprea. Hierarchical multi-agent reinforcement learning for cyber network defense. *Reinforcement Learning Journal*, 6:790–810, 2025.

Natalia Stakhanova, Samik Basu, and Johnny Wong. A taxonomy of intrusion response systems. *Int. J. Inf. Comput. Secur.*, 1(1/2):169–184, January 2007. ISSN 1744-1765. doi: 10.1504/IJICS.2007.012248. URL https://doi.org/10.1504/IJICS.2007.012248.

Naufal Suryanto, Muzammal Naseer, Pengfei Li, Syed Talal Wasim, Jinhui Yi, Juergen Gall, Paolo Ceravolo, and Ernesto Damiani. Redsage: A cybersecurity generalist LLM. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=W4FAenIrQ2.

Microsoft Defender Research Team. Cyberbattlesim, 2021. Created by Christian Seifert, Michael Betser, William Blum, James Bono, Kate Farris, Emily Goren, Justin Grana, Kristian Holsheimer, Brandon Marken, Joshua Neil, Nicole Nichols, Jugal Parikh, Haoran Wei.

Qwen Team. Qwen2.5: A party of foundation models, September 2024. URL https://qwenlm.github.io/blog/qwen2.5/.

Shengye Wan, Cyrus Nikolaidis, Daniel Song, David Molnar, James Crnkovich, Jayson Grace, Manish Bhatt, Sahana Chennabasappa, Spencer Whitman, Stephanie Ding, et al. Cyberseceval 3: Advancing the evaluation of cybersecurity risks and capabilities in large language models. *arXiv preprint arXiv:2408.01605*, 2024.

Tianyi Wang, Yixia Li, Long Li, Yibiao Chen, Shaohan Huang, Yun Chen, Peng Li, Yang Liu, and Guanhua Chen. Sppo: Sequence-level ppo for long-horizon reasoning tasks. *arXiv preprint arXiv:2604.08865*, 2026a.

Yiting Wang, Wanghao Ye, Ping Guo, Yexiao He, Ziyao Wang, Bowei Tian, Shwai He, Guoheng Sun, Zheyu Shen, Sihan Chen, Ankur Srivastava, Qingfu Zhang, Gang Qu, and Ang Li. SymRTLO: Enhancing RTL code optimization with LLMs and neuron-inspired symbolic reasoning. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2026b. URL https://openreview.net/forum?id=XFP5jntOdx.

Zhun Wang, Tianneng Shi, Jingxuan He, Matthew Cai, Jialin Zhang, and Dawn Song. Cybergym: Evaluating AI agents' real-world cybersecurity capabilities at scale. In *The Fourteenth International Conference on Learning Representations*, 2026c. URL https://openreview.net/forum?id=2YvbLQEdYt.

Yuxiang Wei, Olivier Duchenne, Jade Copet, Quentin Carbonneaux, LINGMING ZHANG, Daniel Fried, Gabriel Synnaeve, Rishabh Singh, and Sida Wang. SWE-RL: Advancing LLM reasoning via reinforcement learning on open software evolution. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=ULblO61XZ0.

Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang, Shaokun Zhang, Jiale Liu, Ahmed Hassan Awadallah, Ryen W White, Doug Burger, and Chi Wang. Autogen: Enabling next-gen LLM applications via multi-agent conversations. In *First Conference on Language Modeling*, 2024. URL https://openreview.net/forum?id=BAakY1hNKS.

An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zhihao Fan. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2024.

John Yang, Akshara Prabhakar, Karthik Narasimhan, and Shunyu Yao. Intercode: Standardizing and benchmarking interactive coding with execution feedback. *Advances in Neural Information Processing Systems*, 36:23826–23854, 2023.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. In *The eleventh international conference on learning representations*, 2022.

Yao-Ching Yu, Tsun-Han Chiang, Cheng-Wei Tsai, Chien-Ming Huang, and Wen-Kwang Tsao. Primus: A pioneering collection of open-source datasets for cybersecurity LLM training. In Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (eds.), *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 10391–10413, Suzhou, China, November 2025. Association for Computational Linguistics. ISBN 979-8-89176-332-6. doi: 10.18653/v1/2025.emnlp-main.527. URL https://aclanthology.org/2025.emnlp-main.527/.

Huaye Zeng, Dongfu Jiang, Haozhe Wang, Ping Nie, Xiaotong Chen, and Wenhu Chen. ACECODER: Acing coder RL via automated test-case synthesis. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 12023–12040, Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 979-8-89176-251-0. doi: 10.18653/v1/2025.acl-long.587. URL https://aclanthology.org/2025.acl-long.587/.

Andy K Zhang, Neil Perry, Riya Dulepet, Joey Ji, Celeste Menders, Justin W Lin, Eliot Jones, Gashon Hussein, Samantha Liu, Donovan Julian Jasper, Pura Peetathawatchai, Ari Glenn, Vikram Sivashankar, Daniel Zamoshchin, Leo Glikbarg, Derek Askaryar, Haoxiang Yang, Aolin Zhang, Rishi Alluri, Nathan Tran, Rinnara Sangpisit, Kenny O Oseleononmen, Dan Boneh, Daniel E. Ho, and Percy Liang. Cybench: A framework for evaluating cybersecurity capabilities and risks of language models. In *The Thirteenth International Conference on Learning Representations*, 2025a. URL https://openreview.net/forum?id=tc90LV0yRL.

Chiyu Zhang, Marc-Alexandre Côté, Michael Albada, Anush Sankaran, Jack W Stokes, Tong Wang, Amir H. Abdi, William Blum, and Muhammad Abdul-Mageed. Defenderbench: A toolkit for evaluating language agents in cybersecurity environments. In *Workshop on Scaling Environments for Agents*, 2025b. URL https://openreview.net/forum?id=BtZBwwjO5d.

Ruijie Zhao, Mingwei Zhan, Xianwen Deng, Yanhao Wang, Yijun Wang, Guan Gui, and Zhi Xue. Yet another traffic classifier: A masked autoencoder based traffic transformer with multi-level flow representation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 5420–5427, 2023.

Yaqin Zhou, Shangqing Liu, Jingkai Siow, Xiaoning Du, and Yang Liu. *Devign: effective vulnerability identification by learning comprehensive program semantics via graph neural networks*. Curran Associates Inc., Red Hook, NY, USA, 2019.

Derui Zhu, Dingfan Chen, jinfu chen, Jens Grossklags, Walter Pretschner, and Weiyi Shang. More than just functional: LLM-as-a-critique for efficient code generation. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2026a. URL https://openreview.net/forum?id=0Zri6HSYaK.

Yuxuan Zhu, Antony Kellermann, Dylan Bowman, Philip Li, Akul Gupta, Adarsh Danda, Richard Fang, Conner Jensen, Eric Ihli, Jason Benn, Jet Geronimo, Avi Dhir, Sudhit Rao, Kaicheng Yu, Twm Stone, and Daniel Kang. CVE-bench: A benchmark for AI agents' ability to exploit real-world web application vulnerabilities. In *Forty-second International Conference on Machine Learning*, 2025. URL https://openreview.net/forum?id=3pk0p4NGmQ.

Yuxuan Zhu, Antony Kellermann, Akul Gupta, Philip Li, Richard Fang, Rohan Bindu, and Daniel Kang. Teams of LLM agents can exploit zero-day vulnerabilities. In Vera Demberg, Kentaro Inui, and Lluís Marquez (eds.), *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 23–35, Rabat, Morocco, March 2026b. Association for Computational Linguistics. ISBN 979-8-89176-380-7. doi: 10.18653/v1/2026.eacl-long.2. URL https://aclanthology.org/2026.eacl-long.2/.

Terry Yue Zhuo, Dingmin Wang, Hantian Ding, Varun Kumar, and Zijian Wang. Cyber-zero: Training cybersecurity agents without runtime. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=1gRTeAik4G.

APPENDIX

A RELATED WORKS

A.1 CYBERSECURITY AND LLMS

LLM-based cybersecurity research has evolved from static evaluations, including vulnerability detection (Zhou et al., 2019; Zhu et al., 2025), threat intelligence QA (Alam et al., 2024), and

secure code evaluation (Bhatt et al., 2023; Wan et al., 2024; Bhatt et al., 2024; Guo et al., 2024), toward interactive agentic systems (Happe et al., 2026; Fang et al., 2024a; Zhu et al., 2026b; Fang et al., 2024b; Happe & Cito, 2023), driven by competitions such as DARPA AIxCC (DARPA, 2026; Liu et al., 2024) and evaluated on CTF challenges and sandboxes (Abramovich et al., 2025; Zhang et al., 2025a; Shao et al., 2024a; Wang et al., 2026c; Zhang et al., 2025b). However, these frameworks feature static targets without active defenders, while traditional DRL simulators (Kiely et al., 2025; Team., 2021) operate on numeric vector spaces unsuitable for LLM-native reasoning, and recent cybersecurity LLMs (Zhuo et al., 2026; Yu et al., 2025; Suryanto et al., 2026) rely on offline Supervised Fine Tuning (SFT) datasets that cannot support online RL against adaptive defenders. Trident bridges this gap by providing both an interactive text-native environment and a large-scale hybrid dataset for RLVR training, a need made urgent by the autonomous zero-day exploitation capabilities now demonstrated by frontier models such as Claude Mythos (Anthropic, 2026). The closest concurrent work to ours is (Bullwinkel et al., 2026), which also uses GRPO for adversarial red teaming. However, Trident trains executable code-as-policy agents against active DRL defenders in network-security environments with environment-verifiable rewards, while Bullwinkel et al. (2026) focuses on prompt-based jailbreaks against LLMs.

A.2 Code-as-Policy and RLVR in agentic LLMs

Modern agentic frameworks formulate complex tasks as POMDPs using step-by-step prompting strategies in ReAct (Yao et al., 2022)-based manner (Wu et al., 2024; Chen et al., 2024), where an agent interleaves reasoning traces with discrete tool calls to iteratively progress toward a goal. However, this approach is bottlenecked by the engineering overhead of designing universal tool wrappers across heterogeneous simulators and is vulnerable to error compounding in cybersecurity. The "Code-as-Policies" paradigm (Liang et al., 2023; Wei et al., 2025) overcomes these limitations by generating executable programs that embed complex control flows directly into actions, eliminating brittle tool APIs. Recent works have integrated LLMs into RL pipelines for agentic reward design (Ma et al., 2024) and Jiang et al. (2025) adapted this into cyber defense environments, but these deploy LLMs as auxiliary modules rather than as primary policy generators. In contrast, Trident Agentic adopts the "Code-as-Policies" paradigm with iterative self-repair (Zhu et al., 2026a; Wang et al., 2026b) to generate complete adversarial strategies directly. While RLVR has driven advancements in mathematics (Shao et al., 2024b; He et al., 2026) and competitive programming (Zeng et al., 2025), Trident is, to the best of our knowledge, the first to successfully adapt it for training "Code-as-Policies" adversarial agents in cybersecurity.

B Environment Setup & Reward Design

This section provides a comprehensive breakdown of the experimental testbeds, victim models, and reward formulations used to train and evaluate Trident Agentic. We deploy Trident Agentic across two distinct cybersecurity simulators: CybORG CAGE 4 (Appendix B.1) and CyberWheel (Appendix B.2). For each environment, we detail the underlying network topology, the architecture of the victim DRL defenders (blue agents), and the baseline heuristic attackers (red agents). Crucially, we explain the exact mechanism by which Trident is integrated into these native simulators to dynamically generate and deploy its code-based adversarial policies (highlighted in blue for clarity). Finally, we define the RLVR signals, outlining both the environment-specific tactical reward parsers and the universal system-level penalties (Appendix B.3) used to optimize Trident Agentic.

B.1 CAGE 4

For our experiments utilizing the CybORG CAGE Challenge 4 (Kiely et al., 2025) (CAGE 4) environment, we adopted its official enterprise network scenario, which is designed to emulate a complex, multi-subnet military network. While CyberWheel provides dynamically generated network graphs, CAGE 4 offers a complementary, rigorous testbed by maintaining a structurally fixed macro-topology while introducing severe micro-level variability. By randomizing the number of hosts, active services, and initial vulnerabilities in every episode, CAGE 4 allows us to evaluate the generated red policies for both strategic consistency across a stable network layout and tactical adaptability against shifting host-level configurations.

B.1.1 NETWORK TOPOLOGY

The CAGE 4 environment simulates a highly segmented military enterprise architecture. The network is divided into four isolated enclaves connected via an internet node: two Deployed Networks (Network A and B), a Headquarters (HQ) Network, and an undefended Contractor Network that serves as the red agent's initial breach point. The topology enforces strict routing bottlenecks. For example, each Deployed Network contains a Restricted Zone facing the internet and an air-gapped Operational Zone accessible only through the Restricted Zone's router. Similarly, the HQ Network exposes a Public Access Zone, which internally routes to isolated Admin and Office networks.

To prevent deterministic policy memorization, the environment dynamically randomizes both the network configuration and the attack surface at the start of each episode. Each subnet is procedurally generated with a randomized allocation of 3 to 10 user workstations and 1 to 6 server hosts. Furthermore, while all hosts run baseline services like SSH, the environment probabilistically assigns supplementary services and injects exploitable vulnerabilities. This combination of dynamic node generation, strict hierarchical routing, and stochastic vulnerability distribution ensures that the Trident-generated red policy cannot rely on static killchains. Instead, it demands generalized reconnaissance, adaptive lateral movement, and conditional exploit logic.

B.1.2 AGENT CONFIGURATION AND TRIDENT INTEGRATION

Victim defenders (blue agents): We utilized the official pre-trained weights for the HMARL and standard MARL defensive agents (Lyu et al., 2021) available in the official GitHub repository of (Singh et al., 2025). Additionally, we retrained the Graph Neural Network(GNN) defender using its official repository code (King et al., 2025). This agent represents the network as a temporal attributed graph, processing partial observations through a GNN to output independent action probabilities for each of the five active blue agents. Following the official implementation, the policy was trained via PPO for 500,000 episodes (500 steps per episode) using a batch size of 2,500. The actor and critic networks utilize a 256-unit hidden layer and a 128-dimensional node embedding representation. Optimization was performed using learning rates of 3.0×10^{-4} for the actor and 1.0×10^{-3} for the critic, with a surrogate clipping coefficient of 0.2 over 4 update epochs. All baseline blue policies remain strictly frozen during Trident's training and evaluation.

Baseline attacker (red agent): The default attacker in CAGE4 is the `FiniteStateRedAgent`. It employs an internal Finite State Machine (FSM) to orchestrate its attack killchain. The agent tracks each discovered host across 9 progressive compromise states (e.g., Known, User Shell, Root Shell). During execution, it selects a target host based on heuristic prioritization (e.g., defaulting to a 75% bias toward servers) and selects its next action (e.g., `ExploitRemoteService` or `Impact`) by sampling from a hardcoded state transition probability matrix associated with the target's current state.

Trident integration: Rather than replacing the red agent entirely, Trident patches `FiniteStateRedAgent` by overriding three core methods: `state_transitions_probability`, `set_host_state_priority_list`, and `_process_new_observations`. This narrows the generation space to high-level tactical decisions while delegating low-level killchain execution to the underlying framework, reducing syntax errors.

B.1.3 RLVR REWARD DESIGN AND SPECIFICATIONS

Consistent with our Code-as-Policy framework, we utilized a tactical log-based reward parser to evaluate Trident's performance in CAGE 4. As defined in Equation 5, the episodic reward $R(c, o_{0:T})$ is determined by a conditional branch. If the sandbox sanity check fails or the code triggers unrecoverable runtime errors during the self-repair loop, the episode is immediately terminated with a severe execution penalty ($p_{\text{exec}} = -5.0$, and $p_{\text{empty}} = -10.0$ for empty generation).

For executable policies, the tactical reward is parsed directly from terminal logs:

- **Positive Milestones ($\mathcal{E}^+$):** Rewards scale based on the absolute count of successful actions, breaking the traditional binary success metric. We assign weights of $q_{\text{impact}} = 0.2$, $q_{\text{root}} =$

0.3, $q_{\text{user}} = 0.1$, and $q_{\text{scan}} = 0.05$. To prioritize rapid exploitation, an efficiency multiplier $r = 2.0$ scales the reward based on the achievement-to-step ratio.

- **Negative Events ($\mathcal{E}^-$):** To penalize operational noise, the parser deducts points for invalid actions ($s_{\text{invalid}} = 0.05$) and failed exploits ($s_{\text{exploit_fail}} = 0.1$).

The tactical portion of the reward is clipped within $[R_{\min} = -1.0, R_{\max} = 100.0]$. This separation of syntactic penalties and bounded tactical scores ensures that the GRPO gradient cleanly distinguishes between invalid code and poor strategic planning.

B.2 CyberWheel

For our CyberWheel (Oesch et al., 2024) experiments, we evaluate Trident on a 15-host enterprise network topology. Although CyberWheel natively supports dynamically generated networks of varying scales (capable of simulating over 100,000 hosts) our primary objective was to validate Trident Agentic's generalizability beyond the CAGE 4 environment rather than exhaustively evaluate across topology sizes. Investigating policy performance across diverse network scales remains an compelling direction for future work.

B.2.1 Network topology

The 15-host network is centrally orchestrated around a single `core_router` and partitioned into three functional subnets to replicate a standard enterprise architecture:

- **DMZ Subnet :** Contains 5 outward-facing workstations (`dmz0` to `dmz4`). This serves as the primary entry point and decoy deployment zone for the blue agent.
- **User Subnet :** Contains 5 internal workstations (`host0` to `host4`).
- **Server Subnet :** Contains 5 critical business servers, specifically defining two proxy servers, one mail server, and two SSH jump servers (`server0` to `server4`).

To support complex lateral movement, explicit cross-subnet interfaces are configured, such as bridging `dmz1` directly to `server2`, and `host4` to `dmz2`.

B.2.2 Agent configuration and Trident integration

Victim Defender (blue agent): To establish a rigorous victim model, we utilized a Proximal Policy Optimization (PPO) blue agent following the official CyberWheel implementation. Its policy employs a standard Actor-Critic architecture consisting of Multi-Layer Perceptrons (MLPs) with two hidden layers of 64 units each and ReLU activations. We maintained the official hyperparameters: a discount factor $\gamma = 0.995$, GAE $\lambda = 0.95$, and learning rates of 2.5×10^{-4} with cosine restarts (minimum floor 1×10^{-5}). The surrogate objective was clipped at 0.3, with an entropy coefficient of 0.02, value function coefficient of 0.5, and maximum gradient norm of 0.5. This blue policy remains strictly frozen during Trident's training and evaluation.

Baseline Attacker (red agent): The default `ARTAgent` (Oesch et al., 2024) executes a hardcoded, deterministic killchain sequence (e.g., PingSweep → PortScan → Lateral Movement → Discovery → Privilege Escalation → Impact) grounded in MITRE ATT&CK techniques. Crucially, its target selection logic is modularized into a `strategy` class (defaulting to a greedy `ServerDowntime` breadth-first search).

Trident integration: Trident edits `ARTAgent`'s strategy. The generated Python policy (the Coder's output) dynamically replaces the default target selection module via the `set_strategy()` method. This architecture allows Trident to dictate high-level tactical routing while relying on the underlying `ARTAgent` framework to robustly handle the low-level execution of atomic killchain commands, narrowing the generation space and reducing syntax errors.

B.2.3 RLVR reward design and specifications

To ensure the Trident framework remains versatile and portable across heterogeneous environments, the reward system is decoupled from the simulator's internal engine. Instead, Trident relies on a

tactical log-based reward parser, allowing the agent to learn strictly from observable terminal logs and general environment metadata rather than privileged simulator states. Using this log-centric design, rewards are extracted from trajectory text via regular expressions to identify specific tactical events:

- **Positive Milestones** ($\mathcal{E}^+$): Rewards are allocated for advancing through the killchain, with weights set at $q_{\text{impact}} = 0.2$, $q_{\text{priv_esc}} = 0.1$, and $q_{\text{scan}} = 0.05$. To incentivize direct attack paths, an efficiency multiplier ($r = 2.0$) scales the reward based on Equation 5.
- **Negative Events** ($\mathcal{E}^-$): To penalize operational noise, inefficiency, and misdirection, the parser identifies failed actions ($s_{\text{fail}} = 0.1$), redundant repetitions ($s_{\text{repeat}} = 0.02$), and interactions with honeypots ($s_{\text{decoy}} = 0.5$). Crucially, treating decoy hits as negative events addresses a misalignment in default simulation reward structures (Oesch et al., 2024). In CyberWheel's default deception scenario, a red agent receives a positive reward for impacting a decoy, while the blue agent is simultaneously rewarded for a successful misdirection. By classifying decoy interactions as strict negative penalties, we prevent a mutually beneficial local optimum where the red agent naively farms impact rewards on honeypots, ensuring the policy autonomously learns genuine decoy avoidance. The empirical necessity of this penalty to maintain proper adversarial alignment is explicitly demonstrated in our ablation study (Figure 8 and Appendix D.3).

The final episodic reward is clipped within $R_{\min} = -2.0, R_{\max} = 100.0$ to maintain stability during GRPO gradient optimization.

B.3 System-level penalties

To maintain a stable RLVR training loop and enforce the constraints of the Code-as-Policy paradigm, we define a universal system-level penalty p. Aside from environment-specific tactical violation rules, these penalties are shared across both environments to ensure the emergence of robust algorithmic policies. The value of p is determined by the specific failure mode identified during execution:

$$p = \begin{cases} -10.0 \times N_{\text{leak}} & \text{if is_leak}(c) \\ -10.0 & \text{if is_empty}(c) \\ -5.0 & \text{if exec_error}(c) \\ -2.0 & \text{if is_timeout}(c) \\ -1.0 & \text{if format_error}(c) \vee \text{http_error}(c) \end{cases} \tag{6}$$

where N_{leak} denotes the total count of unauthorized access attempts detected in the generated code. The penalty mapping is defined as follows:

- **Information leak** ($-10.0 \times N_{\text{leak}}$)**:** Applied if the policy attempts to bypass the sanitized observation layer to access restricted simulator internals or unobservable ground-truth data (e.g., direct inspection of decoy flags). The penalty scales with the number of detected leak occurrences in the generated code.
- **Empty code** (-10.0)**:** A discrete penalty applied if the Coder fails to produce any executable Python logic.
- **Execution error** (-5.0)**:** Applied if the generated code triggers unrecoverable runtime errors even after exhausting self-repair attempts.
- **Timeout** (-2.0)**:** Applied if the policy execution exceeds the designated time limit, penalizing computationally inefficient or recursive logic.
- **Format & HTTP error** (-1.0)**:** Applied for JSON schema violations or transient infrastructure communication failures. Treating infrastructure-level issues as a minor penalty encourages the agent to prioritize the generation of deployable, robust policies.

C HEURISTIC RED AGENT DESIGN

C.1 CAGE4 RED AGENT INVENTORY

Each agent operates within the CybORG CAGE 4 environment and selects from nine abstract actions: `DiscoverRemoteSystems`, `AggressiveServiceDiscovery`, `StealthServiceDiscovery`, `DiscoverDeception`, `ExploitRemoteService`, `PrivilegeEscalate`, `Impact`, `DegradeServices`, and `Withdraw`. Host states are encoded as a nine-symbol alphabet {`K`, `KD`, `S`, `SD`, `U`, `UD`, `R`, `RD`, `F`}, representing degrees of knowledge, service discovery, user-level access, root access, and out-of-reach hosts respectively.

C.1.1 FINITESTATEREDAGENT (BASE CLASS)

The base finite-state-machine (FSM) red agent maintains a per-host state dictionary and transitions each host according to separate success and failure state-transition matrices on every step. Actions are sampled from a fixed probability matrix conditioned on the current host state, and decoy services discovered via `DiscoverDeception` are stored so that `ExploitRemoteService` can exclude their PIDs via `PIDSelectiveExploitActionSelector`. All derived agents inherit this observation-processing and state-transition logic.

C.1.2 FSREDAGENTCOMBINED

Extends the base FSM agent with *action-specific* failure tracking: per-host, per-action-type attempt and failure counts are maintained, and the action-selection probabilities are dynamically multiplied by 0.1, 0.5, or 1.5 depending on whether the observed success rate falls below 30%, below 50%, or exceeds 70%, respectively. Host selection is additionally weighted by overall failure count using the formula $w = \max(0.1,\ 1.5^{-f})$, exponentially suppressing hosts with many failures while still allowing occasional re-visits.

C.1.3 FSREDAGENTDEGRADER

A variant that steers the agent toward service-degradation objectives: hosts already in state `R` or `RD` (root access obtained) receive a $\times 10$ selection weight boost, and the `DegradeServices` action probability is multiplied by 4 at action selection time regardless of host zone. The agent otherwise inherits all base FSM state-transition matrices unchanged.

C.1.4 FSREDAGENTFAILURETRACKER

Tracks *all* action failures per host (not only deception failures) and selects hosts using exponentially decaying weights $w = \max(0.1,\ 1.5^{-f})$, biasing toward hosts with fewer cumulative failures. Once a host in state `S` or `SD` has accumulated three or more failures, the `ExploitRemoteService` probability for that host is zeroed out, shifting focus away from repeatedly failing exploitation attempts.

C.1.5 FSREDAGENTIMPACTER

Prioritises hosts whose hostnames contain `operational_zone_a` or `operational_zone_b` by assigning them a $\times 100$ selection weight boost. When acting on a zone-matched host, the `Impact` action probability is multiplied by 4; for unidentified hosts (hostname `None`), `DiscoverDeception` is boosted by $\times 2$; for all other hosts, `Impact` is suppressed to zero.

C.1.6 FSREDAGENTTARGETA

Targets hosts in *both* `restricted_zone_a` and `operational_zone_a` with a $\times 100$ selection weight, excluding hosts already in state `RD` from the boost to avoid re-exploiting already-owned machines. Against zone-A hosts, `ExploitRemoteService` is boosted $\times 4$; for other hosts `Withdraw` probability is slightly increased and `PrivilegeEscalate` halved, reducing unnecessary escalation outside the target zones.

C.1.7 FSRedAgentTargetAOperational

Identical targeting and action-weighting logic to `FSRedAgentTargetA`, but the zone keyword is restricted to `operational_zone_a` only, limiting the agent's focus to the operational (not restricted) subnet of zone A. Hosts already in state `RD` within the zone receive a reduced weight of 2 rather than 100, further narrowing aggression to not-yet-fully-owned targets.

C.1.8 FSRedAgentTargetARestricted

Mirrors `FSRedAgentTargetAOperational` but exclusively targets `restricted_zone_a` hosts, focusing the agent on the restricted (higher-security) subnet of zone A while leaving operational hosts with default uniform weights. The same `RD`-exclusion and action-probability adjustments apply.

C.1.9 FSRedAgentTargetB

The zone-B symmetric counterpart of `FSRedAgentTargetA`: hosts matching `restricted_zone_b` or `operational_zone_b` receive a $\times 100$ weight boost (again excluding `RD` state), and `ExploitRemoteService` is multiplied by 4 when acting on zone-B hosts. Outside zone B, `PrivilegeEscalate` is halved and `Withdraw` is marginally increased ($+0.1$).

C.1.10 FSRedAgentTargetBOperational

Focuses on `operational_zone_b` hosts only, following the same $\times 100$ weight and $\times 4$ exploit-boost pattern as `FSRedAgentTargetB`, but confined to the operational subnet. Hosts already at `RD` are not excluded from the weight boost here (unlike the zone-A operational variant), so the agent continues to visit fully-compromised operational-B hosts.

C.1.11 FSRedAgentTargetBRestricted

Mirrors `FSRedAgentTargetBOperational` but scoped exclusively to `restricted_zone_b`, directing exploitation pressure solely at the restricted subnet of zone B. Action-weighting is identical to the other zone-B variants: exploit boost on target hosts, reduced escalation and slightly raised withdrawal elsewhere.

C.1.12 RandomSelectRedAgent

A baseline agent that uniformly samples a random valid action from the full action space at each step, with no host-state tracking or FSM logic. Parameter values (IP addresses, hostnames, subnets) are also chosen uniformly at random from the valid options in the current action space; an optional `disable_withdraw` flag can remove the `Withdraw` action from consideration.

C.2 CyberWheel red strategy inventory

Each class in this directory implements a *target-selection strategy* for the CyberWheel red agent. All strategies inherit from the abstract base class `RedStrategy` and override a single class method, `select_target(agent_obj) -> Host`, which returns the host the red agent should act on in the current step. The agent's internal history (`agent_obj.history`), known unimpacted servers (`agent_obj.unimpacted_servers`), and the set of unknown hosts (`agent_obj.unknowns`) are the primary inputs to every strategy's decision. The difference between the plain and BFS-prefixed variants lies solely in traversal order: plain strategies draw the next host *randomly* from a pool, whereas BFS strategies take the *front element* (`data_list[0]`) for deterministic breadth-first exploration.

C.2.1 BruteForce

The simplest possible strategy: `select_target` unconditionally returns `agent_obj.current_host`, meaning the agent repeatedly attacks the host it is already on without ever moving. This serves as a lower-bound baseline and is useful for stress-testing a single host's defences or verifying killchain execution in isolation.

C.2.2 EXFILTRATION

Designed to locate and impact a designated *leader* (target) host. The agent stays on the current host if it is still `Unknown` (i.e. not yet fully characterised) or if it has been identified as the leader; otherwise it picks a *random* unknown host from `agent_obj.unknowns`, falling back to the current host if the pool is empty. Stale names that no longer appear in the network graph are silently pruned from the unknown pool during selection.

C.2.3 SERVERDOWNTIME

Aims to impact every server reachable in the network, one by one. The agent remains on the current host if it is `Unknown` or is still listed in `agent_obj.unimpacted_servers`; otherwise it draws a *random* host from the unimpacted-server pool, then falls back to the unknown-host pool, and finally stays put if both are empty. Stale hostnames are pruned from whichever pool is being consumed when they are not found in the live network graph.

C.2.4 DFSIMPACT

Implements a depth-first impact sweep: the agent attacks the current host until it has completed the full killchain (i.e. `last_step == len(killchain) - 1`), then immediately jumps to a *random* unimpacted host drawn from `agent_obj.unimpacted_hosts`. Because no breadth-first bookkeeping is performed, the traversal order across hosts is non-deterministic, giving the strategy a randomised DFS character.

C.2.5 BFSEXFILTRATION

A breadth-first variant of `Exfiltration`: the agent stays on the current host if it is `Unknown` or the leader, but when moving it selects `agent_obj.unknowns.data_list[0]` — the oldest-inserted entry — instead of a random one. This FIFO access pattern ensures unknown hosts are explored in discovery order, producing a systematic breadth-first sweep toward the target leader host.

C.2.6 BFSSERVERDOWNTIME **(BASE CLASS)**

A breadth-first variant of `ServerDowntime`: the agent stays on the current host if it is `Unknown` or unimpacted, then moves to `unimpacted_servers.data_list[0]` (oldest known unimpacted server) if available, or `unknowns.data_list[0]` otherwise. The deterministic FIFO ordering guarantees that servers are impacted in the order they were first discovered, making this strategy more predictable and reproducible than the randomised `ServerDowntime`.

D ABLATION STUDY

D.1 CROSS-ENVIRONMENT VALIDATION: REWARD CORRELATIONS

To validate the efficacy of Trident's reward design, we evaluate the reward formulation at a macro level by analyzing the relationship between the red agent's verifiable red reward and the actual degradation of the defender's blue reward.

This relationship is demonstrated through two distinct perspectives. First, Figure 6 illustrates the learning dynamics over time. Across all evaluated DRL defenders in both CAGE 4 and CyberWheel, as the red agent optimizes its policy and the Red Reward increases, the Blue Reward experiences a mirrored, consistent degradation. Second, Figure 7 formalizes this relationship statistically via scatter plots of episodic outcomes, removing the temporal dimension. In CAGE 4, the Pearson correlation coefficients range from -0.87 to -0.94, demonstrating an exceptionally tight mathematical alignment between our rule-based proxy rewards and the true systemic compromise of the DRL defenders. A similar negative trend is successfully maintained in CyberWheel (Spearman: -0.73), confirming the cross-environment robustness of the reward framework.

Note on micro-level ablation: While a fine-grained ablation of individual tactical weights (e.g., q_{impact} vs. $q_{\text{priv_esc}}$) would be ideal, exhaustively tuning these parameters presents a combinatorial

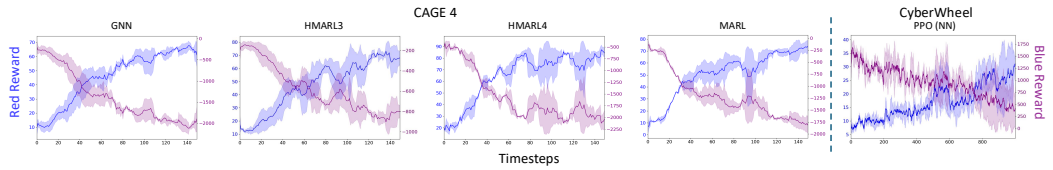

Figure 6: **Learning dynamics of verifiable tactical rewards against defender performance.** As the RLVR process progresses (x-axis: episodes), the red agent successfully maximizes its verifiable reward, consistently driving the blue agent's accumulated reward into a steep decline.

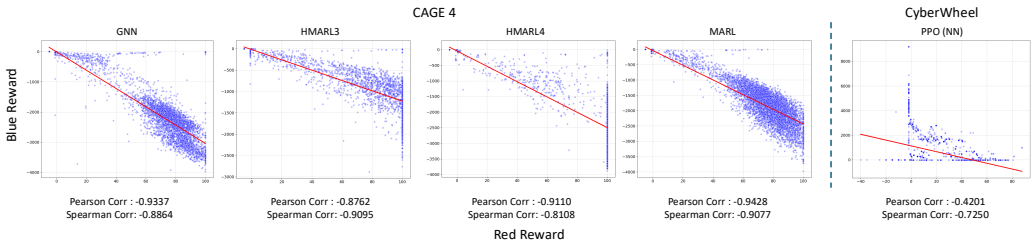

Figure 7: **Correlation between verifiable tactical rewards and defender performance.** Scatter plots illustrating the strong negative correlation between Trident's episodic rewards and the blue agent's rewards. The high correlation coefficients validate that our verifiable reward signal serves as a high-fidelity proxy for actual system compromise across different environments.

explosion that is practically infeasible. Therefore, we rely on heuristic weights based on domain expertise for this study, leaving automated reward optimization as future work.

D.2 Architectural and Performance Ablations in CAGE 4

To isolate the impact of Trident Agentic's core architectural and reward design choices, we conducted further ablation studies within the CAGE 4 environment. For the victim defender, we specifically selected the **HMARL-3** agent since HMARL-3 demonstrated the most consistent learning stability and performance that can provide a reliable and robust baseline for sensitivity analysis. We evaluated two degraded configurations against our full framework: removing the efficiency multiplier (w/o efficiency bonus, $r = 0$) and disabling the Coder's self-repair mechanism (w/o retry, $K = 1$). The results, averaged over three independent training runs, are summarized in Table 3.

Table 3: **Ablation study of Trident's architectural and reward components in CAGE 4.** Lower mean values indicate stronger red agent performance (greater degradation of the blue agent). Displaying individual seeds (Run 1–3) reveals that while $K = 1$ achieves a lower overall mean, it suffers from critical instability driven by a single outlier, whereas $K = 3$ remains consistently robust.

Ablation Configuration	Run 1	Run 2	Run 3	Mean $\pm$ Std. Dev. $\downarrow$
Trident (Ours, $K = 3$)	-843.43	-856.35	-849.76	-849.85 $\pm$ **6.46**
w/o Efficiency Bonus ($r = 0$)	-884.68	-855.80	-767.61	-836.03 $\pm$ 60.99
w/o Retry Loop ($K = 1$)	-646.04	-1358.14	-741.14	**-915.11** $\pm$ 386.62

The necessity of the self-repair Loop ($K = 1$): Restricting the Coder to a single generation attempt ($K = 1$) without the AST-based self-repair loop drastically destabilizes the RLVR training process. Although the $K = 1$ configuration yields a lower mean reward, this metric is deceptive; it is entirely driven by a single highly successful outlier run, masking an extreme generation variance. Without the ability to absorb and correct minor syntactic errors or transient hallucinations, the GRPO signal becomes dominated by execution penalties. This confirms that the $K = 3$ self-repair loop is indispensable for filtering out syntax noise, ensuring that the Planner learns a consistently robust and reproducible strategy.

The role of the efficiency bonus: Removing the efficiency multiplier ($r = 0$ in Equation 5) from the tactical reward signal results in a slightly degraded mean performance and an order-of-magnitude increase in variance. While the agent can still learn effective attack paths without explicit incentives for rapid execution, the absence of this bonus makes the reward signal less focused, leading to higher instability across individual seeds. This indicates that the efficiency bonus contributes positively to both the overall maximization of the reward and the stability of the policy optimization.

D.3 Environment-specific alignment: decoy penalties in CyberWheel

While the macro-level correlations demonstrate general robustness, adapting RLVR to specific environments often requires handling domain-specific simulation mechanics. To validate our handling of such mechanics, we ablate the decoy penalty (s_{decoy}) within the CyberWheel environment.

As illustrated in Figure 8, the presence or absence of this penalty fundamentally alters the agent's learned behavior. When no explicit penalty is applied for interacting with honeypots (left), the Red Reward and Blue Reward paradoxically increase together. This exposes a critical misalignment: the red agent discovers a local optimum by farming impact rewards on decoys, while the blue defender is rewarded for successful misdirection. By enforcing a strict decoy penalty (right), the red agent is forced to autonomously learn behavioral patterns to identify and avoid honeypots. The absolute scale of the Red Reward is significantly lower compared to the unpenalized configuration, reflecting the transition from infinite decoy farming to precise, targeted exploitation of genuine servers.

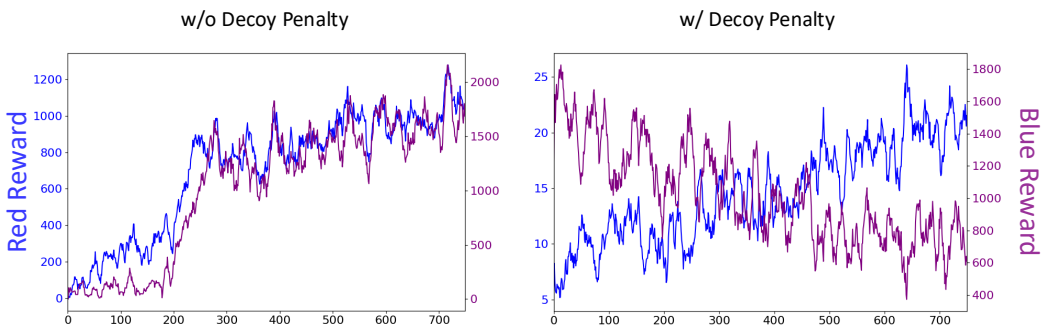

Figure 8: **Impact of decoy interaction penalties on optimization alignment in CyberWheel.** (Left) Without a specific penalty, the agent reaches a misaligned local optimum by repeatedly interacting with decoys, mutually benefiting both agents. (Right) By incorporating a strict negative penalty for decoy hits, Trident learns to avoid honeypots, effectively reducing the blue agent's reward while minimizing operational noise.

E Simulator-Native PPO Baseline

Implementation Details. We implement the DRL red-agent baseline using a hand-written PPO implementation without Stable-Baselines3. A separate PPO agent is trained against each frozen blue defender. The policy and value functions use separate two-layer MLPs with 64 hidden units per layer, ReLU activations, and orthogonal initialization. Training uses generalized advantage estimation (GAE), the clipped PPO objective, entropy regularization, and value-function learning.

For CAGE4, `DRLRedAgent` subclasses `FiniteStateRedAgent` and retains its host-state tracking and action-parameter resolution. PPO replaces its rule-based host and action selection logic. The wrapper converts the dynamically generated network and the red agent's current host knowledge into a fixed-length numerical observation and flattens the available action–host combinations into a discrete action space with an additional no-op action. Invalid action–host combinations fall back to `Sleep()` and receive a small penalty; no hard action masking is applied.

For CyberWheel, the PPO attacker is trained against the frozen PPO blue defender through the simulator-native red-agent interface. In both environments, PPO therefore learns a step-wise numeri-

cal policy directly from simulator transitions, without the Log Summarizer, Planner, Coder, or GRPO components used by Trident Agentic.

Training Budget and Evaluation Protocol. Each PPO red agent is trained for 4 million simulator steps against its corresponding frozen defender. Training trajectories are generated in parallel on CPU using `joblib.Parallel` with process-based execution.

For CAGE4, Trident Agentic at the reported training step 150 uses two contexts per optimization step, $G = 8$ rollouts per context, and 500 simulator steps per rollout, corresponding to

$$150 \times 2 \times 8 \times 500 = 1.2 \text{ million}$$

simulator steps. Thus, the PPO baseline receives approximately $3.33\times$ more simulator interaction than Trident Agentic at the reported checkpoint, reducing the likelihood that weaker PPO performance is caused by insufficient environment exposure.

Both methods are evaluated under the same CAGE4 protocol: 500 episodes of 500 simulator steps, repeated across three independent runs. Exact computational matching is not meaningful because PPO and Trident differ in model size, observation representation, policy parameterization, update frequency, and problem formulation. PPO learns a step-wise policy from numerical transitions, whereas Trident generates a complete executable policy and receives episode-level feedback.

F Dataset Analysis

We characterize the Trident trajectory dataset $\mathcal{D}$ in terms of coverage, diversity, outcome distribution, redundancy, and train–validation consistency. The released dataset contains 9,500 training and 500 validation trajectories for CAGE 4, and 3,556 training and 188 validation trajectories for CyberWheel. Table 4 summarizes the main statistics.

F.1 Dataset Statistics

Each trajectory contains a serialized execution log (`raw_log`), the source heuristic code (`current_code`), the generating heuristic label, and, for CyberWheel, the episodic reward. Since `current_code` is fixed within each heuristic group, trajectory-level statistics are computed from `raw_log`.

Table 4: **Summary of the released Trident dataset.** Action-type sequences are extracted from `raw_log` after removing hostnames, IP addresses, and step indices.

	CyberWheel	**CAGE 4**
Train / validation trajectories	3,556 / 188	9,500 / 500
Source generators	5 strategies	12 agents
Trajectories per generator (train)	672–765	747–852
`raw_log` length (chars), mean $\pm$ std	$18{,}761 \pm 9{,}931$	$110{,}814 \pm 44{,}377$
Documented action vocabulary observed	6 / 6	9 / 9
Action unigram / bigram entropy (bits)	2.49 / 3.57	2.96 / 5.87
Unique action-type sequences	2,570 / 3,556 (72.3%)	8,546 / 9,500 (90.0%)
Exact `raw_log` duplicates	0	0
Train–val consistency (p-values)	χ^2 0.155; KS length 0.733; KS reward 0.586	χ^2 0.667; KS length 0.251

F.2 Coverage

Trajectories are distributed approximately uniformly across the generating heuristics (Figure 9). Training counts range from 672–765 across the five CyberWheel strategies and 747–852 across the twelve CAGE 4 agents, with a maximum-to-minimum ratio of 1.14 in both environments.

F.3 Diversity

The dataset covers all documented strategic actions in both environments (Table 5). Action unigram / bigram entropy is 2.49 / 3.57 bits for CyberWheel and 2.96 / 5.87 bits for CAGE 4. Moreover, 72.3%

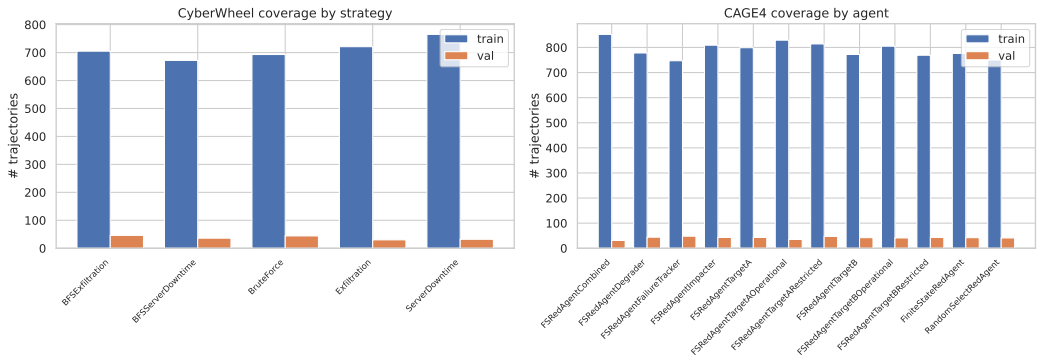

Figure 9: **Coverage by generating heuristic.** Number of training and validation trajectories generated by each CyberWheel strategy (left) and CAGE 4 agent (right).

Table 5: **Action types observed in the training split.**

Environment	Observed action types
CyberWheel (6 / 6)	`pingsweep`, `portscan`, `discovery`, `lateral-movement`, `privilege-escalation`, `impact`
CAGE 4 (9 / 9)	`DiscoverRemoteSystems`, `AggressiveServiceDiscovery`, `StealthServiceDiscovery`, `DiscoverDeception`, `ExploitRemoteService`, `PrivilegeEscalate`, `Impact`, `DegradeServices`, `Withdraw`
CAGE 4 (auxiliary)	`Sleep`, `RedSessionCheck`

of CyberWheel and 90.0% of CAGE 4 trajectories have unique action-type sequences, with no exact `raw_log` duplicates.

Figure 10 further shows substantial variation in serialized trajectory length, reflecting differences in episode progression across randomized environments and defender interactions.

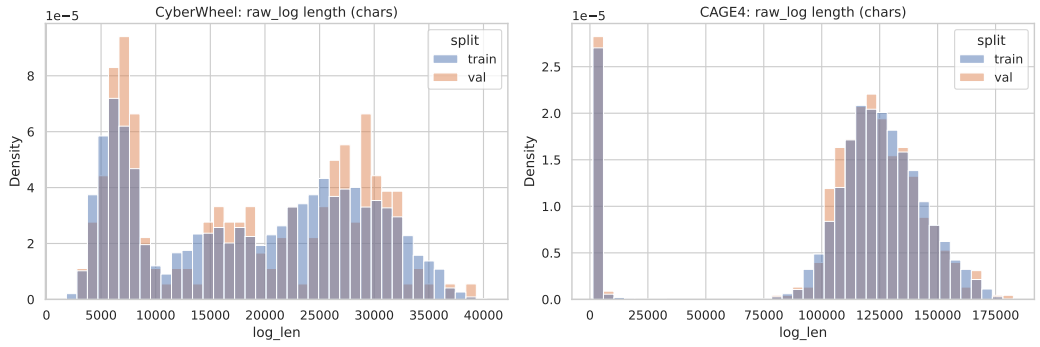

Figure 10: **Log-length distribution.** Distribution of `raw_log` length for the training and validation splits of CyberWheel (left) and CAGE 4 (right).

F.4 Outcome Distribution

Figure 11 shows that the generating heuristics produce a broad range of episode outcomes. CyberWheel uses the released episodic reward. Because CAGE 4 trajectories do not contain a reward field, we use an observable milestone score computed from successful and failed actions in `raw_log`. The CAGE 4 score is used only for relative comparison among generating heuristics.

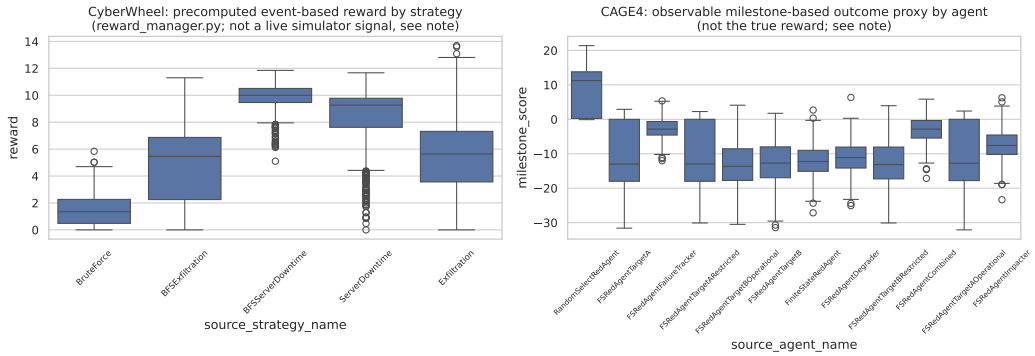

Figure 11: **Outcome distribution by generating heuristic.** CyberWheel episodic reward (left) and CAGE 4 observable milestone score (right).

F.5 Redundancy

We measure within-generator redundancy using the Jaccard similarity of action-type bigram sets over 200 sampled trajectory pairs per generator (Figure 12). Deterministic heuristics show higher similarity, while heuristics affected by randomized network states and defender responses exhibit greater variation.

Action-sequence similarity is a conservative measure of log similarity: `raw_log` also contains hostnames, IP addresses, discovered services, and environment-specific state information. The training split contains 986 CyberWheel and 954 CAGE 4 action-sequence collisions, but no exact `raw_log` duplicates.

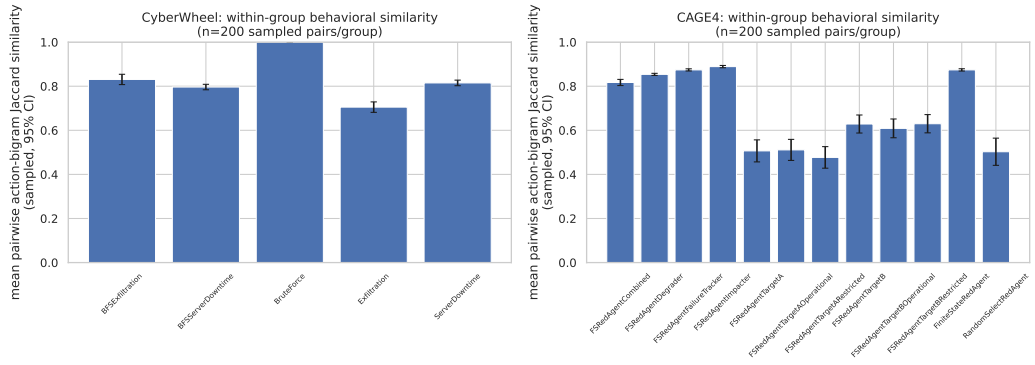

Figure 12: **Within-generator behavioral similarity.** Mean pairwise Jaccard similarity of action-type bigram sets over 200 sampled trajectory pairs per generator.

F.6 Train–Validation Consistency

We compare the training and validation splits using a χ^2 test on generator proportions and two-sample KS tests on log length and CyberWheel reward. All tests yield $p > 0.15$ (Table 6), with closely overlapping log-length distributions (Figure 13).

Table 6: **Train–validation consistency tests.**

Environment	Test	Statistic	p-value
CyberWheel	Strategy proportions (χ^2, df $= 4$)	6.67	0.155
CyberWheel	`raw_log` length (KS)	0.050	0.733
CyberWheel	Reward (KS)	0.057	0.586
CAGE 4	Agent proportions (χ^2, df $= 11$)	8.51	0.667
CAGE 4	`raw_log` length (KS)	0.046	0.251

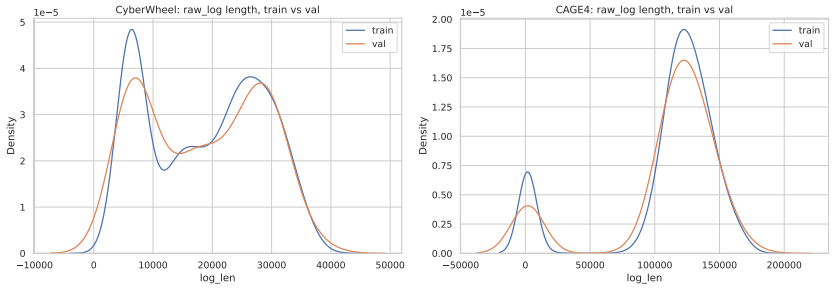

Figure 13: **Training and validation log-length distributions.** Kernel density estimates for CyberWheel (left) and CAGE 4 (right).

G Prompt Details

G.1 Log Summarizer

```
# Mission Trajectory Summary
### Step {t}
Knowledge: {Current Host: {host} | Unimpacted: {S} | Unknowns: {U}}
Host Intel: {host_id}({role}, {phase},)[+NEW?] ...
Action:    {Success|Failed} {action_type} to {target}
(repeated for t = 0,...,T)
## Episode Impact Summary
- Episode {i} achieved the most server impacts: {count}.
```

Figure 14: Example of a serialized trajectory raw log template from the CyberWheel environment.

G.2 CAGE4

G.2.1 General knowledge

```
dDISCOVER_REMOTE_SYSTEMS = """
- **Action**: DiscoverRemoteSystems(subnet, session, agent)
- **Description**: Outputs the IP addresses of all live hosts within
    a specified subnet.
- **Import Path**: `from CybORG.Simulator.Actions.AbstractActions
    import DiscoverRemoteSystems`
- **Input Requirements**:
    - `subnet`: A target IPv4Network (e.g.,
    IPv4Network('10.0.96.0/24')).
    - `session`: The ID of the session executing the action.
    - `agent`: The name of the agent executing the action.
- **Prerequisite**: The agent must identify a known subnet to
    investigate, typically found in the initial observation's
    'Interface' and 'Subnet' dictionaries.
- **Key Observation**:
```

```
    - Gaining a shell on a host within the subnet is NOT necessary to
    execute this action.
    - Success is indicated by `TernaryEnum.TRUE` and a dictionary
    containing the IP addresses of discovered hosts.
- **Common Logic Error**: Attempting to scan a subnet without first
    identifying it from host interface data.
"""

SERVICE_DISCOVERY = """
- **Action**: AggressiveServiceDiscovery(ip_address, session, agent)
    / StealthServiceDiscovery(ip_address, session, agent)
- **Description**: Outputs all processes with associated services
    (local address and port) on a known host.
- **Import Path**: `from CybORG.Simulator.Actions.AbstractActions
    import AggressiveServiceDiscovery, StealthServiceDiscovery`
- **Versions**:
    - **Aggressive**: Faster execution but more likely to alert the
    Blue Team.
    - **Stealth**: Slower execution (requires multiple 'Sleep' steps
    to wait for results) but less likely to alert the Blue Team.
- **Prerequisite**:
    - The target `ip_address` must be known to the agent (e.g., from
    initial observation or DiscoverRemoteSystems).
    - This action MUST be run on a host before attempting exploits or
    DiscoverDeception, as it enables them.
- **Output**: Detailed service information (e.g., port 22, 25, 80)
    found under the 'Processes' key in the observation.
- **Common Logic Error**: Attempting an Exploit without first running
    Service Discovery to identify vulnerable ports.
"""

DISCOVER_DECEPTION = """
- **Action**: DiscoverDeception(ip_address, session, agent)
- **Description**: Probes a remote host to identify decoy services.
    This is a high-level investigative action that relies on a
    specific sequence of prior reconnaissance.
- **Import Path**: `from CybORG.Simulator.Actions.AbstractActions
    import DiscoverDeception`
- **The Mandatory Reasoning Chain (Ordered Combination)**:
    To successfully run this action and interpret the result, the
    agent MUST follow this exact sequence:
    1. **Subnet ID**: Identify the subnet from initial observations
    (e.g., '10.0.96.0/24').
    2. **Host Discovery**: Execute `DiscoverRemoteSystems` on that
    subnet to populate the IP list.
    3. **Service Mapping**: Execute `ServiceDiscovery` (Aggressive or
    Stealth) on the specific `ip_address`. The agent cannot detect a
    "decoy service" if it doesn't first know what "services" exist.
    4. **The Probe**: Execute `DiscoverDeception`.
    5. **The Verification**: Execute `Sleep()` (or wait 1 tick). The
    results of a deception probe take **two ticks** to appear in the
    observation space.
- **Result Interpretation Logic**:
    - **DECOY DETECTED**: Look for `Properties: ['decoy']` within the
    `Processes` list for that host.
    - **NO DECOY**: If the action is `TRUE` but the target host's
    entry is MISSING from the observation after the `Sleep` tick, it
    implies no decoys were found.
- **Logical Check for Analyst**:
    - Flag as **CRITICAL ERROR** if `DiscoverDeception` is called
    without a preceding `ServiceDiscovery` on the same IP.
    - Flag as **INEFFICIENT** if the agent fails to `Sleep`
    immediately after, as it is making decisions based on stale data.
"""
```

```
EXPLOIT_REMOTE_SERVICE = """
- **Action**: ExploitRemoteService(ip_address, session, agent)
- **Description**: Exploits a specific service on a target host to
    gain a user-privileged shell (User Session).
- **Import Path**: `from CybORG.Simulator.Actions.AbstractActions
    import ExploitRemoteService`
- **The Mandatory Reasoning Chain (Ordered Dependency)**:
    For this action to run successfully, the agent MUST follow this
    exact strategic sequence:
    1. **Subnet Discovery**: Identify the target subnet from initial
    observations.
    2. **Host Identification**: Execute `DiscoverRemoteSystems` to
    find live IP addresses (e.g., '10.0.96.108') within that subnet.
    3. **Service Mapping**: Execute `ServiceDiscovery` (Aggressive or
    Stealth) on the target IP. This is a STRICT prerequisite; the
    agent must know which services are running to exploit one of them.
    4. **The Exploit**: Execute `ExploitRemoteService_cc4`.
    5. **The Verification**: Execute TWO `Sleep()` actions (or wait 3
    ticks total). The exploit and session stabilization take three
    ticks to reflect in the observation space.
- **Result Interpretation Logic**:
    - **SUCCESS**: The observation will contain a new `Sessions`
    entry for the target host with `Type: <SessionType.SSH: 2>` (or
    similar) and `username: 'user'`.
    - **HOSTNAME REVEAL**: A successful exploit often reveals the
    target's internal hostname (e.g.,
    'contractor_network_subnet_user_host_5') as a new key in the
    observation dictionary.
- **Key Observation**:
    - Success is indicated by `TernaryEnum.TRUE` in the success key
    after the mandatory wait period.
- **Logical Check for Analyst**:
    - Flag as **CRITICAL FAILURE** if the agent attempts an exploit
    without prior `ServiceDiscovery` on the target IP.
    - Flag as **TIMING ERROR** if the agent attempts to use the new
    session immediately without waiting for the required 3-tick
    stabilization period (Action + 2 Sleeps).
"""

PRIVILEGE_ESCALATE = """
- **Action**: PrivilegeEscalate(hostname, session, agent)
- **Description**: Exploits a specific host where the agent already
    possesses a user-privileged shell to upgrade that shell to root
    privileges.
- **Import Path**: `from CybORG.Simulator.Actions.AbstractActions
    import PrivilegeEscalate`
- **Prerequisites**:
    1. **Target Identification**: The agent must identify the
    target's `hostname` (e.g.,
    'contractor_network_subnet_user_host_4') from observations.
    2. **User Session**: The agent MUST already have an active user
    shell on the target host (visible in the 'Sessions' dictionary).
- **Key Observation**:
    - If a user shell is not already present, this action will fail;
    a user shell must be obtained first (typically via
    `ExploitRemoteService`).
    - Success is indicated by `TernaryEnum.TRUE`.
    - Upon success, the `username` associated with the session ID in
    the observation changes to `root`.
- **Logical Check for Analyst**:
    - Flag as **SEQUENCE ERROR** if the agent attempts
    `PrivilegeEscalate` on a host where no prior user session was
    established.
    - Verify if the agent used the `hostname` string rather than an
    IP address, as this action specifically targets the host entity.
```

```
"""

DEGRADE_SERVICES = """
- **Action**: DegradeServices(hostname, session, agent)
- **Description**: Attempts to reduce the reliability of all services
    on a target host. This makes it significantly less likely for a
    Green agent (the mission user) to successfully complete their
    local work on that host.
- **Import Path**: `from CybORG.Simulator.Actions.AbstractActions
    import DegradeServices`
- **The Mandatory 5-Step Strategic Chain**:
    This is the most complex prerequisite chain in the scenario. The
    agent MUST follow this exact sequence:
    1. **Subnet ID**: Identify the target subnet (e.g.,
    '10.0.96.0/24') from initial interface data.
    2. **Host Discovery**: Execute `DiscoverRemoteSystems` to
    identify live IP addresses on that subnet.
    3. **Service Mapping**: Execute `ServiceDiscovery` (Aggressive or
    Stealth) on the target IP.
    4. **Initial Foothold**: Execute `ExploitRemoteService` to gain a
    User-level shell.
    5. **Administrative Control**: Execute `PrivilegeEscalate` to
    upgrade the session to ROOT privileges.
- **Prerequisites**:
    - The agent MUST possess a ROOT-privileged session on the target
    `hostname`.
- **Key Observation**:
    - The action targets the `hostname` (e.g.,
    'contractor_network_subnet_user_host_5'), not the IP address.
    - Success is indicated by `TernaryEnum.TRUE`.
- **Mission Impact**: Effective use of this action directly disrupts
    the Green agent's ability to fulfill mission requirements by
    lowering service reliability.
- **Logical Check for Analyst**:
    - Flag as **PREREQUISITE FAILURE** if the agent attempts to
    degrade services without a ROOT session (username: 'root').
    - Verify that the agent transitioned correctly from IP-based
    targeting (Steps 1-3) to Hostname-based targeting (Steps 4-5).
"""

IMPACT = """
- **Action**: Impact(hostname, session, agent)
- **Description**: A high-value offensive action that attempts to
    disrupt the mission by stopping Operational Technology (OT)
    services on a target host. This has a significant negative effect
    on the Blue Team's reward.
- **Import Path**: `from CybORG.Simulator.Actions.AbstractActions
    import Impact`
- **The Mandatory Strategic Chain (Prerequisites)**:
    Successful execution requires a strict 5-step progression similar
    to `DegradeServices`:
    1. **Subnet ID**: Discover the target subnet from initial
    observations.
    2. **Host Discovery**: Run `DiscoverRemoteSystems` to find live
    hosts on the target subnet.
    3. **Service Discovery**: Run `AggressiveServiceDiscovery` or
    `StealthServiceDiscovery` to identify services on the target host.
    4. **Initial Foothold**: Run `ExploitRemoteService_cc4` to obtain
    a user-level shell.
    5. **Administrative Control**: Run `PrivilegeEscalate` to secure
    a root-level shell on the host.
- **Target Constraints**:
    - **Privilege Level**: Requires a RED ROOT shell on the target
    host.
```

```
    - **Service Presence**: Only succeeds if an OT service
    (`OTSERVICE`) is present on the host.
    - **Zone Limitation**: OT services are primarily located on
    servers within Operational Zones (Zones A and B).
- **Result Interpretation Logic**:
    - **SUCCESS**: Indicated by `TernaryEnum.TRUE`.
    - **Observation Detail**: The resulting observation will show the
    stopped service in the `Processes` dictionary, often identifying
    `<ProcessName.OTSERVICE: 29>`.
- **Logical Check for Analyst**:
    - Flag as **PREREQUISITE FAILURE** if the agent lacks root
    privileges (`username: 'root'`) at the time of the action.
    - Flag as **TARGETING ERROR** if the agent attempts `Impact` on a
    host that does not contain an OT service (e.g., standard user
    hosts), leading to a failure.
"""

WITHDRAW = """
- **Action**: Withdraw(hostname, session, agent)
- **Description**: Strategically removes ALL of the Red Agent's
    sessions (User and Root) from a specific target host, effectively
    erasing its presence.
- **Import Path**: `from
    CybORG.Simulator.Actions.ConcreteActions.Withdraw import Withdraw`
- **Input Parameters**:
    - `hostname`: The target host to withdraw from (e.g.,
    'contractor_network_subnet_user_host_5').
    - `session`: The parent session ID initiating the withdrawal.
- **Prerequisites**:
    - **Active Presence**: The agent MUST currently possess a valid
    session (User or Root) on the target `hostname`.
    - **Objective Completion**: Typically executed AFTER completing
    objectives (like Impact or Intelligence Gathering) to clean up
    traces.
- **Result Interpretation Logic (The "Vanishing" Act)**:
    - **SUCCESS**: Indicated by `TernaryEnum.TRUE`.
    - **Observation Signature**: Unlike other actions, success is
    defined by **data removal**. The entry for the target host (e.g.,
    `contractor_network_subnet_user_host_5`) will completely
    **disappear** from the observation dictionary.
- **Verification Method (Negative Testing)**:
    - To scientifically verify that `Withdraw` was successful, the
    agent may attempt a privileged action (like `DegradeServices`) on
    the target immediately after.
    - **Expected Outcome**: This subsequent action MUST fail
    (`TernaryEnum.FALSE`), confirming that access has been
    successfully revoked.
- **Logical Check for Analyst**:
    - Flag as **CLEANUP SUCCESS** if the host entry vanishes from the
    log after execution.
    - Flag as **FAILED RETREAT** if the host entry persists in the
    observation despite a `TRUE` result (potential simulation
    artifact).
"""

ACTIONS_LIST = [
    DISCOVER_REMOTE_SYSTEMS,
    SERVICE_DISCOVERY,
    DISCOVER_DECEPTION,
    EXPLOIT_REMOTE_SERVICE,
    PRIVILEGE_ESCALATE,
    DEGRADE_SERVICES,
    IMPACT,
    WITHDRAW
]
```

Listing 1: General knowledge base for all agentic components.

G.2.2 SUMMARIZER

```
LOG_SUMMARIZER_SYSTEM_PROMPT = f"""
You are an expert Cybersecurity Analyst and Code Reviewer.
Your task is to analyze the performance of an autonomous Red Agent
    based on its **Source Code** and **Execution Log**.
Reference the following Knowledge Base to evaluate if the agent uses
    actions correctly according to CybORG specifications.

### CYBORG ACTION KNOWLEDGE BASE
{"".join(ACTIONS_LIST)}
d
### ANALYSIS GUIDELINES
- Compare the agent's logic (Source Code) with its actual performance
    (Log).
- Identify if failures are due to environment constraints or logical
    errors in the code.
- Explain the "Reasoning Gap" between what the code intended and what
    actually happened.
"""

SUMMARY_INSTRUCTION = """
### OBJECTIVE
Compare the agent's **Source Code Logic** with the **Execution Log**
    and provide a strategic analysis.

### OUTPUT FORMAT
1. **Strategy Overview**: Briefly explain the intended logic (e.g.,
    FSM, Weighted Random).
2. **Behavioral Gap Analysis**: Link specific code lines/methods to
    execution results.
3. **Critical Success/Failure**: Highlight the turning points of the
    mission.
4. **Insight for Improvement**: What should be changed in the code to
    improve performance?
"""
```

Listing 2: System Prompt for the Trainable Summarizer.

G.2.3 PLANNER

```
PLANNER_SYSTEM_PROMPT = f"""
You are the Strategic Lead for a Red Team autonomous agent in a
    CybORG CAGE4 environment.

Your job is NOT to write code.
Your job is to output a compact, machine-consumable "StrategySpec" in
    JSON that will later be used to PATCH an existing
    FiniteStateRedAgent.

### CYBORG ACTION KNOWLEDGE BASE (Strict Prerequisites)
{"".join(ACTIONS_LIST)}

### CAGE4 STRATEGIC CONTEXT
- The network is divided into subnets (Contractor, HQ, Restricted,
    Operational Zones A & B).
- Critical targets (OT service) are typically in Operational Zones.
```

```
- Blue Agents are active; stealthy discovery is safer after
    aggressive failures.
- You may only influence the agent by:
  (1) Host-state priority weights (which host-state to focus on),
  (2) State->Action probability matrix (which action to pick given
    host-state),
  (3) A small bugfix flag for decoy PID key consistency.

### OUTPUT FORMAT (STRICT JSON ONLY)
Return ONLY valid JSON (no Markdown fences, no extra commentary).

Schema:
{{
  "intent": {{
    "summary": "<one sentence>",
    "focus": ["<short bullets>"]
  }},
  "host_state_priority": {{
    "K": <0-100>, "KD": <0-100>, "S": <0-100>, "SD": <0-100>,
    "U": <0-100>, "UD": <0-100>, "R": <0-100>, "RD": <0-100>
  }},
  "state_action_probs": {{
    "K":  {{"DiscoverRemoteSystems": <0-1>,
    "AggressiveServiceDiscovery": <0-1>, "StealthServiceDiscovery":
    <0-1>}},
    "KD": {{"AggressiveServiceDiscovery": <0-1>,
    "StealthServiceDiscovery": <0-1>}},
    "S":  {{"DiscoverRemoteSystems": <0-1>, "DiscoverDeception":
    <0-1>, "ExploitRemoteService": <0-1>}},
    "SD": {{"DiscoverDeception": <0-1>, "ExploitRemoteService":
    <0-1>}},
    "U":  {{"DiscoverRemoteSystems": <0-1>, "PrivilegeEscalate":
    <0-1>, "Withdraw": <0-1>}},
    "UD": {{"PrivilegeEscalate": <0-1>, "Withdraw": <0-1>}},
    "R":  {{"DiscoverRemoteSystems": <0-1>, "Impact": <0-1>,
    "DegradeServices": <0-1>, "Withdraw": <0-1>}},
    "RD": {{"Impact": <0-1>, "DegradeServices": <0-1>, "Withdraw":
    <0-1>}}
  }},
  "patch_flags": {{
    "decoy_pid_key": "ip"  // must be exactly "ip"
  }}
}}

Hard constraints:
- host_state_priority values must sum to 100.
- For each state in state_action_probs: probabilities must sum to 1.0
    (within 1e-6).
- Use only the exact action class names used in the schema.
- Keep Withdraw probability at 0.0 unless the summary explicitly
    indicates retreat.
"""

USER_PROMPT = """
### Situation Summary from LogSummarizer:
{summary}

### Task
Produce the StrategySpec JSON for the next iteration.
"""
```

Listing 3: System Prompt for the Trainable Planner (π_θ)

G.2.4 CODER

```
CODER_SYSTEM_PROMPT = f"""
You are a CybORG CAGE4 Python developer.

You are PATCHING an existing implementation: FiniteStateRedAgent.
You will be given BASE_CODE and STRATEGY_SPEC (JSON).

### CYBORG ACTION KNOWLEDGE BASE (Strict Prerequisites)
{"".join(ACTIONS_LIST)}

### HARD RULES
- Output MUST be plain text only (no Markdown fences, no
    explanations).
- Do NOT output a unified diff.
- Do NOT output the full file.
- Do NOT change imports or class names.
- **DO NOT output the same function multiple times. Exactly ONE block
    per requested method.**

### YOUR TASK
Generate UPDATED DEFINITIONS for EXACTLY these 3 methods, in this
    exact order:
1) state_transitions_probability
2) set_host_state_priority_list
3) _process_new_observations

### REQUIRED OUTPUT FORMAT (STRICT TEMPLATE)
You must follow this exact structure and stop generating immediately
    after END_ALL.

BEGIN_FUNC state_transitions_probability
def state_transitions_probability(self):
    # your code here
END_FUNC

BEGIN_FUNC set_host_state_priority_list
def set_host_state_priority_list(self):
    # your code here
END_FUNC

BEGIN_FUNC _process_new_observations
def _process_new_observations(self, observation):
    # your code here
END_FUNC

END_ALL
"""

USER_PROMPT = """
You are given:
(1) BASE_CODE: the current FiniteStateRedAgent source code
(2) STRATEGY_SPEC: a JSON object produced by the Planner

Return ONLY the three updated def blocks in the required format.

BASE_CODE:
{base_code}

STRATEGY_SPEC (JSON):
{strategy_spec}
"""
```

Listing 4: System Prompt for the Frozen Coder

G.3 CyberWheel

G.3.1 General knowledge

```
ROLE_OF_SELECT_TARGET = """
- **Your ONLY job**: Return a single `Host` object  the host you want
    to attack next.
- **You do NOT write attack logic.** `ARTAgent` (art_agent.py)
    automatically executes the full killchain on whichever host you
    return:
    1. PingSweep the target's subnet (if not yet scanned)  discovers
    new hosts into `unknowns`
    2. PortScan the target (if not yet scanned)
    3. LateralMovement to the target (if not already there)  updates
    `current_host`
    4. Discovery  identifies host type; if 'Server', adds to
    `unimpacted_servers`
    5. PrivilegeEscalation  gains root access
    6. Impact  shuts down the server, removes it from
    `unimpacted_servers`
- One step = one action above. The agent progresses through the
    killchain automatically over multiple steps.
- Returning the same host repeatedly is correct ^ the agent keeps
    advancing the killchain on it.
"""

HOST_INTERNAL_LOGIC = ROLE_OF_SELECT_TARGET

STATE_CURRENT_HOST = """
- **`agent_obj.current_host`** (`Host`): The host the Red Agent is
    physically on right now.
  - Use `agent_obj.current_host.name` (str) for comparisons.
"""

STATE_UNIMPACTED_SERVERS = """
- **`agent_obj.unimpacted_servers`** (`HybridSetList`): Known
    'Server'-type hosts not yet impacted. PRIMARY target pool.
  - Only populated after Discovery has confirmed a host's type as
    'Server'.
  - API:
    - `len(agent_obj.unimpacted_servers) > 0`
    - `agent_obj.unimpacted_servers.data_list[0]`  - first element,
    deterministic (BFS-style)
    - `agent_obj.unimpacted_servers.get_random()`   random element
    - `agent_obj.current_host.name in agent_obj.unimpacted_servers`
    membership check
"""

STATE_UNKNOWNS = """
- **`agent_obj.unknowns`** (`HybridSetList`): Hosts discovered via
    PingSweep but not yet typed.
  - Could be Servers, Users, or anything  unknown until Discovery
    runs on them.
  - Explore these when `unimpacted_servers` is empty.
  - Same API as `unimpacted_servers`.
"""

STATE_HOST_HISTORY = """
- **`agent_obj.history.hosts[hostname]`** (`KnownHostInfo`): Per-host
    tracking flags (read-only).
  - `.type`: `'Server'`, `'User'`, or `'Unknown'`  set after
    Discovery runs on the host.
  - `.impacted`: `True` if successfully impacted  never re-target
    these.
```

```
  - `.escalated`: `True` if privilege-escalation has succeeded on
    this host.
  - `.scanned`: `True` if the host has been port-scanned.
  - `.on_host`: `True` if the red agent is currently on this host.
  - `.is_leader`: `True` if this is the designated high-value target.
  - `.service_count`: `int`  number of services discovered via port
    scan (safe to use; does NOT expose service internals).
    A value of `0` after scanning means no exploitable attack surface
    was found  deprioritize these hosts.
  - Tip: The trajectory's **Host Intel** lines encode these states
    compactly as
    `hostname(Type,phase,svc:N)`. Use `.service_count` in your Python
    code to act on this signal.
"""

STATE_NETWORK_LOOKUP = """
- **`agent_obj.network.hosts[hostname]`** (`Host`): Convert a
    hostname string into a Host object to return.
  - Example: `return
    agent_obj.network.hosts[agent_obj.unimpacted_servers.data_list[0]]`
  - **SAFE**: Only `.name` (str) is needed for comparisons or as a
    return value.
  - **FORBIDDEN** on any Host object  these are simulation internals
    invisible to a real red agent.
    Using them is cheating and will be penalized:
    - `.decoy`, `.isolated`, `.services`, `.processes`, `.os`
    - `.host_type.type` (HostTypes enum), `HostTypes.DECOY_SERVER`,
    `HostTypes.DECOY_USER`
    - `agent_obj.network.decoys` (the decoy host registry dict)
"""

ACTIONS_CONTEXT = """
Actions run automatically by art_agent.py in this order on your
    chosen target. You do not call them.
- **PingSweep**  discovers hosts on the subnet into `unknowns`
- **PortScan**  prerequisite for LateralMovement
- **LateralMovement**  moves agent to target, updates `current_host`
- **Discovery**  sets host `.type`; 'Server' hosts added to
    `unimpacted_servers`
- **PrivilegeEscalation**  prerequisite for Impact
- **Impact**  removes host from `unimpacted_servers`; repeating on
    same host = 0 reward
"""

RULE_EFFICIENCY = """
- **10x Early Window**: Rewards are multiplied 10x during the first
    25% of the episode (~25 steps out of 100).
  Reach and impact Servers as fast as possible. Do not wander.
"""

RULE_DECEPTION = """
- **Decoy Detection (Behavioral Only)**: Blue deploys fake
    'Server'-looking hosts mid-episode.
  You CANNOT inspect whether a host is a decoy  that information is
    not observable.
  Use the trajectory's **Host Intel** lines to infer suspicious hosts:
  - `svc:0` after scanning  no discovered attack surface  likely a
    decoy or empty host.
  - Phase stuck at `lateral` or `scanned` despite repeated targeting
    host may be hardened or a decoy.
  - New `Unknown` hosts with UUID-style names appearing mid-episode
    likely blue-deployed decoys.
  Strategy: track a per-host failure count using a class-level dict;
    skip hosts that have
```

```
  failed to advance past `lateral` after N attempts. Rotate to the
    next candidate in
  `unimpacted_servers` or `unknowns` rather than fixating on one host.
"""

RULE_REDUNDANCY = """
- **Avoid Redundant Targets**:
  - Skip hosts where `agent_obj.history.hosts[name].impacted` is True.
  - Do not fixate on a single unknown host - try others if there is
    no progress.
"""

RULE_GENERALIZATION = """
- **Network-Agnostic**: Your strategy runs on networks of any size
    and topology.
  Do NOT hardcode host names, counts, or subnet assumptions.
  Drive all decisions from the observable APIs only.
"""

REFERENCE_STRATEGY = """
The simplest correct baseline (BFSServerDowntime) - improve on this,
    do not copy it:

```python
from cyberwheel.red_agents.strategies.red_strategy import RedStrategy
from cyberwheel.network.host import Host

class BFSServerDowntime(RedStrategy):
    @classmethod
    def select_target(cls, agent_obj) -> Host:
        current_type =
    agent_obj.history.hosts[agent_obj.current_host.name].type

        # Stay on current host if it is unknown (Discovery not run
    yet) or an unimpacted server
        if current_type == "Unknown" or agent_obj.current_host.name
    in agent_obj.unimpacted_servers:
            return agent_obj.current_host

        # Go to the nearest known unimpacted server
        if len(agent_obj.unimpacted_servers) > 0:
            return
    agent_obj.network.hosts[agent_obj.unimpacted_servers.data_list[0]]

        # Explore unknowns when no servers are known
        if len(agent_obj.unknowns) > 0:
            return
    agent_obj.network.hosts[agent_obj.unknowns.data_list[0]]

        return agent_obj.current_host
```
"""

STATE_KNOWLEDGE_LIST = [
    STATE_CURRENT_HOST,
    STATE_UNIMPACTED_SERVERS,
    STATE_UNKNOWNS,
    STATE_HOST_HISTORY,
    STATE_NETWORK_LOOKUP,
]
ACTIONS_KNOWLEDGE_LIST = [ACTIONS_CONTEXT]
RULES_KNOWLEDGE_LIST = [RULE_EFFICIENCY, RULE_DECEPTION,
    RULE_REDUNDANCY, RULE_GENERALIZATION]
```

Listing 5: General Knowledge base among all agents.

G.3.2 SUMMARIZER

```
CYBERWHEEL_MASTER_KNOWLEDGE = f"""
### [DOMAIN KNOWLEDGE]

#### 1. How select_target() Works
{HOST_INTERNAL_LOGIC}

#### 2. Observable Agent State & Legal APIs
{"".join(STATE_KNOWLEDGE_LIST)}

#### 3. Killchain (auto-executed by art_agent.py)
{"".join(ACTIONS_KNOWLEDGE_LIST)}

#### 4. Strategic Rules
{"".join(RULES_KNOWLEDGE_LIST)}

#### 5. Reference Implementation
{REFERENCE_STRATEGY}
"""

LOG_SUMMARIZER_SYSTEM_PROMPT = f"""
You are the **Lead Cyber-AI Strategist**. Your role is to perform a
    post-mission forensic analysis on an autonomous Red Agent.
You are provided with the agent's **Source Code** (its intended
    logic) and the **Execution Log** (the reality of the environment).

{CYBERWHEEL_MASTER_KNOWLEDGE}

### MANDATORY ANALYSIS STEPS:
1. **Legal API Compliance**: Did the code access any FORBIDDEN oracle
    paths?
   FORBIDDEN: `.decoy`, `.isolated`, `.services`, `.processes`, `.os`,
   `host_type.type` (HostTypes enum), `HostTypes.DECOY_SERVER`,
    `HostTypes.DECOY_USER`,
   `agent_obj.network.decoys`
   The red agent cannot observe any of these - any strategy using
    them is cheating and should be flagged.

2. **Killchain Awareness**: Did the code correctly account for the
    automatic killchain?
   (The strategy only picks a target  art_agent.py handles pingsweep
    portscan  lateral  discovery  priv-esc  impact.)

3. **Efficiency**: Did the agent reach `unimpacted_servers` quickly
    within the 10x early window (first 25% of steps)?
   Look for signs of wandering: too many steps on `unknowns`,
    repeated targets with no reward.

4. **Decoy Avoidance (Behavioral)**: Check the **Host Intel** lines
    in the trajectory.
   Did any host show `svc:0` after scanning, or remain stuck at
    `lateral`/`scanned` phase across many steps?
   These are the behavioral decoy signals  no oracle access to
    `.decoy` or HostTypes enum is needed.

5. **Pool Management**: Did the agent correctly exhaust
    `unimpacted_servers` before falling back to `unknowns`?
   Did it avoid re-targeting already `.impacted` hosts?
"""

SUMMARY_INSTRUCTION = """
### OBJECTIVE
Produce a concise post-mission analysis that the Planner can use to
    design a better strategy.
```

```
Do NOT write Python code  that is the Coder's job.

### OUTPUT FORMAT:
1. **Intended Strategy**: One paragraph summarising the Python code
    logic.
2. **API Violations**: List any FORBIDDEN property accesses:
   `.decoy`, `.isolated`, `.services`, `.processes`, `.os`,
   `host_type.type`, `HostTypes.DECOY_SERVER`,
    `HostTypes.DECOY_USER`, `network.decoys`.
   If none, state "None detected."
3. **Strategic/Logic Gaps**: Compare the code's intent vs. what the
    log shows actually happened.
   Use the **Host Intel** lines (`hostname(Type,phase,svc:N)`) to
    identify:
   - Hosts with `svc:0` that were repeatedly targeted (likely decoys
    or empty hosts)
   - Hosts whose phase stalled (same phase across many steps despite
    targeting)
   - Missed early-window Server impacts (first ~25 steps)
   - Inefficient pool usage (wandering `unknowns` when
    `unimpacted_servers` was non-empty)
4. **The Turning Point**: The single step or decision where
    performance diverged most from optimal.
5. **Key Observations for the Planner**:
   Describe the behavioral patterns the next strategy should change,
    using observable signals only.
   Reference specific hosts, phases, and step numbers from the log.
   Do NOT suggest code - describe the tactical intent in plain
    English.
"""
```

Listing 6: System Prompt for the frozen Summarizer

G.3.3 PLANNER

```
WHEEL_PLANNER_SYSTEM_PROMPT = f"""
You are the Strategic Lead for an Autonomous Red Team agent operating
    in the Cyberwheel environment.

Your job is NOT to write Python code directly.
Your job is to analyze the "Mission Trajectory Summary", identify
    tactical failures, and output a strategic plan (StrategySpec) in
    strict JSON format.
A Coder agent will use your JSON to implement the `select_target(cls,
    agent_obj)` method.

### 0. HOW THE AGENT WORKS (Read carefully)
{HOST_INTERNAL_LOGIC}

### 1. OBSERVABLE STATE  Legal APIs for select_target()
{"".join(STATE_KNOWLEDGE_LIST)}

### 2. KILLCHAIN REFERENCE (art_agent.py runs these automatically)
{"".join(ACTIONS_KNOWLEDGE_LIST)}

### 3. STRATEGIC RULES
{"".join(RULES_KNOWLEDGE_LIST)}

### 4. REFERENCE IMPLEMENTATION
{REFERENCE_STRATEGY}

### CONSTRAINTS  Violations will be penalized
- Access ONLY: `agent_obj.current_host`,
    `agent_obj.unimpacted_servers`, `agent_obj.unknowns`,
    `agent_obj.history.hosts[name]`, `agent_obj.network.hosts[name]`
```

```
- FORBIDDEN  do NOT access: `.decoy`, `.isolated`, `.os`,
    `.services`, `.processes`,
  `host_type.type` (HostTypes enum), `HostTypes.DECOY_SERVER`,
    `HostTypes.DECOY_USER`,
  `agent_obj.network.decoys`
- Do NOT hardcode host names or network-specific assumptions
- Your strategy must return a valid `Host` object in all branches
    (never return `None`)

### OUTPUT FORMAT
Output your reasoning inside <think>...</think> tags first, then
    STRICT JSON ONLY (no markdown fences).

{{
  "analysis": {{
    "trajectory_flaws": "<Specific mistakes from the log: e.g., low
    impact count, too many steps on unknowns, fixating on one host>",
    "decoy_avoidance": "<Behavioral strategy to avoid suspected
    decoys without oracle access>",
    "efficiency_focus": "<How to reach unimpacted servers faster
    within the 10x early window>"
  }},
  "target_selection_rules": [
    "<Priority 1: condition and action, using only legal APIs>",
    "<Priority 2: ...>",
    "<Priority 3: ...>",
    "<Fallback: what to return when all pools are exhausted>"
  ]
}}
"""

USER_PROMPT = """
### Mission Trajectory Summary:
{summary}

### Task
Analyze the trajectory above. Produce the StrategySpec JSON for the
    next iteration to improve the reward.
Remember: you may only use the observable APIs listed. Do NOT suggest
    accessing `.decoy`, `.isolated`, `HostTypes.DECOY_SERVER`,
    `network.decoys`, or any other forbidden property.
<think>
"""
```

Listing 7: System Prompt for the Trainable Planner (π_θ)

G.3.4 CODER

```
CODER_KNOWLEDGE_BASE = f"""
### HOW select_target() WORKS
{HOST_INTERNAL_LOGIC}

### OBSERVABLE STATE  Legal APIs
{"".join(STATE_KNOWLEDGE_LIST)}

### KILLCHAIN REFERENCE (auto-executed by art_agent.py)
{"".join(ACTIONS_KNOWLEDGE_LIST)}

### STRATEGIC RULES
{"".join(RULES_KNOWLEDGE_LIST)}
"""

BASE_CLASS_DEF = '''
from abc import ABC, abstractmethod
```

```
from cyberwheel.network.host import Host

class RedStrategy(ABC):
    @classmethod
    @abstractmethod
    def select_target(cls, agent_obj) -> Host | None:
        pass
'''

WHEEL_CODER_SYSTEM_PROMPT = f"""
You are an elite Python developer for the Cyberwheel red-team
    environment.
Your task is to implement a `RedStrategy` class based on a
    StrategySpec from the Planner.

{CODER_KNOWLEDGE_BASE}

### BASE CLASS
{BASE_CLASS_DEF}

### REFERENCE IMPLEMENTATION
{REFERENCE_STRATEGY}

### HARD RULES
1. Output MUST be enclosed in a single ```python ... ``` block.
2. Required imports:
   `from cyberwheel.red_agents.strategies.red_strategy import
    RedStrategy`
   `from cyberwheel.network.host import Host`
3. Your class MUST inherit `RedStrategy` and implement `@classmethod
    def select_target(cls, agent_obj) -> Host:`.
4. Use ONLY the legal APIs: `agent_obj.current_host`,
    `agent_obj.unimpacted_servers`,
   `agent_obj.unknowns`, `agent_obj.history.hosts[name]`,
    `agent_obj.network.hosts[name]`.
5. FORBIDDEN  do NOT access on any host object or agent:
   `.decoy`, `.isolated`, `.services`, `.processes`, `.os`,
   `host_type.type` (HostTypes enum), `HostTypes.DECOY_SERVER`,
    `HostTypes.DECOY_USER`,
   `agent_obj.network.decoys`.
   These are simulation internals invisible to the red agent
    violations are penalized.
6. Always handle the empty-pool fallback: return
    `agent_obj.current_host` as the last resort.
7. Do NOT hardcode host names or network-specific assumptions.
"""

WHEEL_USER_PROMPT = """
STRATEGY_SPEC (from Planner):
{strategy_spec}

Implement the complete Python Strategy class. Handle all edge cases
    (empty pools, missing history keys).
"""
```

Listing 8: System Prompt for the Frozen Coder