An LLM-guided Moving Target Defense Framework For Virtual Power Plants

Juanwei Chen*, Jun Yan*, Anthony Kemmeugne†, Mohsen Ghafouri *, Mourad Debbabi*, Marthe Kassouf †
*Security Research Center, Concordia University
†Hydro-Québec Research Institute (IREQ)
Montréal, Québec, Canada

***Abstract*—To address cybersecurity challenges arising from the large-scale aggregation of distributed energy resources (DERs), this paper introduces a dynamic defense-in-depth and defense-in-breath framework by leveraging moving target defense (MTD), which dynamically alters the system's attack surface to disrupt or delay cyber intrusions in the context of virtual power plants (VPPs). A novel LLM-guided MTD framework is proposed, where a large language model (LLM) adaptively selects MTD strategies—target surfaces, techniques, and movement frequencies — by reasoning over cybersecurity, operational technology (OT), and information technology (IT) data. These insights guide the formulation of an optimization problem to finalize strategy deployment. The framework's effectiveness is validated through analysis of LLM-generated MTD actions and their impact on cyber-physical VPP operations.**

***Index Terms*—Distributed energy resources, virtual power plant, moving target defense, large language model, cybersecurity**

I. Introduction

DERs are increasingly adopted to enhance sustainability, efficiency, and grid resilience, motivating their aggregation through VPPs for coordinated operation. While advances in communication and information technology enable VPPs to improve grid stability and energy optimization, large-scale and heterogeneous DER integration also introduces significant cybersecurity challenges, including expanded attack surfaces and difficulties in enforcing consistent security across diverse devices and stakeholders.

MTD is a proactive cybersecurity strategy that enhances the resilience of VPP-integrated DERs by continuously reconfiguring system attack surfaces to invalidate attackers' reconnaissance. For example, periodically rotating DER communication endpoints or access credentials introduces controlled unpredictability, reducing exploitable attack windows without adding extra defense layers and making MTD well-suited for large-scale, dynamic VPP environments. Existing MTD research largely relies on game-theoretic models that assume rational attackers and defenders with substantial system knowledge, with recent extensions using partially observable stochastic games (POSGs) to capture uncertainty [1]. To further relax these assumptions, reinforcement learning (RL)–based MTD approaches have been proposed, enabling adaptive defense through environment interaction without requiring explicit game models [2]. Despite their adaptability, RL methods typically rely on reasonably accurate models of the environment to map actions to outcomes effectively, hence, limiting their scalability and generalizability across diverse attack surfaces and types, as well as the integration of cyber-physical dynamics.

LLMs appear as powerful tools in addressing the above research gaps based on their strong text-processing and reasoning capabilities. LLMs are increasingly used in cyber threat intelligence (CTI), which focuses on collecting and analyzing threat information to support security decision-making [3]. In particular, LLMs can automatically extract attackers' Tactics, Techniques, and Procedures (TTPs), describing how attacks are carried out, from unstructured sources such as threat reports, enabling faster and more informed threat assessment. These capabilities make LLMs well-suited to support MTD in VPPs by improving scalability and adaptability when securing large and heterogeneous DER aggregations.

Building on recent LLM-based TTP generation research [4], this study proposes an LLM-guided MTD framework that leverages the MITRE ATT&CK knowledge base [5], a standardized taxonomy of adversarial tactics and techniques, to map observed TTPs to defense actions using a predefined MTD library and playbook tailored for VPP cyber–physical environments. Guided by the cyber kill chain model, the LLM infers likely attack targets and intrusion stages, then suggests suitable MTD actions and parameters. These recommendations drive an optimization process that selects final MTD strategies based on real-time context and VPP cybersecurity needs. The main contributions are:

- Leveraging the reasoning and text-processing capabilities of LLMs to infer likely attack targets based on real-time cybersecurity, operational technology (OT), and information technology (IT) data, hence, navigate the MTD action selection.
- Proposing a dynamic MTD framework that adapts defense strategies dynamically to improve scalability and adaptability across diverse attack surfaces and threats, while an optimization process that mitigates hallucinations and enforces guardrails, ensuring each action aligns with system safety and operational objectives.

The rest of this paper is organized as follows: Section II

This research is supported by the Natural Sciences and Engineering Research Council of Canada (NSERC), the Hydro-Québec Hitachi Partnership Research Chair in Smart Grid Security, and PROMPT Québec.

979-8-3315-8138-1

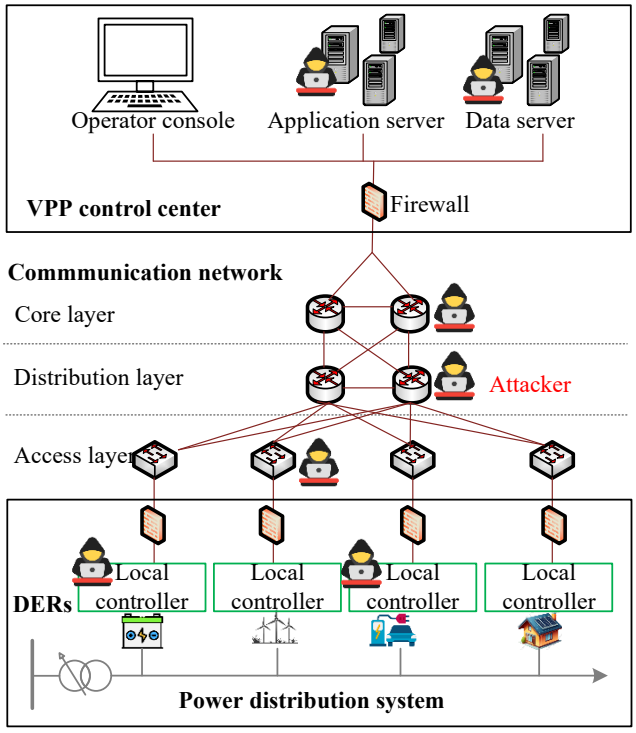

Fig. 1: The overall structure and attack surface of VPPs.

describes the cybersecurity challenges for VPPs and defines the problem addressed in this study. Section III details the proposed LLM-guided MTD framework. Section IV describes experiments and results, and Section V concludes the paper.

II. VPP and Cybersecurity Challenges

In this section, the main components of the VPP and their respective functions are first introduced. This is followed by a discussion of the cybersecurity challenges encountered and the specific problem addressed in this paper.

A. Overall VPP Framework

The VPP consists of DERs connected to the distribution network at the physical layer, a supporting communication network, and a centralized VPP control center, as illustrated in Fig. 1. The DERs encompass generation and storage units, along with embedded controllers responsible for local monitoring and control. These controllers interface with the communication network to report measurement data and receive control commands from the VPP control center. Given the wide geographic distribution, communication is facilitated through a three-tier hierarchical wide area network (WAN) architecture. The VPP control center functions as the central coordinating entity, aggregating and managing DERs to offer services such as energy trading, voltage regulation, and frequency support.

B. Cybersecurity Challenges and Problem Definition

VPPs aggregating large-scale DERs face significant cybersecurity challenges due to heterogeneity in DER types, communication protocols, and deployment locations, often outside secured utility zones. This diversity substantially expands the attack surface across DER devices, communication networks, and control applications [6], while operational variability (e.g., location-dependent voltage support roles) causes uneven system impacts under cyber compromise, motivating adaptive and context-aware defense mechanisms.

This paper investigates the use of MTD to enhance VPP resilience by dynamically reconfiguring system configurations to increase attacker uncertainty. Framed within the unified kill chain (UKC) model, which characterizes cyber attacks as a sequence of stages from reconnaissance to impact, the proposed approach targets early-stage attack disruption before physical grid impacts occur. The central research problem is:

- *How can we design an MTD framework for VPPs that dynamically adapts to diverse attack surfaces and stages of the cyber kill chain, while aligning with real-time operational and service constraints?*

III. LLM-guided MTD Framework

To address this problem, we propose an LLM-guided MTD framework that determines what to move, how to move, and when to move. The framework leverages the LLM's reasoning capability as a decision-support module to handle the complexity of cybersecurity management in cyber–physical environments (Fig. 2). By integrating cyber–physical operational data and partially observed cybersecurity information (e.g., TTPs), the LLM infers effective MTD actions (what and how to move), while an optimization algorithm thereafter determines moving frequency of the MTD actions (when to move).

A. LLM Assistant

To reduce hallucination but avoid reliance on large-scale training data, the proposed assistant adopts a retrieval-augmented generation (RAG) approach, grounding all reasoning in structured, runtime-retrieved cyber-physical system descriptors. The cyber-physical configuration dataset—comprising system assets, their configurations, and application-level parameters—is first incorporated into the LLM to provide system awareness. Subsequently, a comprehensive MTD action library, along with a playbook, is constructed to constrain the LLM generate validated MTD actions, by defining the components and techniques that can serve as moving targets based on industrial guidelines (e.g., NIST [7] and MITRE [8]) or academic research (e.g., [9]). Each MTD entry is characterized by attributes such as *name*, *description*, *scope*, *kill-chain effect*, *service disruption*, and *minimal moving interval*. Finally, a structured prompt is formulated to guide the LLM in MTD action recommendation, as shown in Fig. 3, where the system knowledge and defense objectives are first introduced, followed by the reasoning process and the expected output format.

Upon receiving TTP inputs, the assistant retrieves relevant OT and IT data to support adaptive MTD generation. Different VPP applications impose distinct control and communication requirements—for example, frequency and voltage support demand fast, high-rate interactions, whereas normal scheduling

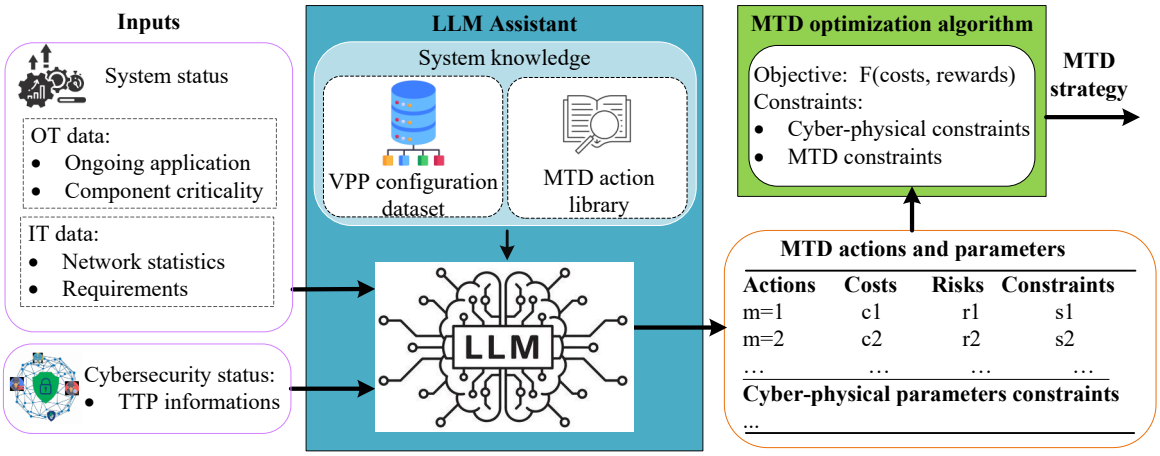

Fig. 2: The overall LLM-guided MTD framework.

is less time-critical—while heterogeneous DER capacities and sensitivities determine their defensive priorities. Accordingly, the LLM is provided with the current VPP application type, DER priority vector, and network latency conditions, and selects candidate MTD actions from a predefined, allowlisted library, parameterizing both the *what* and *how* to move via $\{t_{\mathrm{d}}, t_{\mathrm{min}}\}$ and associated cyber–physical constraints. All LLM outputs are subsequently passed through a deterministic validation layer that enforces feasibility and operational safety constraints, including latency limits, minimal moving intervals, and maximal allowable service disruption; only verified MTD actions are executed, while unsafe outputs are rejected in favor of safe fallback strategies.

B. Development of Optimization Algorithm

The objective of the optimization algorithm is to determine the optimal moving frequency of the MTD actions recommended by the LLM assistant, achieving a balance between the cost of movement and the risk of compromise. The moving cost reflects the disruption introduced to normal VPP operations, which varies depending on the application's timescale. To normalize this effect, the cost is quantified as the ratio between the disruption duration of an MTD action and the average control-loop duration of the corresponding VPP application, expressed as:

$$c_{m,s} = \frac{t_{\mathrm{d}}^{m}}{t_{\mathrm{c}}^{s}}, \tag{1}$$

where $c_{m,s}$ denotes the cost factor of the m-th MTD action under the s-th VPP application, t_{d}^{m} represents the disruption duration of the m-th action, and t_{c}^{s} is the average control-loop duration of the s-th application.

The risk of compromise increases with the frequency of communication between the VPP control center and DERs. Therefore, the risk factor for the s-th VPP application is defined as the reciprocal of its average control-loop duration:

$$r_s = \frac{1}{t_{\mathrm{c}}^{s}}. \tag{2}$$

Given that DERs may have different service priorities due to their different locations, types, and capabilities in voltage support, these priorities are incorporated into the optimization objective. The optimization problem aims to determine the moving interval $t_{i,m}$ of the i-th component for the m-th MTD action by minimizing the overall cost and risk, formulated as:

$$\min\ f(t_{i,m}) = \sum_{i=1}^{N_{\mathrm{D}}} \sum_{m=1}^{M} a_m \left(w_{i,s} r_s t_{i,m} + c_{m,s} \frac{T_{\mathrm{s}}}{t_{i,m}} \right), \tag{3}$$

where $w_{i,s}$ denotes the priority weight of the i-th component under the s-th application, T_s is the total MTD duration for that application, and a_m is a binary indicator representing the feasibility of each LLM-recommended MTD action, defined as:

$$a_m = \begin{cases} 1, & t_{\mathrm{d}}^{m} \le t_{\mathrm{s,max}}, \\ 0, & \text{otherwise}, \end{cases} \quad \forall i \in \mathbb{A}, \tag{4}$$

where $t_{\mathrm{s,max}}$ is the maximum allowable delay for the s-th VPP application and $\mathbb{A}$ denotes the set of MTD actions recommended by the LLM. Finally, each moving interval must satisfy the constraint ensuring it is no shorter than the minimal allowed interval for the corresponding MTD action:

$$t_{\mathrm{min}}^{m} \le t_{i,m}, \quad \forall i \in \mathbb{A}, \forall i \in \mathbb{C}. \tag{5}$$

The above algorithm is solved once the parameter is updated from the LLM assistant. It is convex with respect to the decision variable $t_{i,m}$, as it consists of a linear term and a reciprocal term that is convex for $t_{i,m} > 0$. Therefore, a unique global optimum exists.

You are the **VPP Moving Target Defense (MTD) Assistant**
1. Mission Context
The operating environment is a **Virtual Power Plant** aggregating geographically distributed **Distributed Energy Resources (DERs)**.
Your knowledge base includes:
- Physical and cyber assets, including network topology and accessibility mappings.
- Libraries of **MTD actions**.
2. Objectives
When a TTP or alert is received:
1. **Interpret the TTP** to identify the associated **Unified Kill Chain** stage(s).
2. **Assess affected assets** and potential propagation paths within the cyber-physical VPP.
3. **Generate a runtime MTD policy** based on the current operating state and threat stage.
4. **Recommend safe and effective MTD actions** aligned with the policy and environmental constraints.
3. Reasoning Framework
Follow this reasoning sequence:
Step 1 — Threat Interpretation:...
Step 2 — Context Assessment:...
Step 3 — Dynamic MTD Policy Generation:...
Step 4 — Output Composition:...
4. Required Output Format
kill_chain_stage: ["<Stage_1>", "<Stage_2>", ...]
recommendations:
1. action: <action_name>
params: { <key>: <value>, ... }
rationale: <why this mitigates the threat>
preconditions: <conditions before execution>
...

Fig. 3: Illustration of a simplified prompt template.

IV. Case Study

To validate the proposed MTD framework, the VPP physical layer is simulated using the IEEE 123-test feeder with 28 DERs, including solar PV, wind turbines, and energy storage systems, in OpenDSS to ensure accurate capture of DER impacts on distribution system operation. The communication network is emulated in Mininet to construct realistic virtual networks with 28 access switches, 7 distribution switches, and 3 core switches, integrated with a software-defined network (SDN) for dynamic reconfiguration, of which a real-world SDN controller, Ryu SDN, is deployed. Full platform details are available in [10]. The MTD action library and playbook are built on this cyber-physical setup, incorporating MTD techniques in both the IT and OT layers. An OpenAI model serves as the core of the LLM assistant, which could be entreprise-grade LLMs to avoid risks of, e.g., data leakage. Its generation, specified MTD strategies, and performance analysis are provided thereafter.

A. LLM Recommendations

To evaluate the LLM assistant's MTD recommendations, a DoS attack on distribution switches is analyzed. The attack progresses through initial WAN access, network reconnaissance, and a subsequent phase of bandwidth exhaustion. The LLM receives the following TTP input before the attacker moves forward to the final bandwidth exhaustion stage: ***Tactic:*** *Reconnaissance and discovery;* ***Technique:*** *Network service discovery (T1046);* ***Procedure:*** *Automated scanning of WAN ranges and subnets to identify exposed interfaces, open ports, and service banners for mapping distribution switch locations and potential pivot points.* In the simulated cyber-physical state, the VPP is delivering voltage support in response to a 30% sudden load increase between Buses 78 and 84. The DERs are prioritized for voltage support as follows: {24, 15, 17, 14, 26, 13, 11, 10, 23, 16, 25, 9, 27, 28, 22, 21, 8, 5, 6, 7, 20, 4, 3, 2, 19, 18, 1}.

TABLE I: LLM recommended MTD actions

Actions	Reasoning
Rotate-api-keys	Invalidate potential credential reuse after initial access, enhancing the security posture against unauthorized access.
Randomize-sdn-paths	Disrupt predictable communication paths, thereby impeding reconnaissance and reducing the risk of lateral movement.
Rehome-access-to-ds	Break attacker persistence by moving an access switch to another distribution switch, limiting lateral movement potential.

Based on these inputs, the LLM recommends three MTD actions: *rotate-api-keys* ($m = 1$), *randomize-SDN-path* ($m = 2$), and *rehome-access-to-ds* ($m = 3$). As summarized in Table I, these actions collectively disrupt the reconnaissance process and reduce the likelihood of follow-on FDI, MitM, and DoS attacks by denying consistent access to network services and hindering adversarial mapping of the infrastructure. Note the recommended MTD actions would vary along with attacks targeting different devices or system layer, which is not discussed in details here due to page limits.

B. MTD Strategies

Upon receiving the LLM-recommended MTD actions, an optimization process is executed to determine the moving interval of each action. The cyber-physical simulation platform is built on a Linux-based system running Ubuntu (64-bit) within a virtual machine configured with 6 GB of RAM and an Intel® Core™ i7-8700 processor operating at 3.2 GHz. The average interval for each voltage control loop is set as $t_{\mathrm{c}}^{s} = 0.5$ s (considering the maximal disruption introduced by the three actions) and $t_{s,\max} = 1.5$ s. The delay and minimum interval parameters for the three MTD actions are $t_{\mathrm{d}}^{m} = \{0.2, 0.15, 0.1\}$ and $t_{\min}^{m} = \{1, 0.3, 0.2\}$, respectively.

By incorporating the DER priority defined earlier, the optimization determines the moving interval for each action and DER. As shown in Fig. 4a, *rehome-access-to-ds* exhibits the shortest interval due to its lower cost factors. For the most critical DER group (e.g., DER 11–15 and 24–26), the moving interval is approximately 0.8 s, while less critical DERs (e.g., DER 1) have a longer interval of around 1.82 s. For *randomize-SDN-path*, the minimum and maximum intervals are approximately 1.0 s and 2.2 s, respectively, whereas for *rotate-api-key*, they are about 1.83 s and 4.06 s.

To demonstrate the adaptability of the proposed framework across different VPP applications, another case is evaluated under a normal VPP scheduling mode ($t_{\mathrm{c}}^{s} = 5$ min, $t_{s,\max} =$

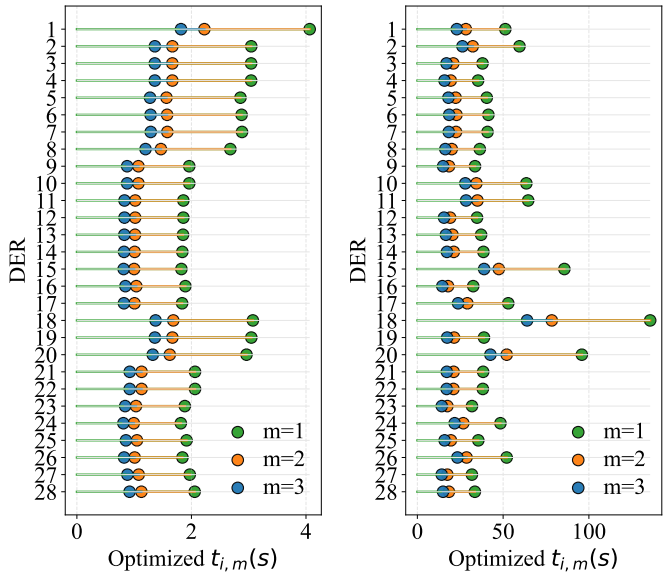

Fig. 4: Voltage support profiles of Bus 85 under different cases.

4 min) with different DER priorities. As illustrated in Fig. 4b, the optimized moving intervals range from 30–136 s for *rotate-api-key*, 17–79 s for *randomize-SDN-path*, and 14–63 s for *rehome-access-to-ds*, demonstrating strong adaptability to varying cyber-physical operating conditions.

C. Performance Analysis

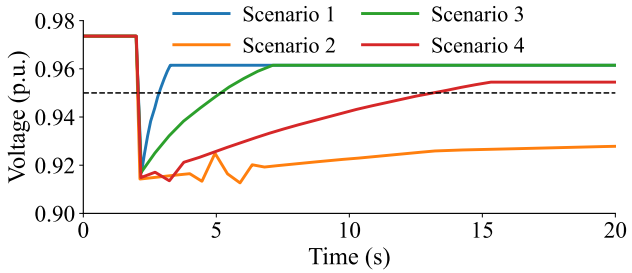

Fig. 5: Voltage support performance analysis

The effectiveness of the proposed MTD strategies against DoS attacks on VPP voltage support is evaluated by comparing four scenarios, specifically, *scenario 1:* normal operation without MTD; *scenario 2:* a DoS attack targeting four distribution switches, compromising DERs 1–4, 9–16, 21–24, without MTD; *scenario 3:* MTD-enabled operation with no successful compromise; and *scenario 4:* a worst-case scenario where partial compromise occurs despite MTD. Voltage responses at Bus 85 under these conditions are shown in Fig. 5. In the normal case, voltage is restored within 2 seconds over 10 control loops. Under attack without MTD, voltage restoration fails entirely. With MTD in place, all defense actions are executed at the start of the voltage support process. This forces the attacker to restart reconnaissance within short moving intervals (e.g., 0.8 s for critical DERs), significantly hindering their ability to proceed. If the attacker fails to execute the DoS attack, the only impact is a slight delay in voltage recovery to 5 seconds, as shown by the green curve. In the worst-case scenario — where the attacker successfully compromises DER communication within a short interval (0.8s) — some control loops are affected. However, after the access switches rehome, the DERs recover, and the voltage is prolonged but still successfully restored, as indicated by the red curve. Overall, the MTD strategies effectively disrupt or delay DoS attacks, preserving voltage support functionality with only minimal degradation under worst-case conditions.

V. Conclusion

This paper investigates the use of LLM to guide MTD strategies by integrating cybersecurity, OT, and IT data into adaptive decision-making for cyber-physical VPP operations. The experimental results demonstrate that LLM-guided MTD can effectively bridge the gap between cyber defense mechanisms and real-time system requirements, thereby enabling more responsive and context-aware protection of VPP services. Future work will focus on further exploiting the reasoning and contextual understanding capabilities of LLMs as human-like decision assistants in MTD, enhancing the scalability and adaptability of the framework across broader VPP applications, larger DER aggregations, and evolving threat landscapes.

References

[1] J. Tan, H. Jin, H. Zhang, Y. Zhang, D. Chang, X. Liu, & H. Zhang. A survey: When moving target defense meets game theory. Computer Science Review, 48, 100544, 2023

[2] T. Zhang et al., "Moving Target Defense Meets Artificial-Intelligence-Driven Network: A Comprehensive Survey," in IEEE Internet of Things Journal, vol. 12, no. 10, pp. 13384-13397, 15 May 2025.

[3] V. Clairoux-Trepanier, I.M. Beauchamp, E. Ruellan, M. Paquet-Clouston, S.O. Paquette, E. Clay. The use of large language models (LLMs) for cyber threat intelligence (CTI) in cybercrime forums. arXiv preprint arXiv:2408.03354. Aug 6, 2024.

[4] Y. Zhang, X. Zhou, H. Wen, W. Niu, J. Liu, H. Wang, & Q. Li. Tactics, techniques, and procedures (ttps) in interpreted malware: A zero-shot generation with large language models. arXiv preprint arXiv:2407.08532, 2024.

[5] MITRE ATT&CK®, 2026, https://attack.mitre.org/. Accessed 3 Feb. 2026.

[6] J. Chen, J. Yan, A. Kemmeugne, M. Kassouf, M. Debbabi. Cybersecurity of distributed energy resource systems in the smart grid: A survey. Applied Energy. 383:125364, April 1, 2025.

[7] R. Ross, V. Pillitteri, R. Graubart, D. Bodeau, and R. McQuaid, Developing Cyber-Resilient Systems: A Systems Security Engineering Approach. NIST Special Publication 800-160, Volume 2, 2021.

[8] D. Bodeau, R. Graubart, W. Heinbockel, E. Laderman, "Cyber Resiliency Engineering Aid – The Updated Cyber Resiliency Engineering Framework and Guidance on Applying Cyber Resiliency Techniques," Mitre, May 2015.].

[9] J. -H. Cho, D. P. Sharma, H. Alavizadeh, S. Yoon, N. Ben-Asher, T. J. Moore, S. K. Dong, H. Lim, F. F. Nelson. "Toward Proactive, Adaptive Defense: A Survey on Moving Target Defense," in IEEE Communications Surveys & Tutorials, vol. 22, no. 1, pp. 709-745, Firstquarter 2020.

[10] Chen J, Kemmeugne A, Yan J, Ghafouri M, Kassouf M, Debbabi M. An SDN-based Framework for Cyber-physically Coordinated Voltage Support of Virtual Power Plants Against DoS Attacks. IEEE Transactions on Smart Grid. Sep 26, 2025.