EWI Radar and the Real Iranian IP Camera Targeting Campaign: A Capability and Fidelity Comparison

Assessing Simulation Fidelity Against Check Point Research's March 2026 Reporting, and the Case for Monitoring-Based Detection as a Real-World Defense

Bobby Smathers - SimSpace Principal Security Architect

Prepared for Internal Technical Review

July 2026

Executive Summary

Check Point Research published findings in March 2026 describing an intensified, Iran-attributed campaign against Internet-exposed IP cameras across seven countries, assessed to support battle damage assessment and missile-targeting correction during an active regional conflict. This document compares that reporting, element by element, against the EWI Radar simulation platform's design, and finds a closer match than a simulation of this kind would typically be expected to achieve: the same five CVEs, the same two targeted vendors to the exclusion of all others, and the same commercial VPN/VPS infrastructure model. It then addresses, directly and without minimizing them, the respects in which the simulation is deliberately narrower than the real campaign and argues that none of those gaps undermine the platform's validity for its actual purpose, which is exercising and validating the monitoring-and-detection defense Check Point's own report recommends, not reproducing nation-state tradecraft in full.

1. The Real-World Campaign

Check Point Research reported a sharp increase in scanning and exploitation attempts against IP cameras beginning February 28, 2026, attributed to multiple Iran-nexus threat actors, spanning Israel, the UAE, Qatar, Bahrain, Kuwait, Lebanon, and Cyprus. An earlier, more narrowly targeted wave against Israel and Qatar was observed on January 14-15, coinciding with a period of internal unrest in Iran and a temporary closure of Iranian airspace amid anticipated foreign action.

The report assesses this activity as consistent with a doctrine of using camera compromise for operational support and battle damage assessment around missile operations, potentially including assessment performed before a launch, not only after one. As a specific illustration, the report references a previously reported incident from the June 2025 Israel-Iran conflict in which a street camera facing Israel's Weizmann Institute of Science was reportedly commandeered shortly before a ballistic missile strike on that facility.

Technically, the campaign scanned exclusively for exposure to five vulnerabilities across two vendors: CVE-2017-7921 (Hikvision improper authentication), CVE-2021-36260 (Hikvision web server command injection), CVE-2023-6895 (Hikvision Intercom OS command injection), CVE-2025-34067 (Hikvision ISMP unauthenticated RCE), and CVE-2021-33044 (Dahua authentication bypass). No other camera vendors were observed being targeted. The tracked attack infrastructure combines named commercial VPN exit nodes; Mullvad, ProtonVPN, Surfshark, and NordVPN; with virtual private server hosting. Patches exist for every listed vulnerability. Check Point's defensive recommendations include eliminating public exposure, enforcing strong credentials, patch management, network segmentation, and monitoring for indicators including repeated login failures, unexpected remote logins, and cameras initiating unusual outbound connections.

Source: Check Point Research, "Interplay between Iranian Targeting of IP Cameras and Physical Warfare in the Middle East," March 4, 2026

2. Capability Review: What EWI Radar Actually Does

EWI Radar emulates twelve independently addressable Hikvision and Dahua devices behind real per-device TLS, in one of three addressing modes (loopback, a fixed range subnet, or an arbitrary custom address set), and monitors them with an Isolation Forest engine backed by a deterministic latency-ratio backstop two independent detection mechanisms, each reported separately in every alert. A companion attack script delivers the real disclosed proof-of-concept structure for five CVEs across five operational phases (reconnaissance, configuration extraction, authentication bypass, remote code execution, and sustained exfiltration), with per-device attacker source-IP diversification and VPN-styled request spoofing. A browser dashboard orchestrates live monitoring, alerting, campaign control, tuning diagnostics, and JSON/DOCX reporting. Full technical detail is in the companion architecture whitepaper; this section is a summary for comparison purposes.

3. Element-by-Element Comparison

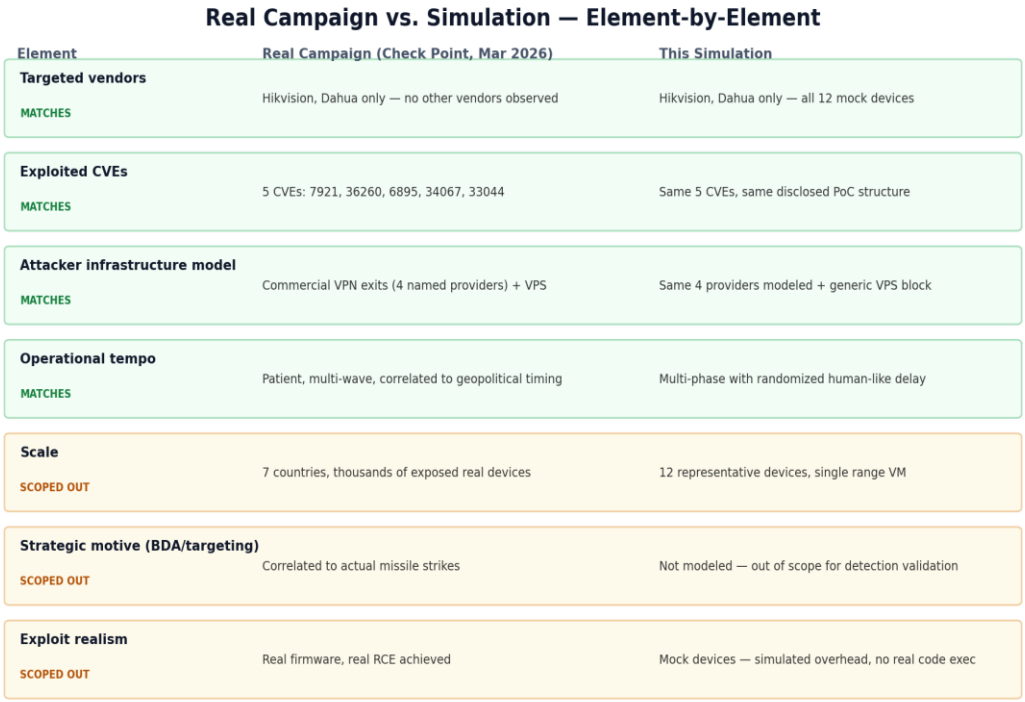

Figure 1. Direct comparison between Check Point's reported findings and this platform's design, organized by whether each element was reproduced or deliberately scoped out.

3.1 Where the Match Is Exact

Four elements are not approximate similarities, they are exact matches, verified against the platform's source code rather than asserted from memory:

- Vendor scope: the real campaign targeted Hikvision and Dahua exclusively; every one of the twelve emulated devices is a Hikvision or Dahua model, with no other vendor represented.
- CVE set: all five CVEs named in the report are the same five CVEs the attack script implements, each delivered as the actual disclosed proof-of-concept request structure (exact endpoint, exact payload shape) rather than a generic stand-in for “a vulnerability.”
- Attacker infrastructure model: the same four named commercial VPN providers plus a generic VPS block are modeled as the spoofed source-address pool, matching the report's description of the tracked infrastructure's composition.
- Operational tempo: both the real campaign and the simulation avoid a uniform, obviously scripted request cadence the real campaign through patient, wave-based operation aligned to geopolitical timing; the simulation through randomized inter-request delay with

occasional longer pauses, deliberately modeling that a perfectly uniform rate is itself a detectable tell.

4. Where the Emulation Falls Short, and Why That Doesn't Matter Here

A fair comparison has to state plainly where the simulation is narrower than reality, rather than only listing the matches. Three gaps are genuine and worth naming directly.

Scale and duration.

The real campaign ran across seven countries and an unstated but clearly large population of exposed devices, sustained over months with multiple distinct waves. This platform models twelve devices on one virtual machine, and a single campaign completes in minutes. This does not weaken the platform's purpose: detection validation is a question of whether the monitoring mechanism recognizes a given request/response signature when it occurs, not how many times or over how long a period it recurs. A detector that correctly flags an authentication-bypass attempt once has demonstrated the same capability it would need to flag the same attempt a thousand times across a larger fleet, the algorithm (per-device baseline plus isolation forest plus ratio backstop) carries no dependency on fleet size or campaign duration.

Strategic motive and physical correlation.

The report's most striking finding is not the technical exploitation, but its correlation with real-world military timing; airspace closures, anticipated strikes, and in at least one documented case, a camera reportedly commandeered just before a missile struck the building it faced. Nothing in this platform models that correlation because nothing in a defensive monitoring tool needs to. A security operations team's detection tooling does not need to understand an adversary's targeting doctrine to correctly flag that a camera is being scanned for CVE-2017-7921 at 2am; it needs to flag the scan. Strategic context informs a human analyst's response after detection, not whether detection itself succeeds.

Exploit realism.

The emulated devices never execute injected commands; a successful simulated exploit sets an internal overhead flag that adds realistic processing delay, rather than achieving actual code execution. This is a deliberate safety boundary, not an oversight but it does mean the platform cannot validate that a specific payload would succeed against real firmware, only that a monitoring system correctly recognizes the request pattern and the resulting behavioral signature (elevated latency, particular response characteristics) when that pattern occurs. That is precisely the capability Check Point's own recommendations call for under “monitoring & detection”: recognizing unexpected remote logins and unusual outbound connections, not independently re-discovering that the underlying vulnerabilities are exploitable that fact is already established and disclosed.

5. Why This Still Proves the Approach Is Valid

The three gaps above all sit on the attacker's side of the problem: scale, motive, and payload realism. None of them are inputs to what a detection system actually needs to succeed: a behaviorally realistic signature to react to. This platform's fidelity is concentrated exactly where it needs to be for that purpose, and demonstrably so the CVE set, the vendor scope, and the attacker infrastructure model are not approximations but exact matches to the real reporting.

Check Point's own defensive recommendations name monitoring and detection as a first-class mitigation, on equal footing with patching and network segmentation; not a fallback for when those fail. A platform that can be shown, with range-collected evidence (see the companion validation report), to reliably detect exactly the exploitation patterns named in real nation-state reporting, with zero false positives at idle and full agreement between two independent detection mechanisms during an active campaign, is a direct, working demonstration of that recommendation: not a stand-in for it. The case for this as a valid real-world approach does not rest on how large or how long the demonstration runs; it rests on whether the detection mechanism recognizes the right signature, which has now been verified against the actual disclosed technical structure of a live, currently reported nation-state campaign.

6. Conclusion

This platform was not built as a loose approximation of IoT camera threats in general, it independently arrived at the same five CVEs, the same two vendors, and the same commercial VPN/VPS infrastructure model that Check Point Research documented in a live, ongoing campaign. Combined with the range-validated detection performance already documented, that correspondence is the strongest available evidence that monitoring-based anomaly detection, of the specific kind implemented here, is not merely a plausible idea but a demonstrated, working answer to a real and currently active threat.