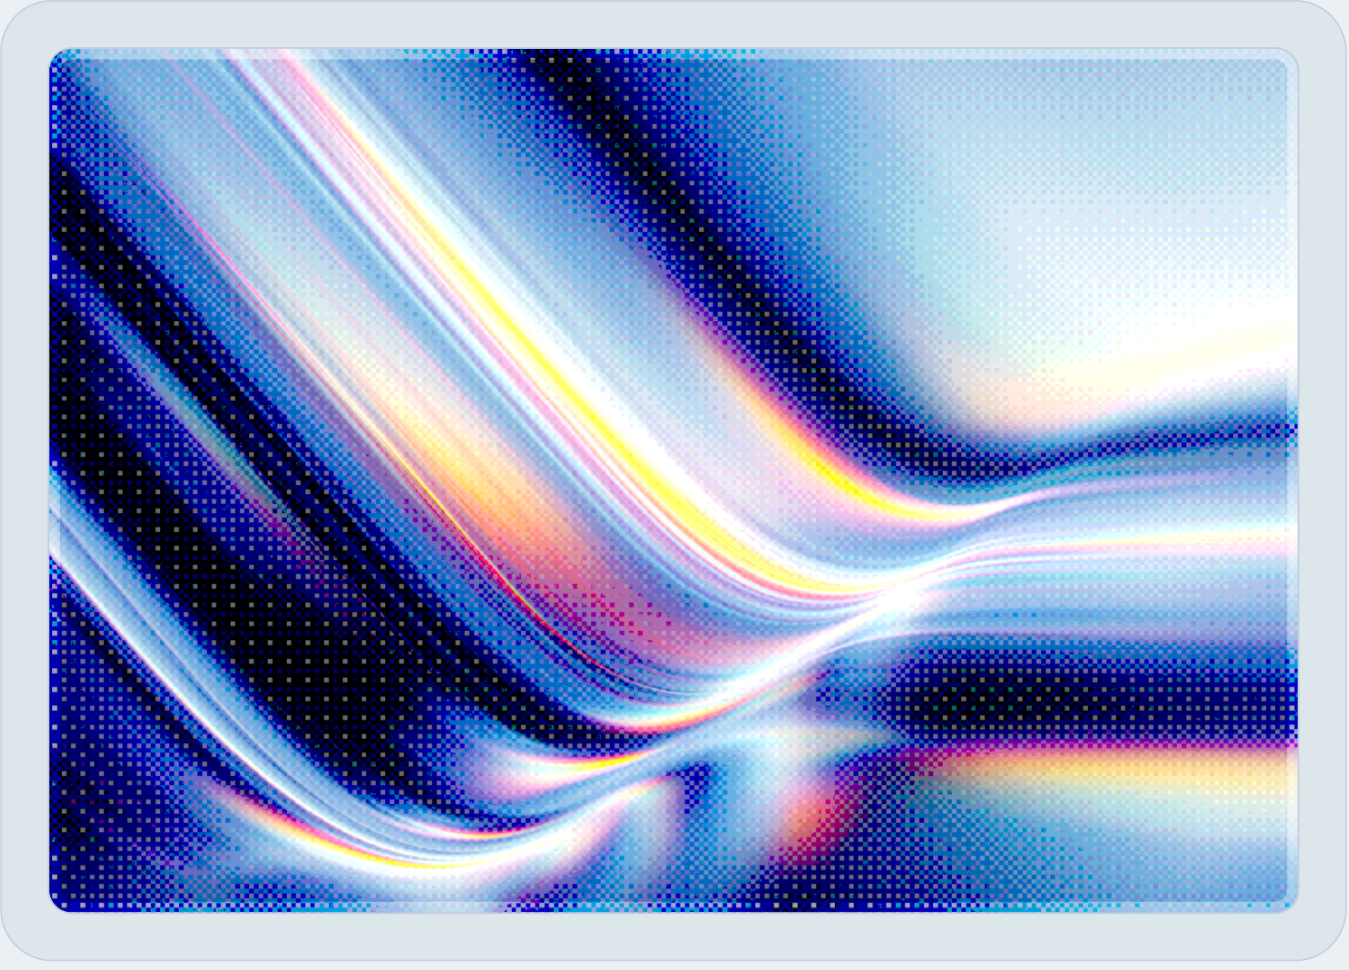
CSAI Foundation | Cloud Security Alliance

GT13: Suspected AI-Driven Intrusions Against Security Management Planes

Lessons from the compromise of 24 Check Point Security Management Servers

2026-11-25

AI-assisted Rapid Research



CSAI

cloud security alliance®

© 2026 Cloud Security Alliance. Some rights reserved.

You may download, store, display, view, print, redistribute, and link to this document in its original, unmodified form, provided that attribution to the Cloud Security Alliance is maintained and all trademark and copyright notices remain intact.

This document may not be modified or altered. You may quote portions of the document as permitted by the Fair Use provisions of the United States Copyright Act, provided that attribution is given to the Cloud Security Alliance.

This document may be shared on professional and social media platforms in its original form with attribution.

This document was generated with AI assistance and has not undergone official CSA review and approval processes.

Key Takeaways

GreyNoise reports that an unattributed actor it tracks as GT13 stole data from 24 Check Point Security Management Servers at no fewer than 22 organizations in 17 countries and territories. Confirmed attack activity ran from 17 October to 2 November 2026 [1]. The victims span finance, energy, government, transportation, mining, engineering, and manufacturing. GreyNoise expects the campaign to continue because the actor appears to keep persistent access to the compromised servers [1][4].

The target is what makes this campaign significant. A security management server is the control plane for an organization's firewall estate. It holds the firewall database, policy, administrator credentials, and the internal certificate authority (CA) that gateways use to trust one another. GreyNoise observed the actor copying firewall databases and internal CA material, minting a new certificate named `pivot.p12`, and unloading security policies. Unloading policy effectively disables the firewall [1].

The actor used an exploit for CVE-2026-16232 to run more than 100 management tasks [1]. That flaw is a SmartConsole authentication bypass, rated CVSS 9.3, which Check Point patched and disclosed as exploited in the wild on 22 July 2026 [2]. GreyNoise says it could not determine how or when the actor first got into the Check Point servers, so the initial access vector is unconfirmed [1].

GreyNoise suspects a "mature AI-enabled offensive apparatus" using separate agentic harnesses for scanning, exploitation, and follow-on interaction, alongside human operators [1]. The evidence for AI involvement is behavioral and inferential: automated check-ins, a bulk task looped across victims in IP order, and troubleshooting failures that GreyNoise attributes to a possible multi-agent architecture problem. It does not show directly that AI was used.

Organizations running Check Point Security Management or Multi-Domain Management should do four things now. Confirm the July 2026 jumbo hotfix is installed. Restrict GUI client and management access to trusted addresses. Hunt for the indicators GreyNoise published. Treat any server that was internet-exposed before patching as potentially compromised, including its internal CA.

Background

The campaign as reported

On 23 November 2026 GreyNoise published "99 IPs but a C2 ain't one," which describes an active intrusion campaign against Check Point Security Management Servers [1]. The Hacker News, SecurityWeek, and Security Affairs covered it over the next two days [3][4][5]. GreyNoise saw most of the activity through a server the actor had compromised and turned into a relay. The relay ran Rejetto HFS, a file server that was not a Check Point product.

The actor exploited CVE-2026-61500 on that server on 2 October 2026. It rebuilt the internal state of the V8 JavaScript engine's `Math.random()` generator from session identifiers, recovered enough to forge a signed administrator cookie, and was making administrator requests within about two minutes [1] [5]. Three days later it added an SSH key to the server's root `authorized_keys` file. From 20 October the server acted as a "hop," running an unauthenticated upload service, Nmap, and a Chisel reverse proxy [1].

The first archive uploaded to the hop contained route-scan output from a Check Point Security Management Server. GreyNoise concludes that the actor had access to the Check Point servers before it compromised the HFS server, although it could not determine when that access began [1]. Between 20 and 26 October, a host inside the victim GreyNoise calls "Victim A" ran a Chisel client. Through that tunnel the actor used a script named `exploit.py`, which GreyNoise describes as "purportedly an adaptation of the Metasploit module" for CVE-2026-16232, to issue more than 100 management tasks [1].

A companion `pull_` script chained two exploit helpers to reach a second Check Point server through the first one. It took credentials, internal CA files, and firewall configuration [1]. The actor then moved on to internal reconnaissance and a mix of older and recent enterprise exploits, summarized in the table below, and ran password-spraying scripts with credentials themed to the victim [1].

Product	CVEs reported by GreyNoise [1]
Windows Netlogon	CVE-2020-1472, CVE-2026-41089
Citrix NetScaler	CVE-2023-4966, CVE-2025-5777
Ivanti	CVE-2024-21893
Microsoft Exchange	CVE-2021-26855, CVE-2021-34473

Product	CVEs reported by GreyNoise [1]
VMware vCenter	CVE-2021-21972, CVE-2023-34048

On 26 October the actor sent a bulk certificate-management task to all 24 compromised servers. GreyNoise did not see the task being sent, which in CSA's reading suggests the actor issued it from other infrastructure. The results arrived on the hop over `scp`, one file per server in increasing IP order. Each file recorded an error, and one showed a preload failure for `/usr/lib/.libproc.so` registered in `/etc/ld.so.preload`. GreyNoise says this "may indicate an implant" [1]. The victim IP addresses in those files allowed GreyNoise, working with government partners and CERTs, to notify affected organizations [1][5].

The vulnerability context

Check Point's July 2026 advisory describes CVE-2026-16232 as an authentication bypass in the SmartConsole login using an application token. It affects Security Management and Multi-Domain Management on R81.10, R81.20, R82, and R82.10, and older versions are also affected [2]. The NVD entry describes the impact: an unauthenticated attacker can obtain a login token and get full administrator access through SmartConsole [6].

At disclosure, Check Point said in-the-wild exploitation affected "a handful of customers" and "only affects a very specific configuration – when Management is exposed directly to the internet without IP restrictions" [2]. It published six attacker IP addresses alongside the fix. In CSA's comparison, none of the 69 ExpressVPN-attributed addresses in GreyNoise's indicator table (some inferred at /24 granularity) match them [1][2]. Responding to the GT13 reporting, Check Point told The Hacker News that the activity involves "a previously disclosed vulnerability with patches already available." It said the activity "does not represent a new vulnerability or exploitation technique," and it recommended updates and hardening that restricts management access [3].

GT13 follows closely on another campaign that used a Check Point component as a platform for further intrusion. On 12 November 2026, Check Point Research reported a separate campaign that exploited unpatched Mobile Access gateways through PHP2Shell. That actor installed appliance-specific backdoors that can tamper with firewall logs and survive backups and snapshots [7]. The two campaigns are not reported as linked, but they fit a longer-running pattern of attackers treating security infrastructure as a target in its own right.

© 2026 Cloud Security Alliance. AI-assisted Rapid Research.

Security Analysis

Why the management plane is a high-value target

Compromising a firewall gateway gives an attacker a foothold. Compromising the management server gives the attacker authority over every gateway it manages. GreyNoise's account describes three actions that follow from that authority, and each carries a distinct consequence [1].

The first is copying the firewall database, which exposes every network object, rule, and trust relationship the organization has defined. That gives the attacker a ready-made map of the network for later lateral movement. The second is stealing the internal CA and minting a certificate with a name such as `pivot.p12`, which may let an attacker impersonate trusted components. Check Point's internal CA underpins Secure Internal Communication (SIC) between management and gateways, so a stolen CA key is not fixed by patching; it requires re-establishing trust. The third is unloading policy, a direct defense-impairment action, because a gateway with no policy loaded no longer enforces segmentation.

Taken together, these actions give an adversary three capabilities. It can see the defensive architecture, it can create credentials that defenders' own infrastructure will trust, and it can switch enforcement off. In CSA's assessment, the theft of CA material is the most consequential long-term risk in this campaign. Its impact lasts beyond any single patch cycle, and in CSA's experience few organizations have rehearsed a procedure for rotating it.

Assessing the "AI-driven" characterization

The claim that GT13 is AI-driven needs careful reading. GreyNoise states that "like most modern attacks, the actor used LLMs and agentic systems to perform vulnerability scanning, exploitation" and other phases. It also says it suspects "some of the foibles in the timeline reflect agentic offensive operations" [1].

Several observations support that reading. A scripted check-in on 8 October ran `uname -a`, `df -h`, and `whoami`, and was followed by eleven days of silence. On 19 October the actor tried to download a Linux implant called ZULT through an HFS feature even though it already had SSH access, then troubleshot the failed download without finding the cause; GreyNoise suggests this was possibly "a problem with the multi-agent architecture." The bulk certificate task looped across 24 victims. Finally, on 2 November the actor quickly checked its own relay for CVE-2026-98432 and attempted to mitigate it with an `nftables` rule [1].

There are limits to this evidence. GreyNoise notes a gap in hands-on-keyboard activity between roughly 17:00 and 00:19 UTC, but says it does not put much weight on time-zone analysis [1]. In CSA's view, automated loops, scripted check-ins, and odd tool behavior are equally consistent with conventional scripted operations. The reported tooling is also the standard offensive toolkit: Nmap, Chisel, Impacket, Paramiko, and at least 30 Python scripts [1][4].

CSA therefore treats AI involvement in GT13 as plausible and consistent with the evidence, but not demonstrated. Defenders should not rest their response on that question, because the defensive implications do not depend on it.

What makes the AI framing relevant is the pattern it fits. Unit 42's November 2026 report on CL-STA-1238 is better documented. In that campaign an actor used a forked open-source agent framework to run 25 operations against at least 19 organizations and gained access or retrieved data at 10 [8]. Unit 42 concluded that such orchestration "raises the volume and coverage of exploitation, but not its complexity" [8].

GT13 appears to fit that pattern. Reaching 24 management servers at 22 organizations in 17 countries and territories looks primarily like a problem of breadth. Every Check Point technique GreyNoise observed after initial access used CVE-2026-16232, a known and patched vulnerability, or a password-spraying script built on common, default, and weak passwords, although the initial access vector remains undetermined [1]. Elsewhere in the chain the actor showed real technical capability, notably the V8 PRNG-state recovery used against the HFS relay. Agentic tooling may lower the operator-hours cost of reaching many such servers at once, so exposures that once sat below a capable actor's priority line may now fall above it.

The exposure and patch-gap problem

Check Point's July statement framed exploitation as limited to management servers exposed to the internet without IP restrictions [2]. GT13 also illustrates a second-order risk of that configuration. The reported activity suggests that some victims remained exploitable months after the July fix. But even organizations that patched promptly may have been compromised during the exposure window, because a patch does not remove access that was already established. GreyNoise's report that the actor maintains persistent access to victim devices [1] underscores why post-patch verification matters.

Patching closes the door. It does not evict an attacker who is already inside, and it does not invalidate CA material or administrator credentials that were already stolen. The Check Point Research PHP2Shell case points the same way: there, implants on Check Point gateways were designed to survive backups and snapshots [7].

The table below summarizes the management-plane actions GreyNoise reported and the defensive implication of each.

Reported GT13 action [1]	Asset affected	Why patching alone is insufficient	Defensive priority
Copying the firewall database	Rulebase, objects, topology	Disclosure of the network map cannot be reversed	Assume topology is known; review segmentation assumptions
Copying internal CA files; minting `pivot.p12`	Internal CA and trust relationships	Stolen keys stay valid until trust is re-established	Plan CA and SIC trust re-establishment with vendor support
Credential theft via `pull_` script	Administrator and service credentials	Credentials stay usable after the fix	Rotate all management and administrator credentials
Unloading security policies	Gateway enforcement	Can recur if access persists	Alert on policy uninstall and unscheduled installs
`ld.so.preload` entry for `/usr/lib/.libproc.so`	Server operating system integrity	Possible implant persists across patching	Integrity check; rebuild from known-good media
Pivoting through Chisel to a second management server	Adjacent management servers	Lateral paths exist independent of the CVE	Isolate management networks; restrict server-to-server access

Recommendations

Immediate Actions

Organizations running Check Point Security Management or Multi-Domain Management should first confirm that the 22 July 2026 jumbo hotfix covering CVE-2026-16232 (sk185169) is installed on every management server. They should then apply Check Point's two published mitigations [2]. The first is to limit Trusted Clients (GUI clients) to trusted IP addresses or subnets. The second is to protect management access with the firewall, restricting it to trusted addresses, and to confirm that the implied rules for control connections are enabled.

Next, compare management server logs and network telemetry against the indicators GreyNoise published [1]. Because CVE-2026-16232 was already being exploited when it was disclosed on 22 July, the review should start from the earliest date the server was internet-reachable, and no later than June 2026. The GreyNoise indicators are:

- the `newpayapp[.]com` domain and its subdomains
- the listed ExpressVPN egress addresses
- the HASSH value `eeca2460550b9ded084ecf2f70a75356`, which GreyNoise says should be combined with other traits
- the user-agent strings ending in "research"

Teams should also check for the July attacker addresses published by Check Point [2]. On the servers themselves, look for an unexpected `/etc/ld.so.preload` entry (particularly `/usr/lib/.libproc.so`), Chisel binaries or processes, unfamiliar SSH keys in root's `authorized_keys`, and certificate objects that administrators did not create. Management audit logs should be reviewed for script executions, certificate operations, and policy uninstalls during the 17 October to 2 November window.

If any management server was reachable from the internet without IP restrictions before it was patched, treat it as suspect regardless of current patch status. Rotate administrator credentials and plan a rebuild from known-good media, following the vendor's guidance.

© 2026 Cloud Security Alliance. AI-assisted Rapid Research.

Short-Term Mitigations

Over the next 30 to 60 days, organizations should place security management interfaces on a dedicated, non-routable management network. They should reach them only through a hardened administrative jump path with phishing-resistant multifactor authentication, and remove any remaining internet exposure. Where a compromise is suspected or cannot be ruled out, work with Check Point support to re-establish trust between management and gateways and to rotate internal CA material. Stolen CA keys are not neutralized by software updates.

Detection should be extended to alert on management-plane actions that rarely happen legitimately. These include policy uninstall events, certificate creation outside change windows, Run Script tasks, and outbound connections from management servers to anything other than managed gateways and vendor update services. Each of these corresponds to an action GreyNoise observed or to the relay behavior that exposed the campaign.

Because GT13 brought exploits for known NetScaler, Ivanti, Exchange, vCenter, and Netlogon vulnerabilities to bear once it was inside [1], organizations should confirm that internal-only systems are patched rather than assuming the perimeter protects them.

Strategic Considerations

GT13 reinforces a broader shift. Security management planes, along with configuration-management, RMM, and virtualization management servers, are high-leverage targets whose compromise turns defensive infrastructure against its owner. Organizations should classify these systems as tier-zero assets, on a par with identity providers, and apply the same standards of isolation, monitoring, and credential hygiene.

Whether or not AI was central to GT13, the economics described in recent reporting [8] and in CSA's guidance on AI-accelerated exploitation [14] suggest that every internet-reachable management interface should be assumed to be found and tested quickly once a relevant flaw is public. Two metrics follow. The first is time from patch to deployment on security infrastructure. The second is time to verify that a server was not compromised before it was patched. CSA recommends that security leaders measure both, alongside the more familiar remediation metrics. Incident response plans should also include a rehearsed procedure for rotating firewall management trust anchors, because in CSA's experience it is a recovery step few organizations have practiced.

CSA Resource Alignment

CSA's July 2026 research note, *CVE-2026-16232: Check Point SmartConsole Authentication Bypass* [9], analyzed, at the time of disclosure, the flaw that GT13 later used. GT13 adds a lesson to that guidance. Patching an affected management server does not settle whether it was compromised before the fix. Post-patch integrity verification and trust-anchor rotation belong in the same remediation workflow.

Hardening Your Network Against AI-Accelerated Attacks, published under the "Preparing Your Networks for the AI Storm" series [10], gives phased guidance on segmentation, patchability, and responsive network operations for adversaries who use AI to accelerate exploitation. Its emphasis on isolating management interfaces and shortening patch cycles for network infrastructure applies directly to the exposure that allowed GT13 to reach Check Point management servers. Security teams can use it as the architectural roadmap behind the short-term mitigations above. At the program level, *The "AI Vulnerability Storm": Building a "Mythos-ready" Security Program* [14] is CSA's briefing for security leaders on AI-accelerated exploitation and autonomous attacks, and it frames the patch-tempo and exposure assumptions behind the Strategic Considerations above.

CSA's *Frontier Ready Maturity Model* [11] is a five-level maturity model across twelve security categories. It measures readiness to defend against AI-accelerated attacks. Organizations can use it to judge whether their exposure management, asset classification, and detection maturity would have surfaced a GT13-style compromise of their management plane. The lab note *State-Level Tempo for Everyone: Commodity Agent Frameworks and Adversary Tiering* [12] explains why the breadth observed in GT13 and CL-STA-1238 changes the threat profile for organizations that previously assumed they were below a capable actor's priority line.

Finally, the *AI Controls Matrix (AICM) v1.1* [13] provides the control baseline for formalizing these actions. Its Identity and Access Management controls cover privileged administrative access, its Cryptography, Encryption and Key Management controls cover CA key rotation, and its Threat and Vulnerability Management controls cover patch verification on security infrastructure.

References

[1] Miller, Steve (GreyNoise). "99 IPs but a C2 ain't one: An Active Intrusion Campaign Against Check Point Security Management Servers." GreyNoise, November 23, 2026.

[2] Finkelstein, Lotem (Check Point). "Security Advisory – Action Required – July 2026 Security Update." Check Point Blog, July 22, 2026.

[3] Lakshmanan, Ravie. "Hackers Compromise 24 Check Point Management Servers in Suspected AI-Driven Attack Campaign." The Hacker News, November 25, 2026.

[4] Arghire, Ionut. "Check Point Security Management Servers Hacked in AI-Assisted Campaign." SecurityWeek, November 24, 2026.

[5] Paganini, Pierluigi. "GT13: AI-Assisted Hackers Target 22 Organizations Through Check Point Servers." Security Affairs, November 25, 2026.

[6] NIST. "CVE-2026-16232 Detail." National Vulnerability Database, July 22, 2026 (last modified October 24, 2026).

[7] Check Point Research. "PHP2Shell Exploitation on Check Point Gateway." Check Point Research, November 12, 2026.

[8] Palo Alto Networks Unit 42. "Attackers Use Autonomous AI to Target Semiconductor, Defense and Civil Society Sectors." Unit 42, November 12, 2026. (Describes cluster CL-STA-1238 and the Cairn framework.)

[9] Cloud Security Alliance AI Safety Initiative. "CVE-2026-16232: Check Point SmartConsole Authentication Bypass." CSA Lab Space, July 25, 2026.

[10] Cloud Security Alliance. "Hardening Your Network Against AI-Accelerated Attacks (Preparing Your Networks for the AI Storm)." Cloud Security Alliance, 2026.

[11] Cloud Security Alliance. "Frontier Ready Maturity Model." Cloud Security Alliance, 2026.

[12] Cloud Security Alliance AI Safety Initiative. "State-Level Tempo for Everyone: Commodity Agent Frameworks and Adversary Tiering." CSA Lab Space, November 13, 2026.

[13] Cloud Security Alliance. "AI Controls Matrix (AICM) v1.1." Cloud Security Alliance, 2026.

[14] Cloud Security Alliance. "The 'AI Vulnerability Storm': Building a 'Mythos-ready' Security Program." Cloud Security Alliance, 2026.