xage SECURITY

WHITEPAPER

CISA's Emergency Directive on Cisco VPNs [CISA ED 25-03]: Short-term and Long-term Response Strategy

Urgent Directive for Cisco ASA/Firepower – A Symptom of Bigger Issues

Late September 2025 saw U.S. federal agencies scrambling under CISA Emergency Directive 25-03, which mandates immediate action on Cisco Adaptive Security Appliance (ASA) and Firepower devices. An advanced threat actor, ArcaneDoor 2, has been exploiting zero day flaws (CVE-2025-20333 and CVE-2025-20362) to achieve unauthenticated remote code execution on ASA VPN endpoints and even modify the device firmware (ROMMON) for boot-persistent malware.

In response, CISA's directive isn't a mere "please patch when convenient" notice, but an urgent call for patching, forensic triage, and reporting within days. CISA is also advising EOL'd Cisco ASA and Firepower devices to be replaced immediately. This whirlwind response highlights a hard truth: our legacy VPN perimeter devices—once the remote access workhorses of enterprise and industrial networks—have become high-risk single points of failure. And while emergency patches will put out this fire, a more structural solution is needed to prevent the next one.

Inside CISA's Emergency Directive 25-03

Emergency Directive (ED) 25-03, issued on September 25, 2025, requires Federal Civilian Executive Branch agencies to "identify and mitigate potential compromise" of all Cisco ASA and Firepower instances. CISA's orders are explicit and time-bound:

Key Dates
Triage/patch: Sept 26, 2025
EOS ASA Removal: Sept 30, 2025
Reporting: Oct 2, 2025

Required Actions

- **Identify All Devices:** Immediately account for every Cisco ASA platform (hardware, virtual, modules) and all Cisco Firepower Threat Defense appliances in use. Nothing can be left untracked, as unseen devices could be backdoors into agency networks.
- **Forensic Triage (Core Dumps):** For every internet-facing Cisco ASA, follow CISA's detailed core dump collection and hunt procedures and upload memory dumps to CISA's malware portal by September 26, 2025, 11:59PM EDT. This allows CISA to analyze devices for the telltale signs of compromise. If any ASA shows "compromise detected," agencies must immediately disconnect it from the network (but not power it off) and report the incident, working with CISA on containment.
- **Patch or Disconnect:** By the same urgent deadline (Sept 26, 2025), apply all available Cisco software updates on devices that show no signs of compromise. Notably, any ASA models that have end-of-support dates on or before Sept 30, 2025 must be permanently disconnected by Sept 30, 2025. This is essentially the forced retirement of legacy boxes that can't be fully secured. Even slightly newer models (supported through 2026) had to be patched by Sept 26 and must install any future patches within 48 hours of release.

- **Fortify Virtual & Firepower Appliances:** Virtual ASAv and physical Firepower appliances weren't spared. They also required updates by Sept 26 and rapid patching of any new fixes within 48 hours.
- **Report Outcomes:** By October 2, 2025, every agency head must report back a complete inventory of affected products and the actions taken (patched, removed, compromised) using CISA's template. This ensures CISA can gauge compliance and remaining exposure across the federal enterprise.

These aggressive timelines (essentially 24 to 48 hours to triage and patch) underscore the directive's urgency. And for good reason. Cisco and government investigators revealed a "widespread" exploitation campaign in the wild. The attackers (identified with the codename ArcaneDoor) have been modifying ASA firmware to plant hidden bootkits that survive reboots and even software upgrades. In other words, an unpatched ASA could remain infected even after a normal update or restart, thanks to malware burrowed in its ROM. This is every network operator's nightmare: a trusted firewall/VPN appliance turned into a persistent threat actor beachhead.

Cisco's analysis of compromised devices found the attackers were highly sophisticated. They disabled logging, intercepted CLI commands, and even intentionally crashed devices to thwart forensic analysis. They targeted ASA 5500-X series units with VPN web services enabled, deploying a malicious ROMMON implant (dubbed "RayInitiator" by UK's NCSC) and an in-memory Linux shellcode loader ("LINE VIPER") to maintain access. In short, this isn't the work of opportunistic script kiddies; it's likely a nation-state or well-resourced group that has turned a critical remote access tool into a foothold for espionage or attack.

CISA's directive is the correct short-term antidote: find the compromises, rip out or fix vulnerable devices, and report the scope of impact. Federal CISOs and network teams have treated it as a four-alarm fire, working around the clock to comply. But as the dust settles, a sobering question remains: how do we stop getting here in the first place? To answer that, we must recognize a pattern that has played out repeatedly in recent years: a pattern of VPN and remote access solutions becoming Achilles' heels.

A Recurring Pattern: VPN Vulnerabilities from Ivanti to Cisco ASA

A recent precedent: In early 2024, CISA's ED 24-01 forced rapid remediation of Ivanti Connect Secure after zero days enabled authentication bypass, device takeovers, and persistence. The pattern mirrors ASA: an internet-facing VPN portal plus implicit network access created a high-value attack path that required emergency disconnection and rebuild—an architectural, not vendor-specific, problem.

Sound familiar? An internet-facing VPN gateway, a zero day used to bypass logins, leading to hidden webshells or implants and long-term stealthy access. The Ivanti incident was almost a preview of

what we're now seeing with Cisco ASA. In both cases, an ostensibly secure remote access solution became an open door.

And it's not just Ivanti or Cisco. We've seen VPN credentials and flaws lead to major breaches before. The Colonial Pipeline ransomware attack in 2021 famously began via a single leaked VPN password. According to the 2025 Verizon DBIR report, stolen or abused credentials are still the top attack vector responsible for breach—and VPN portals are a prime target for password guessing, theft, or exploit.

The common thread is that legacy VPN architectures present a tantalizing combination to attackers: they are both exposed to the internet (often with web-based logins or clients accessible 24/7 from anywhere) and, once breached, they provide deep network access. A single vulnerability or stolen login can potentially let an intruder roam freely inside critical networks.

In Ivanti's case, the attackers who bypassed the VPN could gain administrative access to compromised networks without detection, even reaching domain controllers to attempt full enterprise takeover. With Cisco ASA, the ArcaneDoor group didn't even need valid credentials. The zero days let them get code execution as "unauthenticated" users and from there, they turned the VPN device itself into a beachhead.

This recurring pattern is a big red arrow pointing at an architectural flaw: legacy VPNs and similar remote access concentrators inherently introduce significant risk. ED 25-03 and last year's Ivanti directive are reactive firefights. To get ahead of the next crisis, organizations should be asking how to redesign remote access so that one device or one credential can't expose the whole kingdom.

Why Legacy VPN Architectures Are Inherently Risky

VPNs have been around for decades, but like many legacy technologies, they are showing cracks in today's threat landscape. While legacy VPNs enable convenient access to corporate assets over the internet, which boosts productivity, they also introduce significant risks by granting overly broad and unrestricted access to anyone with valid credentials.

To understand a better path forward, it's important to first understand the challenges with the existing approach. Below is a deep dive into the challenges presented by legacy VPNs and why they pose such a security risk to organizations who use them.

Broad, Implicit Trust

Traditional VPNs operate on a castle-and-moat concept. Once a user authenticates and lands inside the network via the VPN, they are often treated as an insider. The VPN tunnel drops the user onto an internal subnet, effectively placing the remote user's machine on the same network as critical assets. It's an all-or-nothing access model—either you're kept out entirely, or you're let in and can

access pretty much any device. Even if access control lists restrict some segments, users typically still see a broad swath of the network. Attackers love this because a single compromised account or device opens up a buffet of targets behind the firewall.

Lateral Movement and "Flat" Networks

In OT environments (and many IT ones), the VPN often connects to networks where assets have minimal built-in security and trust anyone on the local subnet. Legacy protocols (e.g. industrial control protocols, or even SMB and RDP in IT) assume the network is segregated. When a VPN user connects, those insecure protocols become reachable over the VPN tunnel. If an attacker piggybacks on a VPN connection, they can exploit vulnerable devices or use tooling like ransomware to spread laterally with little resistance. Since microsegmentation is typically absent and the VPN user is essentially an insider, threats can spread unchecked once inside.

Exposure of a High-Value Portal

A VPN gateway (Cisco ASA, Ivanti, etc.) usually exposes some login interface or services on the internet so that legitimate users can connect remotely. This means the device itself is constantly exposed for potential attack from the open internet. And indeed, VPN and firewall appliances have proven to be bug-rich. Cisco's ASA has a long list of CVEs over the years and Ivanti's Pulse Connect Secure was found to contain open-source code modules that hadn't been updated in 20 years, according to researchers.

Attackers have learned that instead of laboriously phishing users, they can often exploit a VPN appliance vulnerability to get in without credentials, as happened here. With VPNs, your network's front door is not only wide open to the internet, but in many cases built on aging code that adversaries continually hammer on for new weaknesses. It's a giant attack surface.

Single Point of Failure

VPN architectures concentrate remote access through a handful of chokepoints. If that VPN server or firewall fails, whether due to a cyber incident or just a configuration error, remote access for an entire organization can grind to a halt. Worse, if an attacker compromises the VPN appliance, they can potentially monitor or tamper with all traffic passing through it, or use it as a springboard to internal systems. The ArcaneDoor attackers demonstrated this by implanting malicious boot code in ASA devices to maintain persistence. A hacked VPN is like handing keys of the castle to the adversary and also giving them a hidden tunnel in and out.

Limited Visibility & Forensics

Traditional VPNs don't provide much insight into what a user does once connected. They might log connection times or how much data was transferred, but they don't record user activity on the target systems. Security teams are often blind to specific commands run or changes made over a VPN session. As a result, detecting misuse or compromises inside a VPN session is difficult,

and investigating incidents after the fact can be a nightmare of piecing together logs from various internal systems. In the ASA campaign, for instance, only by pulling core dumps and using specialized tools could agencies determine if the device was backdoored. The device's own logs were of little help once an attacker disabled them.

Patching and Operational Overhead

As we see with ED 25-03, when a critical vulnerability hits a VPN, organizations face a gut-wrenching choice: patch immediately (often causing downtime) or remain exposed and risk breach. VPNs often terminate user sessions and require reboots on upgrade, so emergency patching can disrupt business continuity. In practice, some organizations delay patching VPN appliances due to uptime requirements, which leaves windows of exposure that attackers exploit. It's a lose-lose: patching a critical remote access system in a hurry is disruptive, but not patching is courting disaster. Moreover, layering compensating controls (firewall rules, IDS monitoring of VPN traffic, jump hosts, etc.) adds complexity and often slows down user access, impacting productivity.

Credential Theft & Abuse

VPNs concentrate a lot of power in a single set of credentials. A valid username/password (or token) lets an actor appear as an authorized user and often go undetected. Phishing and password-stealing malware feed on this. VPN systems themselves sometimes don't enforce strong multi-factor authentication (or it's not uniformly enabled for all users, especially in OT environments with third-party contractors). This makes credential abuse one of the easiest pathways for attackers, certainly far easier than developing a zero day exploit. Yet if those credentials do get compromised, the VPN will happily let the adversary in, since it can't distinguish a legitimate user from an impostor with the right password.

From Firefighting to Forethought: Embracing Zero Trust Access

So, what's the alternative? How can organizations enable secure remote access for admins, engineers, and third-party vendors without exposing themselves to these risks? The answer lies in shifting from a network-centric, perimeter-based model to an identity-centric, Zero Trust model. Instead of pouring more money and effort into patching and protecting VPNs, forward-thinking security leaders are looking to eliminate the need for VPNs entirely by adopting Zero Trust Access solutions.

Zero Trust Access (ZTA) is a fundamental redesign of how remote connectivity works so that even if credentials leak or vulnerabilities lurk, a breach doesn't turn into an organizational catastrophe. As the name implies, "Zero Trust" means embracing the principle of trusting no one and no device by default, even if they are already "inside" the network. Every access request is verified, and privileges are tightly scoped.

A ZTA approach flips the VPN model on its head. Rather than connecting a user to an entire network segment via a tunnel, ZTA connects a user only to a specific application or asset that they explicitly need, and nothing more. Access is granted based on identity and policy, incorporating factors like who the user is, what role they have, the security posture of their device, where they are connecting from, etc. Continuous verification replaces the one-time "login and you're in" approach of VPNs. It's the principle of least privilege applied to remote access: a user gets only the minimum access required, for the minimum time necessary, with the maximum scrutiny.

Crucially, Zero Trust systems don't inherently trust internal network location either. Whether you're at corporate HQ or connected in from home, you face the same authentication and authorization checks for a given resource. In other words, the network boundary matters far less. In this approach, identity and context are the new perimeter. This mitigates insider threats and compromised internal hosts as well, by narrowing what any one principal can do.

Another major difference is how connections are brokered. Instead of opening broad network tunnels, many ZTA solutions (including Xage's) use an application proxy or broker model: remote users connect to a Zero Trust service (often a cloud-based or distributed platform), and that service in turn connects into the target environment on the user's behalf. The user never directly talks to the destination device over the network in the way they would over a VPN. This broker can enforce policy (e.g. check the user's MFA status, device health, authorization rights), translate protocols, and strictly limit the session to the single target asset and operation needed. The result is granular, narrowly scoped connectivity rather than an open pipe into the network.

Example

Imagine the legacy approach: an OT engineer needs to check a substation SCADA system, so they VPN into the corporate network, then maybe RDP from there into the OT network jump box, then from there access the SCADA HMI. Along this path they potentially had access to an entire subnet (or multiple) via VPN, and the jump box had accounts for every user plus broad network reach in OT—lots of exposure points.

Now picture the Zero Trust approach: the engineer opens a secure application in their browser (or client) and requests access to "SCADA-HMI-Substation123." The Zero Trust Access platform verifies their identity (with MFA), checks that they're allowed to access that specific HMI, and then creates a temporary secure session that connects the engineer only to that HMI—effectively a direct application tunnel. When their work is done, the session is closed. The user never "entered" the broader network at all, and could not see or reach any other devices in that substation or data center.

This is exactly the model that Xage Zero Trust Access implements for industrial operations and enterprise alike. Below, we outline how Xage's approach works and how it directly addresses the shortcomings of legacy VPNs.

How Xage Zero Trust Access (ZTA) Mitigates VPN Risks

Xage's identity-first Zero Trust architecture spans cloud, enterprise, and remote sites, enforcing per-user, per-asset access without placing users on internal networks.

Xage Zero-Trust Capability	What it does (technical)	How it enforces control	Resulting risk reduction (vs. legacy VPN/ASA-class issues)
Identity-bound, least-privilege access	Binds every session to a verified user/service identity and role; grants only the specific asset/command needed.	Continuous verification; context-aware policy (user, device posture, asset, protocol, time); automatic session expiry.	Compromised accounts are sandboxed to a narrow slice; no "network ride" to crown jewels; unauthorized assets remain invisible.
Agentless enforcement at the site edge	Distributed enforcement points (gateways) sit at plant/site/data-center/VPC demarc; no agents on PLCs/HMIs.	Policy pulled from Fabric; local autonomy/caching; rolling node updates; no single choke-point appliance.	Removes VPN single point of failure; patch/maintenance doesn't drop all access; resilient during WAN hiccups.
No network placement for users	Users never get an internal IP; access is brokered—Xage connects into the environment on the user's behalf.	Outbound-only connections from sites; reverse-proxy/protocol brokerage; per-session micro-tunnel to one asset.	No subnet exposure or scanning; collapses public attack surface (no web-VPN portal to exploit); blocks arbitrary pivots.
Per-asset granular policy (segmentation)	Enforces identity-based per-device/app authorization (e.g., specific PLC or HMI function, read-only vs. program).	Allow-lists per asset; protocol and command scoping; just-in-time, time-bounded sessions.	Lateral movement is structurally blocked; flat network risks are contained even if an endpoint is compromised.
Multi-factor authentication everywhere	MFA at login and step-up for sensitive operations/assets—even for legacy/OT that lack native MFA.	Integrations: SSO/SAML/AD, PIV/CAC, tokens; broker requires MFA before brokering downstream access.	Stolen passwords alone are useless; contractor access is uniformly strong; reduces credential-abuse paths common on VPNs.
Multi-factor authentication everywhere	MFA at login and step-up for sensitive operations/assets—even for legacy/OT that lack native MFA.	Integrations: SSO/SAML/AD, PIV/CAC, tokens; broker requires MFA before brokering downstream access.	Stolen passwords alone are useless; contractor access is uniformly strong; reduces credential-abuse paths common on VPNs.
Built-in PAM	Privileged workflows without shared creds; downstream creds are ephemeral and rotated.	Per-user accounts, JIT elevation, command policies; integrated vault with automatic credential issuance/rotation; session recording.	Eliminates standing secrets; users never see device passwords; precise audit of admin actions; limits post-compromise blast radius.

xage SECURITY

Xage Zero-Trust Capability	What it does (technical)	How it enforces control	Resulting risk reduction (vs. legacy VPN/ASA-class issues)
Continuous monitoring & audit trails	Captures every access decision and session (incl. RDP/VNC/SSH actions) in a tamper-resistant ledger.	Real-time logs, session shadowing/recording, SIEM/SOAR export, anomaly alerts.	Fast, confident IR; attacker log-tampering on devices is offset by broker-side evidence; simplifies compliance (e.g., NERC CIP).
Overall effect	Zero-trust overlay replaces "VPN > subnet" with identity-aware, per-resource lanes.	dentity broker + distributed enforcement; policy-first access instead of network reachability.	Neutralizes web-VPN exploitation paths (ED 25-03 class); preserves operations during patching/rebuilds; turns remote access from a choke-point into a controlled, auditable service.

In short, Xage Zero Trust Access is designed to structurally eliminate the risks that ED 25-03 and its predecessors keep highlighting. Instead of internet-facing devices with implicit trust, you have a shielded, identity-verifying broker. Instead of broad network tunnels, you have narrow per-resource lanes. Instead of hoping users don't abuse privileges, you enforce least privilege by default. Instead of scrambling to pull device memory in a compromise, you likely prevent the compromise or at least contain its impact to a single asset, and you have robust logs to investigate with. It's a fundamentally different, and safer, way to do remote access.

Deployment: Replace "VPN -> Subnet" with Brokered, Per-Asset/Application Sessions with Xage Zero Trust Access

- **What changes:** Deploy Xage Nodes at site demarc; each Node makes an outbound-only TLS connection to the Xage Fabric and brokers a single, per-asset session on demand—no public portal or inbound listeners.
- **Prerequisites:** Reuse your IdP (SAML/OIDC/AD) and groups; map roles to asset-level policies. Allow Node egress to the Fabric and Node reachability only to listed assets (RDP/SSH/HTTPS/serial). No agents on PLCs/HMIs; no subnet redesign.
- **User flow:** After SSO, users see only approved assets and launch an MFA-protected session to that target. The Node creates a micro-tunnel to that device/port for that session only; anything else is out-of-policy and denied—replacing "VPN in and roam a subnet" with direct-to-asset access.
- **Rollout:** Start with 1–2 Nodes and 10–20 high-value assets; pilot one internal and one vendor workflow; then expand by role/site. Retire public VPN listeners last and keep legacy VPN as break-glass.
- **Day-2 operations:** Update Nodes in rotation; stream decisions and session metadata to SIEM/SOAR; use step-up MFA/approvals for sensitive changes; terminate sessions or quarantine identities/assets centrally—without taking sites offline.

Conclusion: Shift from Short-term to Long-term & Durable Strategies

CISA's ED 25-03 should rightly spur urgent action—patch your Cisco ASA and Firepower devices now, hunt for compromises, and report the all-clear (or the casualties) to CISA by the deadlines. But once those immediate tasks are done, leaders in federal agencies, utilities, and critical infrastructure must take a step back and ask: "How do we make sure we're not the next Ivanti or Cisco VPN emergency case?" The lesson from these directives is that simply maintaining the status quo (with occasional fire drills) is not a winning strategy. The old remote access paradigm—VPN concentrators exposing broad networks – is too risky, as attackers have demonstrated again and again.

The good news is that alternatives exist today.

Xage Zero Trust Access allows organizations to fundamentally transform their security posture without crippling usability. You don't have to choose between enabling remote work and protecting your core systems. You can do both, with modern architectures that were built with today's threats in mind. By implementing identity-bound, granular access with strong authentication and integrated auditing, you turn remote access from a soft underbelly into a robust front line. Instead of dreading the next CISA emergency directive, you can get ahead of it.

For CISOs, OT security architects, and IT managers in the hot seat, the approach should be twofold: tactical and strategic.

Tactically, comply with directives like ED 25-03: patch those devices, collect those dumps, verify your network isn't already breached. But strategically, use this moment to drive change: conduct that VPN exposure risk assessment you've been putting off, pilot a Zero Trust solution in a segment of your environment, and create a roadmap to phase out legacy VPN dependencies. It's fitting that Xage's solution was referenced as a member of CISA's JCDC and as a VPN replacement path— the public and private sector both recognize that our future cannot look like our past in this domain.

In operational technology environments, where downtime is costly and safety is paramount, adopting Zero Trust access can also improve business continuity. It enables maintenance and support to continue securely even during patch cycles or network perturbations, and it provides confidence that one compromised node won't take down an entire production line or grid segment. Moreover, the forensic and oversight capabilities (session recording, tamperproof logs) mean that if something does go wrong, you're not investigating in the dark—a boon for incident response and post-incident learning.

For organizations in critical infrastructure and government, the path forward is clear: patch urgently today, but design for security tomorrow. By the time the next zero day comes knocking, your goal should be that your attack surface is so minimized, and your controls so granular, that an emergency directive becomes a footnote rather than a front-page headline.

Zero Trust Access, as delivered by solutions like Xage, is how you achieve that resilience. Not by trusting less, but by architecting smartly. The result is safer operations, easier compliance, and fewer 3 a.m.. incident response calls. In the end, moving beyond legacy VPNs isn't just a network change, it's a mindset change: from implicit trust to explicit verification; from perimeter defense to identity-centered defense; from patch-and-pray to plan-and-protect.

About Xage

Xage is a global leader in zero trust access and protection on a mission to pioneer a secure tomorrow. Control access and prevent attacks in the cloud, in the data center, at the remote operational edge anywhere on Earth, and even in orbit with the Xage Fabric Platform. Xage is easy to manage and can be deployed in a day, giving users easy and secure access to the assets they need from anywhere, while preventing advanced adversaries and insider threats at every stage of the attack chain. Learn why organizations like the U.S. Space Force, PETRONAS, and Kinder Morgan choose Xage at xage.com.