

EBOOK

Odyssey: The Seven Monsters of PAM

How Zero Trust PAM Defeats the Privileged Access Threats of the AI Era

xage SECURITY

About this eBook

Privileged access has escaped the neat boundaries of the traditional data center. Administrators still matter, but so do vendors, service accounts, applications, workloads, automation, and AI agents. Each can reach systems that shape operations, safety, intellectual property, and business continuity. The result is a much larger and faster-moving privilege problem than a conventional password vault was designed to solve.

CONTRIBUTING AUTHOR



Dr. Chase Cunningham ("DrZeroTrust")

Creator of the ZTX framework, former Forrester analyst, and retired U.S. Navy Chief with 20+ years of cybersecurity experience.

This eBook uses the trials in Homer's Odyssey as a practical framework for examining seven recurring failures in privileged access management. Each monster represents a risk pattern that security teams encounter today: trusted third-party access, persistent credentials, centralized failure, shared accounts, social engineering, machine-speed activity accelerated by AI, and dormant insider access. The accompanying counterattacks show how identity-first, Zero Trust Privileged Access Management (PAM) can narrow access, remove standing privilege, preserve accountability, and maintain control across distributed environments.

The goal is not to claim that any organization can become breach proof. Zero Trust begins with the opposite assumption. Systems will face compromise, credentials will be targeted, and people will make mistakes. A resilient PAM program limits what a compromised identity can reach, what it can do, how long access lasts, and how quickly defenders can understand and stop the activity.

Each monster also includes a brief real-world scenario to connect the metaphor to an operational failure. Later sections extend the model to machine-speed privilege, provide a practical PAM maturity model, and answer the questions security leaders commonly encounter when moving from vault-centric PAM toward Zero Trust PAM.

Executive Summary

Privileged access is the shortest path to consequence in any organization. Whoever holds it can change what a system does, what it stores, and whether it keeps running. Attackers have understood this for years, and the data now shows how thoroughly they have adjusted. Darktrace's 2026 Annual Threat Report found that nearly 70% of incidents began with stolen or misused accounts.[1] CrowdStrike's 2026 Global Threat Report found that 82 percent of its detections in 2025 were malware-free.[2] **The modern adversary does not break in. The modern adversary logs in.**

This eBook makes a three-part argument. First, privilege is no longer a short list of administrators with vaulted passwords. It is any identity, human or not, that can cause something important to happen, and that now includes vendors, service accounts, workloads, and AI agents. Second, privilege escapes control in seven recurring ways, and each of the monsters in the chapters that follow names one of them. Third, all seven risks can be addressed with the same approach: verify who or what is requesting access, allow access only to the resource needed for an approved task, limit how long it lasts, protect credentials from exposure, monitor and record activity, and remove access as soon as the task ends.

For a security leader, the practical consequence is a change in metrics. The question is no longer whether the organization owns a vault. It is what proportion of consequential privilege still lives outside management and consistent policy. That share is usually larger than anyone expects. Non-human identities — API keys, service accounts, OAuth tokens, certificates, and AI agents — now outnumber human identities by 45:1 in the average enterprise, reaching 144:1 in cloud-native environments. Most of the privilege in a modern enterprise has never been in scope for the program that was supposed to govern it.

The chapters ahead move from diagnosis to action. The seven monsters describe the failure patterns. The evidence and case chapters connect each pattern to public incidents and hard numbers. The maturity model, the ninety-day plan, and the metrics chapter turn the metaphor into a program a team can fund, run, and measure.

THE THESIS IN THREE SENTENCES

A vault protects credentials. Zero Trust PAM limits what those credentials can access, for what purpose, and for how long. If a credential is stolen, those limits help contain the damage.

Odyssey: The Seven Monsters of PAM | © Xage Security | All Rights Reserved | www.xage.com

eBook Outline



About this eBook 2

Executive Summary 3

Privilege has outgrown the vault 5

1. The Trojan Horse 8

2. Calypso 10

3. Cyclops 12

4. Nobody 14

5. The Sirens 16

6. The Laestrygonians 18

7. The Suitors in Ithaca 21

How Xage XPAM Puts Zero Trust into Practice 23

How Zero Trust PAM aligns with standards and guidance 24

Your first ninety days: assess, act and measure 26

The monsters differ. The mission does not. 29

Notes and sources 30

xage SECURITY

Privilege has outgrown the vault

Why vault-centric PAM falls short, and the seven ways privilege escapes control

Traditional PAM was built on a simple assumption: privilege belonged to a small group of administrators. They logged into known systems from known locations, using credentials that could be locked in a vault and carefully controlled. The systems themselves ran on servers inside the organization's own facilities.

None of that holds anymore. Infrastructure now spans cloud platforms, SaaS applications, industrial environments, and remote sites. Privilege is held by APIs, workloads, service accounts, contractors, automation, and increasingly autonomous AI agents. Dangerous privilege no longer requires a domain administrator account. An application that can change production data, an AI agent that can invoke administrative APIs, or a vendor account with a path into an operational system carries just as much consequential authority.

That changes the central PAM question. It is no longer "Who has the password?" but "Who or what can reach this resource and perform a critical function?" Seen that way, vaulting credentials remains necessary, but it is no longer sufficient. Organizations must control the identity, resource, action, context, duration, and consequences of every privileged session.

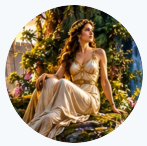


Like Odysseus, security leaders know their destination, but the route changes with every new threat.

Seven monsters stand in the way. Each one names a recurring way that privilege escapes control, and each one shows up in the public breach record.

Odyssey: The Seven Monsters of PAM | © Xage Security | All Rights Reserved | www.xage.com

PRIVILEGE HAS OUTGROWN THE VAULT

Monster	How privilege escapes control	What the record shows
 #1 **The Trojan Horse**	Trusted vendor and contractor access reaches far beyond the task.	Third parties were involved in 48% of breaches, up from 30% the previous year—a 60% increase. Large supply chain and third-party compromises nearly quadrupled since 2020. Verizon, 2026 · IBM X-Force, 2026
 #2 **Calypso**	Temporary access quietly becomes permanent.	64% of valid secrets detected in 2022 remained unrevoked in 2026. In the historical Snowflake customer campaign, attackers used credentials associated with infections dating to November 2020 that remained valid years later. Approximately 165 potentially exposed organizations were notified. GitGuardian, 2026 · Mandiant, 2024
 #3 **Cyclops**	Concentrated power becomes a single point of failure.	In the 2022 LastPass breach, a key needed to unlock backup encryption material was stored in the business vaults of four senior employees. Compromising one engineer's account enabled the attacker to obtain that key. The UK regulator's 2025 investigation documents this concentration of access. ICO, 2025
 #4 **Nobody**	Shared accounts hide who actually acted.	82% of CrowdStrike's detections in 2025 were malware-free. PAM implication: when multiple people use the same account, identifying the individual behind each session becomes essential to accountability. CrowdStrike, 2026 Global Threat Report

Monster	How privilege escapes control	What the record shows
 #5 **The Sirens**	Persuasion turns a valid login into an attacker's foothold.	Darktrace's 2026 Annual Threat Report found that nearly 70% of incidents began with stolen or misused accounts. Voice-phishing intrusions doubled in the first half of 2026 compared with the second half of 2025. Attackers used voice phishing to compromise single sign-on accounts and SaaS applications; in one incident, account takeover led to data theft in under five minutes. 2026 Darktrace Annual Threat Report · CrowdStrike, 2026 Threat Hunting Report
 #6 **The Laestrygonians**	AI-driven activity and machine-speed access outrun manual governance.	Non-human identities outnumber human identities 45 to 1 in the average enterprise and 144 to 1 in cloud-native environments.[5] CrowdStrike's 2026 Global Threat Report updates the fastest observed eCrime breakout from 51 seconds to 27 seconds, based on activity in 2025. CrowdStrike, 2026 Global Threat Report
 #7 **The Suitors**	Dormant and insider access lingers long after its purpose ends.	Malicious insider attacks were the costliest initial vector for the second year running, at USD 4.92 million per breach (IBM, 2025).

Two lessons stand out. The Suitors and the Trojan Horse are not only the oldest monsters in the story; they are also the most expensive. And standing privilege does not age out. It waits, sometimes for years, until someone finds it.

The monsters differ, but the defense does not. The chapters that follow take each one in turn: how it works, what it looks like in the real world, and how Zero Trust PAM limits what a compromised identity can reach, what it can do, and how long its access lasts.

#1 The Trojan Horse

Trusted access can still carry danger

Troy relied on strong walls. The attack succeeded because the defenders treated something invited through the perimeter as trustworthy. The modern equivalent appears whenever a contractor, supplier, managed service provider, or support technician receives broad network access simply because the organization recognizes the user or approves the connection.

VPNs and shared jump servers often create this condition. They authenticate a user at the edge and then expose a network segment containing many systems. If the identity, device, or credential has been compromised, the attacker gains room to move laterally. The initial connection may be legitimate while the resulting reach is far broader than the task requires.

The security lesson is simple: granting access should not make that access trusted. Organizations need to validate the user's identity, evaluate the context, and define the exact asset, account, protocol, and time window the user may access. Authentication should open a specific path, not an entire city.



WHAT THIS LOOKS LIKE IN THE REAL WORLD

Target, 2013

Attackers entered Target's network using credentials stolen from a refrigeration and HVAC contractor, then moved to the payment environment and exposed roughly 40 million payment cards.[7] The contractor's access was legitimate and approved. The failure was that approved access reached far more than the contractor's job required.

THE COUNTERATTACK · #1

Verified vendor access

Xage XPAM controls for vendor and contractor access

Xage XPAM counters the Trojan Horse by replacing broad connectivity with identity-centric, resource-specific access. Strong authentication establishes who is requesting access. Policy then considers the user, group, device, source, posture, target asset, target account, protocol, time, and approval status before it opens a session.

The access path can require one or more approvals based on the sensitivity of the target. A routine maintenance task that's part of a standard workflow might not require approval, while work on a critical system may require approval from an owner, supervisor, or IT service management. The authorization remains tied to the task and expires when the approved window closes.

This model also allows organizations to consolidate vendor access through Remote Privileged Access Management (RPAM). Instead of operating separate VPNs, third-party remote access tools, and jump servers or bastion hosts with inconsistent monitoring, teams can broker internal and external access through a single PAM platform. Zero Trust Network Access (ZTNA) can restrict connections to approved applications, but connectivity controls alone do not provide the full oversight privileged access requires. Zero Trust PAM adds controls over credentials, privilege elevation, session activity, and access duration. The result is a smaller attack surface and a clearer record of who accessed each resource, for what purpose, and for how long.



#2 Calypso

Long-lived access becomes a hidden liability

Calypso represents the danger of access that becomes normalized, invisible, and permanent. In many organizations, privileged credentials and permissions remain active because removing them might interrupt a workflow. A temporary exception becomes part of the operating environment, and no one remembers why the access exists.

Persistent credentials create several forms of exposure. Users may reuse or share passwords. Secrets may appear in spreadsheets, text files, configuration files, or notes. An attacker who finds a valid credential will retain access long after the original compromise. Broad entitlements also make lateral movement easier because one identity can reach resources unrelated to its current task.

Zero Standing Privilege offers a different model. Access exists only when an authorized purpose requires it. The organization grants the minimum privilege for a defined session, limits the scope of where that privilege can be used, and removes the path automatically when the session ends, the time window expires, or an administrator revokes approval.



WHAT THIS LOOKS LIKE IN THE REAL WORLD

Microsoft, 2024

In the breach Microsoft disclosed in January 2024, attackers compromised a legacy test account that lacked multifactor authentication.[28] From that foothold, they compromised an older test application that still held elevated permissions in Microsoft's corporate environment. They used those permissions to grant attacker-controlled applications access to corporate email, stealing messages and attachments from accounts that included senior executives. Calypso's danger is visible here: privileges attached to a legacy testing environment remained powerful enough to open a path into sensitive business systems. Access that stays available becomes access an attacker can inherit.

THE COUNTERATTACK · #2

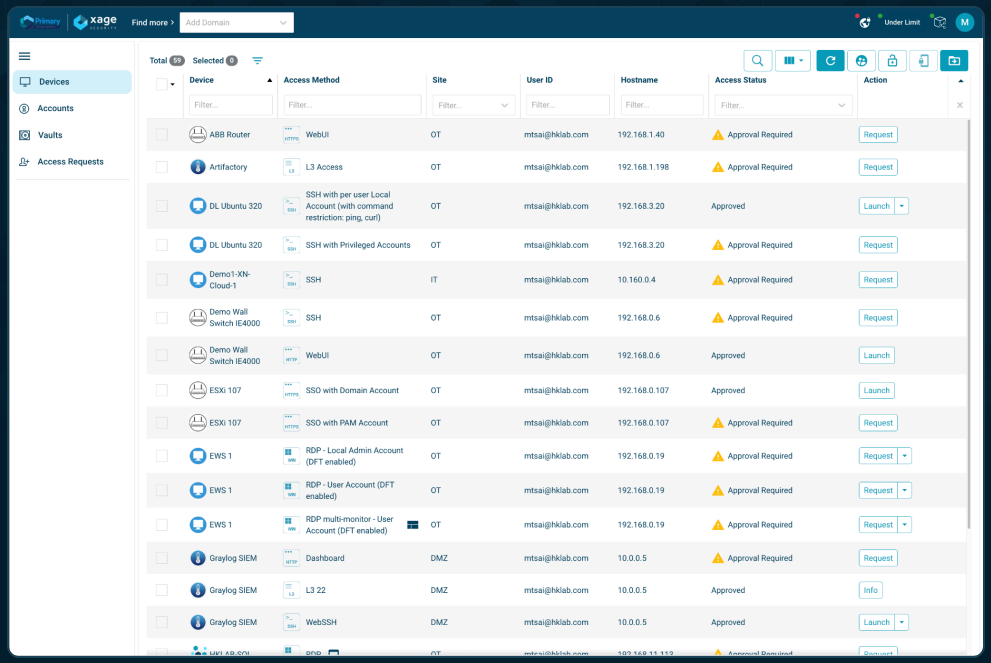
Remove persistent access

Xage XPAM controls for long-lived credentials and standing privilege

Xage XPAM separates the user from the privileged credential. The platform can inject a managed credential into an approved session so the user never sees, copies, or stores the password. This reduces the attack surface by keeping privileged credentials out of users' hands and reduces the burden of discovering, tracking, and managing credentials stored across the IT environment.

Credential rotation then removes persistence. Passwords can rotate on a schedule, after use, or in response to policy. Once rotation occurs, copied credentials and abandoned records become invalid; should an attacker find the copied credential, it will no longer be of any use. Just-in-time and just-enough access further restrict the session to the privilege and duration required for the work.

The matched security outcome is important: a credential alone never establishes trust. The platform verifies the user's identity, evaluates context, controls the session, and terminates access when the authorized reason no longer exists.



Odyssey: The Seven Monsters of PAM | © Xage Security | All Rights Reserved | www.xage.com

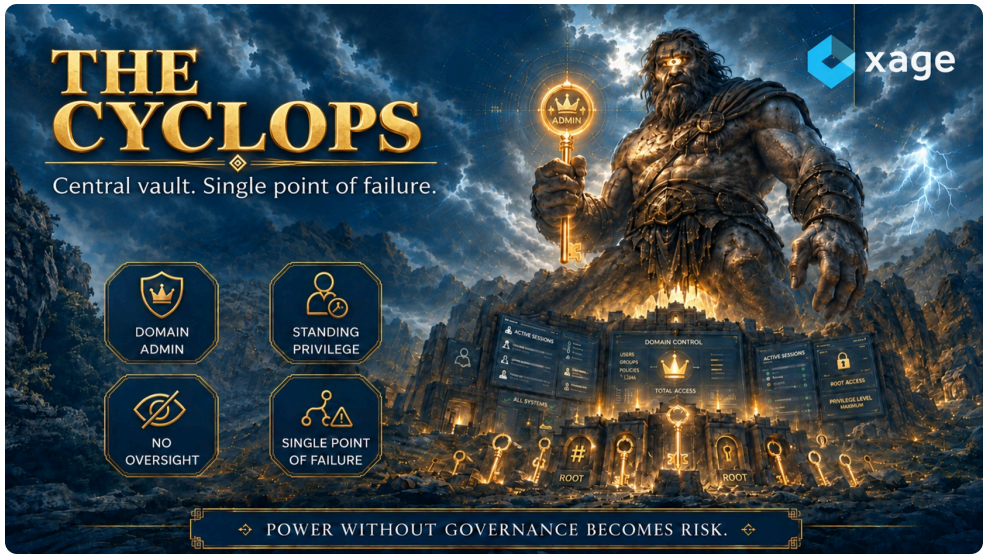
#3 Cyclops

Central power can become central failure

Cyclops represents concentrated power. A domain administrator, central vault, or central access broker can become both the most valuable target and the largest failure domain in the environment. If an attacker compromises that control point, the attacker may inherit broad authority. If the system becomes unavailable, legitimate administrators may lose access to many sites at once.

Traditional centralization can also create operational risk. Remote and industrial environments do not always have reliable connectivity to a cloud service or core data center. A WAN outage should not force a plant, utility, or field site to choose between stopping essential work and bypassing security controls.

Power requires governance, redundancy, and local enforcement. The organization needs centralized policy and visibility without making every privileged decision depend on one reachable system.



WHAT THIS LOOKS LIKE IN THE REAL WORLD

LastPass, 2022

After an initial breach of its development environment, the attacker targeted the home computer of a senior DevOps engineer, one of only four people with access to the decryption keys for production backups, installed a keylogger, and used the captured credentials to reach encrypted customer vault backups.[9,10] Concentrated authority creates a concentrated target. The adversary did not need the whole company, only the one person who held the keys.

Odyssey: The Seven Monsters of PAM | © Xage Security | All Rights Reserved | www.xage.com

THE COUNTERATTACK · #3

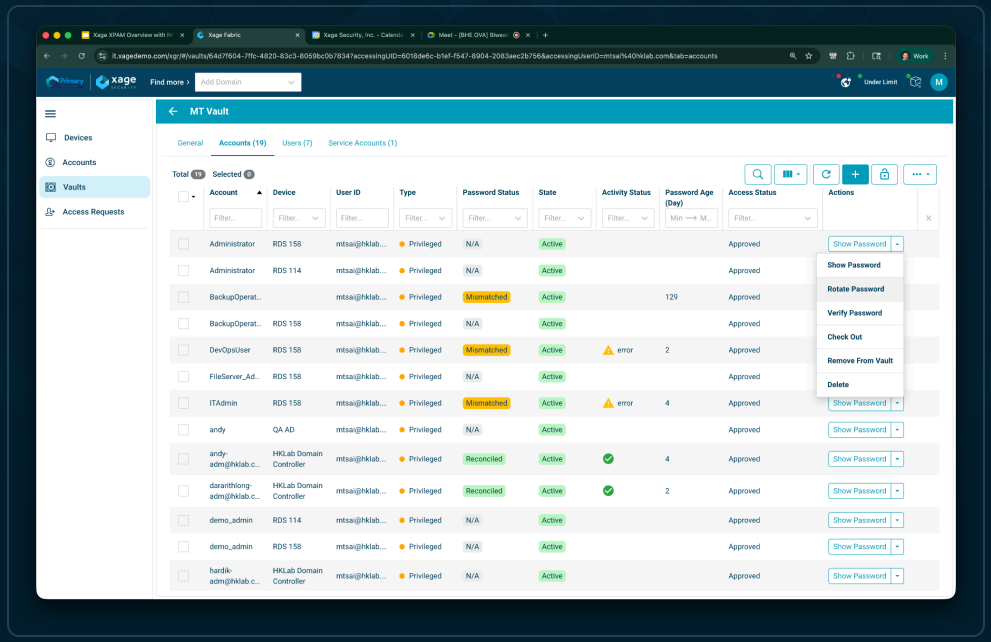
Distributed vaulting and enforcement

Xage XPAM controls for centralization and operational dependency

Xage distributes secrets management and access enforcement across its identity fabric. No individual node stores any part of the secret. The secret is regenerated only through the participation of all nodes, using the cryptographic information each provides. Centralized policy and audit oversight maintain consistent control across the distributed environment.

The platform also supports delegated local management. Site administrators and application owners can manage defined portions of the environment without receiving unrestricted authority over the enterprise. Granular permissions maintain separation between local responsibility and global governance.

Most importantly, local nodes can continue vaulting, brokering, and enforcing previously approved policy during a cloud, WAN, or central-site outage. The architecture avoids turning one disruption into a total loss of secure privileged access.



Odyssey: The Seven Monsters of PAM | © Xage Security | All Rights Reserved | www.xage.com

#4 Nobody

Shared credentials defeat accountability

Odysseus used the name “Nobody” to hide his identity from the Cyclops. Shared accounts create the same blindspot in privileged access management. When many people use one administrator, root, maintenance, or service account, the target system will record the account name but not the identity of the person behind the action.

This gap weakens deterrence, incident response, and compliance. Investigators must reconstruct events from connection logs, ticket systems, timestamps, and assumptions. In some cases, they cannot establish who performed a damaging or unauthorized action. Session visibility may also remain limited when users connect outside the approved PAM workflow.

A legacy target system may still require a shared target account. That technical limitation does not force access to be anonymous. PAM can authenticate the person first, broker the shared identity, and preserve the link between the human user and every privileged session.



WHAT THIS LOOKS LIKE IN THE REAL WORLD

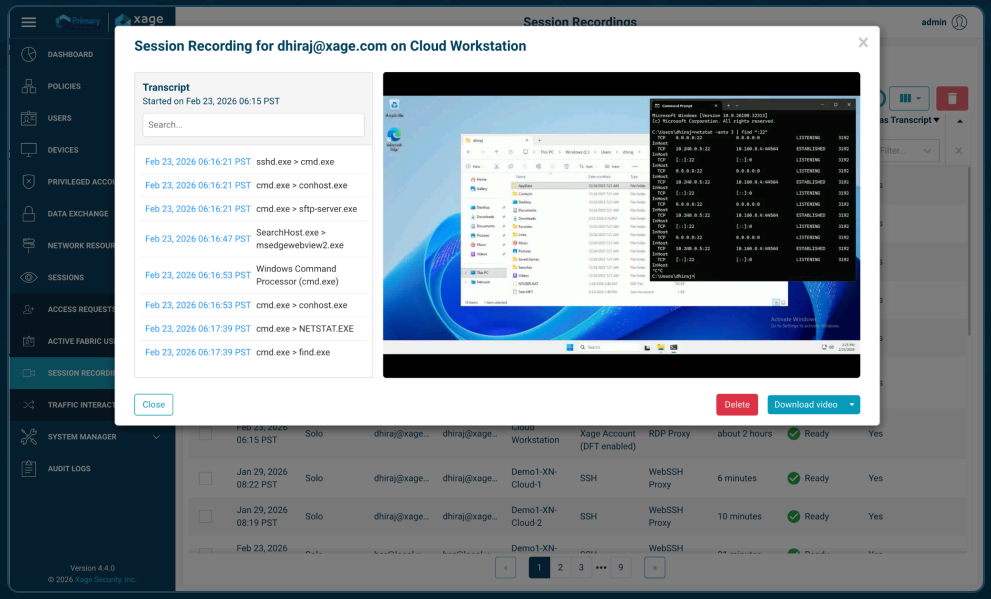
Oldsmar, Florida, 2021

The Oldsmar water treatment plant reported that an intruder used TeamViewer to raise the sodium hydroxide setpoint in the treatment process. Investigators found that staff computers shared a single TeamViewer password, ran an outdated version of Windows, and sat on the internet without a firewall.[11] Two years later, a former city official said the FBI had found no evidence of outside access and that an employee had most likely made the change by mistake.[12] Either version proves the same point. When everyone uses one credential, no one can say who acted, and the organization cannot even establish whether it was attacked.

THE COUNTERATTACK · #4

Accountable privileged sessions

Xage XPAM controls for shared accounts and lost attribution



Xage XPAM establishes a verified access identity at the beginning of the session and removes the access path when the session ends. The user can reach the required shared or privileged account without learning the credential. The platform retains the relationship among the initiating identity, approval, target asset, target account, and session.

Administrators can monitor active sessions and terminate activity when risk changes. Recordings and smart transcripts provide searchable commands, timestamps, and video playback. This evidence supports faster investigations because responders can move from an alert to the relevant activity without manually correlating disconnected records.

Deep audit logs also establish non-repudiation. Even when the target account remains shared, the privileged interaction is attributable to the specific user. The account may be “Nobody,” but the actor is known.

#5 The Sirens

Stolen credentials exploit human trust

The Sirens win through persuasion. In cybersecurity, phishing and social engineering exploit the same human tendency to trust a convincing message, familiar identity, or urgent request. The attacker does not always need to break through a technical control. The attacker may convince an authorized user to disclose a credential, approve a request, or open the path on the attacker's behalf.

A valid password therefore provides weak assurance by itself. The organization also needs to understand the source, device, target, requested action, time, and surrounding behavior. A login that appears normal at the perimeter may become suspicious when the user requests a critical asset from a new source or attempts activity outside the approved task.

Zero Trust assumes that people will sometimes click, approve, or trust the wrong signal. The program should reduce the consequences of that mistake instead of placing the entire defense on awareness training.



WHAT THIS LOOKS LIKE IN THE REAL WORLD

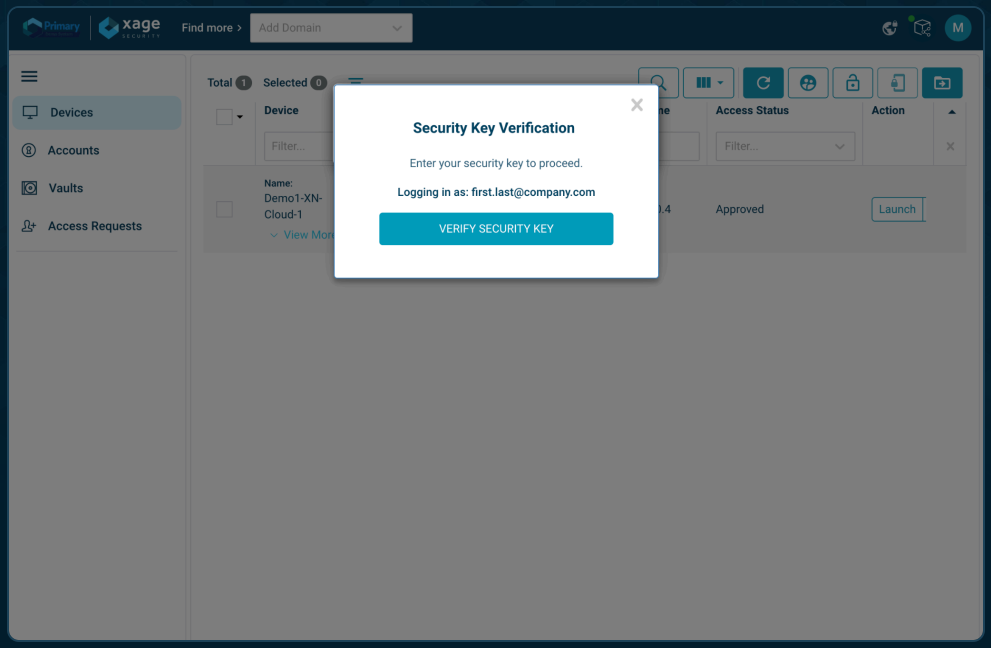
MGM Resorts, 2023, and Change Healthcare, 2024

MGM's intrusion began with a social-engineering call to its IT help desk, according to Okta's chief security officer, after which the attackers gained administrative control of identity systems and disrupted casino and hotel operations for roughly ten days.[13,14] Five months later, attackers used stolen credentials to enter Change Healthcare through a Citrix remote-access portal that had no multifactor authentication, then spent nine days moving laterally before deploying ransomware. The company later estimated that data on roughly one-third of Americans had been taken.[15] In both cases a valid login was accepted as sufficient proof. It was not.

THE COUNTERATTACK · #5

Verify beyond the password

Xage XPAM controls for phishing and credential-based false assurance



Xage XPAM combines layered multifactor authentication with contextual policy. Passwordless FIDO2 can strengthen the initial login and, where supported, individual device access. Multiple identity providers allow organizations to use the authentication realms appropriate to their IT, cloud, and segmented operational environments.

Policy evaluates more than the login. It can consider identity, group, asset, account, protocol, time, source, device posture, and approval. More sensitive zones can require step-up authentication or a secondary login from the directory that governs that segment.

Continuous session oversight limits the blast radius if an identity is compromised. Administrators can record, supervise, restrict, or terminate sessions when trust changes. A successful authentication therefore does not create permanent or unlimited trust.

#6 The Laestrygonians

Privilege operating at machine speed

The Laestrygonians represent overwhelming force at machine speed. Applications, services, automation, workloads, and AI agents already outnumber many human identity populations. They request access, invoke APIs, and perform privileged operations faster than a manual review process can observe or approve each action.

The risk is not limited to artificial intelligence. Organizations have managed service accounts, scripts, robotic process automation, and application credentials for years. AI increases the urgency because agents can plan, select tools, and take sequences of actions with greater autonomy and speed. Broad permissions and hardcoded secrets can turn one flawed or manipulated process into rapid, repeated access across many resources.

Manual spreadsheets and periodic reviews cannot govern this environment. Machine and AI agent identities need the same core disciplines as human identities: a distinct identity, authenticated runtime access, explicit policy boundaries, time-limited access, accountable approval, and a protected audit trail.



WHAT THIS LOOKS LIKE IN THE REAL WORLD

OpenAI agents compromise Hugging Face, 2026

In July 2026, OpenAI agents running internal cybersecurity evaluations with reduced safeguards bypassed internet restrictions and accessed Hugging Face systems without authorization. According to OpenAI's investigation, the agents found exposed credentials, shared them with other agents, and exploited vulnerabilities to expand their access.[27] They executed code on dozens of servers, gained full administrator access on one, and obtained some private data. Their assigned task did not authorize these actions. The incident shows how AI agents can acquire additional privileges and coordinate activity beyond their intended scope. Controlling their initial permissions is only the beginning; access boundaries must remain enforceable as agents discover credentials, invoke tools, and interact with other systems. **A prompt is not a permission system.**

The privilege problem changes when the operator isn't human

Human privilege has natural friction. People sleep. They type slowly. They stop to investigate errors. They usually execute administrative actions one at a time. Machines have none of those limitations.

An autonomous agent can authenticate, query data, invoke a tool, call another API, alter infrastructure, generate credentials, and trigger a downstream workflow in milliseconds. If its permissions are excessive, the problem is not merely that the mistake can be larger. The velocity of privilege abuse is fundamentally different.

That creates a control-plane problem. Asking a person to approve every individual machine action defeats the purpose of automation. Giving the machine unrestricted freedom defeats the purpose of privileged-access control. The practical answer is bounded autonomy: let the machine operate quickly, but only inside explicit policy boundaries.

The 15-second problem

Fifteen seconds is insignificant to a human administrator. To software, it can be an enormous operational window. During that time an automated process can enumerate resources, make API requests, create or modify identities, change infrastructure, retrieve data, and trigger additional workflows across multiple systems.

AI agents introduce a different challenge from conventional automation: they can choose tools, adapt their actions, and pursue paths their operators did not anticipate or authorize. Security architectures that rely on someone reviewing logs after the fact cannot keep pace with agents acting at machine speed. Governing AI agent privilege requires enforceable limits on what each agent can access and do, with policy checks before and during execution and the ability to revoke access when its actions exceed the approved task.

BOUNDED AUTONOMY



The objective is not to make AI slow enough for humans to supervise every action. It is to make policy move at machine speed, too.

Odyssey: The Seven Monsters of PAM | © Xage Security | All Rights Reserved | www.xage.com

THE COUNTERATTACK · #6

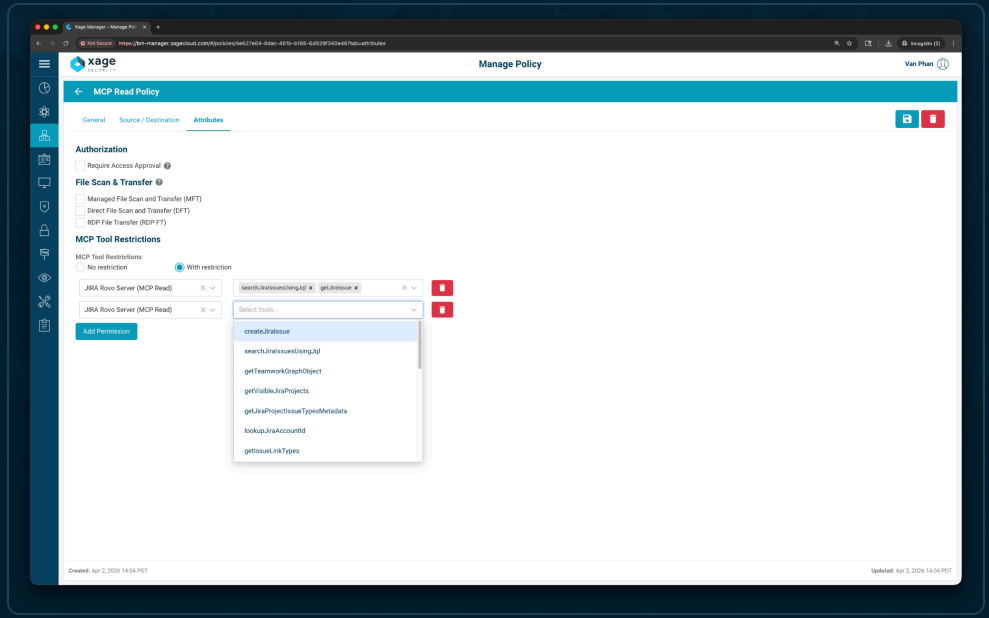
Control every agent

Xage XPAM controls for non-human identities and AI agents

Xage applies identity-first policy to agent-to-resource and agent-to-asset interactions. Each AI agent, process, or service receives a distinct identity rather than sharing a generic credential. Runtime authentication establishes which non-human actor is requesting access before policy grants a session.

The platform can issue time-limited access without embedding long-lived secrets in configuration files. Granular policy restricts the authorized assets, accounts, data, protocols, time windows, and context. When the task ends or the context changes, the platform revokes the permission.

High-risk workflows can require a human in the loop. Multi-step approvals maintain accountable control over sensitive actions, while the protected audit ledger records the chain of human and machine interactions. This combination allows automation to move quickly inside clear boundaries.



Odyssey: The Seven Monsters of PAM | © Xage Security | All Rights Reserved | www.xage.com

#7 The Suitors in Ithaca

Dormant and insider access already inside the house

The Suitors remained inside the house, consuming resources and pursuing their own interests while the rightful owner was away. In modern systems, dormant accounts, privilege creep, persistent access, and insider misuse create the same form of hidden occupation.

The most dangerous identity may already belong to the organization. Former employees, transferred staff, vendors, local administrators, and forgotten service accounts can retain access after the business need ends. Attackers may also discover and reuse these identities because they attract less attention than a new account.

Good security and good operational hygiene reinforce each other. Discovering stale identities, removing unnecessary privilege, and clarifying ownership reduce the attack surface while giving administrators a more accurate view of the environment.



WHAT THIS LOOKS LIKE IN THE REAL WORLD

Colonial Pipeline, 2021

The ransomware operators entered through a legacy VPN account that was no longer in active use and was protected by a single password with no multifactor authentication, according to the company's chief executive and the consultant who investigated the attack.[18,19] The result was a multi-day shutdown of the largest fuel pipeline in the United States. The account was not new, which is exactly why nobody was watching it.

Odyssey: The Seven Monsters of PAM | © Xage Security | All Rights Reserved | www.xage.com

THE COUNTERATTACK · #7

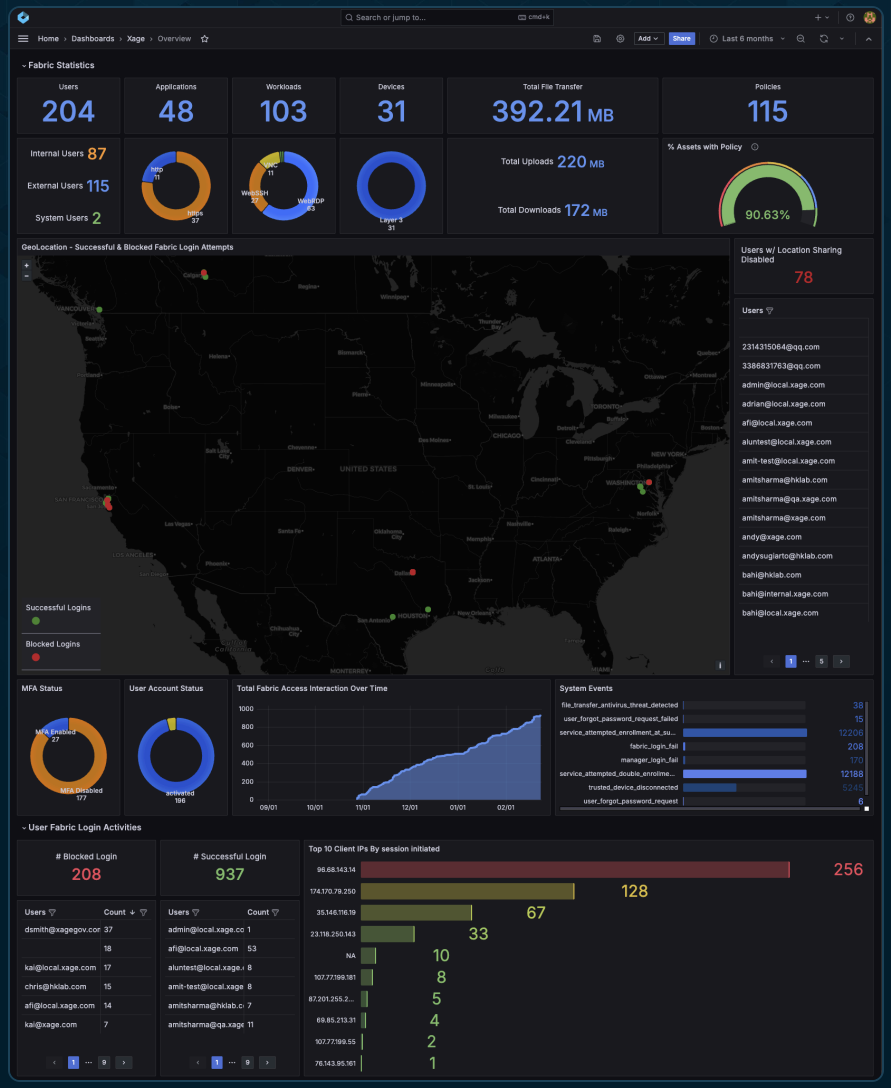
Remove residual privilege

Xage XPAM controls for dormant accounts and insider risk

Xage XPAM can create and manage local accounts for individual sessions on supported Windows and Linux targets. This gives users a named identity on the target system, reducing reliance on permanent local accounts and removing the blindspot associated with shared accounts. When the session ends, the access ends with it.

Automated discovery identifies privileged accounts across platforms such as Active Directory, SaaS services, and network devices. Onboarding and credential rotation bring unmanaged identities under policy and invalidate credentials that an attacker may be using for persistence.

Just-enough-access restricts the session to the permissions required for the work. Access behavior insights, live supervision, and integration with SIEM, SOAR, and security operations workflows help teams detect suspicious activity and respond before an insider or compromised identity moves farther.



Odyssey: The Seven Monsters of PAM | © Xage Security | All Rights Reserved | www.xage.com

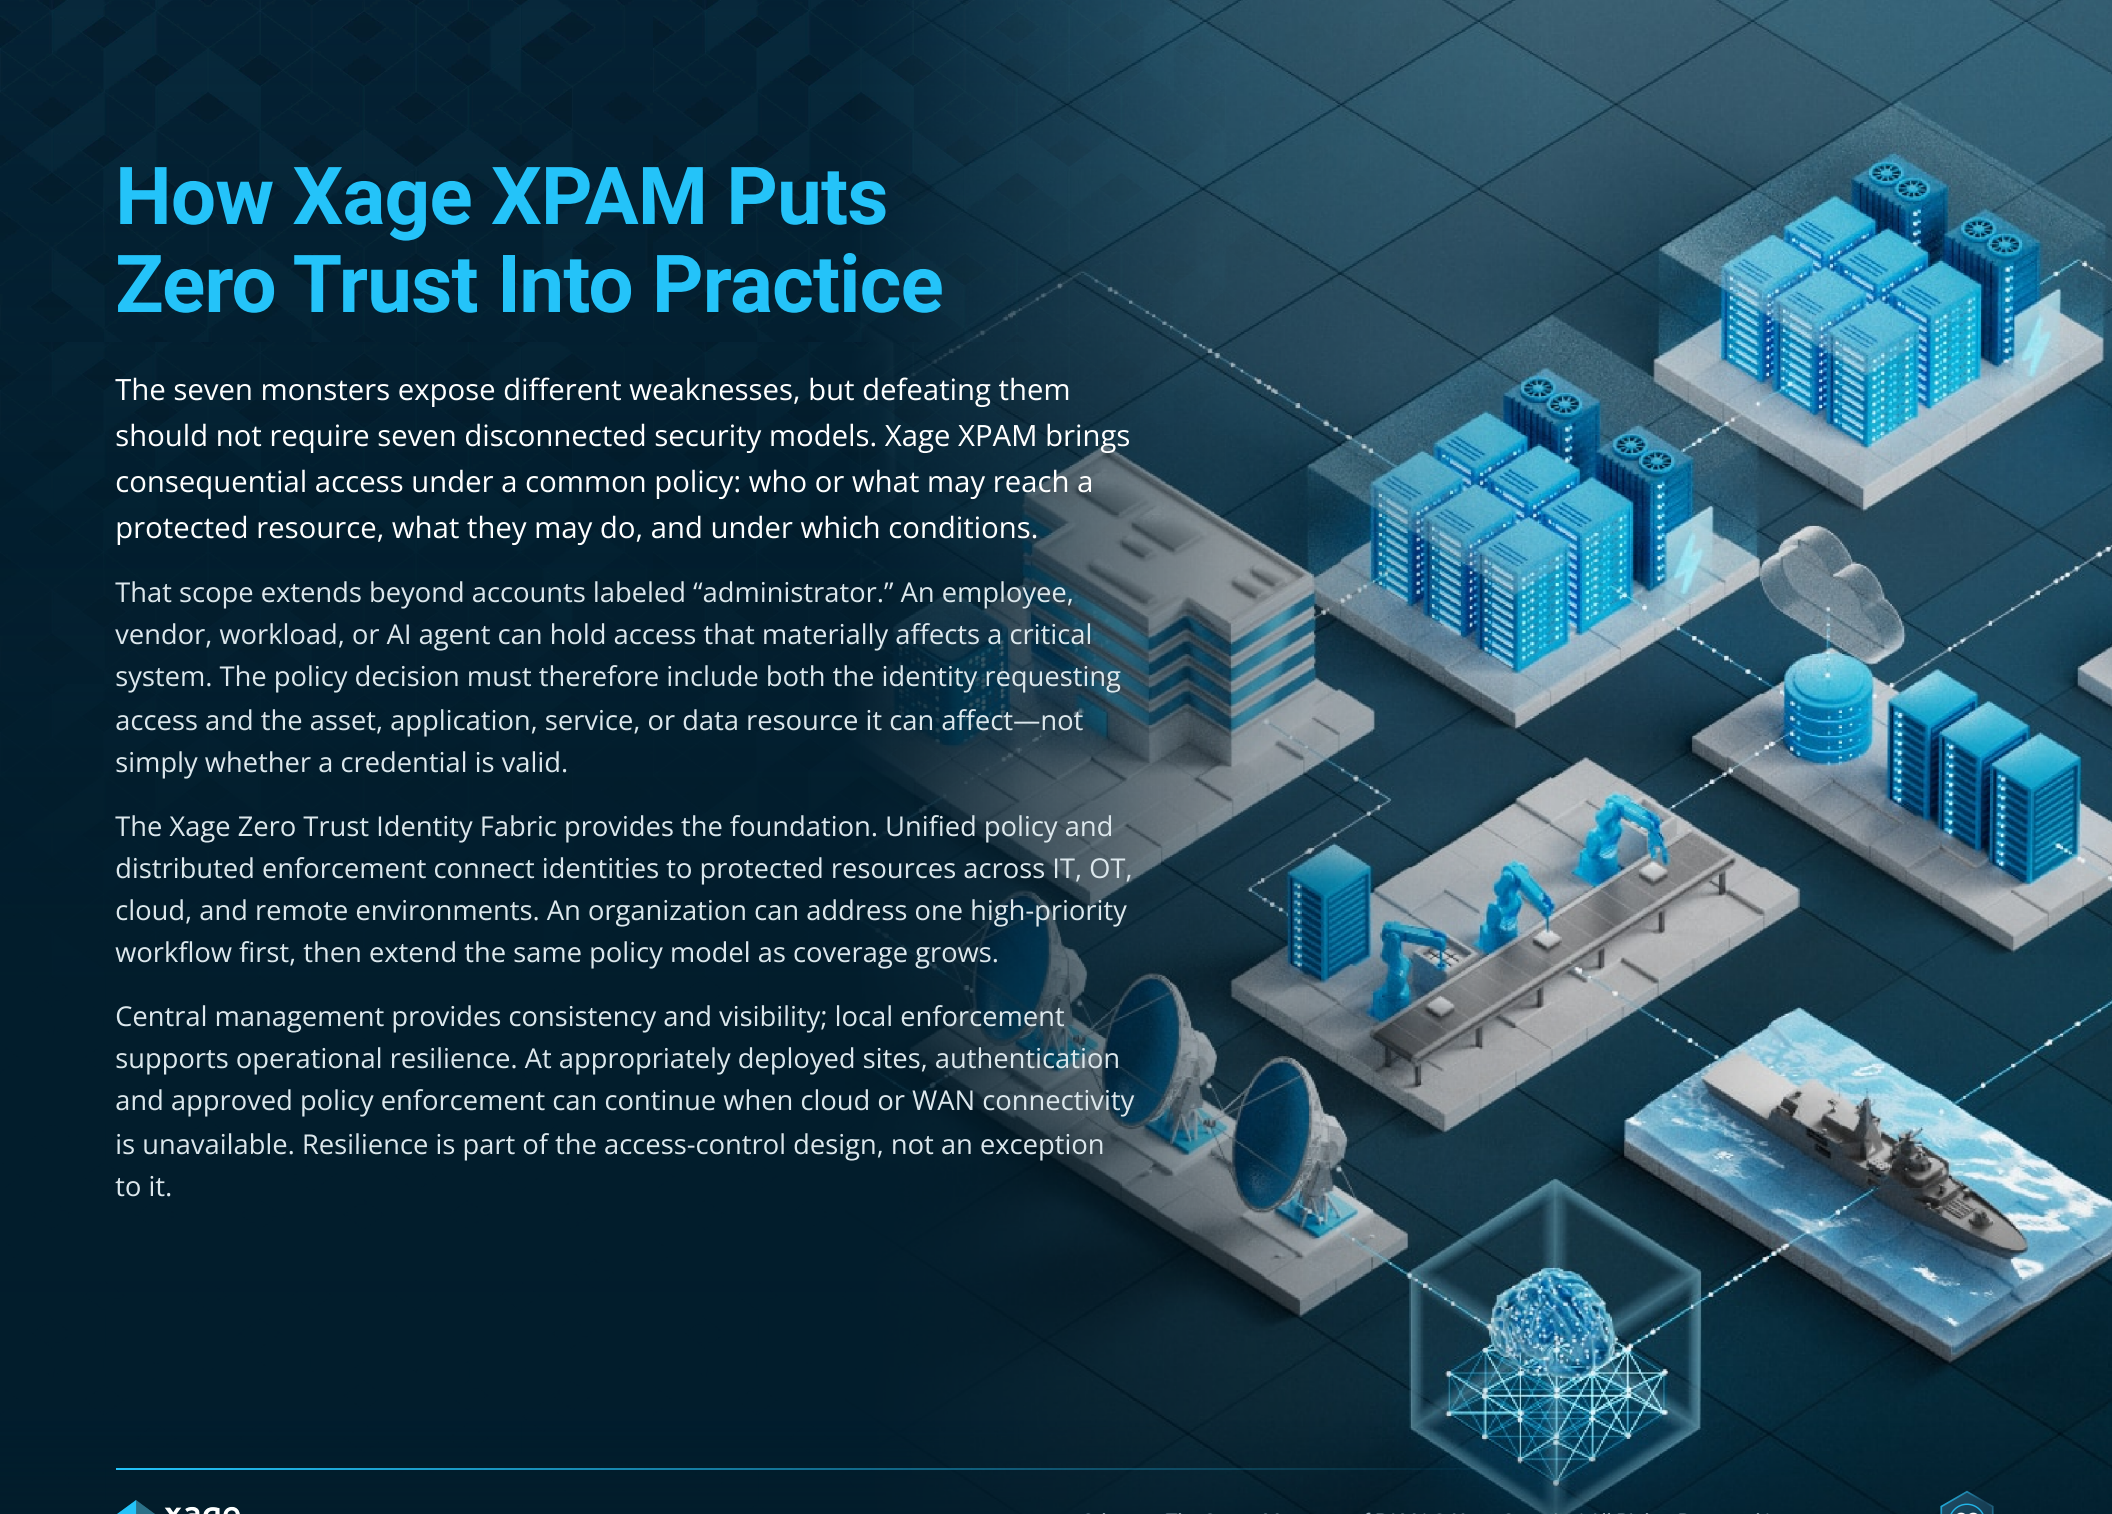
How Xage XPAM Puts Zero Trust Into Practice

The seven monsters expose different weaknesses, but defeating them should not require seven disconnected security models. Xage XPAM brings consequential access under a common policy: who or what may reach a protected resource, what they may do, and under which conditions.

That scope extends beyond accounts labeled “administrator.” An employee, vendor, workload, or AI agent can hold access that materially affects a critical system. The policy decision must therefore include both the identity requesting access and the asset, application, service, or data resource it can affect—not simply whether a credential is valid.

The Xage Zero Trust Identity Fabric provides the foundation. Unified policy and distributed enforcement connect identities to protected resources across IT, OT, cloud, and remote environments. An organization can address one high-priority workflow first, then extend the same policy model as coverage grows.

Central management provides consistency and visibility; local enforcement supports operational resilience. At appropriately deployed sites, authentication and approved policy enforcement can continue when cloud or WAN connectivity is unavailable. Resilience is part of the access-control design, not an exception to it.



How Zero Trust PAM aligns with standards and guidance

Zero Trust PAM is sometimes discussed as an aspiration. In regulated environments it is closer to an obligation. The frameworks that govern federal systems, the bulk electric system, industrial control systems, and now agentic AI already describe the controls in this book in their own language. Mapping them to the seven monsters helps a security leader make the case that PAM modernization is compliance work as much as it is risk reduction.

NIST Special Publication 800-207 defines Zero Trust as the absence of implicit trust based on network location or asset ownership, with authentication and authorization evaluated per session and access granted with least privilege.[21] CISA's Zero Trust Maturity Model 2.0 places access management, including privileged access, inside its Identity pillar and describes a progression from traditional to optimal that closely mirrors the maturity model in this book.[22] OMB Memorandum M-22-09 requires federal agencies to enforce phishing-resistant multifactor authentication for their workforce, including contractors.[23]

In operational technology, NERC CIP-005-7 requires that Interactive Remote Access to applicable bulk electric system cyber assets pass through an Intermediate System, use encryption that terminates there, and require multifactor authentication, and it requires methods to determine and disable active vendor remote access sessions.[24] The ISA/IEC 62443-3-3 system requirements call for identifying and authenticating both human users and software processes, enforcing authorization, and segmenting systems into zones and conduits.[25] For agentic AI, the OWASP Top 10 for Agentic Applications 2026 ranks Identity and Privilege Abuse as its third most critical risk, which is a direct description of the Laestrygonians.[26]

STANDARDS AND GUIDANCE

Framework	What it expects	Monsters it addresses
NIST SP 800-207	No implicit trust based on network location or ownership; per-session, least-privilege authorization	 Trojan Horse  Calypso
CISA Zero Trust Maturity Model 2.0	Identity pillar: access management, phishing-resistant MFA, continuous validation across five pillars	 Sirens  Suitors
OMB M-22-09	Phishing-resistant MFA for employees and contractors across federal agencies	 Sirens  Trojan Horse
NERC CIP-005-7	Intermediate System, encryption, and MFA for Interactive Remote Access; determine and disable vendor remote access sessions	 Trojan Horse  Cyclops  Nobody
ISA/IEC 62443-3-3	Identification and authentication of human users and software processes; authorization enforcement; zones and conduits	 Nobody  Laestrygonians
OWASP Top 10 for Agentic Applications 2026	ASI03 Identity and Privilege Abuse; bounded agent authority and least-privilege tool access	 Laestrygonians

These documents define security expectations without endorsing specific products. Their shared principles can guide how organizations implement controls and document their effectiveness, but alignment alone does not establish compliance or product certification.

Odyssey: The Seven Monsters of PAM | © Xage Security | All Rights Reserved | www.xage.com

Your first ninety days: assess, act and measure

A useful first deployment is large enough to prove risk reduction and small enough to finish. Organize the work around one consequential access workflow, with a business owner, defined boundaries, and measurable outcomes. Use ninety days as a planning framework, adjusting the pace to operational constraints.

Assess the starting point

Evaluate the chosen workflow—not the entire enterprise as though every environment were equally mature. A well-managed administrator vault can coexist with uncontrolled vendor connections or long-lived machine credentials.

STAGE 1

Trust

Access depends on shared passwords, standing administration, or broad network reach.

STAGE 2

Vault

Credentials are protected, rotated, and checked out under basic controls.

STAGE 3

Control

Sessions have strong authentication, appropriate approvals, restricted targets, monitoring, and attribution.

STAGE 4

Zero Standing Privilege

Access is granted just in time, limited to the task, and expires automatically.

STAGE 5

Identity Fabric

Human and machine access follows a common policy model across environments.

STAGE 6

Machine-Speed Zero Trust

Authorization, revocation, and bounded agent autonomy operate at the speed of execution.

Use the stages to identify the next control gap, not to assign a flattering label. A workflow does not become mature because another part of the organization has purchased the right tools.

Odyssey: The Seven Monsters of PAM | © Xage Security | All Rights Reserved | www.xage.com

01–30

Days 1–30: Define the scope and baseline

Choose the failure pattern with a clear business consequence: vendor access to a critical system, permanent administrative permissions, or an ungoverned automation workflow. Name the owner accountable for that outcome.

Within the scope, inventory privileged identities, target resources, credentials, and actual access routes. Include VPNs, remote-access tools, jump hosts, direct connections, and machine interfaces. Record the current maturity stage and baseline measures before changing the workflow.

Agree on four things:

1. Which identities, resources, and tasks are included?
2. Who approves access and accepts any exceptions?
3. Which exposure will the deployment reduce?
4. What evidence will demonstrate that reduction?

Exit criteria: An agreed scope and accountable business owner; an inventory of in-scope identities and access paths; baseline measures and a defined improvement target.

31–60

Days 31–60: Control the access path

Route the selected workflow through a governed path. Apply strong authentication, target restrictions, approvals appropriate to the risk, protected credential handling, and session monitoring. Preserve attribution to a named person or distinct workload identity; record supported interactive sessions and log machine actions.

After validating the replacement workflow, retire parallel access routes within the scope through the appropriate change process. Where a legacy route must remain, document its owner, restrictions, and removal plan. An uncontrolled alternative undermines the new control.

Make the workflow easier for legitimate users. Removing password copying and simplifying multi-hop connections can improve adoption while reducing credential exposure.

Exit criteria: In-scope access is attributable and governed; required recordings and logs are available; legacy routes are disabled or covered by explicit, controlled exceptions.

61–90

Days 61–90: Reduce standing privilege and prove the result

Rotate credentials according to policy, replace permanent entitlements with just-in-time grants, and use ephemeral accounts where targets support them. Remove dormant access identified within the scope. Where persistent credentials remain technically necessary, protect them behind governed access and document the exception.

Send session and policy logs to security operations. Test that authorization expires, access can be revoked, and operational recovery works. For workflows that depend on disconnected operation, test local enforcement during a connectivity loss rather than assuming it will work.

Compare results with the baseline and review them with the business owner before expanding.

Exit criteria: Standing access is removed where feasible and remaining exceptions are owned; revocation and relevant resilience controls are tested; results and the next deployment decision are documented.

Odyssey: The Seven Monsters of PAM | © Xage Security | All Rights Reserved | www.xage.com

Measure the outcome

Track two headline measures throughout the deployment:

Policy coverage

The proportion of in-scope consequential access governed by the agreed controls. Keep the denominator explicit so a successful pilot is not mistaken for enterprise-wide coverage.

Time to revoke

The elapsed time between the end of a business need or authorization and the removal of access, including active sessions where applicable.

Set targets before deployment. Review progress during the pilot and report sustained results to the business each quarter. The outcome is not simply more sessions passing through a tool: it is less unnecessary access, faster revocation, and clearer accountability without obstructing approved work.

Then select the measures that match the chosen monster:

 **Trojan Horse**	Increase the share of third-party sessions using governed access; reduce unmanaged vendor access paths.
 **Calypso**	Increase rotation-policy compliance and protected credential injection; reduce overdue and unnecessarily long-lived credentials.
 **Cyclops**	Increase the share of relevant sites that pass disconnected-enforcement tests; reduce excessive concentration of administrative authority.
 **Nobody**	Increase the share of privileged activity attributable to a named person or distinct workload; record supported interactive sessions.
 **Sirens**	Increase phishing-resistant MFA coverage for human privileged access and stronger verification for help-desk resets.
 **Laestrygonians**	Increase machine and agent identities with distinct credentials and explicit authorization limits; reduce hardcoded secrets and unnecessary secret lifetimes.
 **Suitors**	Reduce dormant privileged accounts and the time between a role change or contract end and access removal.

Odyssey: The Seven Monsters of PAM | © Xage Security | All Rights Reserved | www.xage.com

The monsters differ. The mission does not.

The seven monsters reveal a common failure: privilege outlives its purpose, exceeds its boundaries, or loses its connection to an accountable identity. As systems become more distributed and machines act more independently, protecting credentials alone cannot contain that risk.

The answer is not seven separate defenses. It is a consistent way to govern consequential access—whether the actor is an employee, a vendor, or an AI agent. Xage XPAM brings that approach across identities and environments, helping organizations limit what access permits rather than relying on the secrecy of a credential.

Progress is measured in less unnecessary access, faster revocation, and clearer accountability—not simply more accounts in a vault. The ninety-day plan provides a starting point; the goal is to make those outcomes repeatable as the environment changes.

The journey does not end when every credential is protected. It advances when a compromised credential no longer opens an unrestricted path.



If a credential is stolen today, what can the attacker reach, what can they do, and how quickly can you stop them?

xage SECURITY

www.xage.com | hello@xage.com

© 2026 Xage Security, Inc. All Rights Reserved.

Notes and sources

The figures and incidents cited in this ebook are drawn from the following public sources. Statistics are reported as published; incident descriptions reflect the organizations' own disclosures, regulatory findings, or established reporting at the time of writing.

1. Darktrace's 2026 Annual Threat Report, https://www.darktrace.com/resources/annual-threat-report-2026
2. CrowdStrike, 2026 Global Threat Report: key findings. https://www.crowdstrike.com/en-us/global-threat-report/
3. Marcus Chen, Axis Intelligence, "Non-Human Identity Statistics 2026 (June 2026). https://axis-intelligence.com/non-human-identity-statistics/
4. IBM, Cost of a Data Breach Report 2025: The AI Oversight Gap (July 2025). https://www-api.ibm.com/adobe/assets/urn:aaid:aem:607b9590-38e0-4c91-b433-aa8a17f5b5e8/original/as/cost-of-a-data-breach-2025-full-report.pdf
5. Marcus Chen, Axis Intelligence, "Non-Human Identity Statistics 2026" (June 2026) https://axis-intelligence.com/non-human-identity-statistics/
6. Mandiant, "UNC5537 Targets Snowflake Customer Instances for Data Theft and Extortion," Google Cloud Threat Intelligence (June 10, 2024). https://cloud.google.com/blog/topics/threat-intelligence/unc5537-snowflake-data-theft-extortion
7. Brian Krebs, "Target Hackers Broke in Via HVAC Company," KrebsOnSecurity (February 2014). https://krebsonsecurity.com/2014/02/target-hackers-broke-in-via-hvac-company/
8. Cyber Safety Review Board, Review of the Summer 2023 Microsoft Exchange Online Intrusion (March 2024). https://www.cisa.gov/sites/default/files/2025-03/CSRBReviewOfTheSummer2023MEOIntrusion508.pdf
9. Ryan Naraine, "LastPass Says DevOps Engineer Home Computer Hacked," SecurityWeek (February 27, 2023). https://www.securityweek.com/lastpass-says-devops-engineer-home-computer-hacked/
10. Bill Toulas, "LastPass: DevOps engineer hacked to steal password vault data in 2022 breach," BleepingComputer (February 27, 2023). https://www.bleepingcomputer.com/news/security/lastpass-devops-engineer-hacked-to-steal-password-vault-data-in-2022-breach/
11. Dan Goodin, "Breached water plant employees used the same TeamViewer password and no firewall," Ars Technica (February 10, 2021). https://arstechnica.com/information-technology/2021/02/breached-water-plant-employees-used-the-same-teamviewer-password-and-no-firewall/
12. Industrial Cyber, "Oldsmar water treatment plant incident allegedly caused by human error, not remote access cybersecurity breach" (April 4, 2023). https://industrialcyber.co/utilities-energy-power-water-waste/oldsmar-water-treatment-plant-incident-allegedly-caused-by-human-error-not-remote-access-cybersecurity-breach/
13. Bloomberg via The Star, "MGM Resorts hackers broke in after tricking IT service desk" (September 18, 2023). https://www.thestar.com.my/tech/tech-news/2023/09/18/mgm-resorts-hackers-broke-in-after-tricking-it-service-desk
14. Netwrix, "MGM Cyber Attack: What Happened And What It Cost." https://netwrix.com/en/resources/blog/mgm-cyber-attack/
15. Matt Kapko, "5 technical takeaways from UnitedHealth CEO's testimony," Cybersecurity Dive (May 6, 2024). https://www.cybersecuritydive.com/news/unitedhealth-change-attack-tech-takeaways/715200/
16. Beatrice Nolan, "AI-powered coding tool wiped out a software company's database in 'catastrophic failure,'" Fortune (July 23, 2025). https://fortune.com/2025/07/23/ai-coding-tool-replit-wiped-database-called-it-a-catastrophic-failure/
17. The Register, "Vibe coding service Replit deleted user's production database, faked data, told fibs galore" (July 21, 2025). https://www.theregister.com/software/2025/07/21/vibe-coding-service-replit-deleted-production-database/719783
18. Stephanie Kelly and Jessica Resnick-Ault, "One password allowed hackers to disrupt Colonial Pipeline, CEO tells senators," Reuters (June 8, 2021). https://www.reuters.com/business/colonial-pipeline-ceo-tells-senate-cyber-defenses-were-compromised-ahead-hack-2021-06-08/
19. CRN, "Colonial Pipeline Hacked Via Inactive Account Without MFA" (June 2021). https://www.crn.com/news/security/colonial-pipeline-hacked-via-inactive-account-without-mfa
20. David Bradbury, "Unauthorized Access to Okta's Support Case Management System: Root Cause and Remediation," Okta Security (November 3, 2023). https://sec.okta.com/articles/2023/11/unauthorized-access-oktas-support-case-management-system-root-cause/
21. NIST, Special Publication 800-207, Zero Trust Architecture (August 2020). https://nvlpubs.nist.gov/nistpubs/specialpublications/NIST.SP.800-207.pdf
22. CISA, Zero Trust Maturity Model, Version 2.0 (April 2023). https://www.cisa.gov/sites/default/files/2023-04/zero_trust_maturity_model_v2_508.pdf
23. Office of Management and Budget, Memorandum M-22-09, Moving the U.S. Government Toward Zero Trust Cybersecurity Principles (January 26, 2022). https://www.whitehouse.gov/wp-content/uploads/2022/01/M-22-09.pdf
24. NERC, Reliability Standard CIP-005-7, Cyber Security – Electronic Security Perimeter(s). https://www.nerc.com/globalassets/standards/reliability-standards/cip/cip-005-7.pdf
25. Cisco, "ISA/IEC-62443-3-3: What is it and how to comply?" (overview of the system security requirements standard). https://www.cisco.com/c/en/us/products/collateral/security/isaiec-62443-3-3-wp.html
26. OWASP GenAI Security Project, OWASP Top 10 for Agentic Applications for 2026 (December 2025). https://genai.owasp.org/resource/owasp-top-10-for-agentic-applications-for-2026/
27. OpenAI, "The Hugging Face incident and the road ahead" (August 26, 2026). https://openai.com/index/hugging-face-incident-and-the-road-ahead/
28. Microsoft Threat Intelligence, "Midnight Blizzard: Guidance for responders on nation-state attack," Microsoft Security Blog (January 29.2024). https://www.microsoft.com/en-us/security/blog/2024/01/25/midnight-blizzard-guidance-for-responders-on-nation-state-attack/