xage SECURITY

Identity Security
for the Agentic Enterprise

Control Privileged Access Across Human, Machine, and AI Identities

A practical guide to reducing exposure, governing delegated authority, and containing identity-driven threats before they become business impact.



Table of Contents

Executive Summary 03

Introduction: **Identity Is Now the Control Plane** 04

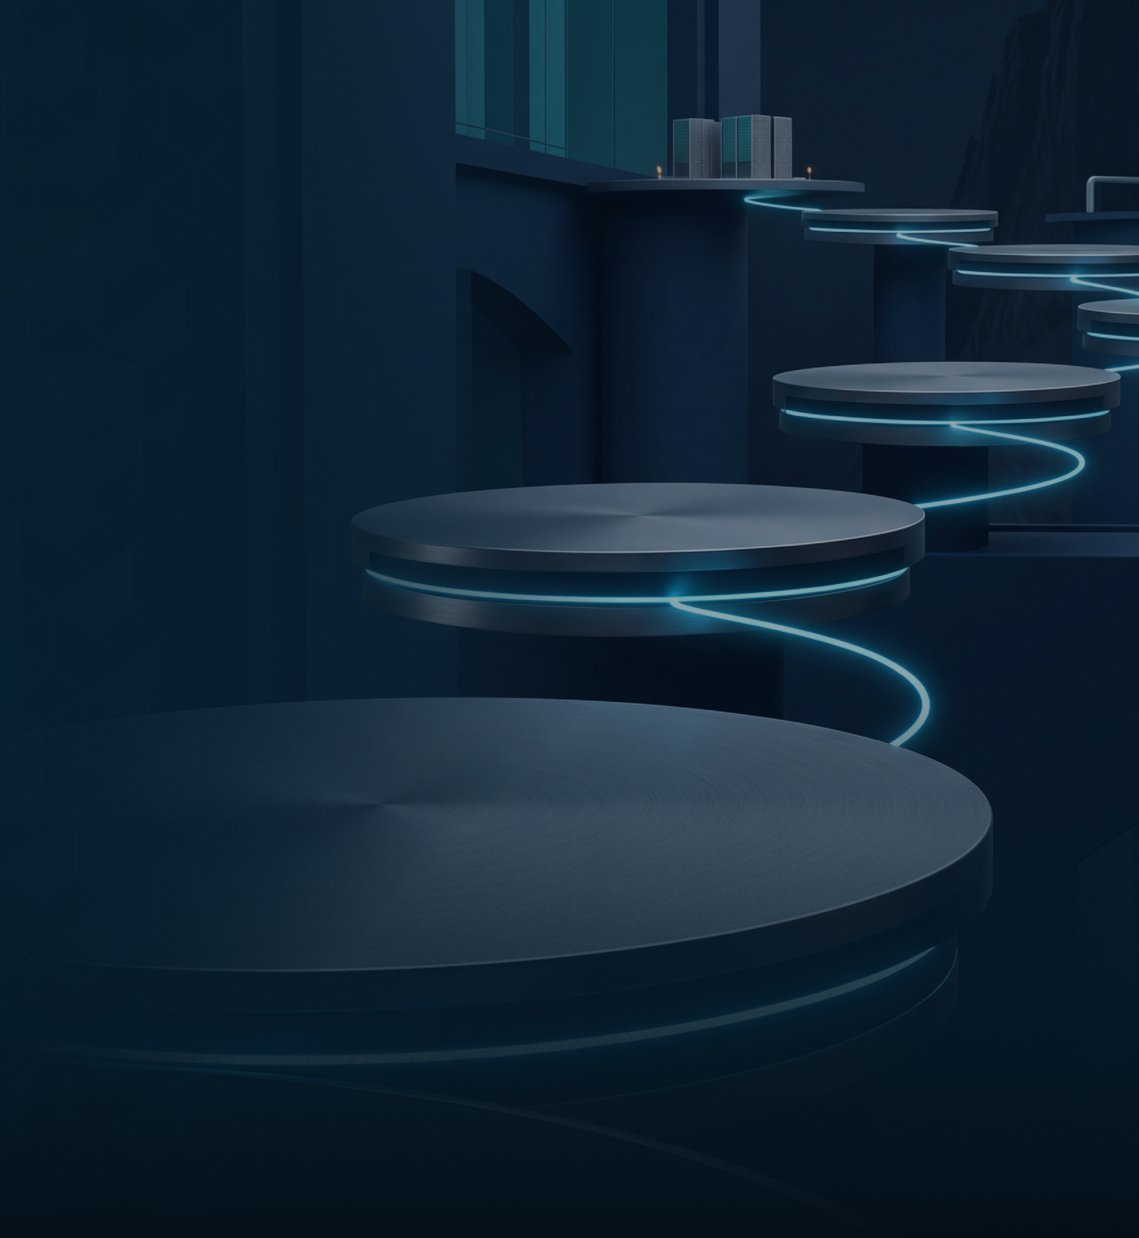
PART I: **The Expanding Identity Security Challenge** 05

PART II: **Building the Foundation** 09

PART III: **Governing Access and Action** 11

PART IV: **From Strategy to Execution** 15

Conclusion 17



Executive Summary

The security problem is shifting from proving identity to governing authority. Enterprises now depend on employees, vendors, service accounts, workloads, machines, automation, and AI agents, all of which can hold or inherit permissions. Traditional PAM remains essential, but approaches built primarily around human administrators, credentials, and login events cannot by themselves govern every identity and action across IT, OT, cloud, edge, and AI.

The response is an expanded model of privileged access: discover the actor, understand the authority it is using, minimize that authority, broker credentials where possible, evaluate the requested action, and revoke or contain access immediately when risk changes.

THE PRESSURE IS INCREASING ON BOTH SIDES.

76%

of organizations report growth in non-human identities

(SANS, 2026)



74%

already use AI agents or automations that require credentials

(SANS, 2026)



56%

YOY increase in AI-driven attacks

(IBM, 2026)



$1M

added to malicious breach costs on average

(IBM, 2026)



Where Xage fits

Xage extends these controls through a Zero Trust identity fabric. It complements IAM and IGA systems that establish identity, authentication, ownership, provisioning, and lifecycle governance, while Xage focuses on runtime enforcement: JIT/JEA, credential brokering, identity-based segmentation, session and action control, and rapid containment across distributed environments.

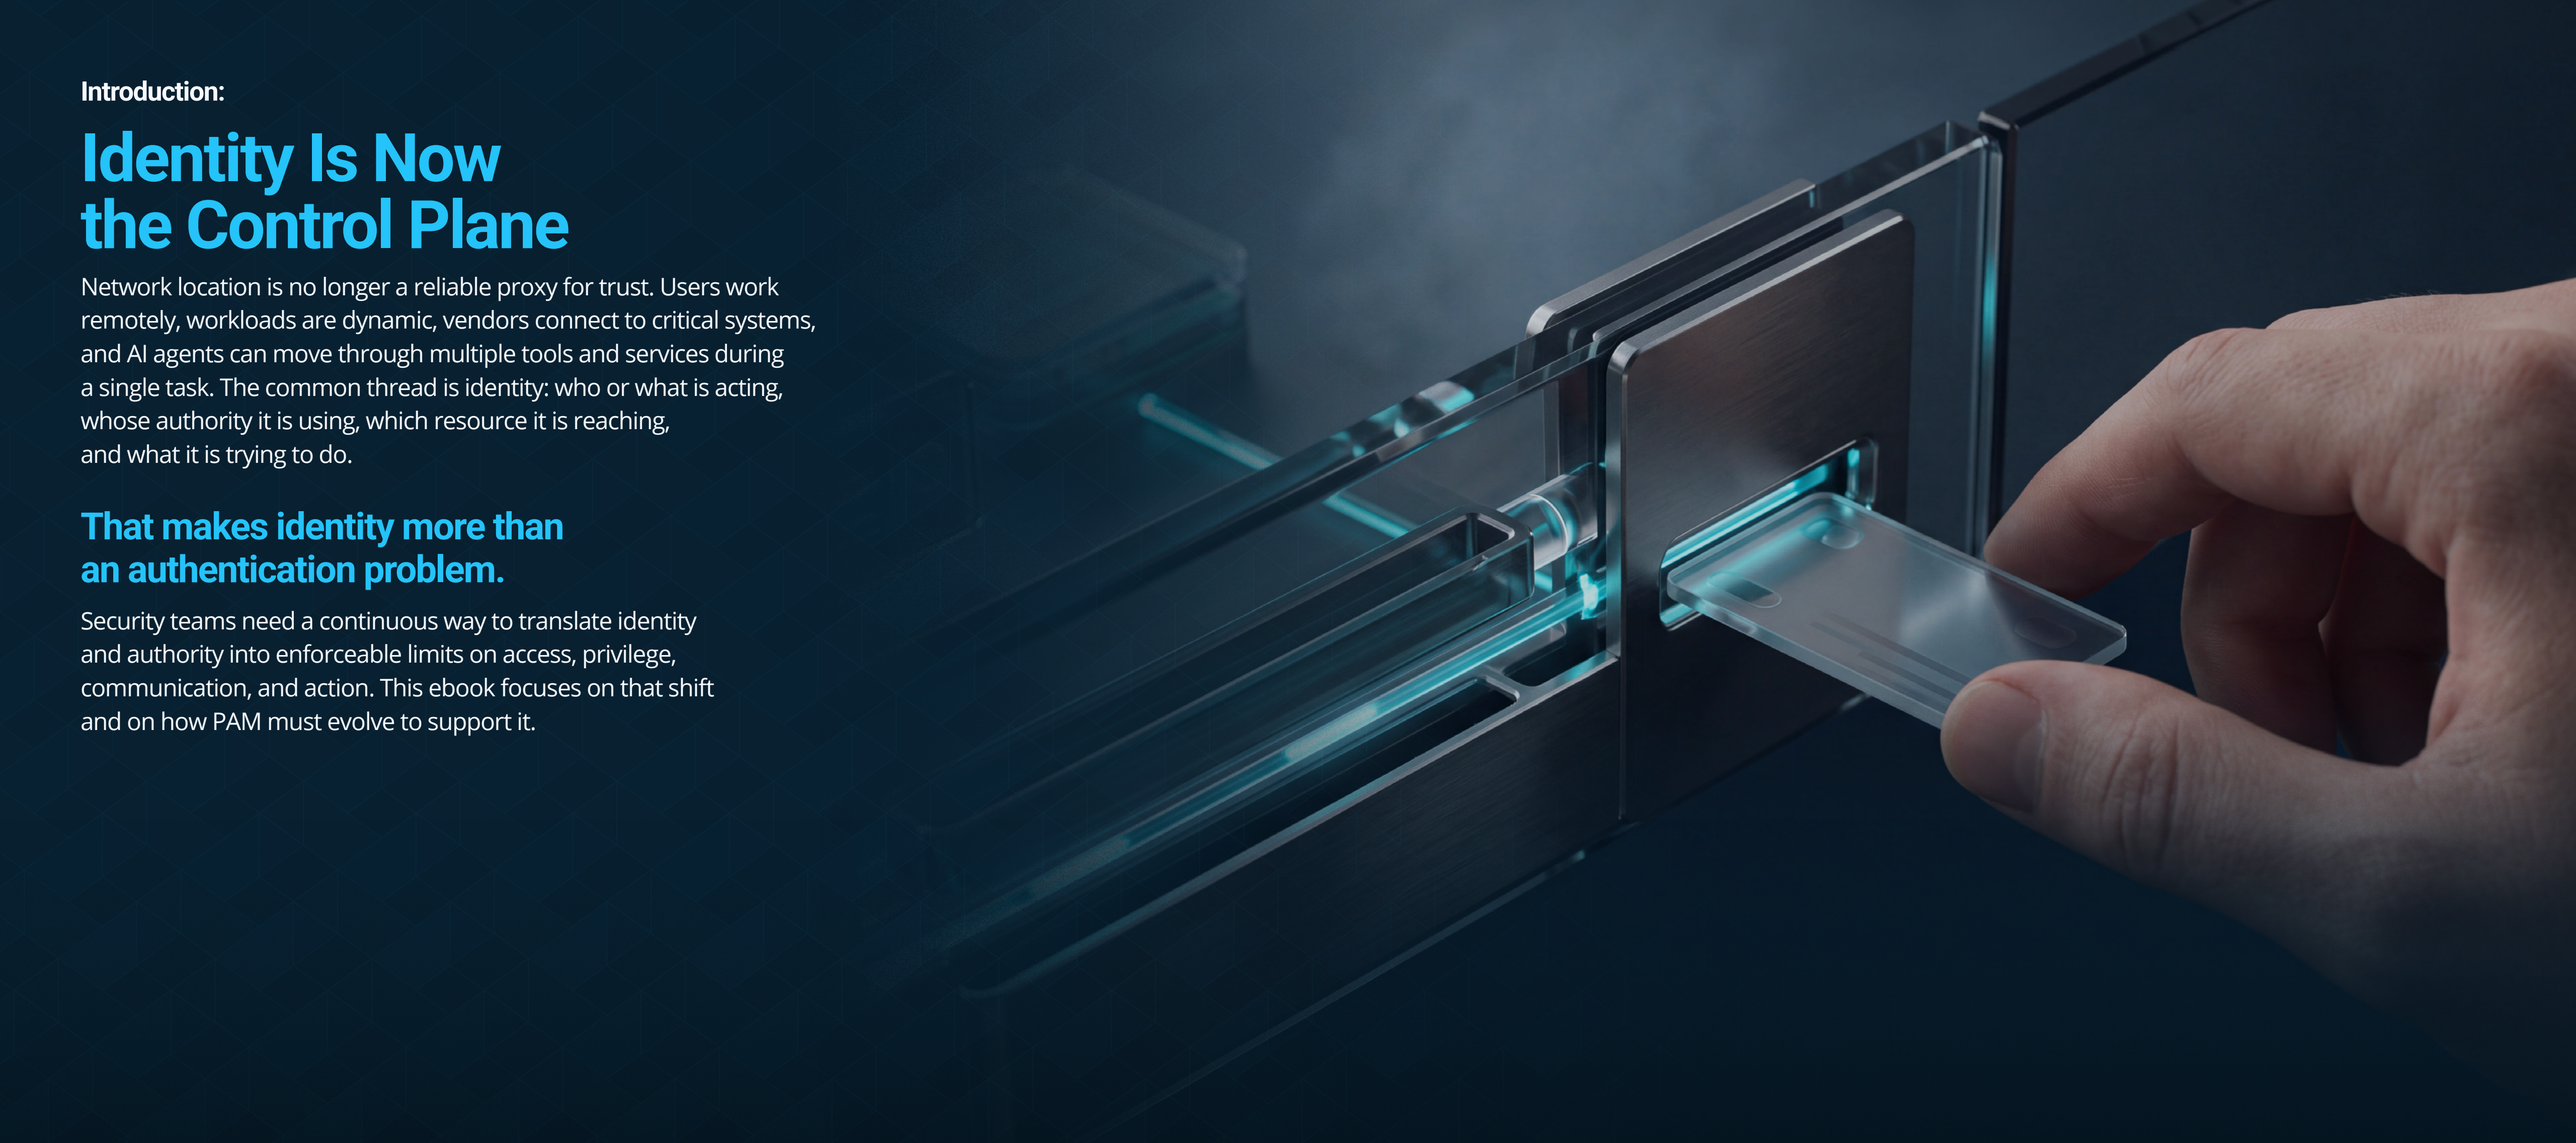
Introduction:

Identity Is Now the Control Plane

Network location is no longer a reliable proxy for trust. Users work remotely, workloads are dynamic, vendors connect to critical systems, and AI agents can move through multiple tools and services during a single task. The common thread is identity: who or what is acting, whose authority it is using, which resource it is reaching, and what it is trying to do.

That makes identity more than an authentication problem.

Security teams need a continuous way to translate identity and authority into enforceable limits on access, privilege, communication, and action. This ebook focuses on that shift and on how PAM must evolve to support it.



PART I

The Expanding Identity Security Challenge

The attack surface is growing faster than human-centric privilege models were designed to handle.

Chapter 1

Identity Explosion Is Expanding the Attack Surface

Why are traditional PAM approaches struggling to cover the modern enterprise?

Privileged access used to be easier to define: administrators logged in to known systems with privileged accounts. Today, authority is distributed across people, service accounts, cloud roles, applications, workloads, devices, API tokens, automations, and AI agents. Many of these identities are non-human, dynamic, and created outside HR-driven lifecycle processes.

AI compounds the problem. Agents can select tools, invoke APIs, use delegated credentials, and execute multistep workflows at machine speed. A single task may traverse a user identity, an agent identity, a workload identity, multiple credentials, and several resources. The attack surface expands with every identity, permission, relationship, and executable action introduced into the environment.

THE EXPANSION IS MEASURABLE

76%

of organizations report growth in non-human identities

(SANS, 2026)

74%

already use AI agents or automations that require credentials

(SANS, 2026)

21%

of organizations experienced an incident involving an AI model or application

(IBM, 2026)



www.xage.com | hello@xage.com | © Xage Security | All Rights Reserved

Why legacy PAM coverage becomes difficult

Legacy PAM architectures commonly center on credential vaulting, administrator access, and session control around a defined set of systems. Those functions remain critical, but the operating model becomes harder to extend when identities and resources are distributed across cloud, SaaS, OT, edge, and AI environments, or when systems cannot depend on a single centralized enforcement point.

The result can be uneven coverage across remote access, secrets, cloud privilege, segmentation, OT access, and agent controls. That gap is becoming more consequential as LLM capabilities advance. Attackers can use AI to identify exposed systems, accelerate reconnaissance, generate exploits, and move more quickly across portions of the environment that lack consistent controls. As the attack surface expands, every uncovered identity, resource, or access path increases the opportunity for compromise.

Each additional tool can also create policy handoffs, operational complexity, and integration work. The core challenge is not that PAM is obsolete. It is that the definition of privileged access has expanded beyond the scope of traditional deployments.

Privilege is now contextual

An identity may be privileged because it can modify production, access sensitive data, create users, issue credentials, approve a transaction, invoke a high-impact API, change an industrial process, or delegate authority to another actor. The security question is therefore broader than “Does this account have admin rights?” It is “What authority can this identity exercise right now, against this resource, for this action?”

IBM’s 2026 report shows why this matters: valid account abuse produced an average breach cost of $5.07 million and took 243 days to identify and contain. Excessive privileges and poor role management added $177,313 to average breach costs (IBM, 2026).

Key Takeaways



Identity growth expands the attack surface through permissions and relationships, not just account count.



Traditional PAM remains necessary, but human- and credential-centric coverage is no longer sufficient.



Privilege should be defined by potential impact, including actions taken by workloads, applications, vendors, and agents.



Reducing complexity requires a control model that can enforce policy consistently across distributed environments.

www.xage.com | hello@xage.com | © Xage Security | All Rights Reserved

Chapter 2

What Identity-Based Control Means

How does identity-based control turn identity security policy into enforceable limits?

Identity security is the broader discipline for governing human and non-human identities across access, authority, behavior, and lifecycle. PAM is a critical subset focused on privileged accounts, credentials, sessions, and elevated access.

Identity-based control is the enforcement mechanism within that broader discipline. It applies policy using identity, ownership, purpose, privilege, context, risk, resource sensitivity, and the requested action to determine what an identity may access, what it may do, and how quickly that authority can be reduced or revoked.

Authentication is only the starting point. A verified identity may still have excessive privilege, request access to a sensitive resource, or attempt an action that falls outside its purpose. Identity-based control carries policy into the interaction itself.

PAM must cover more types of privileged actors

PAM capabilities such as vaulting, credential rotation, session monitoring, JIT/JEA, and revocation remain foundational. But privileged authority is now exercised by more than human administrators. Applications assume cloud roles, workloads use tokens and secrets, vendors access critical systems, and AI agents invoke tools and APIs under delegated authority.

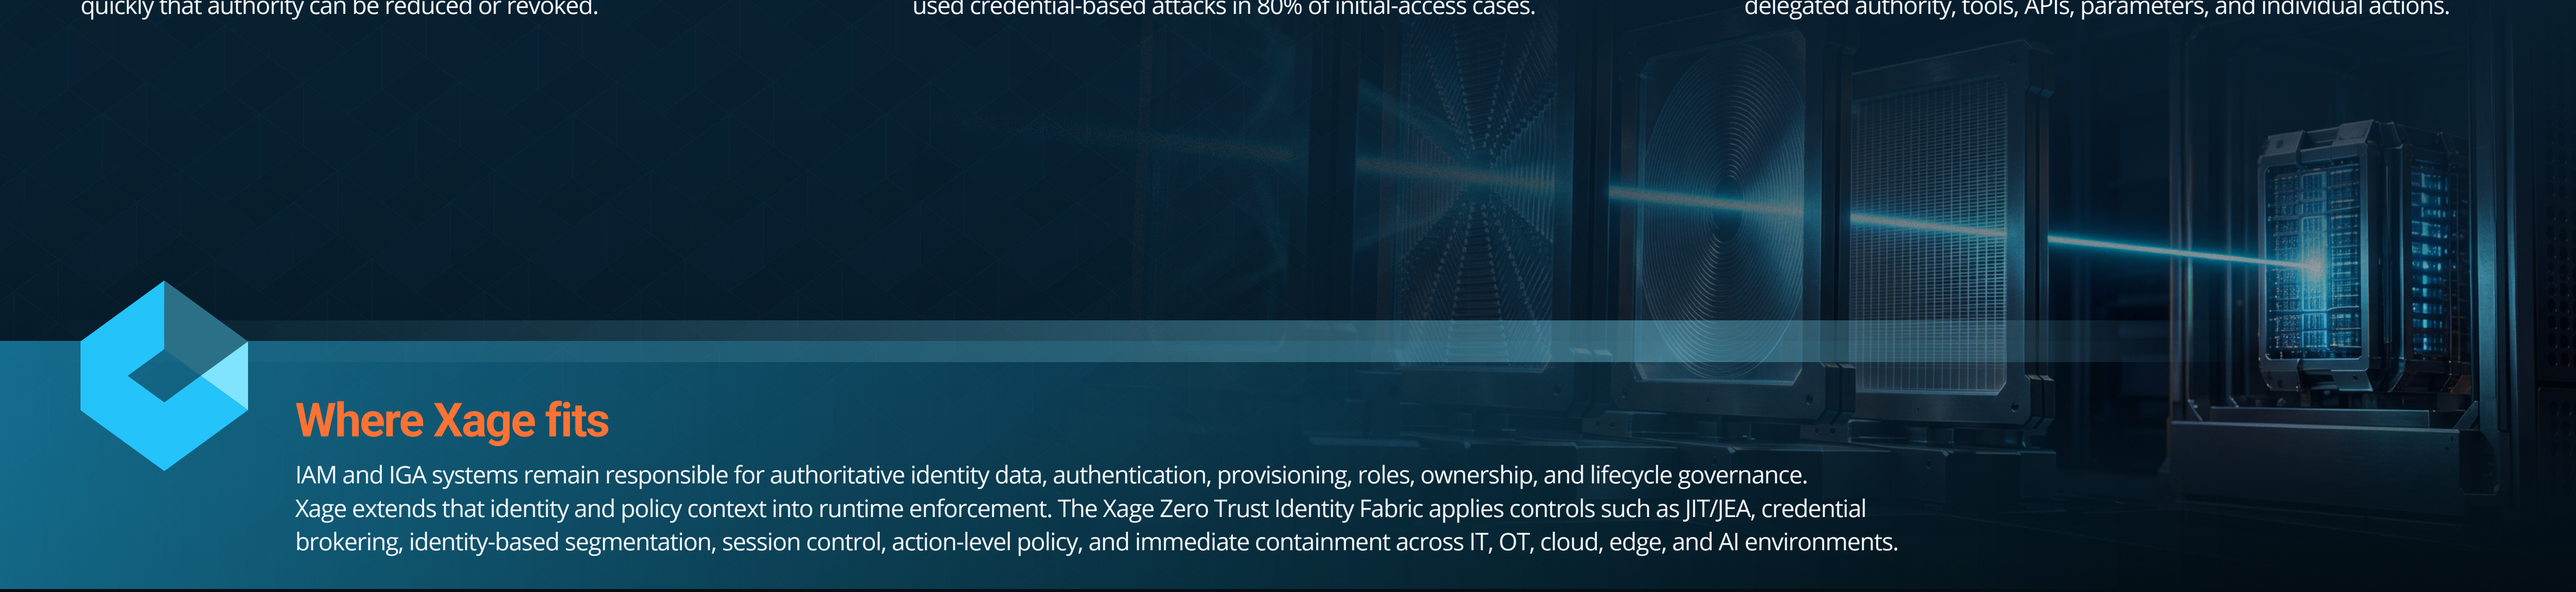
IBM's 2026 Cost of a Data Breach Report found that breaches involving valid account abuse cost an average of $5.07 million and took 243 days to identify and contain. It also found that excessive privileges and poor role management increased average breach costs by $177,313. Microsoft's 2025 Digital Defense Report found that access brokers used credential-based attacks in 80% of initial-access cases.

The issue is therefore not only protecting credentials, but controlling the authority attached to every identity. From access to action.

Identity-based control asks:

- Who or what is acting?
- What authority does it have, and where did that authority come from?
- Which resource and action are involved?
- What context or risk should affect the decision?
- Should access continue, be restricted, require approval, or be revoked?

This becomes especially important for AI agents. Many PAM principles still apply, including least privilege, JIT/JEA, credential brokering, session control, auditing, and revocation. But agents also require control over delegated authority, tools, APIs, parameters, and individual actions.



Where Xage fits

IAM and IGA systems remain responsible for authoritative identity data, authentication, provisioning, roles, ownership, and lifecycle governance. Xage extends that identity and policy context into runtime enforcement. The Xage Zero Trust Identity Fabric applies controls such as JIT/JEA, credential brokering, identity-based segmentation, session control, action-level policy, and immediate containment across IT, OT, cloud, edge, and AI environments.

Chapter 2

Key Takeaways



Identity-based control turns identity policy into enforceable limits on access and action.



PAM principles must extend to non-human identities and AI agents where they exercise privileged authority.



Effective control must govern authority continuously, not only at login.



www.xage.com | hello@xage.com | © Xage Security | All Rights Reserved

PART II

Building the Foundation

Establish visibility, ownership, and governance across human and non-human identities.



Chapter 3

The Identity Foundation

What must organizations know about human and non-human identities before they can control them?

Every identity needs a clear purpose, accountable ownership, an appropriate assurance level, and a lifecycle. The required signals differ by principal type, but the governance objective is consistent: know what the identity is, why it exists, which authority it should possess, and how that authority will be removed.

Identity type	Assurance/ governance signals	Common risk
Humans	Authoritative source, MFA, device posture, role, session risk	Phishing, session theft, excessive access
Applications/workloads	Deployment identity, code provenance, workload identity, owner	Broad roles, token abuse, hidden ownership
Devices/machines	Registration, certificate trust, configuration, expected behavior	Unmanaged trust, lateral paths
AI agents	Unique identity, purpose, owner, model/version, tools, authority source	Delegated authority, tool misuse, action chaining

Non-human identities are especially difficult because they often lack a single lifecycle event. A service account can outlive an application; a cloud role can remain after a project ends; an API key can outlive its developer; an agent can remain active after its purpose or tool set changes. Discovery and review therefore need to be continuous rather than periodic.

Xage does not replace the enterprise systems that own workforce identity lifecycle or access governance. Instead, it can use their identity and entitlement context as inputs to enforcement, while extending control to resources and environments that are difficult to cover with centralized identity systems alone.

www.xage.com | hello@xage.com | © Xage Security | All Rights Reserved

Chapter 4

Minimize and Protect Authority

How can organizations reduce standing privilege and prevent credentials from becoming reusable attack paths?

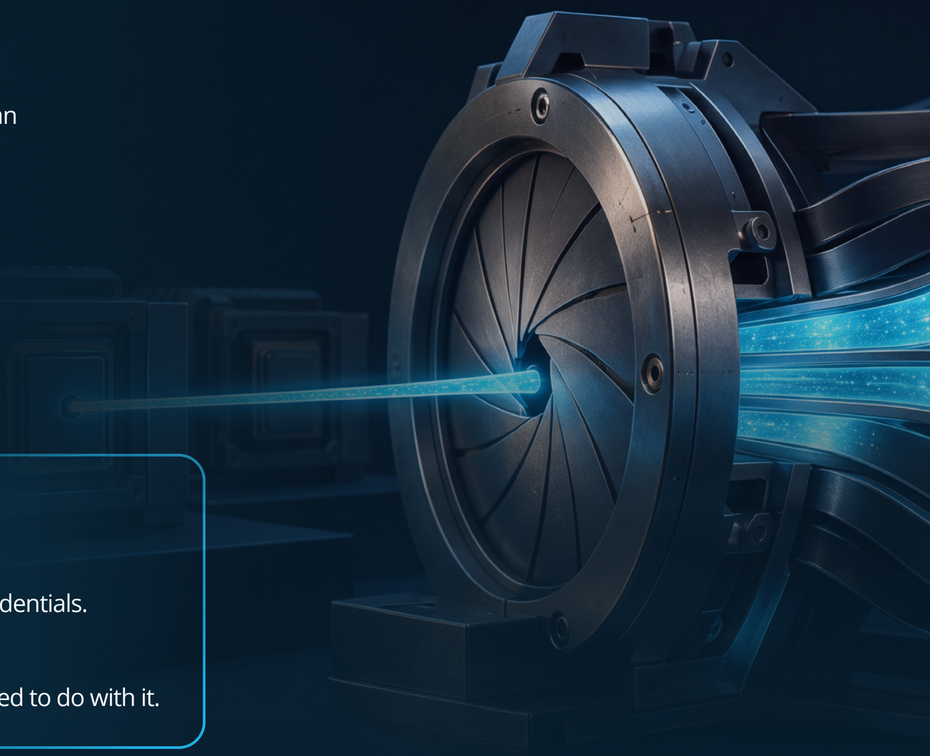
Entitlements accumulate because granting access is usually easier than removing it. Organizations should identify unused permissions, inherited privilege, toxic combinations, over-scoped service accounts, broad cloud roles, and unnecessary delegation, then replace persistent authority with narrower, temporary access wherever possible.

Standing privilege should give way to just-in-time and just-enough access, temporary role assumption, task-bound elevation, and automatic expiration. Delegated authority should be explicit, purpose-bound, scope-limited, time-bound, auditable, and revocable.

Broker credentials instead of exposing them

Credentials, tokens, secrets, and sessions are mechanisms identities use to exercise authority. Where possible, users, workloads, vendors, and agents should not retrieve persistent secrets directly. A broker can inject or issue temporary credentials tied to the requesting identity, task, target, and duration, then revoke them when the task ends.

This matters particularly for non-human identities. 92% of organizations fail to rotate most NHI credentials every 90 days, while only 8% rotate the majority within that period (SANS, 2026). Short-lived credentials, automated rotation, credential injection, certificate lifecycle management, and rapid revocation reduce the amount of reusable authority available to an attacker.



PAM principles apply directly to AI agents:

- Do not let the agent inherit unlimited user permissions.
- Use task-bound, short-lived authority rather than standing credentials.
- Keep secrets outside prompts and model memory.
- Preserve who delegated authority and what the agent is allowed to do with it.

Key Takeaways



Every identity needs ownership, purpose, appropriate assurance, and a revocation path.



Least privilege must address standing entitlements and delegated authority, not only administrator accounts.



Credential brokering and short-lived authority reduce the value of stolen secrets.



AI agents should receive bounded authority tied to a specific task, tool, resource, and duration.

www.xage.com | hello@xage.com | © Xage Security | All Rights Reserved

PART III

Governing Access and Action

Enforce least privilege continuously and connect risk signals directly to containment.

Chapter 5

Continuous Authorization, Segmentation, and Containment

How can policy remain active after authentication?

Traditional access controls often make a decision at login and assume the interaction should continue. Continuous authorization reevaluates that decision as identity, privilege, device or workload posture, resource sensitivity, approvals, behavioral risk, and changes in the requested action.

This enables three controls that static access models struggle to provide consistently:



Action-level authorization:
Govern a command, configuration change, API operation, transaction, tool call, or parameter - not only access to the resource.

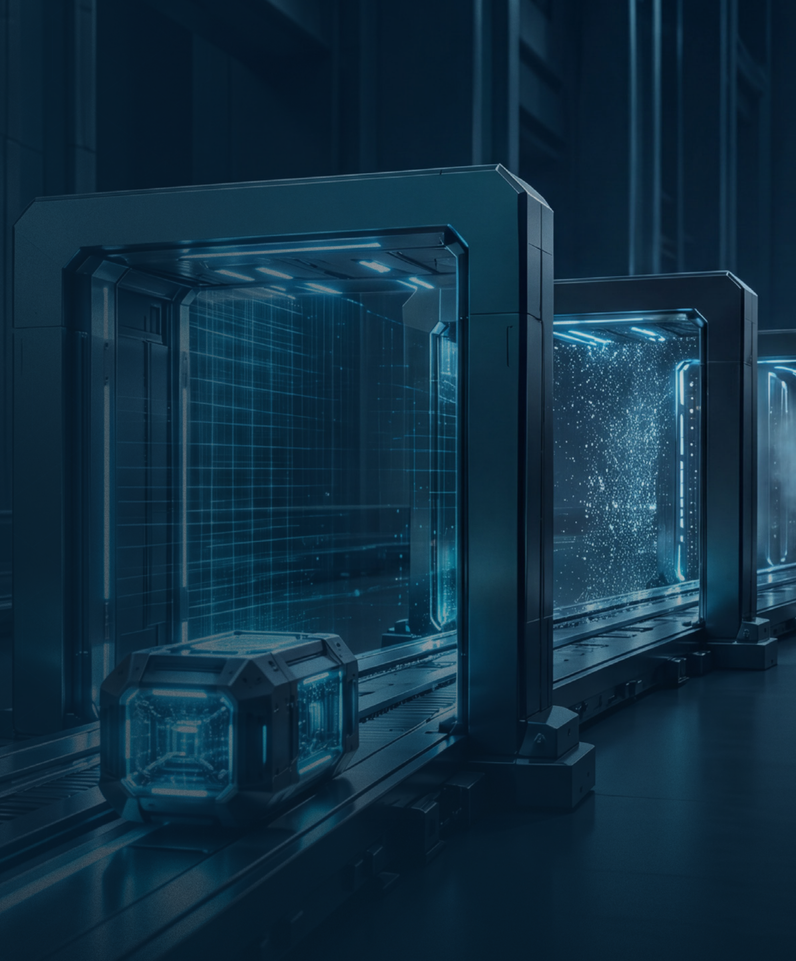


Identity-based segmentation:
Decide which identities may communicate with which systems, services, applications, and data sources, limiting broad network trust and lateral movement.



Immediate containment:
Revoke authority, terminate sessions or actions, block tools, quarantine workloads, or isolate an identity when risk changes.

These controls should also extend across organizational boundaries. Breaches involving third parties accounted for 48% of all breaches in Verizon's 2026 analysis, while IBM found supply-chain compromises took 258 days to identify and contain and averaged $4.96 million (Verizon, 2026; IBM, 2026).



Detection is valuable only when it can change authority

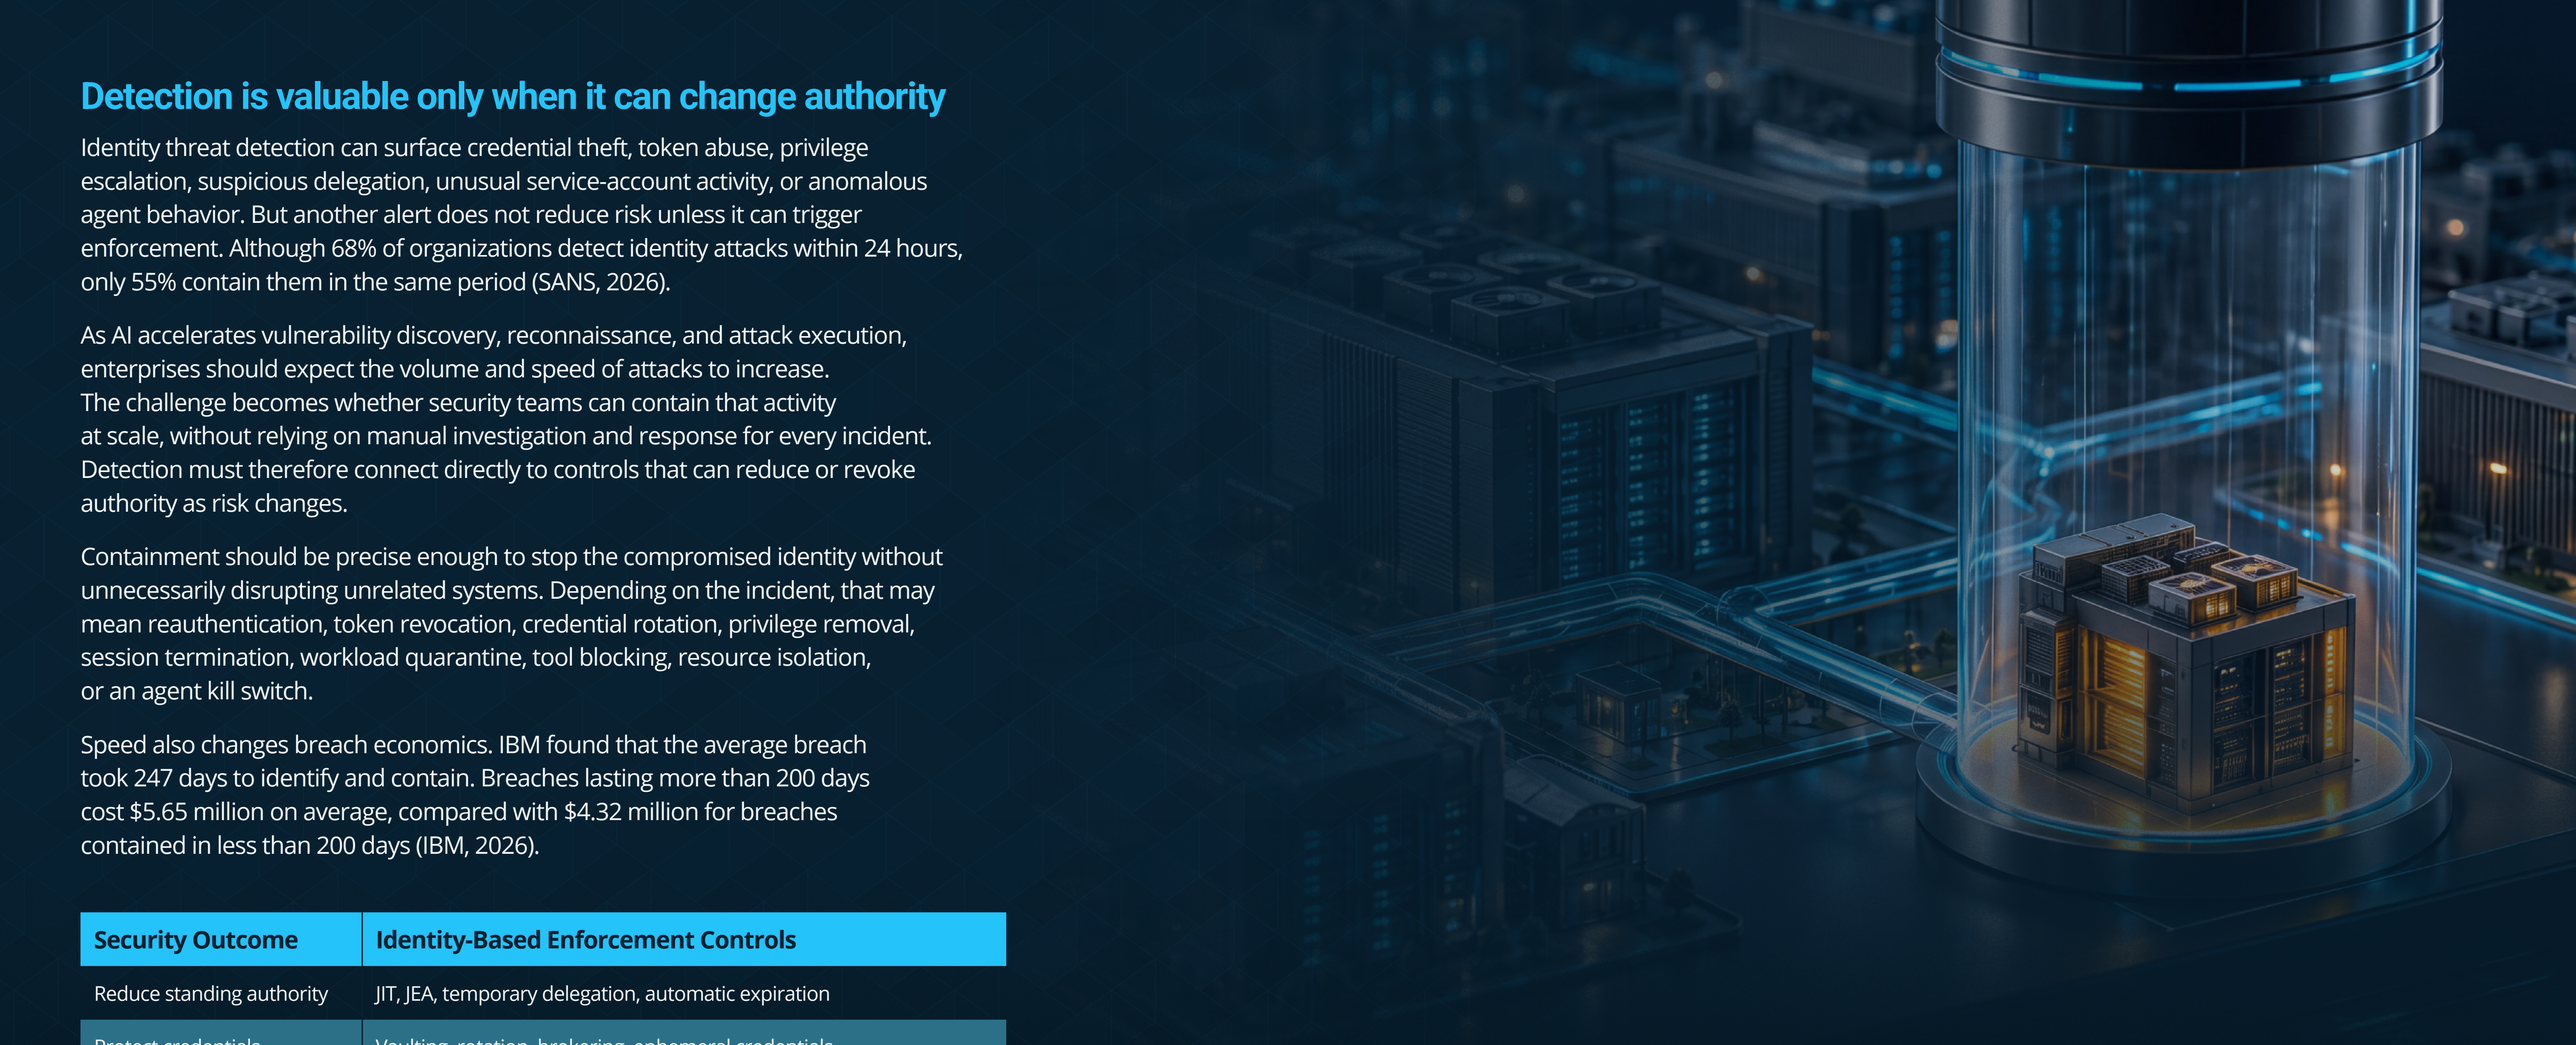
Identity threat detection can surface credential theft, token abuse, privilege escalation, suspicious delegation, unusual service-account activity, or anomalous agent behavior. But another alert does not reduce risk unless it can trigger enforcement. Although 68% of organizations detect identity attacks within 24 hours, only 55% contain them in the same period (SANS, 2026).

As AI accelerates vulnerability discovery, reconnaissance, and attack execution, enterprises should expect the volume and speed of attacks to increase. The challenge becomes whether security teams can contain that activity at scale, without relying on manual investigation and response for every incident. Detection must therefore connect directly to controls that can reduce or revoke authority as risk changes.

Containment should be precise enough to stop the compromised identity without unnecessarily disrupting unrelated systems. Depending on the incident, that may mean reauthentication, token revocation, credential rotation, privilege removal, session termination, workload quarantine, tool blocking, resource isolation, or an agent kill switch.

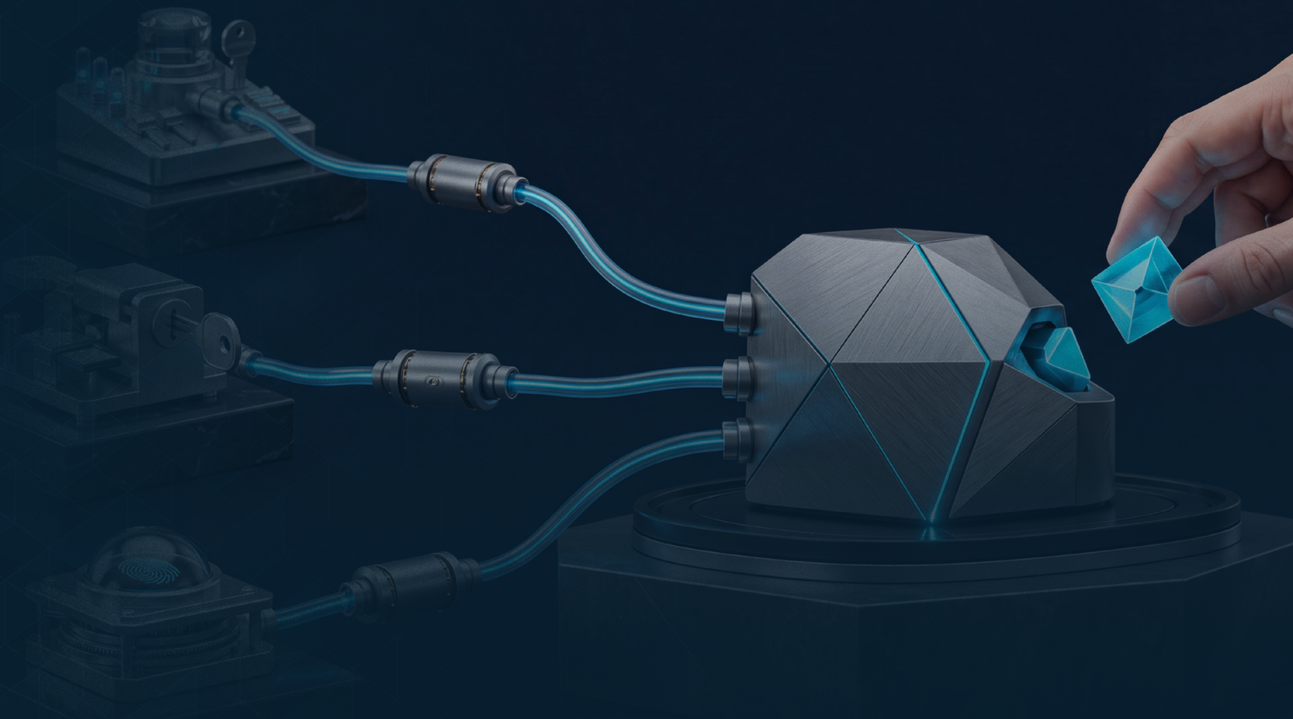
Speed also changes breach economics. IBM found that the average breach took 247 days to identify and contain. Breaches lasting more than 200 days cost $5.65 million on average, compared with $4.32 million for breaches contained in less than 200 days (IBM, 2026).

Security Outcome	Identity-Based Enforcement Controls
Reduce standing authority	JIT, JEA, temporary delegation, automatic expiration
Protect credentials	Vaulting, rotation, brokering, ephemeral credentials
Govern execution	Runtime authorization, action and parameter restrictions, approvals
Limit movement	**Identity-based segmentation and resource isolation**
Contain compromise	Revocation, termination, quarantine, tool blocking, kill switches



Chapter 6

Governing AI Agents and Delegated Authority



What changes when software can choose tools and execute multistep actions on behalf of people and systems?

AI agents turn delegated authority into active execution. They can select tools, call APIs, access data, use credentials, and complete workflows without a human making each individual decision. That makes the agent a first-class identity and, when it can take consequential action, a potentially privileged actor.

Every production agent should have a unique identity, registered purpose, business and technical owner, approved tools, known authority source, risk classification, monitoring, and a defined revocation path. The delegation chain must remain visible so the agent is distinguishable from the user, application, workload role, OAuth grant, or other agent whose authority it uses.

Prompt controls are not an enterprise authorization boundary

Model safeguards and prompt filtering can reduce unsafe outputs, but they cannot reliably determine which enterprise actions to allow. Agents may also ignore, circumvent, or be manipulated into bypassing intended guardrails. The model may propose an action; an independent policy and enforcement layer should determine whether the agent is authorized to execute it.

Agent runtime controls

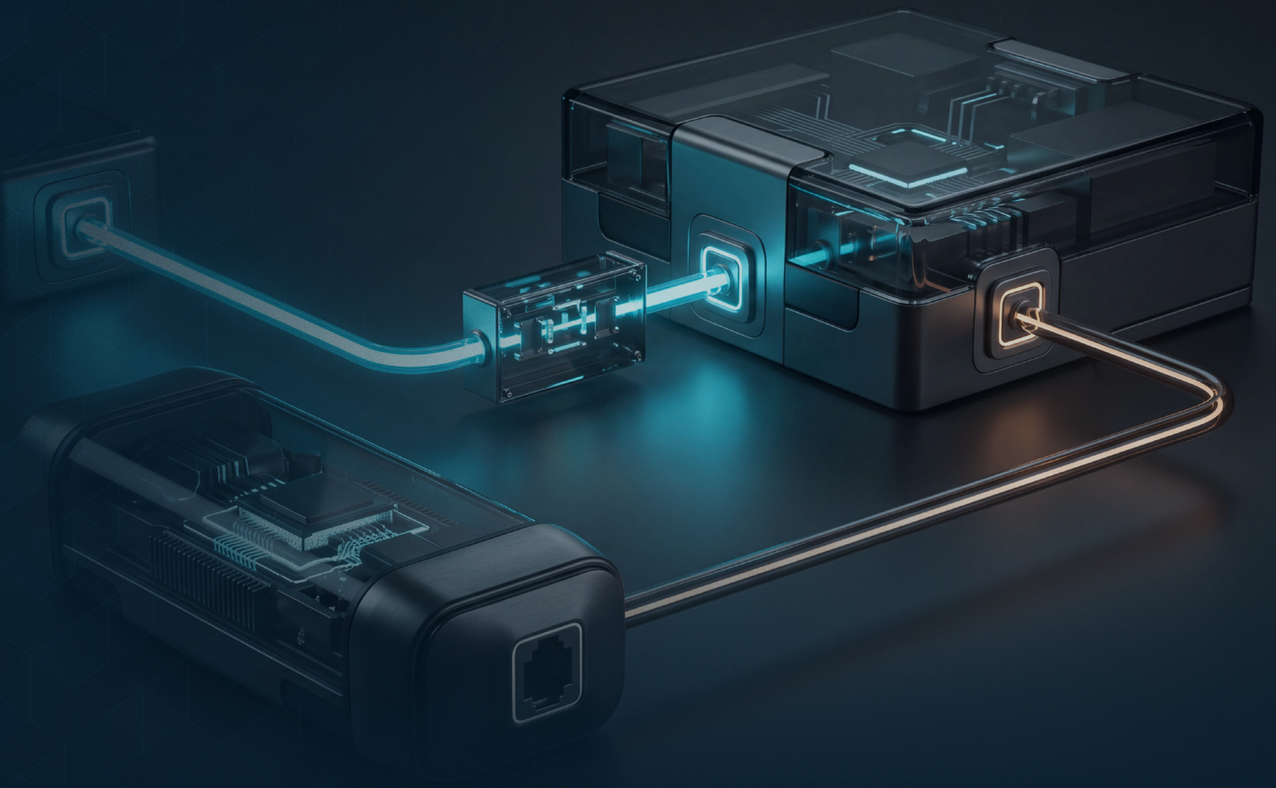
1. Explicit on-behalf-of delegation and task-bound authority
2. Tool, command, API, and parameter restrictions
3. Read/write separation and approval for high-impact actions
4. Limits on sub-agent delegation and lateral movement
5. Complete action logging, anomaly detection, and kill switches

The governance gap is already visible

IBM found that 21% of organizations experienced a breach involving an AI model or application in 2026, up from 13% the prior year. Among organizations with an AI-related breach, 92% lacked proper AI access controls, and those incidents cost an average of $5.33 million (IBM, 2026).

Shadow AI widens the problem. Security incidents involving unsanctioned AI more than doubled to 43%, and those incidents averaged $5.39 million (IBM, 2026). When agents and AI tools appear outside formal governance, security teams may not know who owns them, which credentials they use, what data they can reach, or what actions they can perform.

The security model for agents should therefore look familiar to PAM teams: establish identity, minimize privilege, broker credentials, control sessions, restrict actions, preserve accountability, and revoke authority quickly. The difference is that these controls must operate at machine speed and remain active across every step of the workflow.



Key Takeaways



Authorization should continue throughout the interaction, not stop at login.



Identity-based segmentation and action-level policy help limit lateral movement and high-impact operations.



Detection creates business value when it can trigger immediate, targeted containment.



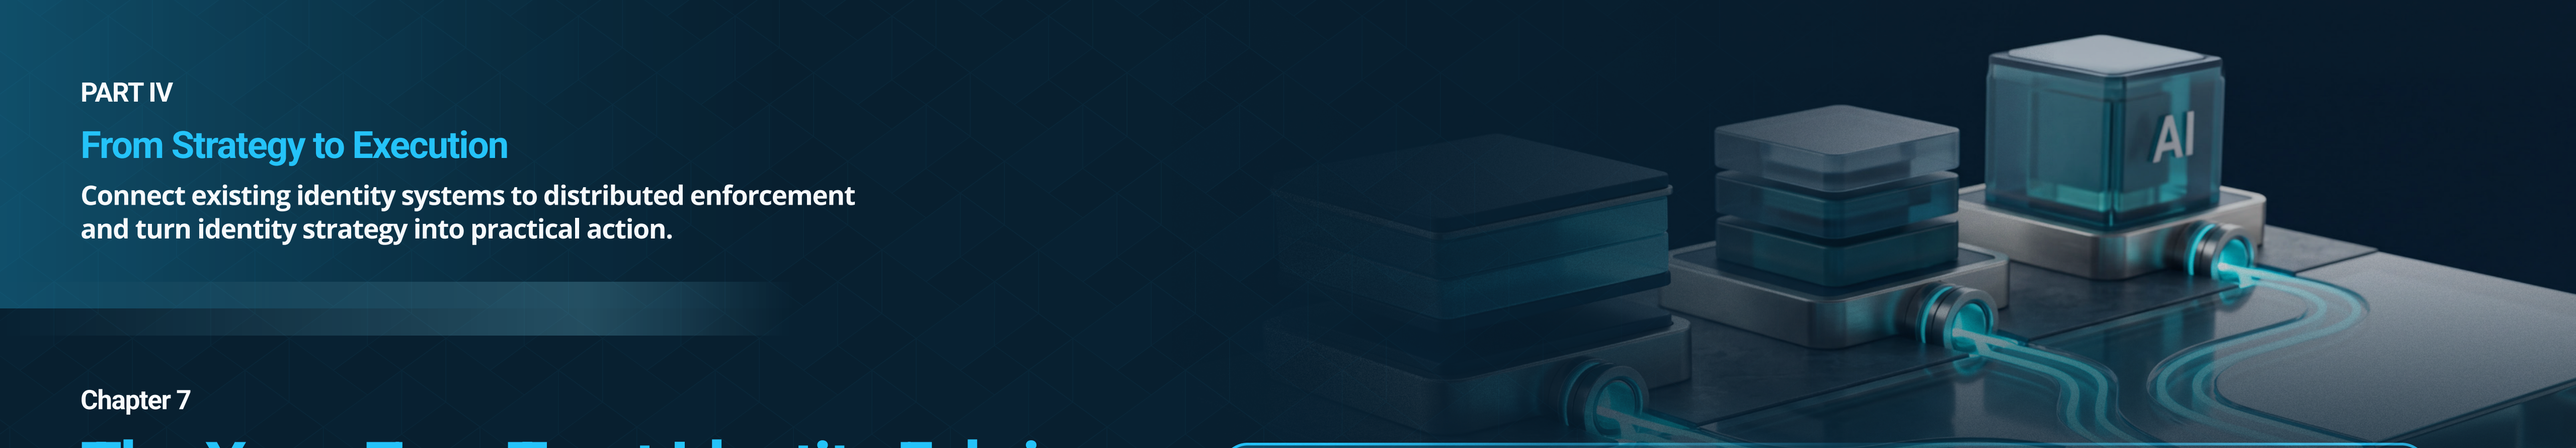
AI agents should be governed as identities, with explicit delegated authority, bounded permissions, and independent runtime enforcement.

www.xage.com | hello@xage.com | © Xage Security | All Rights Reserved

PART IV

From Strategy to Execution

Connect existing identity systems to distributed enforcement and turn identity strategy into practical action.



Chapter 7

The Xage Zero Trust Identity Fabric

How does distributed enforcement reduce PAM complexity while extending control across the enterprise?

The Xage Zero Trust Identity Fabric provides a distributed enforcement layer across IT, OT, cloud, data center, edge, and AI environments. Instead of requiring every protected interaction to depend on network location or a single centralized enforcement point, policy can be applied closer to the resource and continue operating across distributed and constrained environments.

The Fabric brings together capabilities that are often deployed separately:

- JIT/JEA and task-bound privilege to reduce standing authority.
- Credential brokering and injection so users, vendors, workloads, and agents do not need direct access to persistent secrets.
- Identity-based segmentation to restrict which identities can communicate with specific systems and services.
- Session monitoring and recording for privileged interactions.
- Action-level policy for commands, tools, APIs, parameters, and high-impact operations.
- Immediate revocation, isolation, and containment when identity risk changes.

For AI agents, the same fabric can evaluate the agent identity, delegated authority, tool, target resource, requested operation, and relevant conditions before allowing execution. That puts the final authorization decision outside the model and creates boundaries

What Xage does (and does not) own

Xage enforces runtime access, privilege, communication, and action. It integrates with authoritative identity and governance systems for identity proofing, authentication, workforce lifecycle, ownership, provisioning, and access-review decisions. The goal is a coherent control plane, not a new identity system of record.

Key Takeaways



Distributed enforcement extends privileged-access control across environments that centralized PAM can struggle to cover consistently.



Xage complements IAM and IGA by enforcing their identity and governance context at runtime.



JIT/JEA, credential brokering, segmentation, and action-level policy reduce both standing authority and blast radius.



The same enforcement model can govern human users, non-human identities, and AI agents.

Chapter 8

A 90-Day Path to Broader Identity Control

Where should security leaders start?

Start with the identities and actions that can create the most business impact. The goal is not to replace every identity system. It is to close the highest-risk gaps between identity, privilege, and enforcement.

Days	Priority	Actions
1-30	Map high-risk authority	Inventory privileged users, vendors, service accounts, cloud roles, applications, and agents. Identify owners, standing access, long-lived credentials, and critical access paths.
31-60	Reduce persistent privilege	Remove unused access. Introduce JIT/JEA for selected workflows. Broker or rotate high-risk credentials. Register production agents and review delegated permissions.
61-90	Connect risk to enforcement	Apply continuous authorization to a sensitive workflow. Add identity-based segmentation. Define action and agent-tool policies. Establish token revocation, session termination, and containment playbooks.

Measure progress with a small set of control-oriented metrics:

- JIT/JEA and task-bound privilege to reduce standing authority.
- Percentage of high-risk human and non-human identities with accountable owners
- Percentage of privileged access that remains standing versus JIT/JEA
- Percentage of sensitive sessions and actions covered by runtime policy
- Percentage of high-risk credentials brokered, rotated automatically, or issued ephemerally
- Mean time to revoke a compromised identity or terminate a risky interaction
- Percentage of production AI agents with identity, owner, authority source, tool policy, and kill switch

Maturity should be measured by how consistently the organization can control authority, not simply how many identities it can inventory or authenticate.



Conclusion

From Access Management to Continuous Control

The next stage of identity security is not a rejection of PAM. It is an expansion of what PAM must protect. Human administrators remain important, but privileged authority now belongs to vendors, applications, workloads, service accounts, machines, automations, and AI agents. The control point has moved from the account and login event to the identity, authority, interaction, and action.

That changes the security objective. Organizations need to know who or what is acting, minimize the permissions available to it, protect the credentials through which those permissions are exercised, continuously evaluate whether access should continue, restrict high-impact actions, limit lateral movement, and contain compromised authority immediately.

Xage applies this model through an Zero Trust identity fabric, a distributed enforcement layer that works with existing IAM, IGA, and identity sources while extending runtime policy across IT, OT, cloud, edge, and AI. JIT/JEA, credential brokering, identity-based segmentation, session controls, action-level authorization, and immediate containment provide a common set of controls for both traditional privileged users and the emerging population of non-human and agentic identities.



The strategic progression

Traditional PAM protects privileged access. Identity-based control makes that protection continuous and action-aware. Agent-aware control extends the same discipline through delegated and autonomous execution.

Learn more about Xage Extended Privileged Access Management (XPAM)

Explore the Product

Talk to an Expert

Sources:

IBM. 2026 Cost of a Data Breach Report.
Microsoft. 2025 Digital Defense Report.
SANS Institute. 2026 State of Identity Threats & Defenses Survey.
Verizon. 2026 Data Breach Investigations Report.

xage SECURITY

| www.xage.com
| hello@xage.com

| © Xage Security
| All Rights Reserved