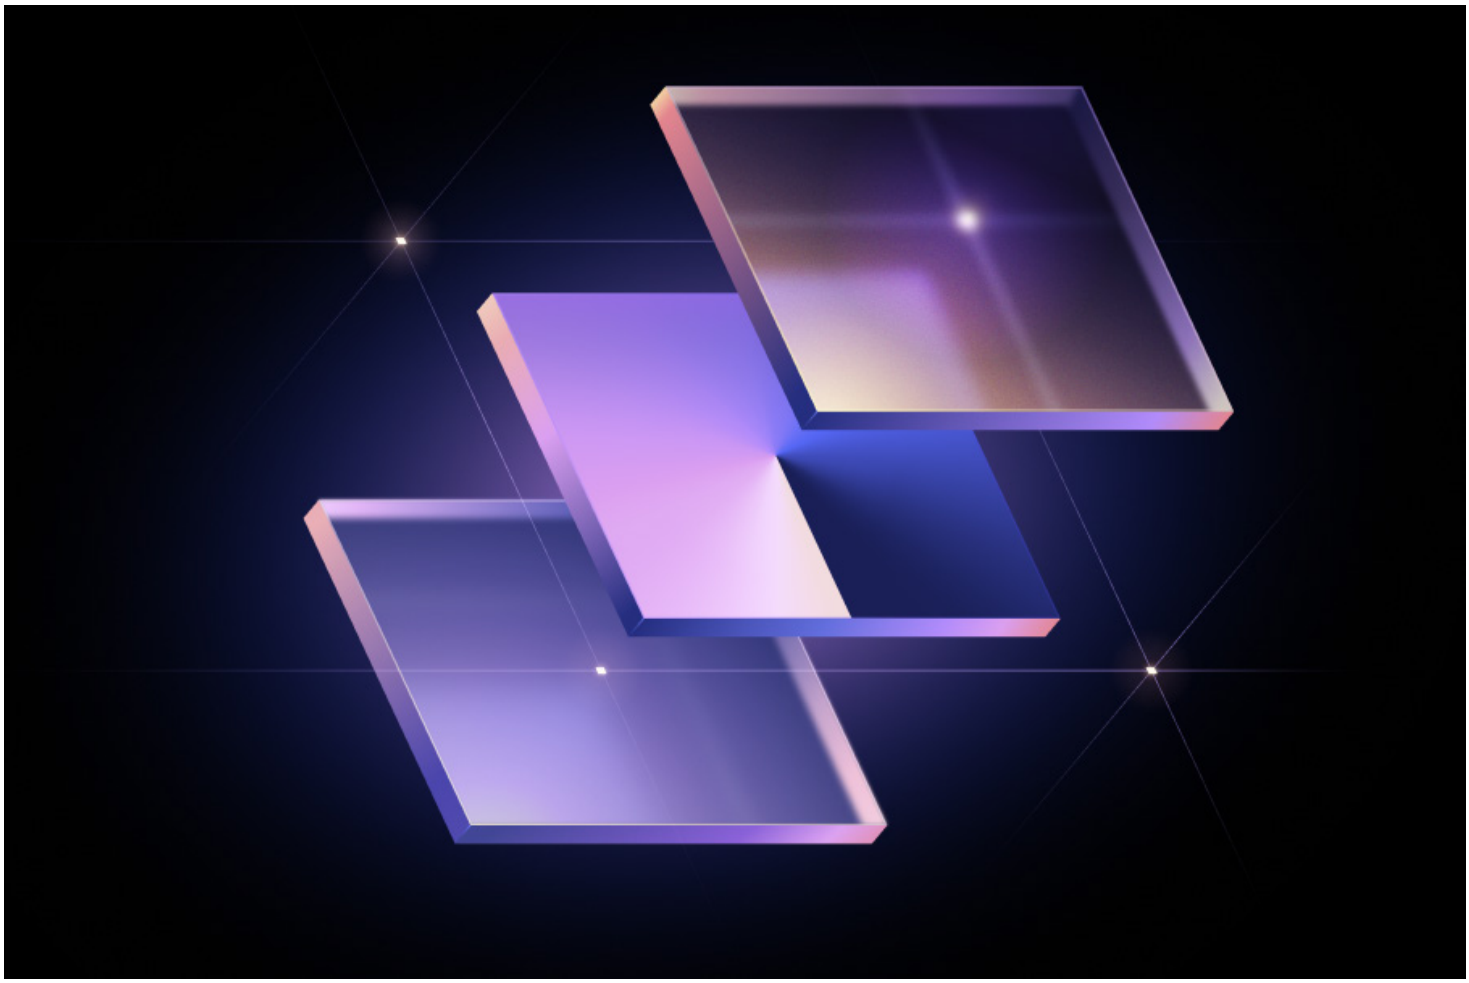

# Agentic security

Managing risk in autonomous workflows

HashiCorp
an IBM Company

# Executive summary

As enterprises adopt cloud, automation, and AI, digital identity has shifted beyond humans. Non-human identities (NHIs) such as service accounts, devices, and AI agents outnumber human users and operate at machine speed, creating new security, operational, and compliance challenges. The growth in AI agents is especially concerning to security professionals. These NHIs are ephemeral, high-privileged, fast-moving, and invisible to legacy IAM. They constitute a huge risk if left unmanaged.

This white paper discusses the approach, use cases, and design principles for agentic security and agentic identity and access management (AIAM) for human and non-human identities, secrets, certificates, and privileged access.

# What is agentic security?

Agentic security addresses the risks of AI agent use and threats to agentic applications by securing environments where humans and autonomous systems coexist. In practice, organizations achieve this by managing identities, access, and secrets for every agent.

To secure these environments, organizations need a framework that addresses the foundational controls required when agents act alongside humans inside production systems. The following five categories represent the core concerns of an agentic security strategy:

- **Non-human identity + governance:** Coordinating human and non-human identities under consistent policies and lifecycle controls without compromising velocity.
- **Certificate lifecycle management:** Issuing short-lived certificates to reduce exposure, prevent outages from expired certificates, and eliminate reliance on long-lived access.
- **Agentic secrets management:** Centrally managing the full lifecycle of secrets for autonomous and ephemeral workloads, including databases, APIs, cloud services, and AI/ML pipelines.
- **Privileged access management for autonomous systems and humans:** Enforcing just-in-time, zero trust access across ephemeral workloads.
- **Workforce user identity + governance:** Eliminates the gap between human identity governance and infrastructure access control to prevent unauthorized, over-privileged, or untraceable access.

Together, these principles form the foundation for agentic security, helping organizations confidently scale AI, automation, and cloud operations while reducing exposure, operational friction, and regulatory risk.

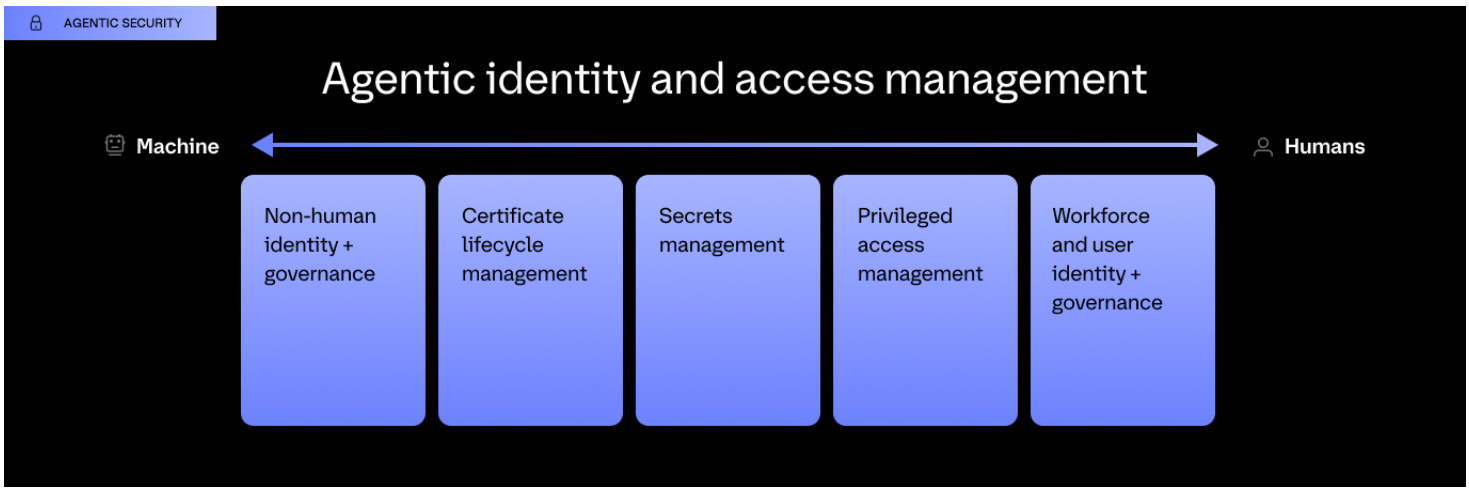

# The identity landscape has shifted

**Explosion of NHIs**
NHIs tied to cloud workloads, microservices, CI/CD pipelines, MLOps components, automation, and agentic AI now outnumber human users and drive enterprise risk. They are ephemeral, fast-moving, highly privileged, and often invisible to legacy IAM, creating a major security blind spot.

**Identity fragmentation across environments**
Hybrid and multi-cloud architectures fragment identity across cloud, on-prem, SaaS, and AI platforms, each with different authentication and authorization models. This results in inconsistent policies, duplicated controls, and limited enterprise-wide visibility, increasing operational complexity and risk.

**Compliance and governance pressures**
Regulations and industry frameworks demand clear identity provenance, least-privilege access, credential rotation, and auditable decisions. Without centralized control across humans and machines, organizations face compliance gaps, regulatory exposure, and reputational risk.

**Establishing trust for human users with IBM Verify**
Verify secures critical systems through MFA, risk-based adaptive access, and passwordless authentication. Policies dynamically adjust based on user behavior and device posture, ensuring role-appropriate access. Integrated with HashiCorp Vault, Verify enables SSO, maps enterprise roles to Vault policies, and centralizes onboarding/offboarding, automatically revoking access when users change roles. Continuous logging and audit trails provide clear visibility for CISOs and compliance teams.

**Managing workload identities with Vault**
Vault secures NHIs using workload identity federation (JWT/OIDC, cloud metadata, SPIFFE/SPIRE), eliminating static credentials and “secret zero” risks. It issues dynamic, short-lived credentials for databases, APIs, PKI certificates, and SSH, with automated rotation, revocation, and least-privilege enforcement. Policies are identity-based, enabling consistent zero trust access across hybrid and multi-cloud environments.

Vault integrates with enterprise IdPs such as Okta, Ping, Azure AD, and Google Workspace, leveraging existing SSO, MFA, and governance frameworks while simplifying onboarding, offboarding, and auditing for all identities.

# Part I: Managing NHIs and reducing risk with HashiCorp Vault and IBM Verify

As organizations expand their use of cloud services, microservices, automation frameworks, and machine learning pipelines, the identity landscape has shifted dramatically. Today, the majority of authentication events within an enterprise originate not from employees, but from NHIs that drive critical automation, yet they often operate with less governance and visibility than human users. The result is a rapidly growing attack surface and a hidden source of operational risk.

**The challenge of NHI sprawl**
Traditional identity systems were designed to manage human users, not ephemeral services. As organizations adopt distributed architectures and agentic AI systems, NHIs proliferate across clouds, CI/CD systems, Kubernetes clusters, edge devices, and third-party SaaS platforms. Unlike human identities, NHIs often:

- Outnumber human accounts by orders of magnitude
- Are created automatically by pipelines and tools
- Live far longer than intended or have undefined expiry
- Have excessive and unbounded permissions
- Are stored insecurely in repo files, config maps, or automation scripts
- Lack clear ownership and lifecycle policies

Attackers have learned to exploit these gaps. Compromised API keys, service account tokens, and misconfigured machine identities are now involved in many major breaches, enabling lateral movement, privilege escalation, or persistent access. Managing NHI sprawl has therefore become a top priority for CISOs and zero trust architects.

**Why centralized identity and secrets management matters**
To reduce NHI-related risk, organizations need a unified way to authenticate workloads, securely distribute credentials, and govern their lifecycle. This requires two complementary capabilities:

1. A trust anchor and secrets authority that can issue, rotate, revoke, and audit machine credentials
2. An identity governance layer that ensures every workload identity is verified, mapped to policy, and aligned to enterprise context (teams, applications, environments)

This is where Vault and Verify create a powerful joint solution. Together, they help organizations modernize identity, enforce least privilege, and eliminate unmanaged or long-lived NHIs across hybrid and multi-cloud environments

## How Vault reduces NHI risk

NHIs introduce risk primarily through unmanaged secrets, excessive privileges, and lack of lifecycle control. Vault directly addresses these challenges by acting as the central identity and secrets authority for workloads, applications, and autonomous systems.

**Centralized secrets and credential control:** Certificates are foundational to secure digital operations, but fragmented and manual certificate management continues to drive outages, security risk, and compliance gaps. As enterprises adopt agentic AI and autonomous workloads, Vault and IBM Guardium Cryptography Manager enable secure scale by automating certificate lifecycles and providing short-lived, identity-based authentication for NHIs.

**Identity-based access for workloads:** Vault authenticates workloads using trusted identity signals such as cloud-native identities, Kubernetes tokens, enterprise IdP assertions, and SPIFFE/SPIRE. By federating with existing identity systems, Vault eliminates "secret zero" risk and enables secure, portable identity for ephemeral and autonomous workloads.

**Governed lifecycle for NHI credentials:** ault enforces strong lifecycle controls for NHIs through short-lived credentials, automated rotation and revocation, and fine-grained, identity-based policies. This approach applies the same rigor to non-human access as human access, reducing blast radius and limiting the impact of compromise.

## How Verify strengthens governance and identity assurance

While Vault secures and issues machine credentials, Verify provides the governance, policy context, and visibility required to manage identities at enterprise scale.

**Unified governance for human and non-human identities**
Verify extends consistent governance across humans, service accounts, AI agents, ML pipelines, APIs, and edge workloads, applying uniform standards for access, lifecycle management, and approvals.

It enforces ownership, purpose, and scope for every NHI, automates provisioning/deprovisioning, and aligns with risk and policy requirements — closing common blind spots in machine identity management.

## A combined solution: Vault + Verify

Together, Vault and Verify provide a modern, integrated solution for NHI security:

- **Strong identity + strong secrets management:** Vault creates and governs short-lived machine credentials. Verify ensures those workloads are validated, known, and authorized.
- **Enterprise governance + zero trust enforcement:** Verify governs identity lifecycles and maps them to frameworks like NIST, CIS, or internal policies. Vault enforces granular, identity-based authorization for all secrets and data access.
- **Consistency across hybrid and multi-cloud:** Both platforms eliminate silos across AWS, Azure, GCP, on-prem, VMware, Kubernetes, and SaaS.
- **Reduced risk from AI and automation:** The combined stack provides workload identity for agentic AI processes, ML pipelines, autonomous systems, and code-driven infrastructure, ensuring their credentials are short-lived, verified, and auditable.

As NHIs continue to multiply, organizations need a modern approach to authentication, authorization, and secrets lifecycle management. Vault and Verify provide a powerful foundation for establishing trust, governing identity, and reducing the risk associated with NHI's across hybrid environments.

Together, they enable security leaders to implement zero trust for workloads, enforce consistent controls across all identity types, and protect the rapidly expanding digital ecosystem in which automation and AI now play a central role.

# Part II: Modernizing certificate lifecycle management with HashiCorp Vault and IBM Guardium Cryptography Manager

Certificates are foundational to securing digital communications and modern applications, yet many enterprises still manage them through fragmented, manual processes that increase the risk of outages, breaches, and compliance failures. Expired certificates, hardcoded keys, and unmanaged certificate authorities remain common causes of operational disruption.

As organizations adopt agentic AI, autonomous workloads, and machine-to-machine automation, certificate lifecycle management becomes a strategic enabler of secure agentic identity access. Short-lived, dynamically issued certificates allow NHIs to authenticate safely without relying on long-lived secrets or manual distribution.

## The challenges of traditional certificate management

Modern enterprises face four main challenges that directly impact AIAM initiatives:

1. Fragmented and manual processes
2. Poorly automated certificate lifecycle management
3. Security and compliance risk
4. Operational complexity in hybrid and multi-cloud environments

Without modern certificate lifecycle management, agentic identities and autonomous workloads are forced to rely on long-lived or manually provisioned credentials, increasing both operational risk and exposure to cyber threats.

## How Vault modernizes certificate lifecycle management

Traditional certificate management is often built around manual requests, ticketing systems, and periodic renewal cycles that cannot scale with cloud-native infrastructure or autonomous workloads. Security teams act as gatekeepers for issuance and renewal, creating operational bottlenecks while increasing the risk of expired or unmanaged certificates.

Vault modernizes this model by transforming certificate lifecycle management from a manual, human-driven workflow into a fully automated, API-driven control plane embedded directly into infrastructure and application pipelines.

Instead of relying on tickets and static provisioning, certificates are issued, renewed, and revoked programmatically at machine speed, aligned to workload identity and policy. This modernized approach includes:

**Dynamic, short-lived certificates:** Certificates are generated on demand and automatically expire, reducing exposure and eliminating reliance on long-lived credentials.

**Automated renewal and revocation:** Rotation is continuous and policy-driven rather than calendar-based or manually triggered.

**Workload identity federation (WIF):** Certificates are tied to trusted workload identity signals, eliminating bootstrap credential risk and ensuring identity-bound authentication.

**Centralized audit and compliance logging:** Every issuance, renewal, and revocation event is logged and traceable, supporting governance and regulatory requirements.

**Policy-driven access controls:** Certificate issuance and usage are enforced through fine-grained policies aligned to enterprise security standards.

By shifting from manual ticketing processes to automated, identity-driven certificate workflows, Vault transforms certificate management into a foundational enabler of secure agentic identity access. Autonomous systems can authenticate dynamically and operate reliably, without operational bottlenecks or unmanaged cryptographic risk.

## Guardium Cryptography Manager strengthens oversight and control

While Vault modernizes certificate issuance and automation, Guardium Cryptography Manager provides the enterprise-wide visibility, governance, and cryptographic control required to operate certificate lifecycle management at scale.

As certificate issuance becomes automated and API-driven, organizations face a new challenge. They must maintain centralized oversight of what has been issued, where it is deployed, who owns it, and whether it complies with enterprise cryptographic standards. Without this visibility, automation can increase speed but reduce control.

Guardium Cryptography Manager addresses this challenge by providing a centralized view and governance layer across certificate authorities, keys, and cryptographic assets.

Key capabilities include:

**Centralized governance and visibility:** A unified dashboard provides real-time insight into issued certificates, active and inactive assets, upcoming expirations, and ownership context. Security teams gain clear visibility into certificate inventory across hybrid and multi-cloud environments.

**Lifecycle monitoring and proactive risk reduction:** Guardium Cryptography Manager identifies certificates nearing expiration, weak cryptographic configurations, and unmanaged or shadow certificate authorities. This reduces outage risk and prevents compliance gaps before they become incidents.

**Hardware-backed key protection:** Integration with hardware security modules (HSMs) and hardened key storage protects root and intermediate certificate authorities, ensuring a strong root of trust and reducing the risk of key compromise.

**Policy enforcement and compliance reporting:** Cryptographic policies can be centrally defined and enforced across environments. Built-in reporting supports audit readiness and alignment with regulatory frameworks and internal standards.

**User-centric certificate lifecycle management experience:** Security and operations teams can manage certificate workflows through a clear interface that balances automation with oversight, reducing reliance on manual spreadsheets or disconnected tracking systems.

By combining automation with centralized governance and real-time visibility, Guardium Cryptography Manager ensures that certificate lifecycle management does not become a blind spot in an automated environment. It provides the oversight required to safely scale agentic identity initiatives while reducing outage risk, strengthening compliance posture, and protecting critical cryptographic assets.

## The combined value of Vault + Guardium for AIAM

When deployed together, Vault and Guardium provide a comprehensive solution for agentic identity access management:

1. End-to-end lifecycle management
2. Reduced risk
3. Operational efficiency
4. Compliance and audit readiness

Vault and Guardium Cryptography Manager transform certificate lifecycle management (CLM) from a compliance requirement into a strategic enabler of secure, scalable AI and autonomous operations, empowering executives to support innovation while reducing enterprise risk.

Vault and Guardium Cryptography Manager together provide a policy-driven, automated, and auditable approach to CLM to enable:

- Secure operation of agentic identities and ephemeral workloads
- Reduction of cyber risk and operational outages
- Streamlined compliance and audit readiness
- Operational efficiency that supports innovation

By modernizing certificate lifecycle management, enterprises can safely scale AI-driven and automated processes, ensuring that agentic identities operate securely and without disruption.

# Part III: Modern secrets management for cloud, agentic AI, and MLOps

As enterprises accelerate cloud adoption, automation, and agentic AI initiatives, secrets management has become one of the most critical foundations of security and operational resilience.

For agentic AI systems, autonomous workloads, and machine learning pipelines, this risk is amplified. Secrets are used not by humans, but by ephemeral, automated identities that operate at machine speed across distributed environments. Without a centralized, automated approach, organizations face increased exposure to breaches, misconfigurations, and compliance failures.

Modern enterprises require secrets management that is designed for non-human identity scale, continuous automation, and zero trust enforcement.

## The business risk of unmanaged secrets

Unmanaged secrets increase the risk of credential leakage, excessive privileges, compliance gaps, and operational disruption. As non-human identities scale across cloud services, CI/CD systems, Kubernetes clusters, and AI pipelines, the absence of centralized secrets governance significantly amplifies enterprise risk.

Common failure patterns include:

- Hardcoded credentials embedded in code repositories and configuration files
- Long-lived API keys and service account tokens that are rarely rotated
- Overprivileged machine credentials with unclear ownership
- Limited visibility into where secrets exist and how they are being used
- Manual distribution processes that cannot keep pace with automation

In agentic environments, these weaknesses become more dangerous because autonomous workflows can propagate compromised credentials rapidly across systems.

## Centralized secrets management as a foundation for agentic security

A modern secrets strategy requires a centralized platform that can secure, generate, distribute, and govern credentials for both human and non-human identities.

In an ideal model, organizations implement secrets management that provides:

**Dynamic, short-lived certificates:** Access to databases, cloud services, APIs, and infrastructure is granted through just-in-time credentials that automatically expire, reducing blast radius.

**Automated rotation and revocation:** Secrets lifecycle processes are continuous and policy-driven, eliminating reliance on manual updates or periodic audits.

**Identity-based access controls:** Secret access is governed by workload identity and enterprise context, ensuring least privilege enforcement across environments.

**Workload identity federation:** Secrets management platforms integrate with trusted workload identity signals such as Kubernetes service accounts, cloud-native identity providers, OIDC/JWT assertions, or SPIFFE/SPIRE. This eliminates "secret zero" risk and enables secure authentication without static bootstrap credentials.

**Centralized auditing and compliance visibility:** Every request, issuance, and access event is logged and traceable, supporting regulatory requirements and continuous risk monitoring.

This approach transforms secrets from unmanaged liabilities into governed, automated security controls that scale with cloud and AI adoption.

## Secrets management for agentic AI and MLOps pipelines

Agentic AI systems and MLOps workflows introduce new challenges for secrets management because pipelines operate continuously across distributed infrastructure. Model training, inference services, orchestration frameworks, and autonomous agents all require secure access to sensitive resources such as:

- Feature stores and training datasets
- Model registries and deployment systems
- Cloud infrastructure and GPU environments
- Third-party APIs and enterprise applications
- Databases and messaging systems

Without centralized secret governance, these pipelines often rely on embedded credentials or overly broad permissions that create significant security risk.

A modern secrets foundation enables AI agents and ML pipelines to operate with just-in-time, least-privilege credentials that are dynamically issued, automatically rotated, and tightly scoped to workload identity.

This ensures that autonomous systems can move quickly without bypassing security controls.

## Measurable outcomes for security and business leaders

By adopting modern secrets lifecycle management, organizations achieve measurable improvements in both security posture and operational efficiency:

- Reduced breach risk through short-lived, identity-bound credentials
- Stronger compliance through centralized auditing and lifecycle governance
- Improved operational resilience through automation and reduced credential outages
- Scalable security foundations for AI, cloud, and DevOps workloads

In the autonomous enterprise, secrets management is no longer a tactical security tool. It is a strategic control plane that determines whether automation and AI can scale safely.

Modern, identity-driven secrets management provides the trust foundation required for agentic systems to operate securely, continuously, and with confidence.

# Part IV: Secure remote access with HashiCorp Boundary

Privileged access remains one of the most significant attack surfaces in modern enterprises, spanning administrators, operators, and automated workflows that require elevated access to critical systems. Traditional privileged access management approaches, which depend on static credentials, VPNs, and network-based trust, are increasingly ineffective in hybrid, cloud, and ephemeral environments.

Boundary delivers a modern, identity-based approach to privileged access by enforcing authentication and authorization through trusted identity providers rather than network location or long-lived credentials. This model applies consistently to both human administrators and agentic workflows, ensuring that access decisions are based on verified identity and policy rather than implicit trust within a privileged network.

By providing just-in-time, least-privilege access to infrastructure and applications across cloud, on-premises, and hybrid environments, Boundary minimizes credential exposure and reduces the risk of autonomous systems operating with excessive or persistent privileges. Human users and AI-driven automation receive scoped, time-bound access aligned to enterprise policy, with full session visibility and auditability.

This approach enables zero trust access for both interactive and machine-driven operations, reducing the risk of agents operating unchecked inside sensitive environments while maintaining operational efficiency.

## Modern PAM for AI, automation, and MLOps

Privileged access risks are amplified in automated and AI workflows. Boundary delivers secure, just-in-time access using identities, enabling CI/CD, MLOps, and agentic automation to operate on behalf of authorized users while enforcing least privileged principles and standard access.

Session logging, identity verification, and automated revocation ensure audit readiness and regulatory compliance while enabling operational efficiency, reduced cyber risk, and scalable zero trust enforcement. Traditional PAM approaches can no longer meet the demands of cloud-native, ephemeral, and automated environments.

# Part V: Workforce identity and governance

As organizations scale AI, automation, and cloud-native infrastructure, human identity remains the root of accountability. Autonomous systems may execute actions at machine speed, but those systems are ultimately authorized, configured, and governed by people.

A secure autonomous enterprise therefore requires strong workforce identity assurance combined with infrastructure-level enforcement. Identity governance cannot stop at login. It must extend into the systems, secrets, certificates, and privileged environments where critical operations occur.

## Establishing trust in human identity with IBM Verify

Verify provides the foundation for workforce identity assurance and lifecycle governance. Modern enterprises require more than username and password authentication. They require:

- Multi-factor and passwordless authentication to reduce credential-based attacks
- Risk-based adaptive access that evaluates device posture, behavior, and context
- Centralized onboarding and offboarding processes tied to enterprise roles
- Automated deprovisioning when users change roles or leave the organization
- Continuous logging and audit trails for compliance and regulatory reporting

By governing identity at the user level, organizations ensure that access decisions are contextual, traceable, and aligned to enterprise policy.

## Extending workforce governance into infrastructure and secrets with HashiCorp Vault

Authentication alone does not secure the enterprise. Once a user is verified, their access to infrastructure, secrets, and production systems must be enforced with least privilege and lifecycle controls. Vault extends workforce governance into the infrastructure layer by:

- Mapping enterprise roles and identity claims to fine-grained access policies
- Issuing dynamic, short-lived credentials tied to a verified identity

– Enforcing identity-based authorization across databases, APIs, PKI, and SSH
– Providing centralized audit visibility into credential issuance and access events

This ensures that human access to sensitive systems is not based on static credentials or implicit trust, but on time-bound, policy-driven authorization.

## A unified identity fabric for humans and machines

Together, Verify and Vault create a unified identity fabric:

– Verify establishes who the user is and governs their lifecycle
– Vault enforces what that user, or the systems acting on their behalf, can access
– Privileged sessions are controlled through just-in-time, identity-based access
– All identity activity remains auditable across hybrid and multi-cloud environments

This integrated model ensures that human identity governance extends seamlessly into non-human identity control. AI agents, automation pipelines, and workloads operate only within the boundaries defined by verified human authority and enterprise policy.

In the autonomous enterprise, identity is not a front-door control. It is the continuous control plane that governs both human decisions and machine execution. By unifying workforce identity assurance with workload-level enforcement, organizations can enable innovation while maintaining security, compliance, and operational resilience.

## Maturity roadmap

The following table illustrates how Vault, HashiCorp Boundary, Verify, and IBM Guardium Cryptography Manager work together to reduce risk and increase compliance readiness and automation levels.

| Maturity roadmap | **Foundational**<br>Ad hoc & reactive | **Standardize**<br>Proactive & policy-driven | **Advanced**<br>Identity & integrity-aware | **Optimized**<br>Adaptive & flexible |
|---|---|---|---|---|
| | **STAGE 1:** Organizations at this stage rely on manual processes and fragmented systems, creating security gaps and operational inefficiencies. | **STAGE 2:** The adoption of centralized platforms and policy-based controls reduces security exposure and establish governance frameworks. | **STAGE 3:** Platform integration creates an identity fabric that adapts to changing threats. Every access request is verified based on real-time context. | **STAGE 4:** Automation handles routine operations to achieve continuous compliance and proactive risk mitigation. |
| **HashiCorp Vault** | - Credentials stored insecurely<br>- Certificates tracked manually<br>- Fragmented identity practices<br>- Shared or privileged access accounts | - Vaulted secrets<br>- Automated certificate management for select workloads | - Workload identity federation<br>- Dynamic secrets<br>- Automated certificate rotation<br>- Policy enforcement for NHIs | - Fully automated dynamic secrets and credential rotation<br>- Federated workload authentication across multi-cloud and AI/automation pipelines |
| **HashiCorp Boundary** | - Shared admin and root accounts<br>- Uncontrolled root access<br>- Limited sessions visibility | | - Just-in-time privileged access<br>- Session logging<br>- Zero trust enforcement | - Adaptive, policy-driven privileged access for humans and machines |
| **IBM Verify** | - Limited policy enforcement<br>- Poor visibility for compliance<br>- Manual on/offboarding | - MFA<br>- Adaptive access for humans<br>- Centralized onboarding/offboarding | - Lifecycle governance for humans<br>- Context-aware authentication<br>- Federated identity with Vault | - Continuous lifecycle management<br>- Automated provisioning/deprovisioning<br>- Continuous certificate lifecycle management for ephemeral and agentic workloads |
| **IBM Guardium** | - Hard-coded or shared certificates<br>- Fragmented CAs<br>- Manual renewal and deployment<br>- Limited compliance and audit visibility | | | - Continuous certificate lifecycle management for ephemeral and agentic workloads |
| **Capabilities** | - Static credentials<br>- Manual secrets<br>- Shared privileged accounts<br>- Ad hoc certificates | - Centralized IAM<br>- Role-based access<br>- Policy-driven secret and certificate management | - Federated identities<br>- Short-lived credentials<br>- Zero trust access & automated PKI<br>- Auditable access for humans and NHIs | - Adaptive policy-driven privileged access for humans and machines |
| **Risks & outcomes** | - High risk<br>- Limited visibility<br>- Compliance gaps | - Reduced exposure<br>- Improved visibility<br>- Policy enforcement begins<br>- Partial regulatory alignment | - Enterprise-wide zero trust enforcement<br>- Reduced risk from AI/automation<br>- Consistent governance | - Minimal exposure<br>- Scalable automation<br>- Enterprise-wide audit readiness<br>- Secure AI and automated workflows |

# Maturity comparison: Risk and capabilities

As organizations progress through maturity stages, security risk decreases dramatically while automation and compliance capabilities increase exponentially. According to the 2025 Thales Data Threat Report, organizations that consistently pass compliance audits are 57% less likely to experience a breach compared to those that fail audits. Additionally, IBM research shows that breaches contained in under 200 days result in 25.9% lower costs, a capability that mature security postures enable through faster threat detection and response.

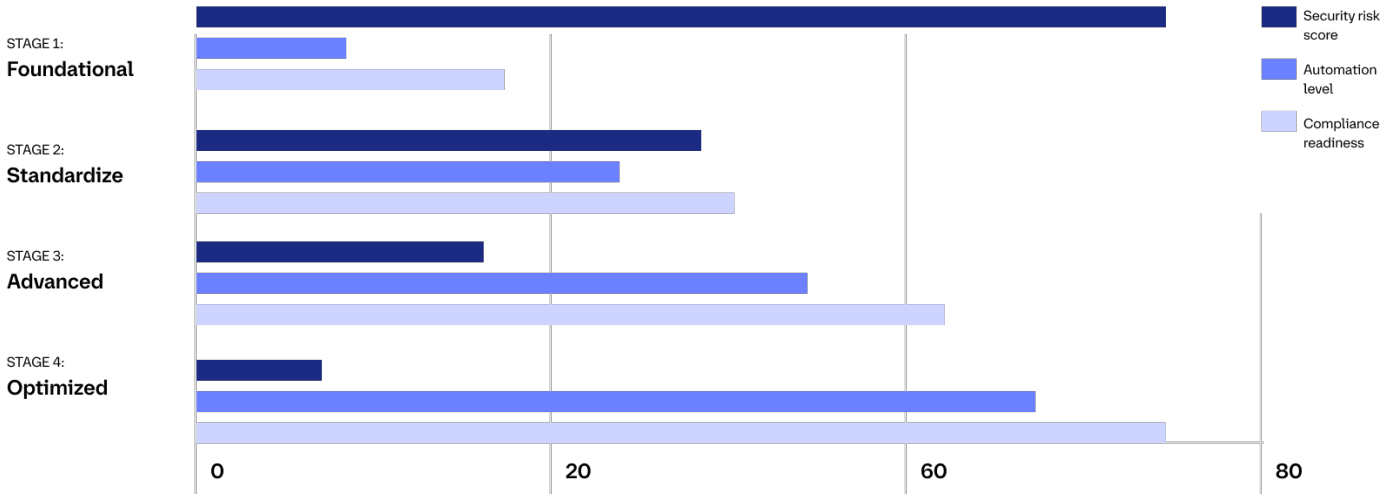

# Conclusion: Building a resilient identity foundation for the autonomous enterprise

Enterprise identity has fundamentally changed. As organizations accelerate cloud adoption, embrace agentic AI, and automate critical business processes, identity becomes the control plane that determines whether innovation happens securely or at unacceptable risk.

HashiCorp Vault, HashiCorp Boundary, and IBM Verify together provide a cohesive, zero trust-aligned identity architecture designed for the modern enterprise:

- **Verify** establishes trust, governance, and strong authentication for human users
- **Vault** secures NHIs through workload identity federation, dynamic secrets, and automated certificate lifecycle management
- **Boundary** enforces just-in-time, identity-based privileged access without exposing credentials

Together, these capabilities enable organizations to secure every identity — human and non-human — across hybrid, multi-cloud, and AI-driven environments, while reducing operational overhead and improving compliance readiness.

By modernizing identity with Vault, Boundary, and Verify, organizations can:

- **Reduce enterprise risk** by eliminating static credentials and enforcing least privilege
- **Improve resilience** by preventing outages caused by expired certificates or mismanaged access
- **Strengthen compliance posture** with centralized auditing and identity governance
- **Enable AI and automation** by providing secure, ephemeral identity to agentic systems
- **Align security with business velocity**, allowing teams to innovate without bypassing controls

## The path forward

The question for today's executives is no longer whether identity modernization is required, but how quickly it can be achieved without disrupting the business. A unified identity foundation that treats humans and non-humans as first-class citizens under a zero trust model is essential to sustaining growth, resilience, and trust in an increasingly autonomous enterprise.

By adopting a modern, integrated approach to identity, organizations position themselves to secure what matters most today, while remaining prepared for what comes next. Learn more about how HashiCorp and IBM enable agentic identity and access management.

© Copyright IBM Corporation 2026

Produced in the
United States of America
April 2026

IBM, the IBM logo, Guardium, HashiCorp, Boundary, Terraform, and Vault are trademarks or registered trademarks of International Business Machines Corporation, in the United States and/or other countries. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on ibm.com/trademark.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT.

IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

