CYBERARK®
THE IDENTITY SECURITY COMPANY®

Buyer's Guide

How to Choose an Identity Security Platform in 2026

What to consider in the new era of identity risk

Table of contents

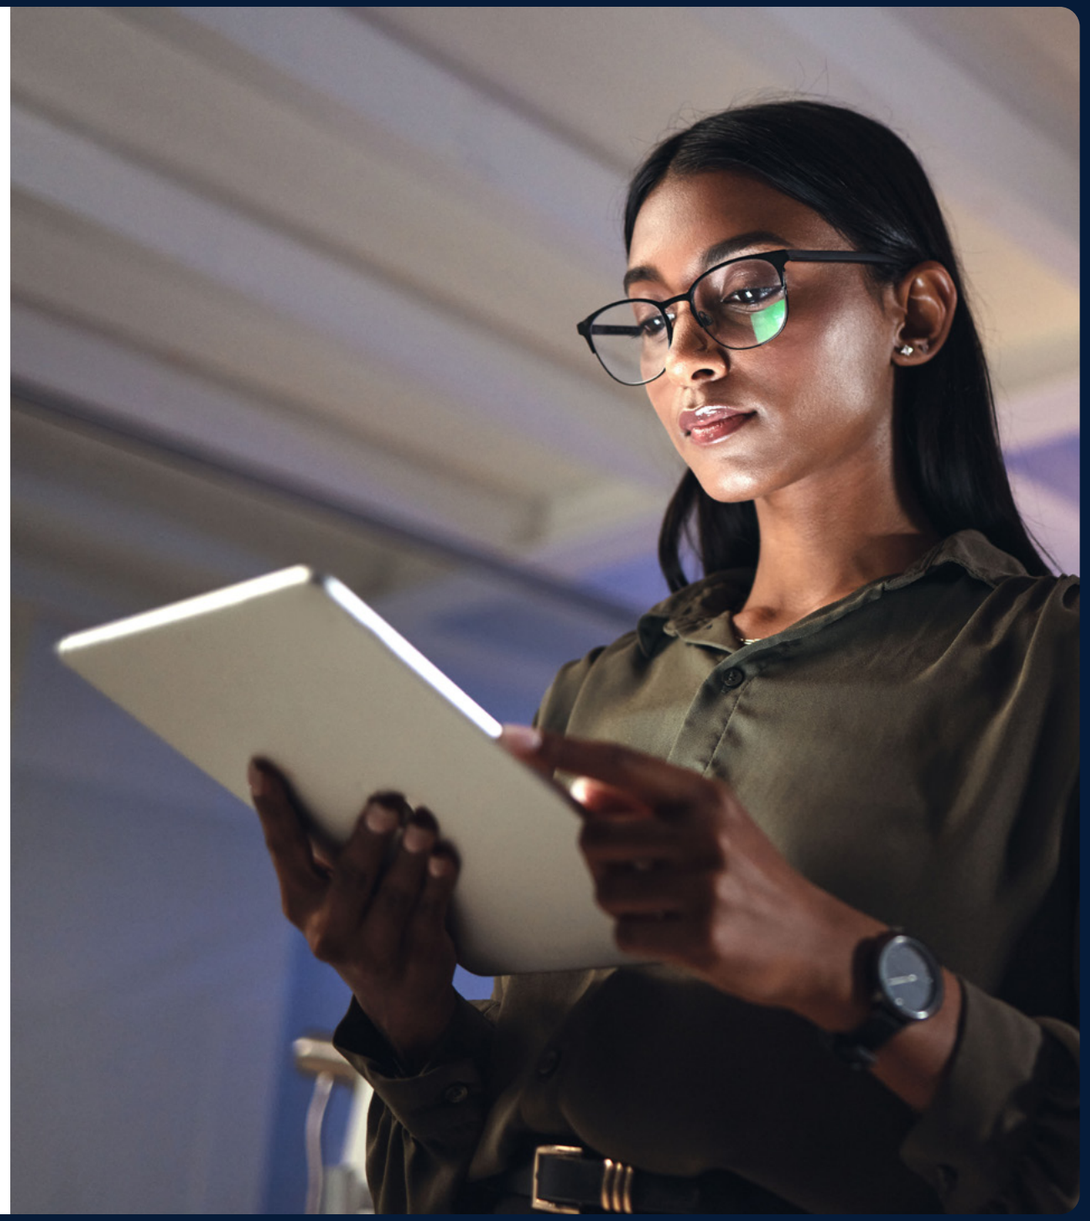

Introduction

If you're reading this, you've probably been to at least one meeting where someone asks:

"Have we locked down every privileged identity?"
It's a hard question to answer.

Privileged access now hides in plain sight: in users, in service accounts, in SaaS apps, and in browsers. The more security tools you add, the less likely they are to work together, and the harder it gets to gauge what's exposed.

Today, identity isn't just the center of cybersecurity investment. It's one of the most complex risk areas to govern, thanks in part to the sheer volume and velocity of identity growth. If you're responsible for choosing a defense, you're facing three converging challenges:

1. A rising wave of machine and AI identities that have privileged access to your sensitive resources.
2. Increasingly sophisticated AI-driven phishing and impersonation attacks.
3. A fractured solution landscape that makes it hard to see the full picture for identities.

What is Identity Security

Identity security is the cybersecurity discipline concerned with reducing all aspects of identity-related risk. This discipline requires governing, identifying, and protecting all identities used within an organization, including those used to access applications, endpoints, infrastructure, and data stores. It also covers processes for detecting and responding to identity-related threats.

Source: CyberArk. The Identity Security Imperative. 2025.

Modern attackers exploit identity seams between tools for identity and access management (IAM), privileged access management (PAM), cloud and SaaS session controls—where no single tool sees the full picture. If your security only thinks about "user accounts," you miss the hidden privileges living in packages, bots, pipelines, and developer workflows.

Case in point: the Shai-Hulud npm attack disclosed in fall 2025. Rather than "hack" code, attackers exploited trusted developer and machine identities to spread like a worm across supply chains, using stolen credentials and automation tokens to infiltrate CI/CD pipelines, cloud resources, and downstream dependencies.

This wasn't a flaw in one application. It was a collapse of identity trust, impacting human and machine identities, developer accounts, build tools, and production infrastructure. When identity-based attacks can leap from dev credentials to production systems with a single npm install, the risk surface explodes. Identity security should be designed to cover every identity—human, machine, developer—across code, cloud, and runtime environments.

In this sophisticated threat landscape, technology choices matter.

The right identity security platform can strengthen your security posture. The wrong one leaves gaps that attackers may walk straight through. Organizations need a fully integrated identity security platform that applies intelligent privilege controls, enforces zero standing privileges, and gives teams a complete picture of identity risk across users, machines, bots, and applications.

This Buyer's Guide is designed to help you cut through the noise—and choose a platform that's purpose-built for the identity threats of today and tomorrow.

The difference between a platform and a pile of products

You likely already have parts of an identity strategy in place—a PAM tool to protect admin credentials, a single sign-on (SSO) or multi-factor authentication (MFA) solution for workforce access, maybe a cloud identity store or secrets vault.

But chances are, you've also noticed the cracks.

These tools weren't designed to work together or to manage the exploding number of human, machine, and AI identities you're dealing with today. They overlap in places, and more importantly, they leave gaps.

You might find yourself wondering:

Shouldn't session monitoring apply to a cloud engineer, too?

Shouldn't bots get the same protections as users?

That instinct is right. But siloed tools:

- can't share controls
- can't share intelligence, and
- can't give you a complete picture of identity risk.

That's where most organizations feel the pain: in the seams between tools. That's often where attackers strike.

A modern identity security platform seals these gaps because it does more than integrate tools—it correlates data, applies shared controls, and delivers a unified view of identity risk. This connective tissue is what separates a collection of tools from a truly connected platform.

Fragmentation drives identity risk

70%

say lack of integration between identity and security tools weakens detection.

49%

of respondents say identity silos are a root cause of cybersecurity risk.

68%

lack complete visibility into cloud entitlements.

Source: CyberArk. 2025 Identity Security Landscape Report. April 2025.

Defining a modern identity security platform

With that foundation in mind, a modern identity security platform should be designed to:

1. **Secure human, machine, and AI identities** across physical, virtual, and cloud environments.
2. **Layer and coordinate controls** like MFA, SSO, privileged access management (PAM), and identity lifecycle management across users and machines.
3. Provide **threat detection, session monitoring, and access reviews.**
4. **Harness AI to analyze data** for proactive threat detection.
5. **Manage and secure diverse environments**—from endpoints to cloud platforms and SaaS applications—across every industry.
6. **Extend intelligent privilege controls™** to all identities across all environments.

Core capabilities of a modern identity security platform

Now that we've defined what an identity security platform should be designed to do, let's look at the core capabilities you need to ensure it can deliver.

1. **Endpoint privilege security:** Most ransomware starts here. A modern identity security platform should help you take control of unmanaged privileges on endpoints by removing local admin rights, enforcing role-specific least privilege, and improving audit readiness.
2. **Privileged access management:** PAM should be applied in a way admins won't hate—inside the tools they already use. To drive real adoption, low-friction privilege controls should fit naturally into native IT workflows, so users stay both secure and efficient.
3. **Secrets management:** Hard-coded and unmanaged secrets can be ticking time bombs. Secure and manage the credentials used by applications, workloads, and other machine identities across both enterprise and external IT environments.
4. **Cloud security:** A modern identity security platform should extend PAM to the cloud and handle complexity and multi-cloud sprawl at scale. It should continuously analyze permissions, grant just-in-time (JIT) access, and monitor actions across accounts and regions—aligned to major cloud frameworks and audit requirements.
5. **Identity management, governance, and administration:** Clean up entitlement sprawl, remove orphaned accounts, and automate reviews by handling provisioning/deprovisioning, policy-based access, approvals, certifications, and risk-based reviews—so access stays aligned with roles, compliance, and least privilege.
6. **Access controls across browsers, SaaS, and modern environments:** If work happens in the browser, so do attacks. A platform should include secure browser or browser-based session protection that can enforce policy, isolate risky actions, and record high-risk SaaS sessions without agents. It should also support zero trust access models with just-in-time privileged access across operating systems, endpoints, and cloud infrastructure.
7. **Integration with existing tools:** The right platform should connect identity signals to the rest of your SOC stack, so you don't have to fly blind during a breach. Integration between the platform and other security and IT tools improves visibility into identity-related vulnerabilities, speeds up attack detection, and expedites containment.

Why intelligent privilege controls are foundational to identity security

Privilege controls are identity security measures that adjust in real time to govern how identities access and interact with enterprise resources. They are intelligent because they modify dynamically based on factors like risk and the requester's level of authorization.

As identities, access paths, and workloads multiply, privilege is no longer a static concept. Today, any human or machine identity can become privileged the moment it touches sensitive data, a risky API, or a production system.

You can see the industry breakdown shows what larger percentages of the workforce are given access to sensitive corporate data, and why it's essential to ensure that access is secure with privilege controls.

The average percentage of workforce given access to sensitive corporate data.

37%
Education

38%
Finance

39%
Healthcare & Life sciences

42%
Government

39%
Technology

34%
Manufacturing

34%
Retail

36%
Other

Source: CyberArk. 2025 Identity Security Landscape Report. April 2025.

A modern platform should be able to dynamically apply least privilege controls based on risk, context, and behavior. Here are the essential controls you should expect.

Privilege control	What it does	Why it matters
Zero standing privileges	Grants access only when needed, then removes it automatically.	Prevents privilege creep, misconfigurations, and long-lived admin rights that attackers exploit.
Vaulting and credential management	Protects passwords, keys, tokens, and secrets.	Machines and AI identities now outnumber humans, and each carries credentials attackers can steal or abuse.
Session protection, isolation, and monitoring	Provides oversight and control of high-risk actions.	Visibility into privileged behavior is crucial for detecting misuse, meeting compliance requirements, and expediting incident response.
Endpoint identity security	Removes local admin rights and enforces role-based least privilege.	Many attacks begin with endpoint privilege escalation. Controlling workstation and server privileges cuts off one of the most common escalation paths.
Identity threat detection and response (ITDR)	Spots privilege misuse, lateral movement, and identity anomalies as they happen.	Early detection of identity threats can help reduce dwell time and prevent escalation before critical systems are exposed.

While these capabilities set the standard for protecting privilege, don't be swayed by the identity platform with the longest feature list. What matters most is whether those features work together—across identity types, environments, and control points.

In the past, organizations relied on fragmented tools: IAM for workforce access, PAM for IT admins, cloud security posture management (CSPM) for cloud permissions, and identity governance and administration (IGA) for access governance. But today, identity risk doesn't stay in one lane. Attackers move laterally between SaaS, endpoints, cloud infrastructure, and machine identities—often slipping through the seams.

MFA now reaches deep into endpoint protection. PAM needs to extend into the browser. Cloud permissions must be governed as tightly as workforce credentials. Your platform must remove the silos between domains and apply controls cohesively—wherever identity risk emerges.

Shared services that actually make a platform a platform

Ultimately, the value of a platform depends on how well it helps your team work. A modern platform should be designed to eliminate friction, not create it.

It should NOT require your team to:

- Switch between different security solutions and user interfaces.
- Cross-reference varying sources of threat intelligence data.
- Assemble reports manually during an audit.

Platform vendors love to talk about features. But shared services are the real test of whether those features actually work together.

How to pressure test a vendor's coverage

Based on your most top-of-mind need (perhaps one you presented to your board) you might go into vendor meetings asking if they have a solution for one specific purpose (for example, protecting IT admins' credentials).

But why stop there? Ask for more.

You know attackers will take advantage of any vulnerable credential: a network admin's credential leading to critical infrastructure or an HR admin's password to an app containing employee data. Make sure the solution you are evaluating applies the same tight security controls across all credentials and identities.

Ask the vendor:

- Can your platform protect every form of credential for human, machine, and AI identities?
- What types of privilege controls do you have in place for securing identities of all users?
- How do each of these controls generate measurable cyber risk reduction?
- How does your platform handle AI identities, machine identities, and ephemeral cloud identities?

Shared services for a modern identity security platform

These are the features that help ensure your team has a user experience that supports the most important job: protecting your organization.

United user interface

A “single pane of glass” gives administrators one place to manage every identity, access type, privilege level, and session.

Why this matters: A shared UI helps reduce cognitive load, speed decision-making, and ensure identity decisions are made with the full picture in view.

Centralized identity administration

A unified function for managing roles, policies, app access, and access workflows across all users and environments.

Why this matters: Teams can unintentionally create inconsistent rules across tools, resulting in overprivileged access and audit findings that are expensive to correct. A central function enforces consistency—and dramatically reduces manual work.

Identity security intelligence

Shared analytics across identity, privilege, cloud, and behavioral signals, all processed by a single threat engine.

Why this matters: Most breaches exploit what one tool misses. When signals are correlated in one place, you detect risk before escalation, not after. Correlation turns noise into early warnings.

Automated connector management

Connector deployment, configuration, and maintenance handled automatically across cloud and on-prem environments.

Why this matters: Teams should not waste time downloading agents, installing connectors, or manually configuring integrations across different systems. Automation can help eliminate brittle integrations and keep security controls consistently enforced. This gives teams back time and helps reduce operational fatigue.

Real-time intelligence that drives action

Signals alone aren’t enough. Platforms should surface risk scores, suggest next steps, and prioritize dangerous behavior automatically.

Why this matters: Security teams don’t have time to dig through logs or session recordings. AI-driven insight and prioritization accelerate investigations and help ensure the most dangerous identity behaviors don’t get buried in a long queue of alerts.

The role of AI in an identity security platform

AI can feel like a buzzword that vendors use to sell you on features that aren't fully realized. But the reality is that AI empowers attackers to analyze public data at scale, mimic tone and formatting, and adapt its messaging in real time.

Your identity platform needs AI to adapt just as quickly. AI should do more than add convenience. It should multiply the impact of your defense. AI should quickly surface risky behavior, anomalies, privilege misuse, and suspicious session activity—without forcing analysts to sift through logs. **How can an identity security platform support this?**

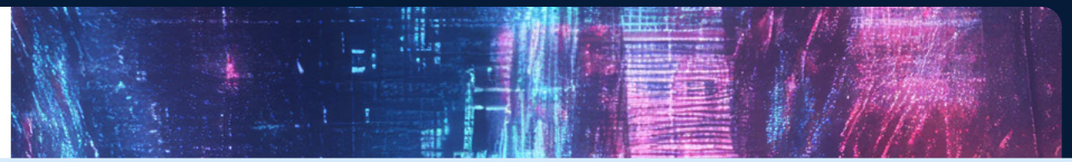

How a modern identity platform should use AI

Streamline onboarding and lifecycle decisions.

AI should identify newly discovered identities, surface pending access needs, and make policy recommendations based on similar users. Admins simply review and approve, reducing the risk of overprivileged accounts.

Accelerate threat detection and investigation.

AI should surface high-risk behaviors and generate concise session summaries, so analysts don't have to sift through logs or recordings. This shortens investigation time and improves response.

Detect machine identity anomalies.

AI should flag unusual access to secrets or token misuse—especially in automated workflows where deviations are subtle but dangerous.

Next steps for selecting an identity security platform

Evaluating identity security in 2026 is more complex than comparing feature lists. You're defending against AI-powered phishing, machine-identity sprawl, and attackers who exploit the blind spots created by fragmented tools. Leading industry guidance favors cloud-first platforms that deploy quickly, update continuously, scale across any environment, and simplify policy management for modern applications.

You've likely heard the phrase "security is a team sport." When vetting vendors, look beyond the product pages and ask about the people behind the solution.

As you move from evaluation to action, the most important question is whether your identity security strategy is helping your team stay ahead—or forcing them to react. That's where the true differences emerge. At CyberArk, we build our platform, intelligence, and expertise around our customers. We believe the right platform is one that helps teams lead with confidence instead of chasing alerts, incidents, and attackers.

Ask:

- Can you detect and respond to identity-based attacks in real time?
- Do they have an integrated team of threat researchers studying the latest attack methods?
- Do they have a red team that can evaluate and spot your vulnerabilities before attackers do?
- Do they have leadership that has been on the front lines of security and cyber warfare?
- Do these players come together, collaborate and continuously improve their offerings?

Learn more about our platform

Centered on intelligent privilege controls, the CyberArk Identity Security Platform seamlessly secures human, machine, and AI identities accessing workloads from hybrid to multi-cloud, and flexibly automates the identity lifecycle, all with continuous threat detection and prevention to enable zero trust and enforce least privilege.

Find out how CyberArk helps the world's leading organizations build resilient identity security programs designed for today's threats.

Schedule a Meeting

About CyberArk

CyberArk (NASDAQ: CYBR) is the global leader in identity security. Centered on intelligent privilege controls, CyberArk provides the most comprehensive security offering for any identity — human or machine — across business applications, distributed workforces, hybrid cloud environments and throughout the DevOps lifecycle. The world's leading organizations trust CyberArk to help secure their most critical assets. To learn more about CyberArk, visit www.cyberark.com, read the CyberArk blogs or follow on LinkedIn, X, Facebook or YouTube.