WRITTEN BY

Francis Odum

Software Analyst®
Cyber Research

REPORT

The Evolution of the Privileged Access Management (PAM) Market & The New Competitive Landscape

The Future of PAM post PANW & CyberArk; How Zero standing Privileges, Cloud, and Agentic AI Are Redefining the Privileged Identity Access Control Plane

britive

As Identity becomes central to agentic security, Privileged Access Management (PAM) is undergoing a structural shift. It is currently an underlooked category which deserves more attention as we prepare for an identity centric agentic stack.

Our opinion is that privileged access will be central to how human, agents and machines (NHIs) evolve within the future identity stack. The acquisitions we have witnessed over the past 12 months have been reinforced by decisive market activity. Palo Alto Networks' $25B acquisition of CyberArk reflects a clear recognition that identity and privilege are now foundational to platform agentic security. Palo Alto Networks had many opportunities across this ecosystem, but chose to go with the PAM route.

Subsequently, we've seen other minor acquisitions such as Okta's acquisition of Axiom. We also saw Delinea's acquisition of StrongDM two weeks ago signal a move toward just-in-time, runtime-aware access for cloud and developer environments. We've recently seen vendors such as Silverfort (PAS), 1Password vaulting and Keeper PAS continue to push for more privilege access products.

See market breakdown below:

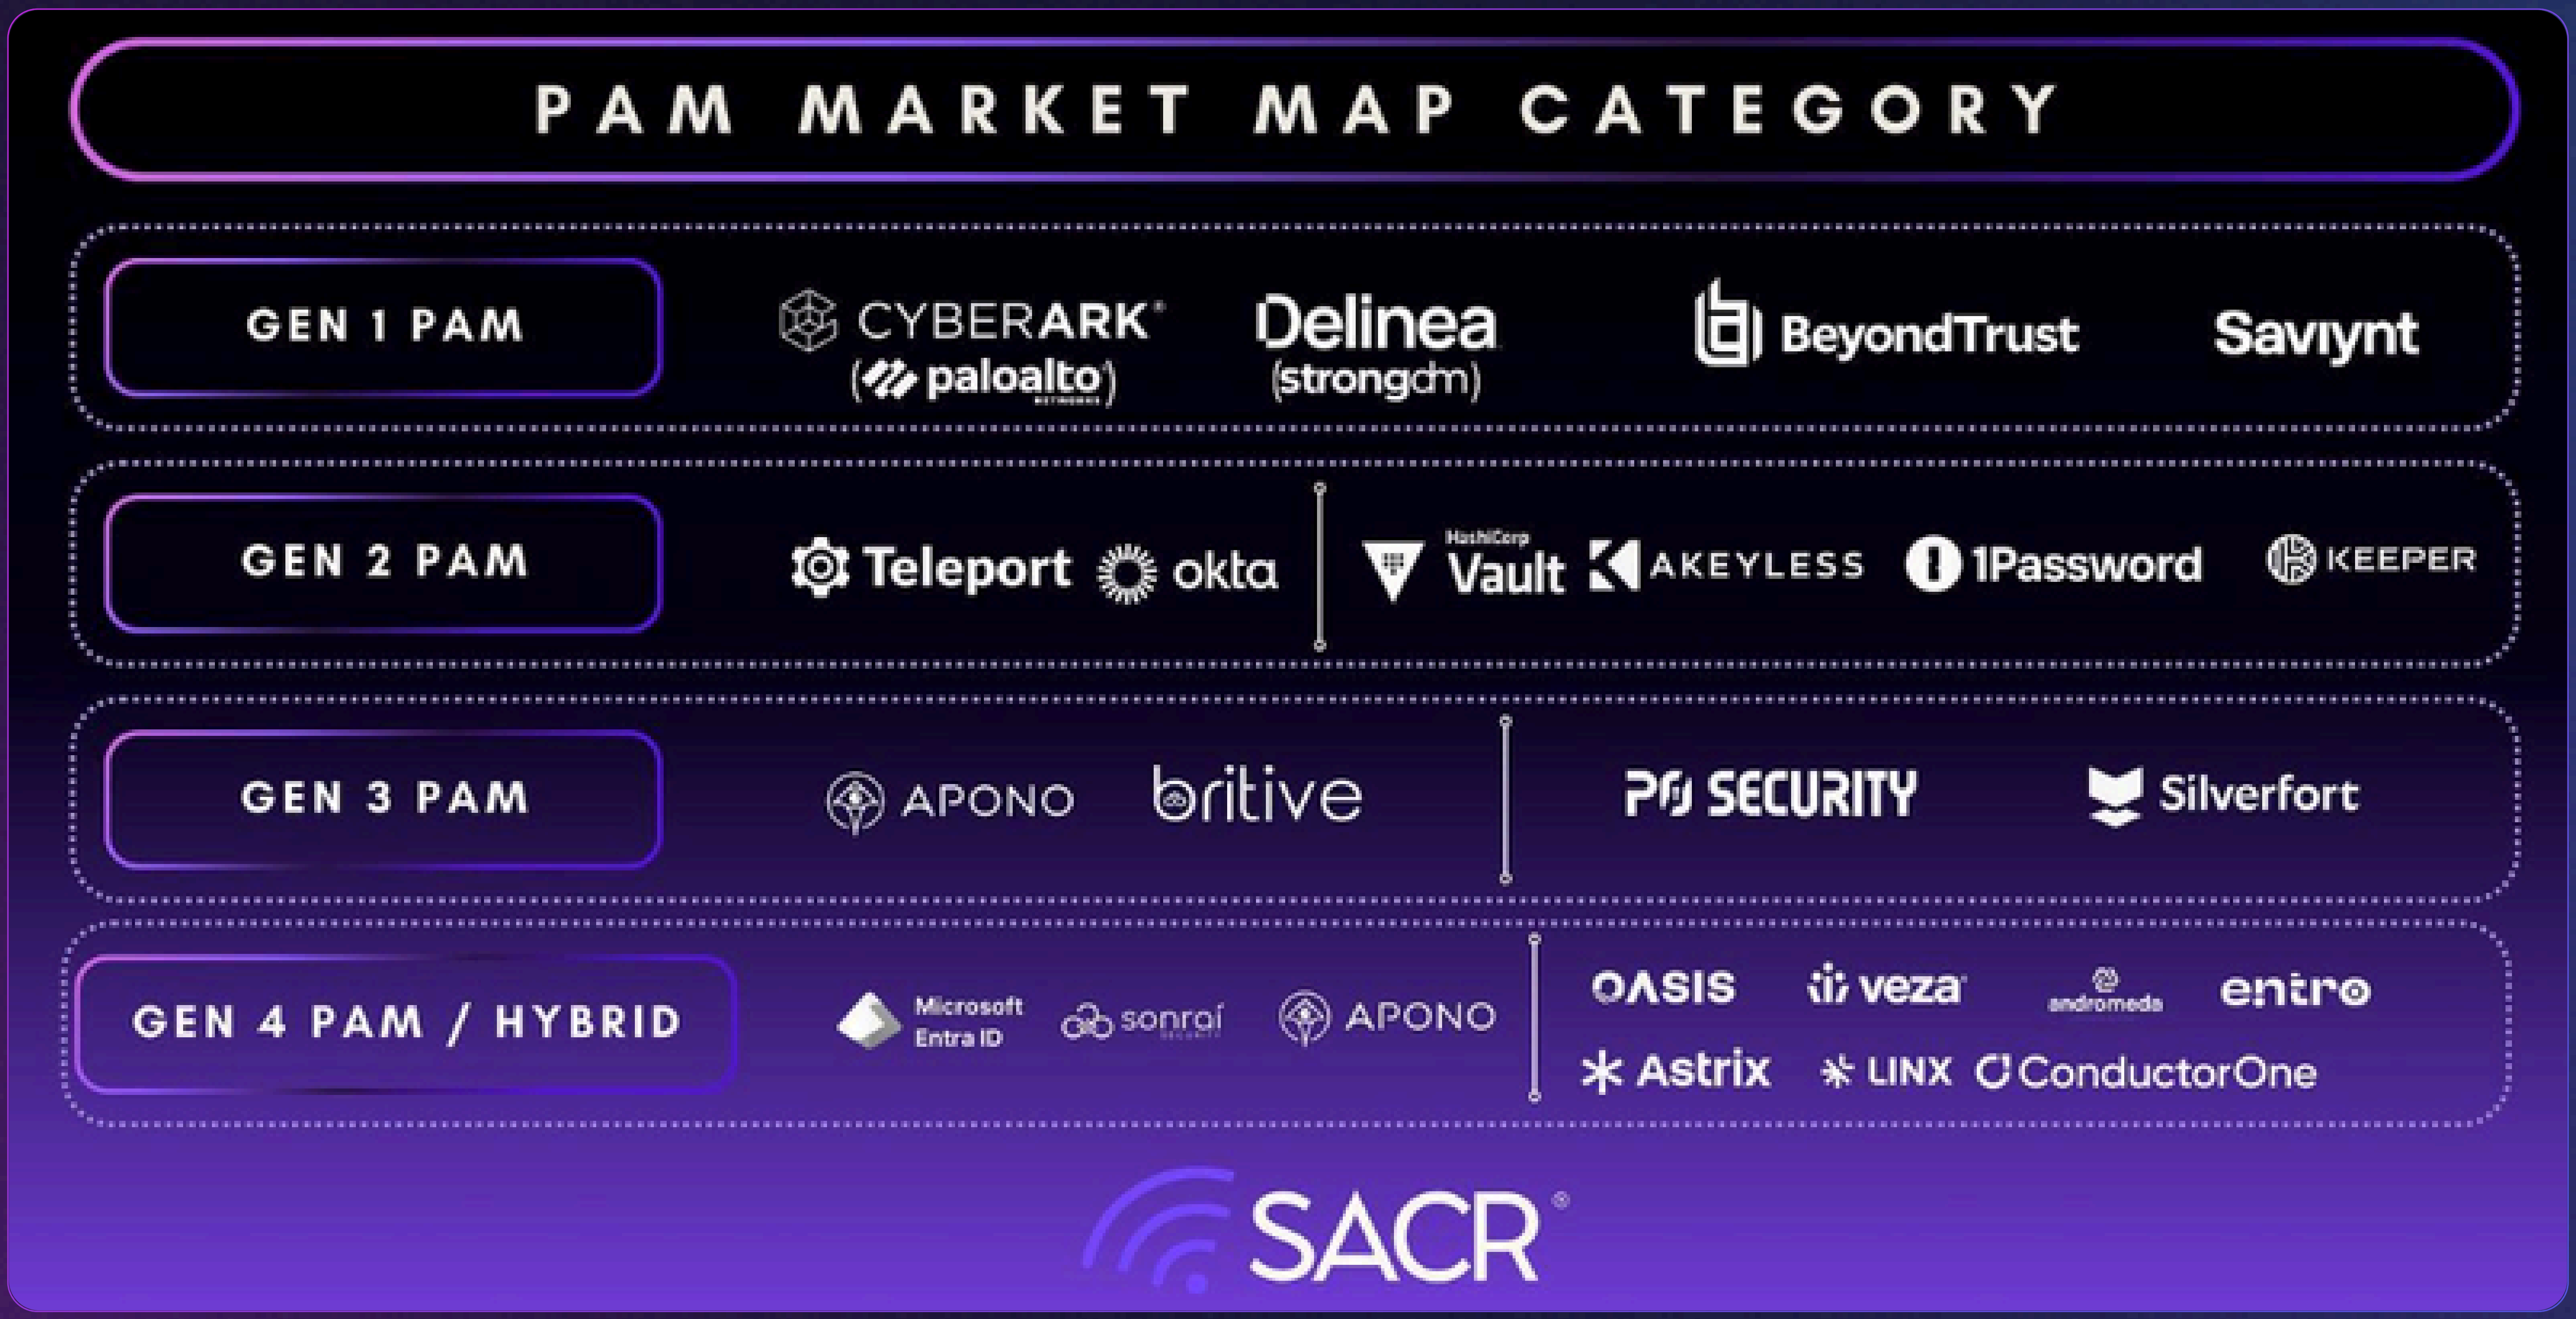

PAM MARKET MAP CATEGORY

GEN 1 PAM: CYBERARK (paloalto), Delinea (strongdm), BeyondTrust, Saviynt

GEN 2 PAM: Teleport, okta | HashiCorp Vault, AKEYLESS, 1Password, KEEPER

GEN 3 PAM: APONO, britive | P0 SECURITY, Silverfort

GEN 4 PAM / HYBRID: Microsoft Entra ID, sonrai, APONO | OASIS, veza, andromeda, entro, Astrix, LINX, ConductorOne

SACR

SACR'S VIEW OF THE MODERN IDENTITY SECURITY STACK

To better understand this framework, it's crucial to understand how SACR thinks about the identity ecosystem. In today's cloud-first and identity-driven environments, identity security has become the new perimeter. The image below outlines the core pillars of a modern Identity Security Framework, illustrating how organizations must govern access across both human and machine users.

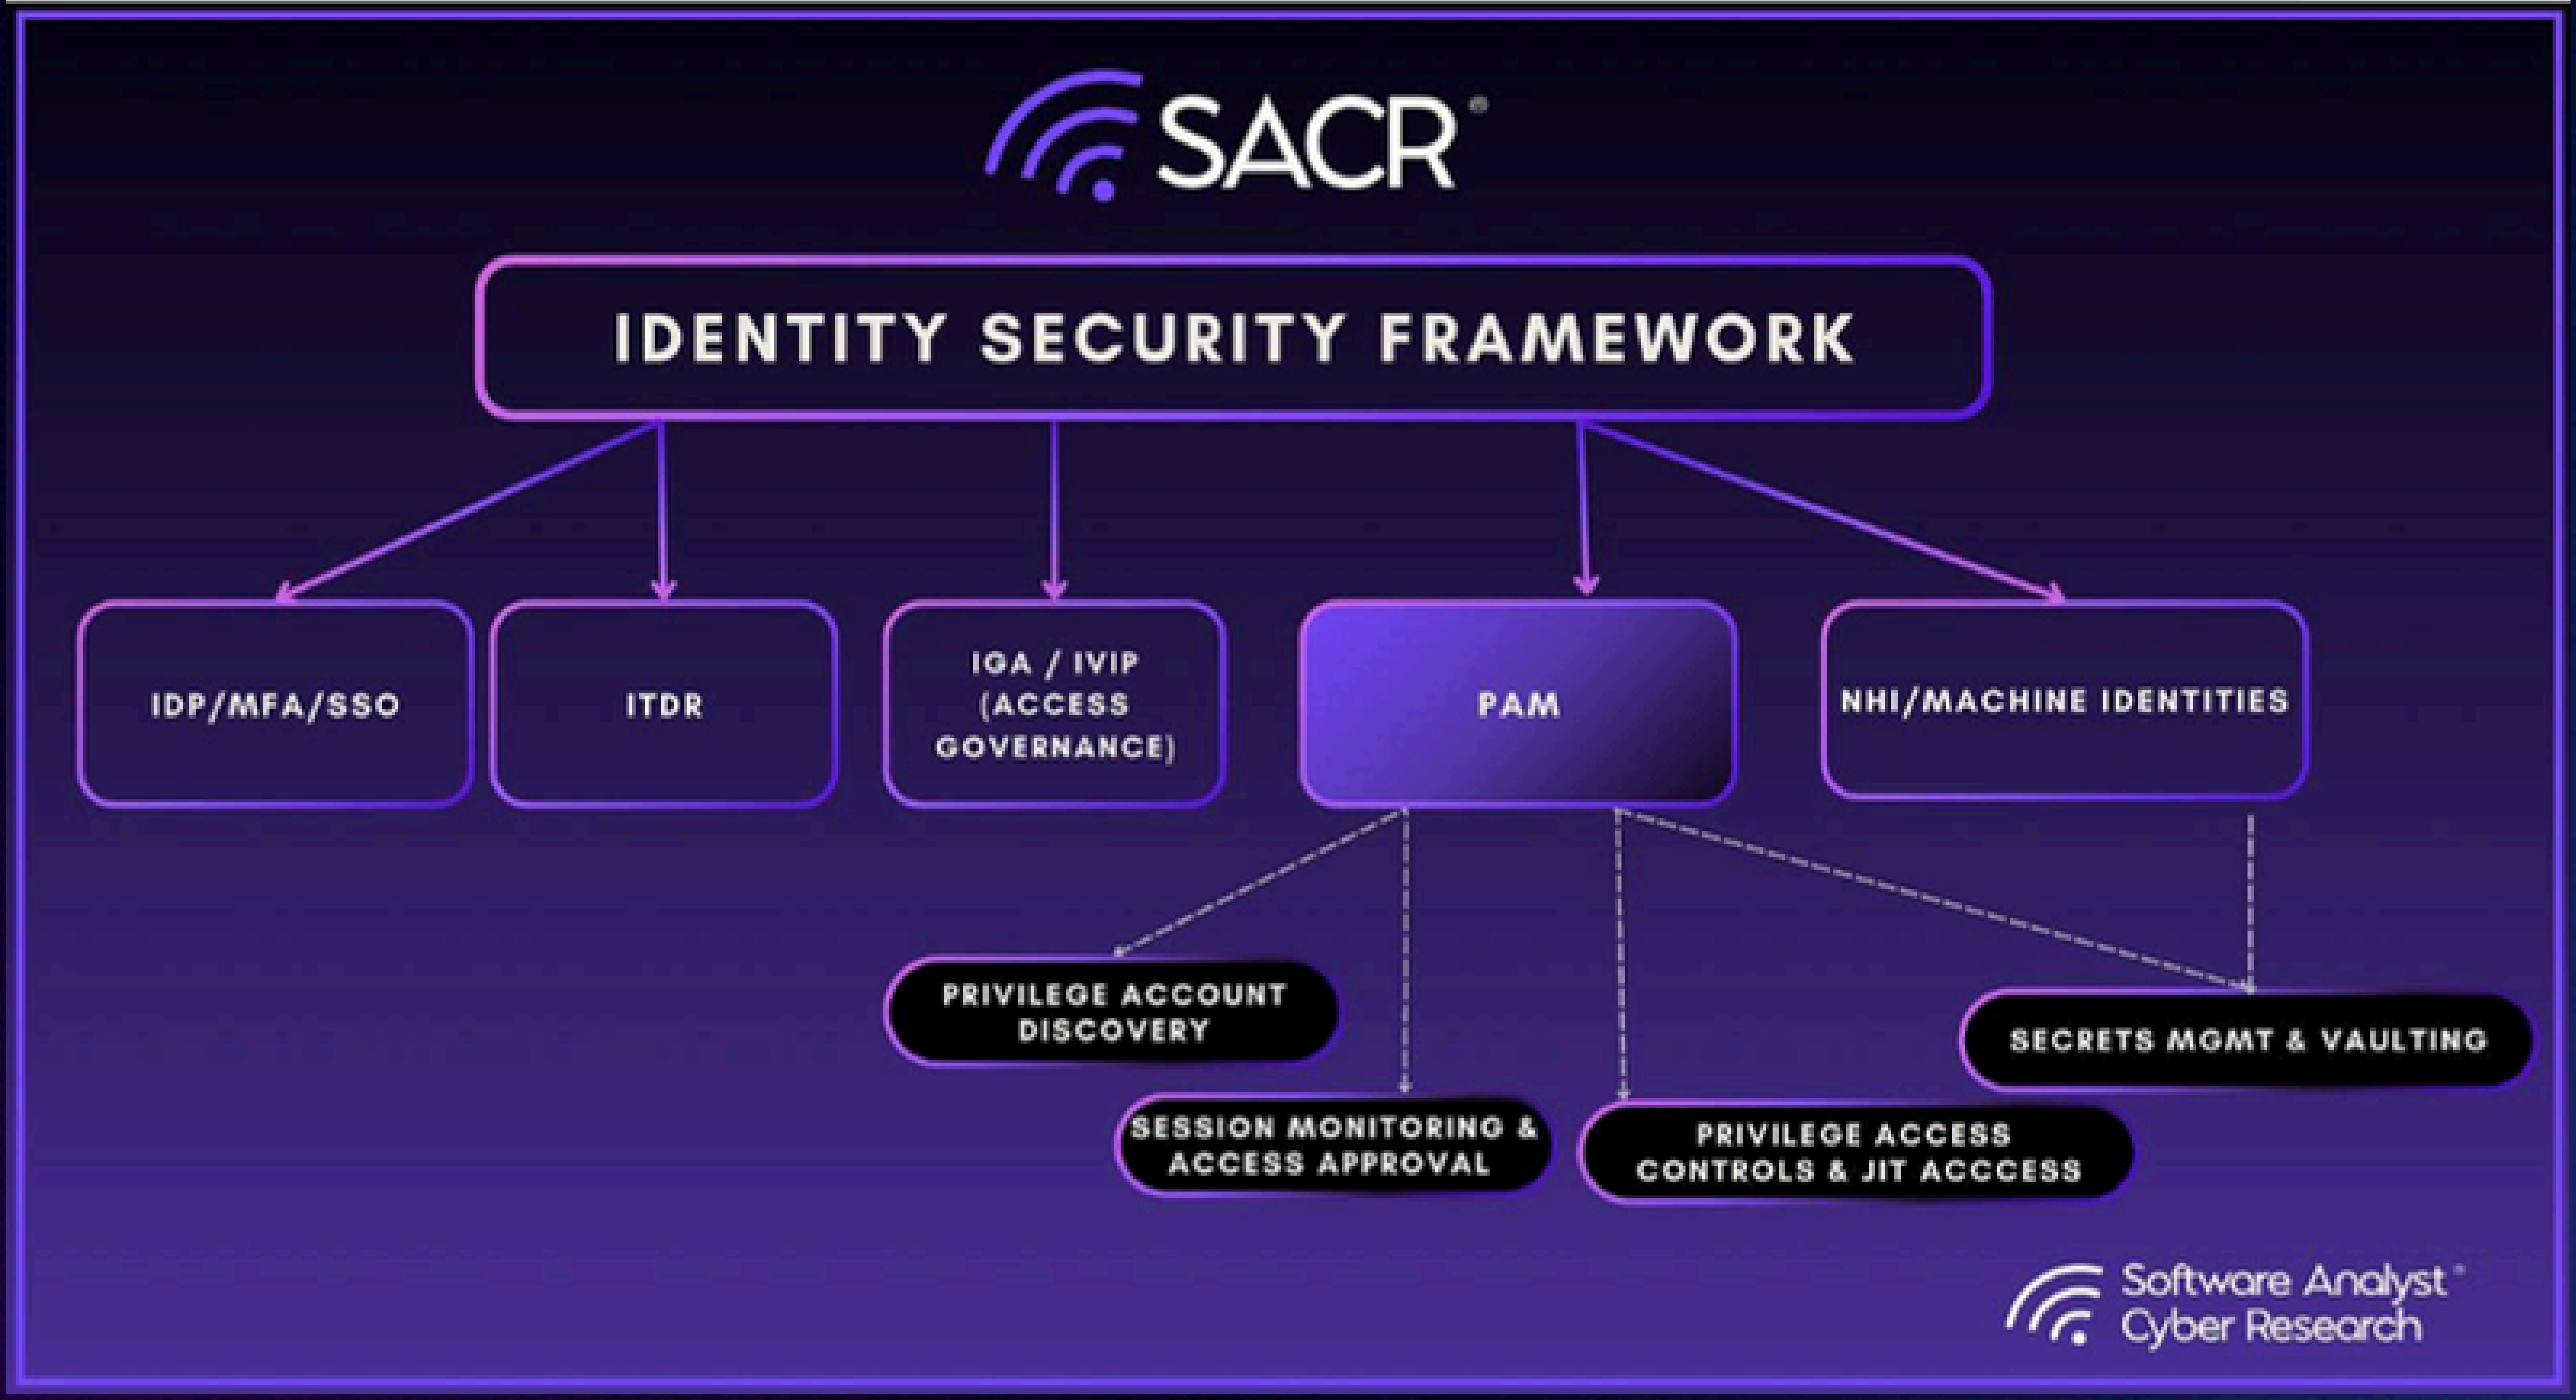
SACR
IDENTITY SECURITY FRAMEWORK
IDP/MFA/SSO
ITDR
IGA / IVIP (ACCESS GOVERNANCE)
PAM
NHI/MACHINE IDENTITIES
PRIVILEGE ACCOUNT DISCOVERY
SESSION MONITORING & ACCESS APPROVAL
PRIVILEGE ACCESS CONTROLS & JIT ACCCESS
SECRETS MGMT & VAULTING
Software Analyst Cyber Research

At the center is Privileged Access Management (PAM). Its the hardest capability to build in identity security relative to others. We believe PAM will be central to managing agent identity. The function responsible for securing the most sensitive and high-impact permissions across your environment. Surrounding PAM are adjacent pillars:

↳ IDP/MFA/SSO, which authenticate users and enable secure logins.

↳ ITDR (Identity Threat Detection & Response), focused on detecting identity-based threats in real time.

↳ IGA/IVIP, which handles visibility into all identities (IVIP) and IGA focuses on governance, access reviews, and joiner/mover/leaver flows.

↳ NHI (Non-Human Identities), which includes service accounts, workloads, bots, and API keys. It is important to realize there is a separation or distinction of NHIs vs agents.

Within PAM itself, we break SACR down into **four essential components**:

Privileged Account Discovery	Find what powerful access exists
Secrets Management & Vaulting	Store credentials safely
Privileged Access Controls & JIT Access	Enforce who gets access, when, and how
Session Monitoring & Access Approval	Observe usage and add checks before granting access

These building blocks work together to enforce least privilege, prevent unauthorized escalation, and contain the blast radius of breaches. The rest of this report will dive deeper into this PAM pillar and, hopefully, provide context for all readers. I provide a foundation breakdown much more in the report.

DEFINING PRIVILEGED ACCESS MANAGEMENT (PAM) IN THE MODERN ENTERPRISE

Privileged Access Management is the identity security discipline focused on protecting, governing, and monitoring access to the most sensitive systems and actions within an environment.

A privileged account is any account that can

- ✓ Change systems
- ✓ Access sensitive data
- ✓ Create or delete users
- ✓ Shut things down
- ✓ Override security controls

Privileged accounts span

- Domain controllers and directory services
- Administrative access to databases, applications, and operating systems
- Public and private cloud IAM permissions
- Network and infrastructure devices
- DevOps secrets, API keys, tokens, and service accounts

Historically, PAM was synonymous with IT administrators and shared root credentials. Today, privilege has expanded dramatically in breadth (more identities) and depth (more powerful actions). Broadly, the move to the cloud has expanded access across the enterprise. We have seen an expansion in developers becoming more admins of critical infrastructure.

Foundational components

The core components of a PAM suite

Privileged Account Discovery

The goal is to identify all privileged accounts across systems, networks, applications, and cloud platforms. There is also another audit logging and compliance where before you can protect privileged access, you need to find it and track it. This pillar handles account discovery, continuous inventory, detailed activity logs, and compliance reporting. It ensures that security teams have the insights to detect risks and prove controls are working.

Credential Vaulting / Rotation

It stores privileged credentials in a secure and encrypted vault accessible only to authorized users or systems. This pillar focuses on protecting the actual secrets: passwords, SSH keys, tokens, and certificates. Credentials are stored in secure vaults, rotated frequently, and retrieved securely without exposing them to users. This reduces the risk of theft, reuse, or unmanaged sprawl.

Access Control & Least Privilege

This category covers Access control, Just-in-Time (JIT) Access and Approval Workflows.
It enforce who gets access, when, and under what conditions. The goal is to eliminate standing privileges and instead issue temporary access based on need and context, often requiring manager or peer approval. It's the core of Zero Trust and minimizes exposure.

Session Management and Monitoring

Once privileged access is granted, this pillar ensures real-time visibility and oversight of what users do with that power. This includes logging, monitoring, session recording, and if needed termination of live sessions. It's crucial for incident response and auditability.

CORE ACTIONABLE SUMMARY FOR READERS

Setting the context for PAM, if you only had a few minutes to read the report. Here are the core takeaways:

Consolidation Has Made Privileged Access a Board-Level Control, not a Point Solution

We highlighted this in our report last year.Palo Alto Networks' acquisition of CyberArk is not simply another large security deal; it is a signal that privileged access has moved to a foundational layer of enterprise security architecture. Platform security vendors do not spend $25B to fill feature gaps, they do it to control a control plane. By embedding privileged identity telemetry into network, endpoint, and SOC workflows, Palo Alto is effectively asserting that identity-driven privilege enforcement must operate at the same level of priority as threat detection and response.

For CISOs, the implication is clear: privileged access is no longer a standalone IAM decision or a compliance checkbox. It is becoming inseparable from how organizations detect, contain, and respond to breaches. This acquisition has created a gap, opening the door for new vendors like Britive, Apono, Teleport and P0 Security.

Machine and Agentic Identities Are Now the Fastest-Growing Privileged Users and the Least Governed

The definition of identities is expanding more and more. The most significant expansion of privilege in modern environments is no longer human administrators, but non-human identities: service accounts, API keys, workloads, automation, and increasingly autonomous AI agents. These identities already outnumber humans by orders of magnitude, and unlike people, they operate continuously, at machine speed, and often with broad, implicit permissions. AI agents amplify this risk further by introducing non-deterministic behavior, an agent granted access to "optimize infrastructure" may legitimately modify or delete production systems if guardrails are weak or misinterpreted.

The Evolution of Privileged Identity

Enterprise security has undergone a structural inversion over the last two decades. Where trust was once anchored to a hardened network perimeter, modern environments have dissolved those boundaries through cloud computing, SaaS adoption, automation scripts, and an API-driven infrastructure. The perimeter no longer meaningfully exists. Identity is now the only consistent boundary. Across any organization, certain users and accounts hold elevated permissions over systems, infrastructure, data, and configurations effectively the "keys to the kingdom." These privileged identities now include not only IT administrators, but developers, cloud engineers, service accounts, APIs, workloads, and increasingly, autonomous AI agents. This is a crucial contrast that is changing PAM.

Cloud and Ephemeral Infrastructure Have Made Standing Privilege Structurally Unsustainable

Cloud has fundamentally broken the assumptions that traditional PAM was built on. Infrastructure is now created and destroyed in minutes, access requirements change continuously, and policy defined through static roles cannot keep pace without creating excessive risk or operational drag. Standing privilege like long-lived permissions granted has become one of the most common root causes of cloud security incidents. In this environment, the question is no longer whether an organization will experience privilege misuse, but whether it can limit the blast radius when it happens. Modern PAM is shifting from credential storage to real-time authorization: provisioning access only when needed, enforcing it in context, and revoking it automatically. Organizations that fail to make this transition are not merely behind the curve; they are operating with an access model that is incompatible with the speed and volatility of their own infrastructure.

HOW PAM IS SLOWLY EVOLVING IN 2026

This new parameter shift, where identity security becomes the "parameter, " is new for many practitioners. What was once a discipline centred on on-prem session control and credential vaulting is now evolving into a core access control layer responsible for enforcing contextual access decisions, continuous verification, and eliminating standing privileges across increasingly dynamic environments.

As a result, the baseline for PAM has materially changed. Just-in-time (JIT) access, remote privileged access, secrets management, and automation are no longer differentiators; they have become stakes. Buyers are increasingly evaluating PAM platforms based on their ability to enforce least privilege in real time, integrate with identity and infrastructure signals, and operate effectively across hybrid, cloud-native, and SaaS environments.

However, the most consequential force reshaping PAM is not cloud adoption or Zero Trust alone, as we've seen in recent years. Agentic AI is now redefining the nature of privileged access itself. The scope of identity has expanded well beyond human administrators to include non-human identities like service accounts and agents that operate with elevated permissions. This expansion in identities and privileged access is exposing a critical gap. Traditional PAM architectures, which were designed for human-centric access and static infrastructure, are increasingly misaligned with these new identities.

For CISOs and investors alike, the next phase of PAM will be determined by which platforms can govern human, machine, and AI identities at runtime, enforce privilege dynamically, and scale trust decisions in systems where access is transient and constantly changing.

This report examines how the evolving PAM market has hit an inflection point for 2026, what it signals about the future of identity security, and which architectural approaches are likely to define the next generation of privileged access control. We selected the following players based on their next-gen PAM criterion and partnered with them to produce this research report for the community. The four vendors do not represent the entire market for next-gen players but they have representative key use-cases that help us illustrate how PAM is evolving for the cloud and agentic world.

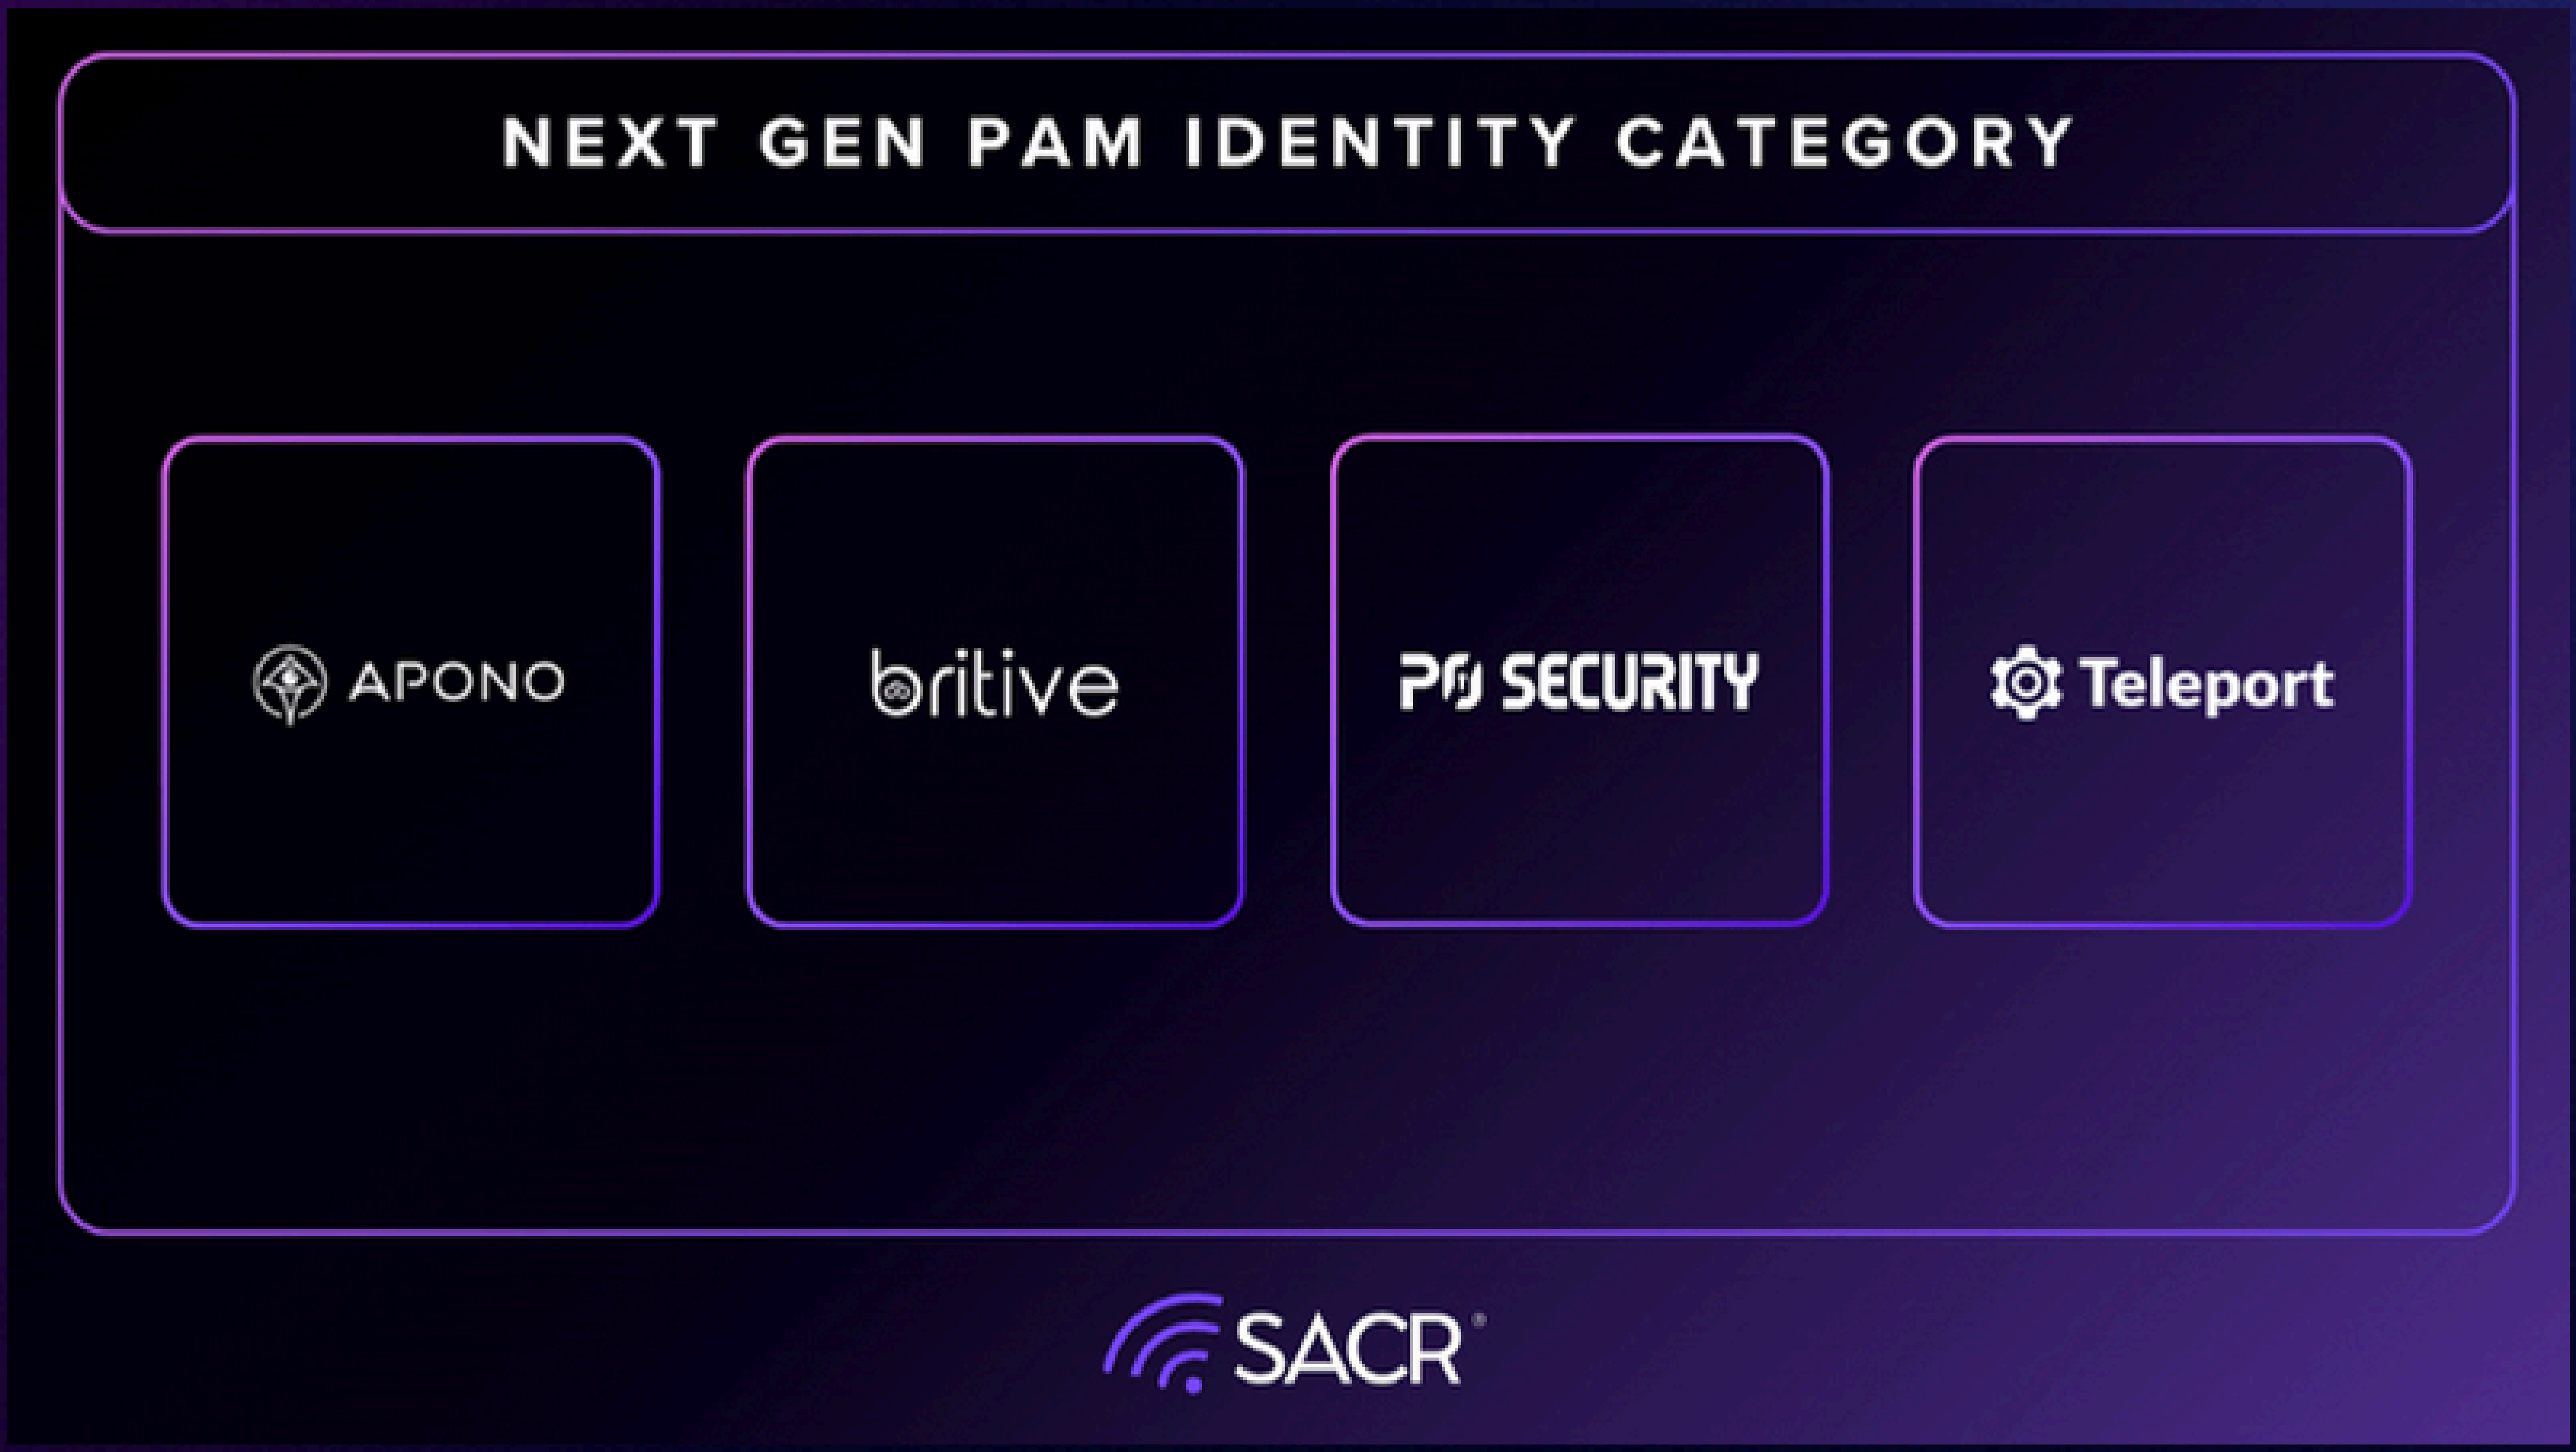

THE HISTORY & EVOLUTION OF PAM

PAM suites then evolved over the years as more applications and infrastructure shifted from on-premise environments to public cloud environments. As we think about PAM, it has changed over the years from an on-prem centric ecosystem to cloud-centric view. More readings can be found here.

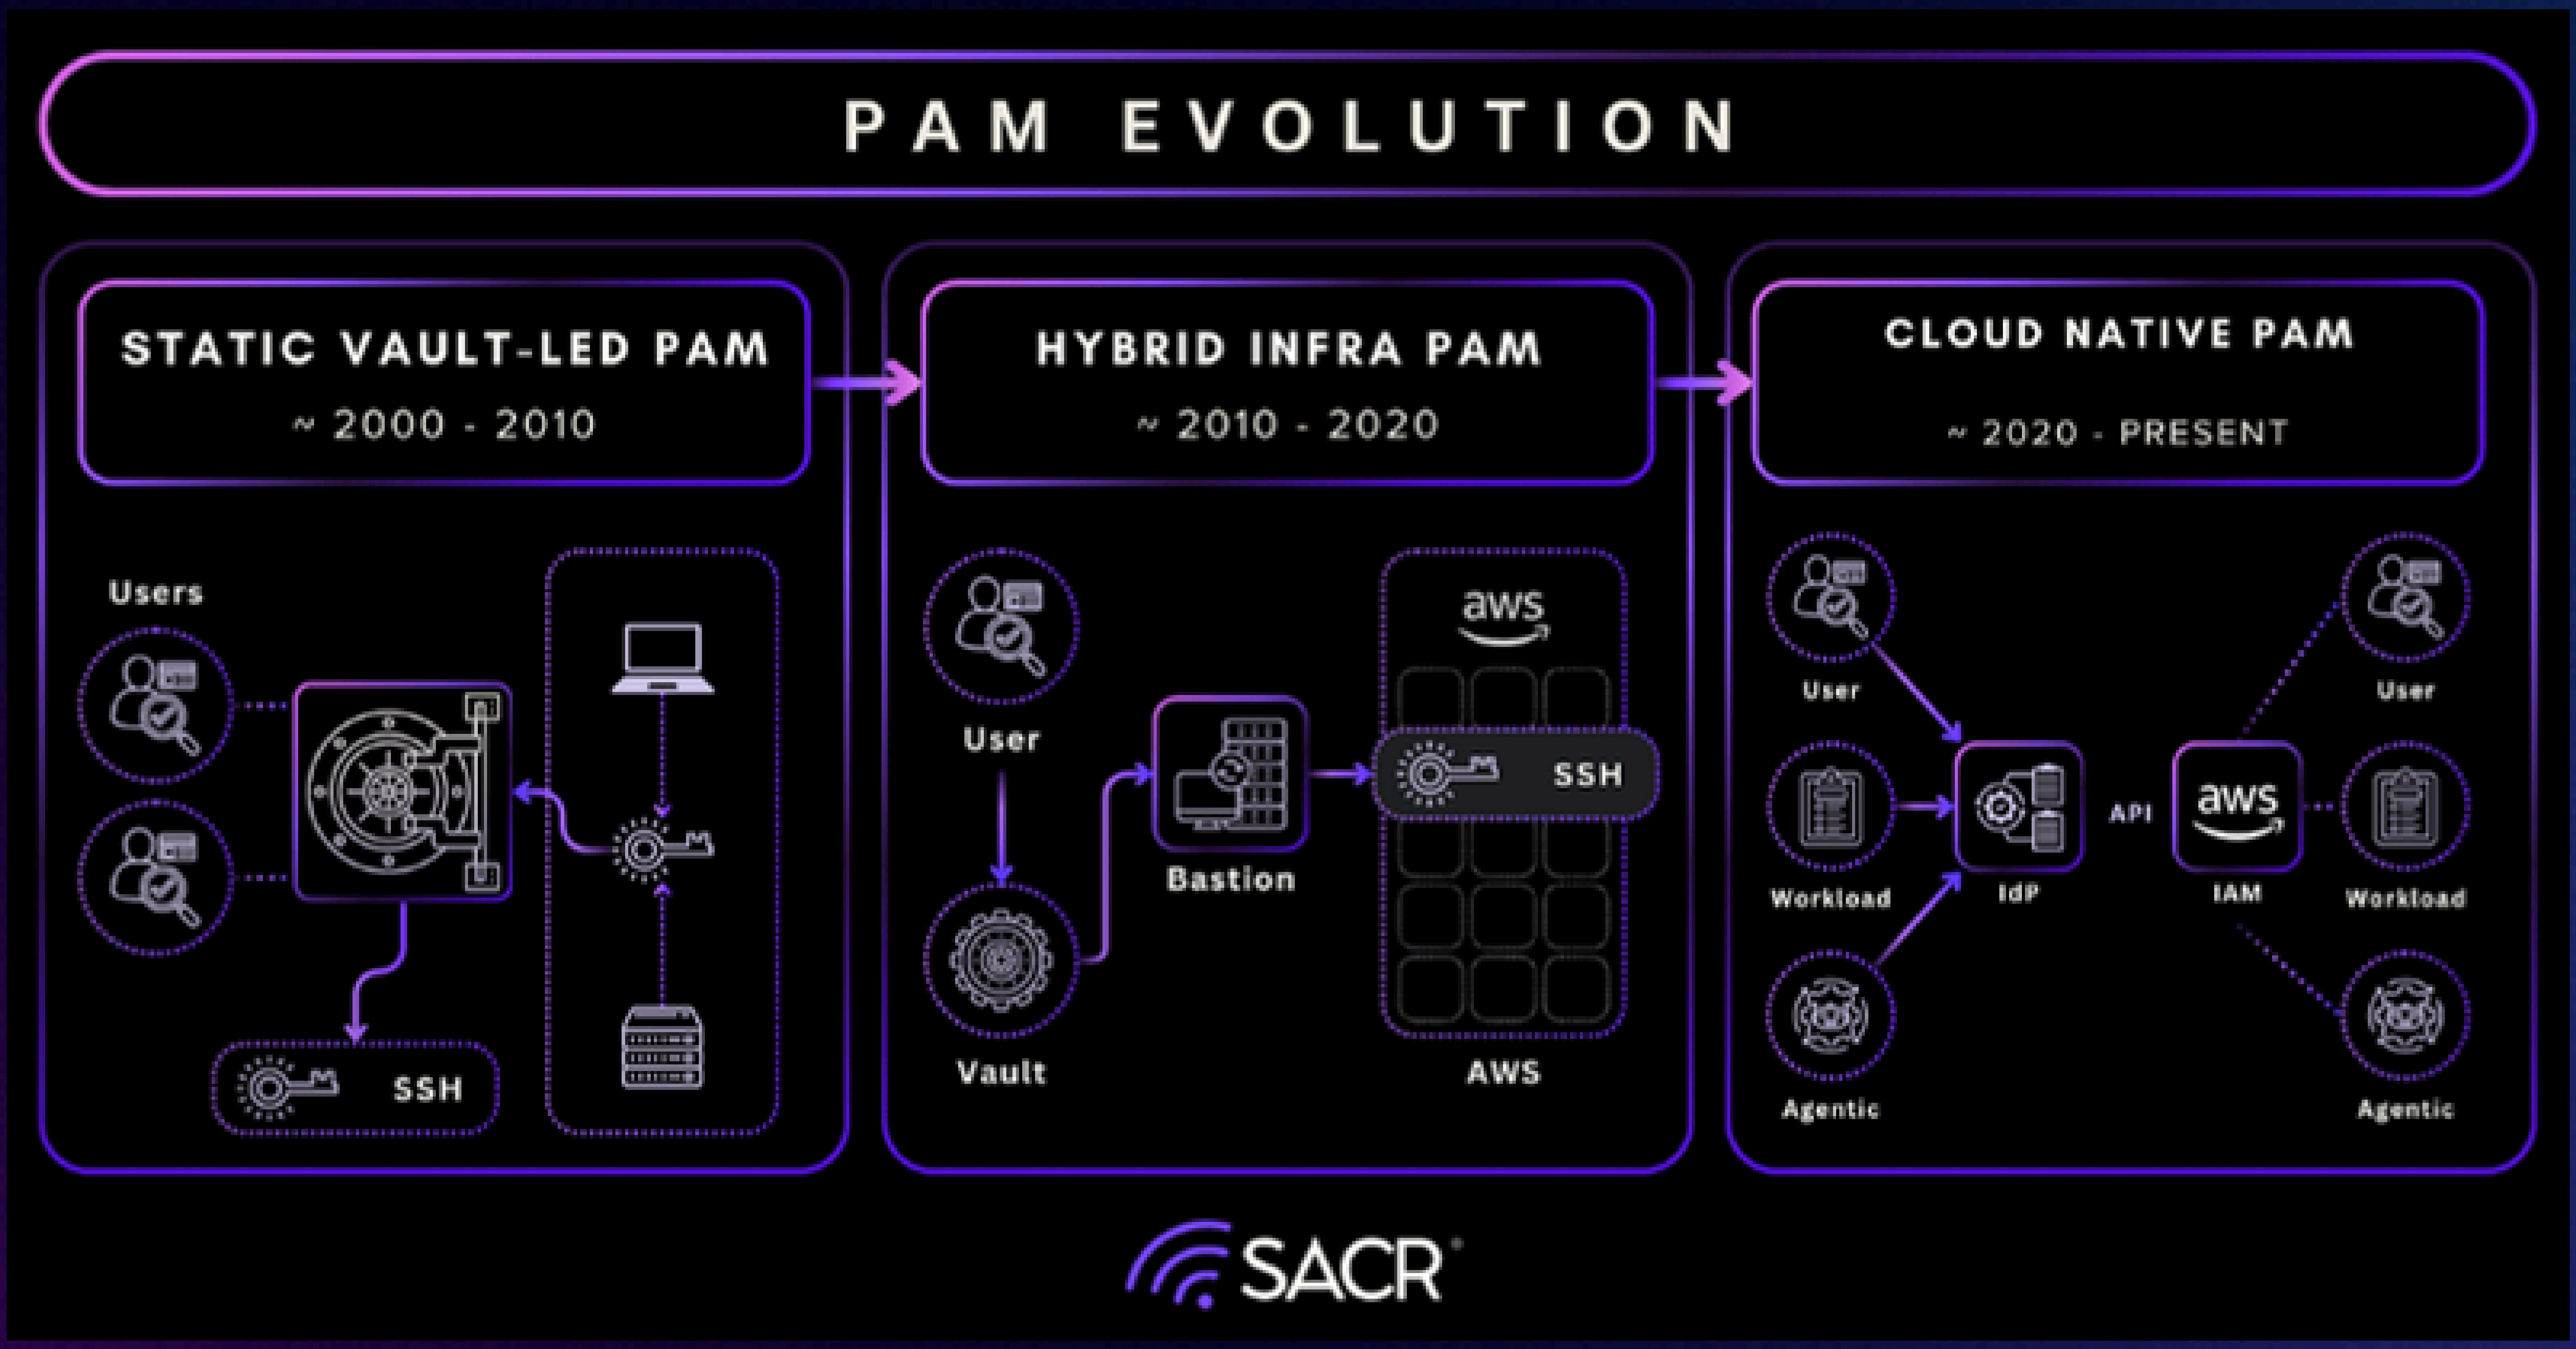

PHASE 1

The Vault-Centric Era (Static Infrastructure, Static Secrets)

The PAM market emerged in the early 2000s alongside large on-premise data centers. Infrastructure was static, credentials were long-lived, and administrative access was often shared and poorly documented.

High-profile breaches exposed the fragility of this model. The 2014 Sony Pictures breach where attackers discovered a literal folder named “Passwords” containing privileged credentials became a defining moment for the category. Adoption accelerated rapidly, coinciding with CyberArk’s IPO and the market’s transition from niche to mission-critical.

Core architectural components of this era included:

- Secure password vaults
- Privileged session management and recording
- Application identity and service account credential storage
- Threat analytics focused on anomalous privileged behavior

This model dramatically reduced risk in static environments but introduced friction. Manual password checkout workflows often failed to scale, leading to shadow accounts and policy bypasses.

PHASE 2

The Hybrid Cloud PAM Era (Session Governance at Scale)

As enterprises migrated to public cloud and virtualized infrastructure, the assumptions underlying vault-centric PAM began to break down. Ephemeral virtual machines, autoscaling workloads, CI/CD pipelines, and SaaS platforms caused an explosion in privileged credentials which many teams refer to as secret sprawl.

PAM platforms expanded to address this hybrid reality:

- ✓ DevOps and secrets management
- ✓ Endpoint privilege management
- ✓ Cloud privilege and entitlement management
- ✓ Bastion-based session mediation for SSH, RDP, and web access

Despite architectural change, market leadership remained stable. CyberArk and BeyondTrust retained dominance, while Centrify and Thycotic merged to form Delinea. At the same time, DevOps-native players like HashiCorp carved out a strong position around secrets management. The core limitation persisted: static roles and long-lived permissions do not align with infrastructure that changes by the minute.

Cloud Introduced Complexities Around Managing Ephemeral and Dynamic Privileged Access

The scale and volatility of modern cloud environments have fundamentally changed the requirements for privileged access management. Infrastructure is no longer composed of long-lived assets with predictable access patterns; instead, thousands of resources, virtual machines, containers, serverless functions, and cloud services are created, modified, and destroyed on a continuous basis.

In this context, access models built on static roles and pre-defined permissions become operationally unmanageable. Roles require constant updates, new resources must be manually onboarded, and permissions frequently lag behind the actual state of the environment, creating both security gaps and administrative overhead.

To be effective in cloud-native environments, privileged access must shift from static entitlement management to dynamic, context-aware authorization. This requires continuous, real-time discovery of infrastructure and identities, coupled with policies that evaluate access requests based on attributes such as workload context, environment, risk signals, and business intent at the time of use. Privilege must be provisioned just-in-time, scoped narrowly to the specific task or resource, and automatically revoked once the task is complete. Solutions that cannot adapt to infrastructure as it comes online without manual intervention; do not reduce risk; they simply introduce friction and complexity that teams will eventually work around.

As a result, defining privileged access policy through Infrastructure as Code (IaC) is becoming a practical requirement rather than a best practice for organizations operating at scale. Security and access controls must be versioned, automated, and deployed alongside infrastructure changes to remain effective. These requirements are increasingly driven not only by security teams, but also by platform, cloud, and DevOps engineers who are responsible for day-to-day operations and uptime.
In environments where velocity is a competitive necessity, privileged access solutions must enforce control without impeding delivery: otherwise, they will be bypassed, undermining both security and governance.

PHASE 3

The Zero Standing Privilege for AI Agents Era (Authorization Over Authentication)

The most significant shift underway is the transition from standing privilege to ephemeral, just-in-time authorization continuously in cloud environments and increasingly for AI agents.

In cloud-native environments, privileged access is less about logging into servers and more about executing API calls that mutate infrastructure state. Modern PAM architectures integrate directly with identity providers and cloud control planes, provisioning temporary credentials only when required and revoking them automatically upon task completion.

This Zero Standing Privilege (ZSP) model reduces blast radius, eliminates permanent secrets, and aligns access with real-time context. Newer entrants argue that managing secrets indefinitely is an anti-pattern, the real solution is eliminating the need for secrets altogether. Importantly, this is not yet the norm. Most enterprises remain hybrid, and vault-based PAM will remain essential for years. The market is not replacing legacy PAM, it is layering dynamic authorization on top of it.

THE EXPANSION OF PRIVILEGE FROM HUMANS TO MACHINES

The complexity of human identity roles requires newer PAM solutions

The breadth and depth of access that engineers need to complete their work in the cloud have created a new attack surface for security teams to manage. This new class of privileged account presents two critical challenges for PAM: unprecedented scale, with exponentially more privileged accounts to manage; and heightened expectations regarding user experience. Unlike traditional privileged users, today's developers represent a large population whose productivity directly impacts business outcomes, making friction intolerable.As a result, we regularly speak with CISOs who are as concerned with their team being viewed internally as business enablers as they are with addressing risk: an outcome that is often incompatible with traditional approaches to privileged access.We've also seen the scope of privileged access projects shift to include many stakeholders not typically associated with the category, including cloud security, platform security, DevOps, and even engineering leadership.These dynamics have redefined what organizations need and expect from PAM solutions: robust self-service capabilities, developer-friendly workflows, and the ability to secure access without impeding velocity.

Machines & Agents Evolving The Market

Identity security is increasingly bifurcated into human and non-human (machine) identities. While the number of human users has stabilized, machine identities: service accounts, API keys, tokens, certificates, workloads are growing exponentially. Most organizations already manage 40–50x more non-human identities than human ones, yet visibility and governance lag far behind. High-profile breaches at Okta, JumpCloud, and others have demonstrated that poorly managed machine identities represent a systemic risk. Machine identity security has evolved from SSH key management into a broader discipline encompassing:

- Discovery and inventory
- Lifecycle management
- Secret rotation and certificate management
- Behavioral detection and response

As PAM expanded into DevOps and secrets management, its boundaries increasingly overlapped with machine identity platforms such as Venafi and emerging startups focused exclusively on this problem.

THE FUTURE OF PAM: THE RISE OF AGENTIC AI & THE REDEFINITION OF PRIVILEGE WITH AGENTS

AI agents represent a new class of privileged identity. When we compare them to traditional service accounts, agentic systems can reason, plan, and execute multi-step actions across domains. If compromised via prompt injection or model manipulation, these agents become high-speed insider threats. Organizations that have already operationalized just-in-time access, session auditing, and dynamic policy enforcement for humans are structurally better positioned to govern AI agents.

However, effective deployment of AI and agents require access to sensitive internal systems, data, and infrastructure, yet a majority of organizations remain uncertain about how to enable that access safely. Recent research indicates that a significant portion of IT leaders lack confidence in their ability to govern AI interactions with proprietary data, highlighting a growing gap between AI ambition and access control maturity. For many enterprises, even achieving zero standing privilege for human users remains an aspirational goal; extending privilege safely to non-deterministic AI agents introduces a materially higher level of risk.

Agentic identity is still an emerging domain, but customer sentiment is converging around a clear conclusion: privileged access maturity is a prerequisite for agentic AI adoption, not a downstream enhancement. AI agents operate continuously, execute multi-step actions, and can affect production environments at machine speed. Without just-in-time access, real-time authorization, session-level auditing, and anomaly detection, these agents effectively function as high-velocity privileged insiders. Organizations that have already operationalized these controls for human users are structurally better positioned to apply the same governance patterns to autonomous systems.

Additionally, AI co-pilots that inherit or "piggyback" on human access rights significantly expand the attack surface in environments with weak privilege boundaries. In such cases, compromise of a single identity can cascade across human and machine workflows.

Forward-looking CISOs are responding by treating PAM as a foundational control for AI readiness. We broadly believe that investing now in deep visibility into sensitive resources, enforcing automated and context-aware permissions, and aligning privileged access decisions with business intent. This groundwork will ultimately determine whether AI becomes a controlled force multiplier or an ungoverned source of systemic risk.

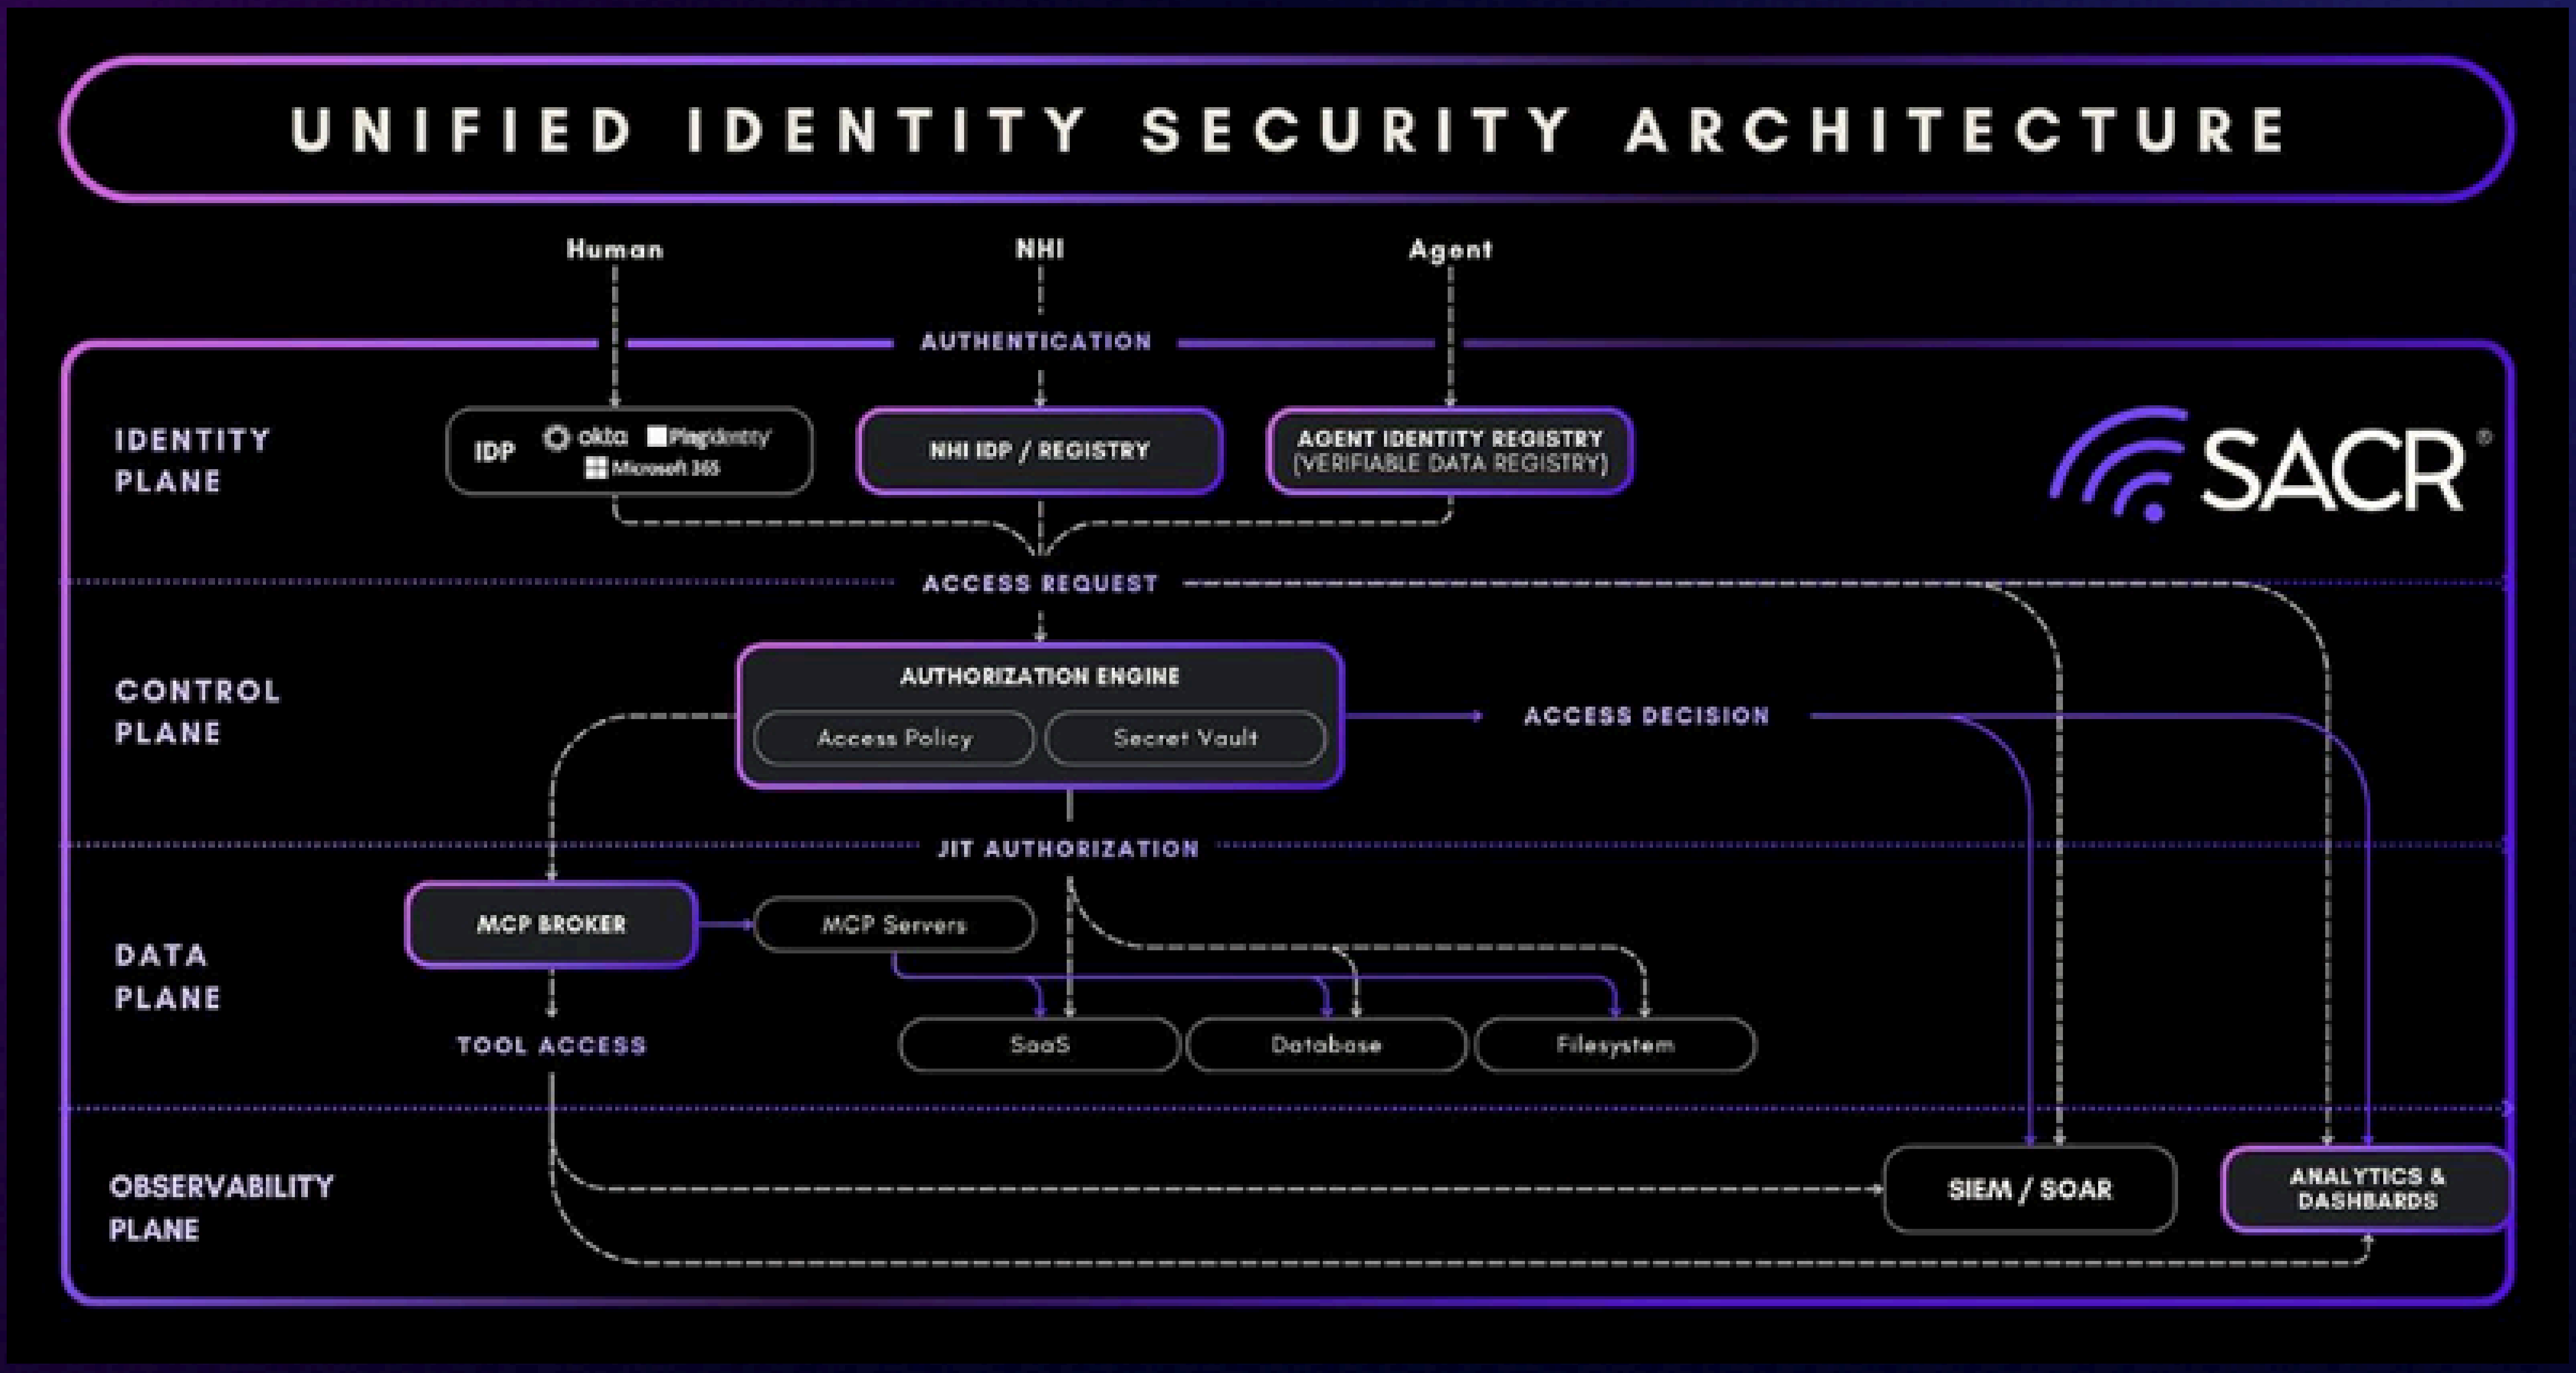

Source: britive.com/use-cases/cloud-pam

FUTURE THEMES TO WATCH

The Rise of Just-in-Time Trust (JIT Trust)

Just-in-Time Trust (JIT Trust) represents the next evolutionary step beyond Zero Trust Architecture (ZTA). It emerges as a unified identity control layer designed for the speed of modern threats and the rise of autonomous, agentic systems and AI agents. Where Zero Trust redefined who can access what, JIT Trust redefines how long, under what conditions, and for what exact purpose access exists. Unified identity systems, stronger authentication, and just-in-time access must be contextually aware of intent.

At its core, JIT Trust treats access as a continuously evaluated, ephemeral resource rather than a static entitlement. Long-lived credentials and standing privileges are replaced with temporary, self-destructing Ephemeral Access Grants (EAGs) and short-lived, certificate-based authentication. These grants are narrowly scoped to the precise resources and actions required for a specific human, machine, or agentic task.

JIT Trust moves beyond traditional authentication toward continuous authorization grounded in behavioral and intent-based signals. Rather than validating identity once at login, the system continuously monitors an entity's intent and behavior: the digital signals generated through AI prompts, tool usage, API calls, and execution patterns to dynamically assess risk and derive trust in real time. When behavior deviates from an established baseline or intent shifts unexpectedly, the system initiates a graduated response: privileges can be reduced, entitlements constrained, or access suspended entirely.

This model establishes Continuous Adaptive Trust (CAT): a control framework where trust is not assumed, but continuously earned and recalibrated. The result is a dramatically reduced attack surface, tighter blast-radius containment, and an access model that aligns with environments defined by automation, ephemerality, and machine-speed execution.

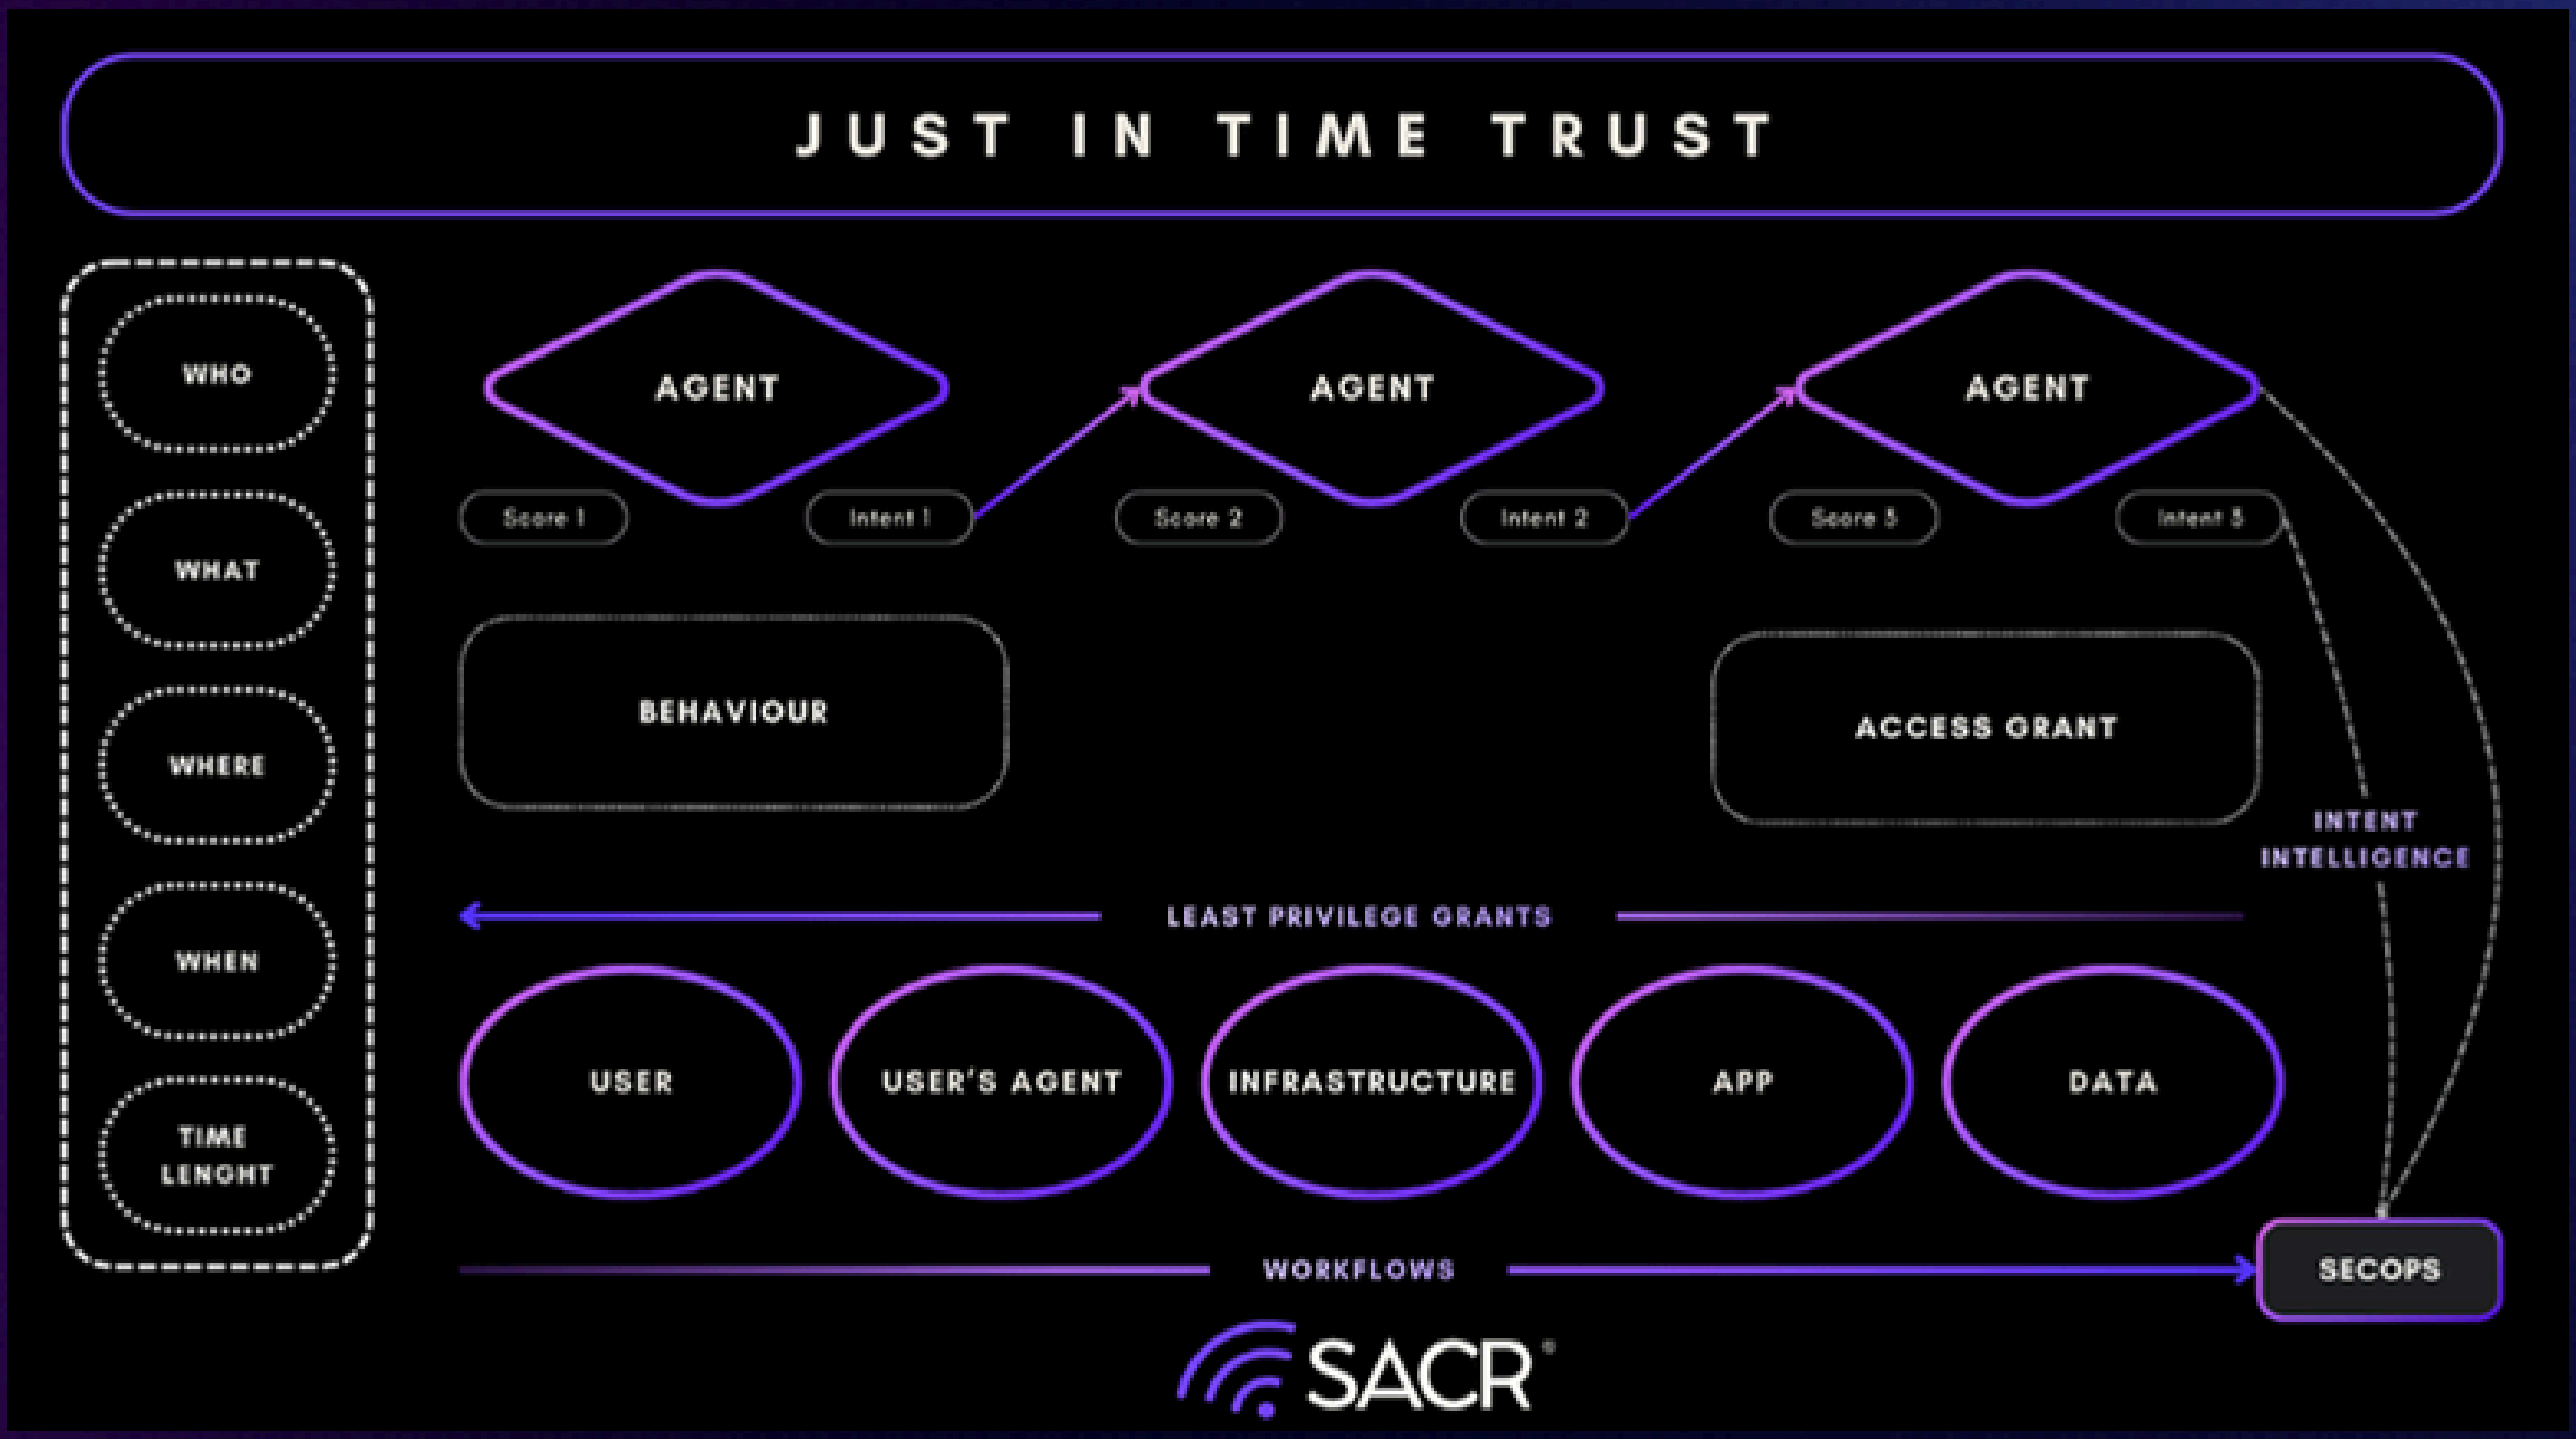

PAM Market Ecosystem

This landscape illustrates a Privileged Access Management (PAM) market that has expanded well beyond its historical roots in password vaulting and session control. What was once a narrow category dominated by a handful of incumbents has evolved into a broad ecosystem spanning identity providers, secrets management, cloud-native access platforms, developer-centric tools, and emerging authorization layers.

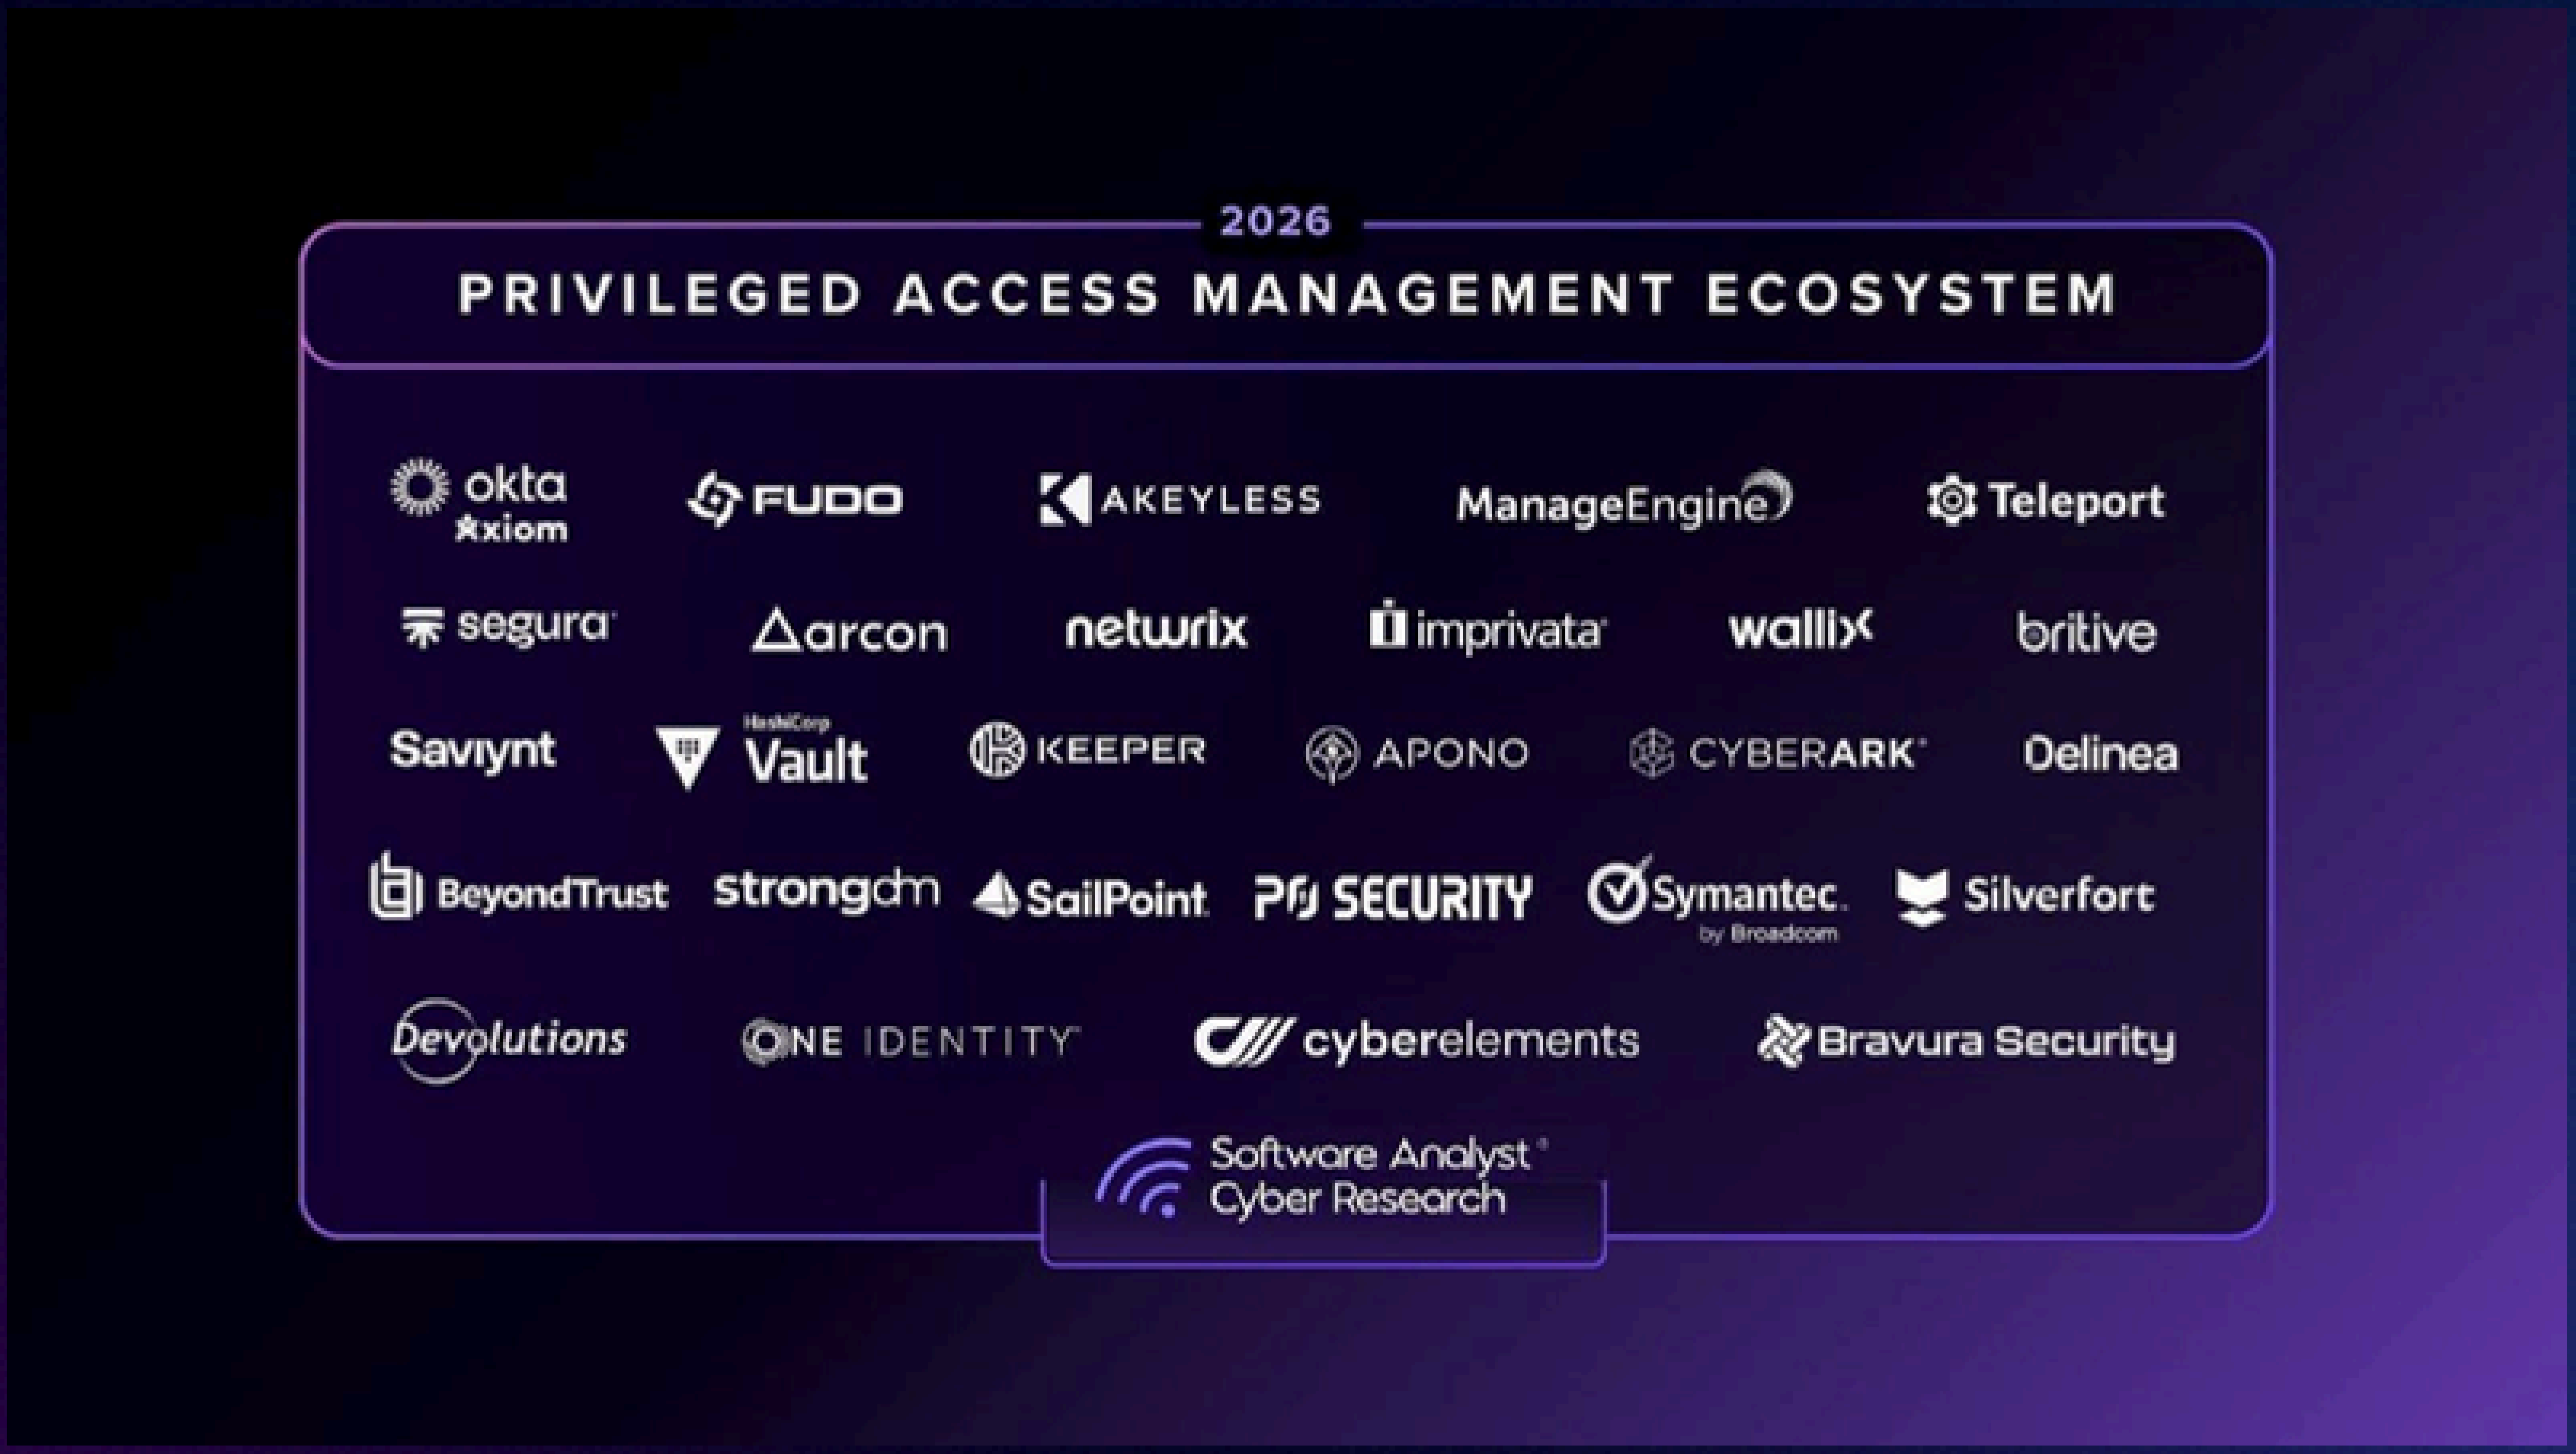

The presence of legacy leaders such as CyberArk, Delinea, and BeyondTrust alongside newer cloud- and API-native players like Apono, Britive, Teleport, and StrongDM reflects a market in transition rather than replacement. Incumbents continue to anchor regulated and hybrid environments, while newer entrants are redefining PAM around just-in-time access, zero standing privilege, and runtime authorization aligned with cloud and DevOps workflows. At the same time, adjacent identity vendors including Okta, SailPoint, and One Identity underscore the growing convergence between PAM, IAM, IGA, and identity threat detection.

Collectively, this ecosystem signals that PAM is no longer a point solution but a foundational control layer within the modern identity stack. As non-human identities, automation, and agentic systems proliferate, the market is shifting toward platforms that can enforce privilege dynamically, at scale, and in context positioning PAM as a central pillar of enterprise security architecture heading into 2026.

Based on our extensive work, we want to dive into next-gen platform that are set to capitalize on the next evolution of cloud and agentic PAM architecture.

ANALYST ASSESSMENT

EXECUTIVE SUMMARY

britive

Britive has emerged as a next-generation, cloud-native Privileged Access Management (PAM) platform architected around runtime-centric authorization rather than a vault-centric credential store. By enforcing access decisions at runtime using native control planes and APIs, the platform eliminates the need for persistent privileged accounts or credential checkout. This design enables a Zero Standing Privilege (ZSP) operating model, where permissions are created only when needed, scoped to the task, and automatically revoked once execution completes. In contrast, legacy Phase 1 PAM providers (e.g., CyberArk, Delinea) rely on persistent accounts and vaulted credentials that continue to exist even when access is not actively in use.

Broadly, Britive supports three distinct identity classes under a unified authorization model: human identities, non-human identities (NHIs) such as those used by workloads and pipelines, and agentic AI identities. In cloud-native environments, this model also extends to Kubernetes, where privileged access often spans clusters, namespaces, and platform services and benefits from the same time-boxed, policy-governed authorization approach.

Core Architectural Philosophy

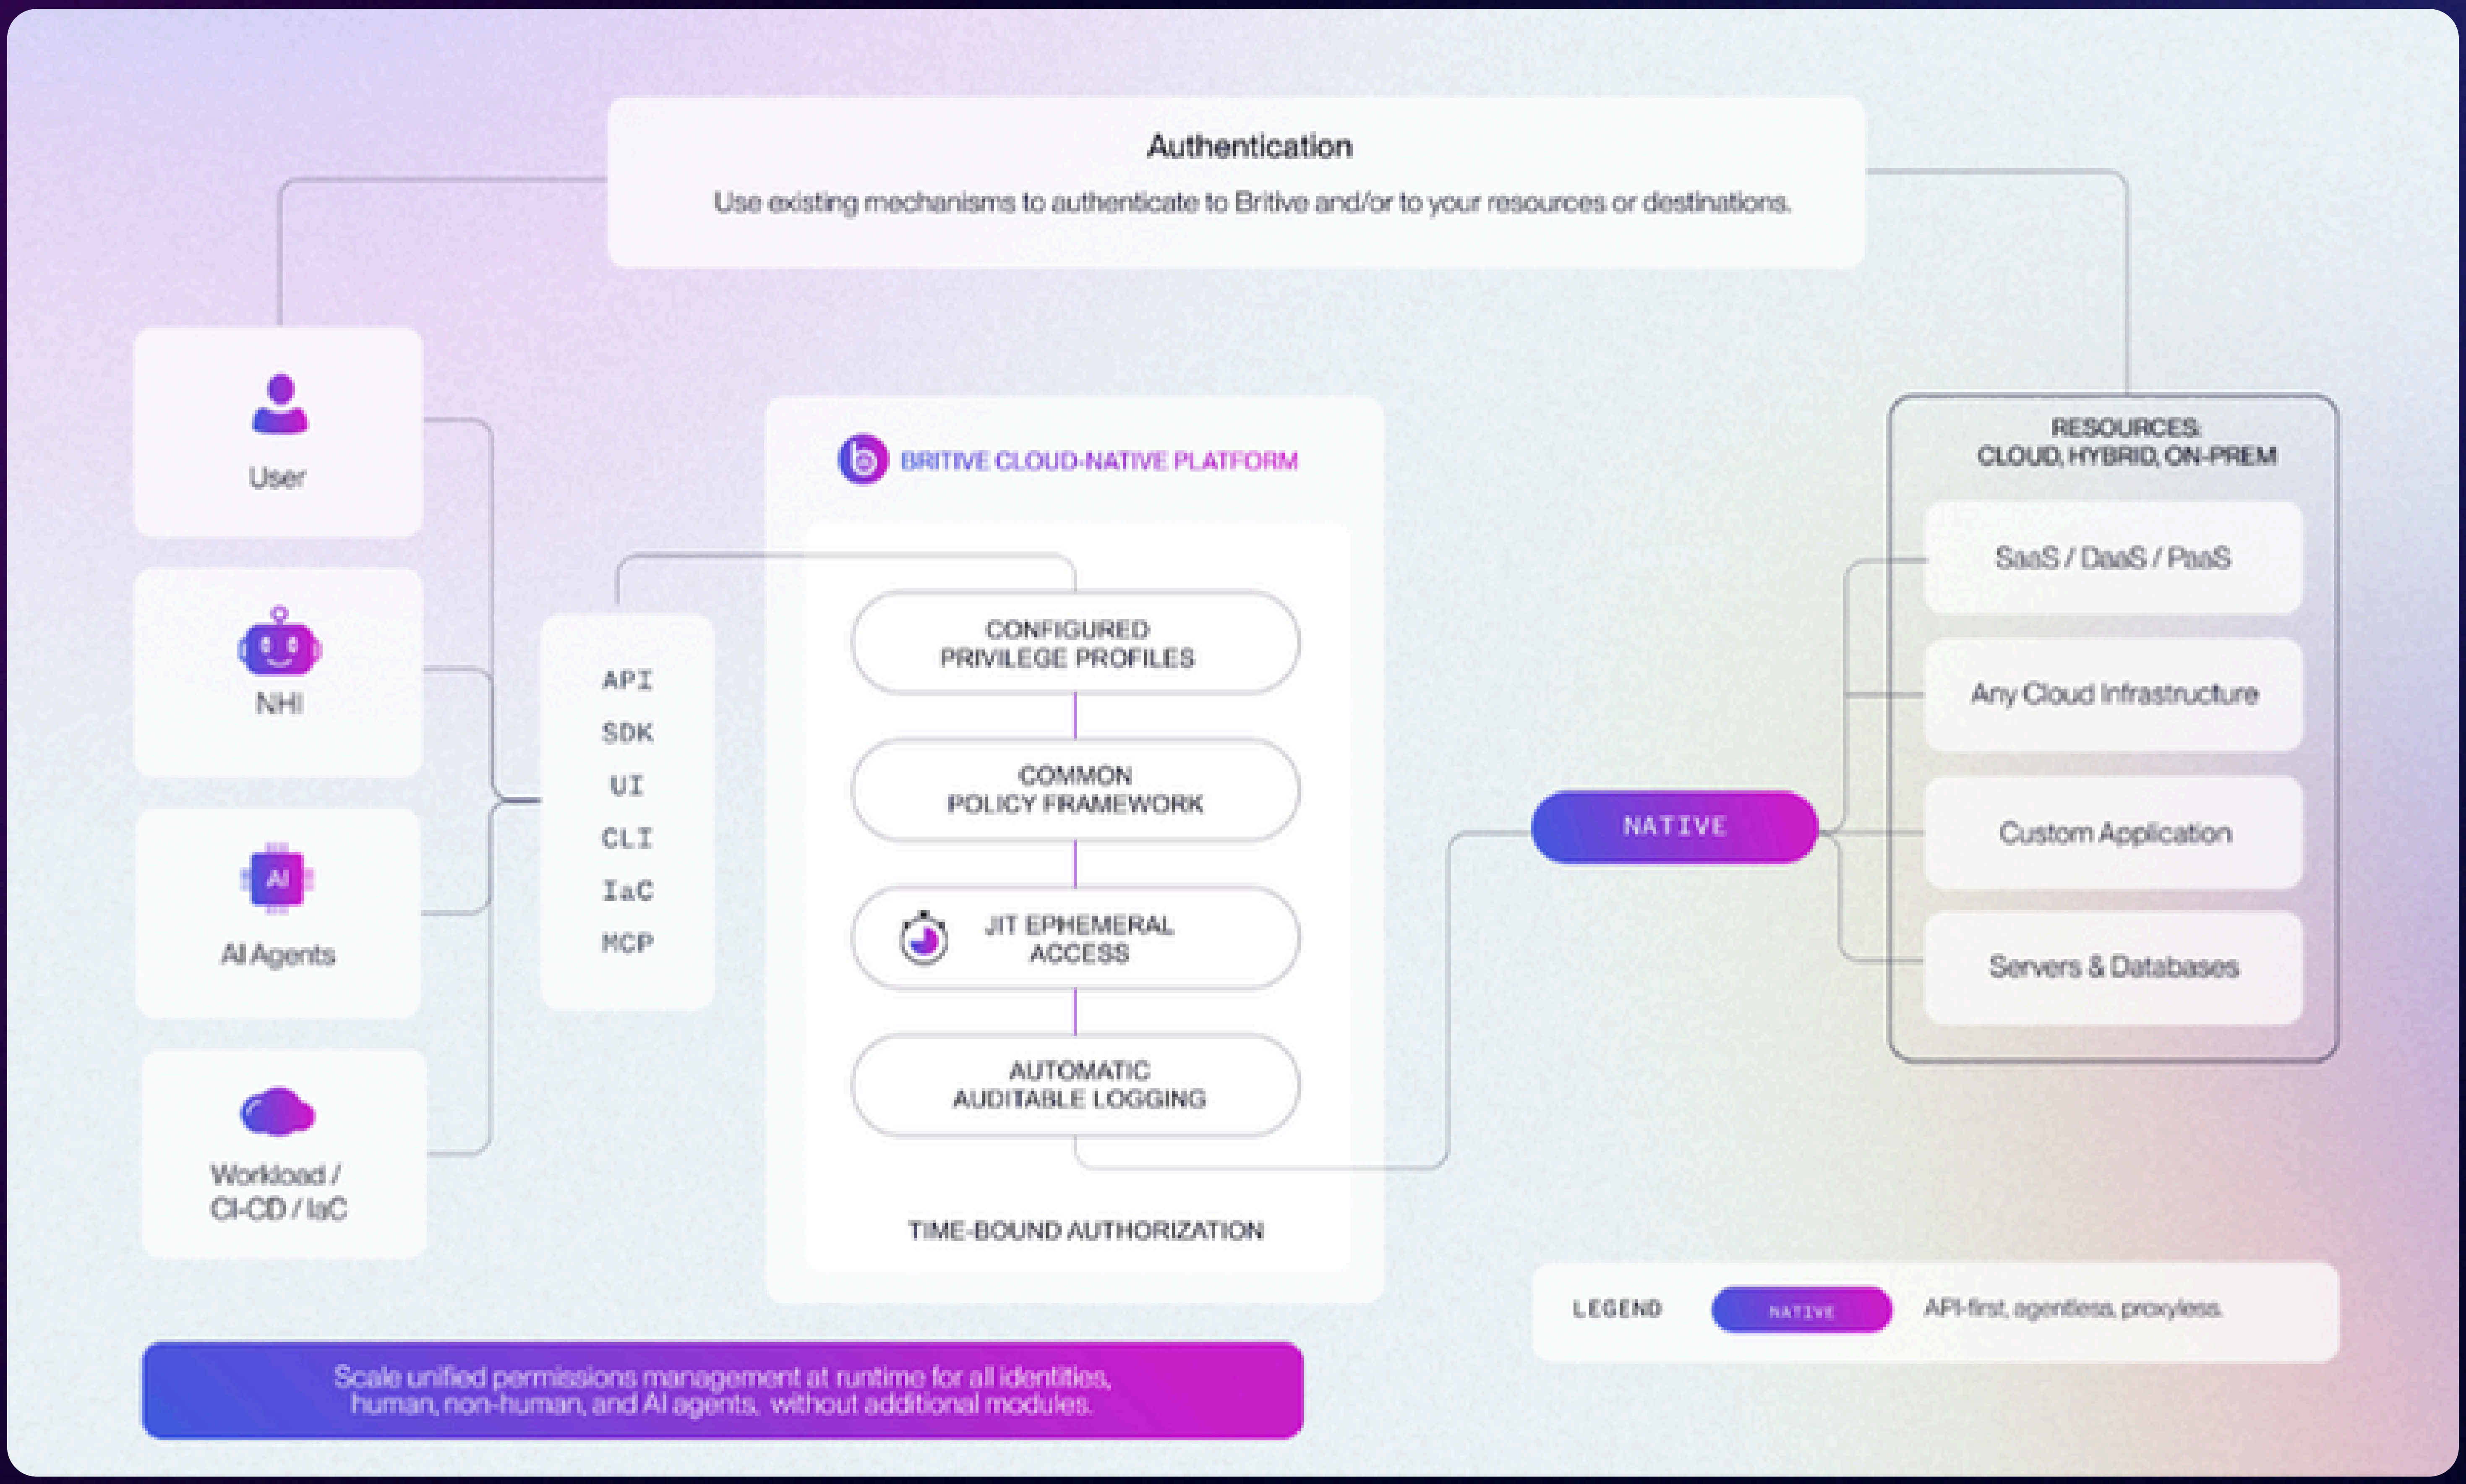
Authentication
Use existing mechanisms to authenticate to Britive and/or to your resources or destinations.
User
NHI
AI
AI Agents
Workload / CI-CD / IaC
API
SDK
UI
CLI
IaC
MCP
BRITIVE CLOUD-NATIVE PLATFORM
CONFIGURED PRIVILEGE PROFILES
COMMON POLICY FRAMEWORK
JIT EPHEMERAL ACCESS
AUTOMATIC AUDITABLE LOGGING
TIME-BOUND AUTHORIZATION
NATIVE
RESOURCES: CLOUD, HYBRID, ON-PREM
SaaS / DaaS / PaaS
Any Cloud Infrastructure
Custom Application
Servers & Databases
LEGEND
NATIVE
API-first, agentless, proxyless.
Scale unified permissions management at runtime for all identities, human, non-human, and AI agents, without additional modules.

From our analysis, Britive's primary technical differentiator is its agentless and proxyless approach. While many modern competitors still utilize connectors or jump boxes to gate access, Britive integrates directly into the native identity control planes of major Cloud Service Providers (CSPs). We observed that, as a SaaS-native platform with an agentless, proxyless design for cloud-native targets, Britive can typically be stood up quickly without the overhead of deploying and maintaining proxy-heavy architectures. Onboarding can be accelerated through standard identity integrations, including SSO (SAML 2.0) and automated provisioning (SCIM 2.0), enabling teams to sync users and groups from providers such as Okta, Ping, and Microsoft Entra ID and begin building a common policy layer rapidly across environments.

Secondly, we believe Britive has a strong common policy model and access profiles system. We were impressed to see that Britive has an abstracted policy layer that spans all target systems, including cloud, SaaS, and on-premises environments (depending on system requirements). This provides unified security controls (MFA, ServiceNow integration, approval workflows) regardless of whether individual systems natively support them. This is an important differentiator that few vendors provide consistently across the ecosystem.

Britive's architecture separates the authorization control plane from the underlying entitlement systems of target resources. Rather than simply placing users into native groups or roles for a fixed time window (the industry's standard JIT model), Britive constructs access profiles as collections of permissions wrapped in a unified policy framework that applies consistently regardless of backend system capabilities.

Britive intentionally abstracts each target system's entitlement model. This abstraction provides a common layer of security controls, including MFA challenges, approval workflows, ServiceNow integration, time-of-day restrictions, and IP constraints, that persist whether AWS supports those features natively or not, whether Google does or does not, or whether an on-premises Active Directory domain has step-up verification built in.

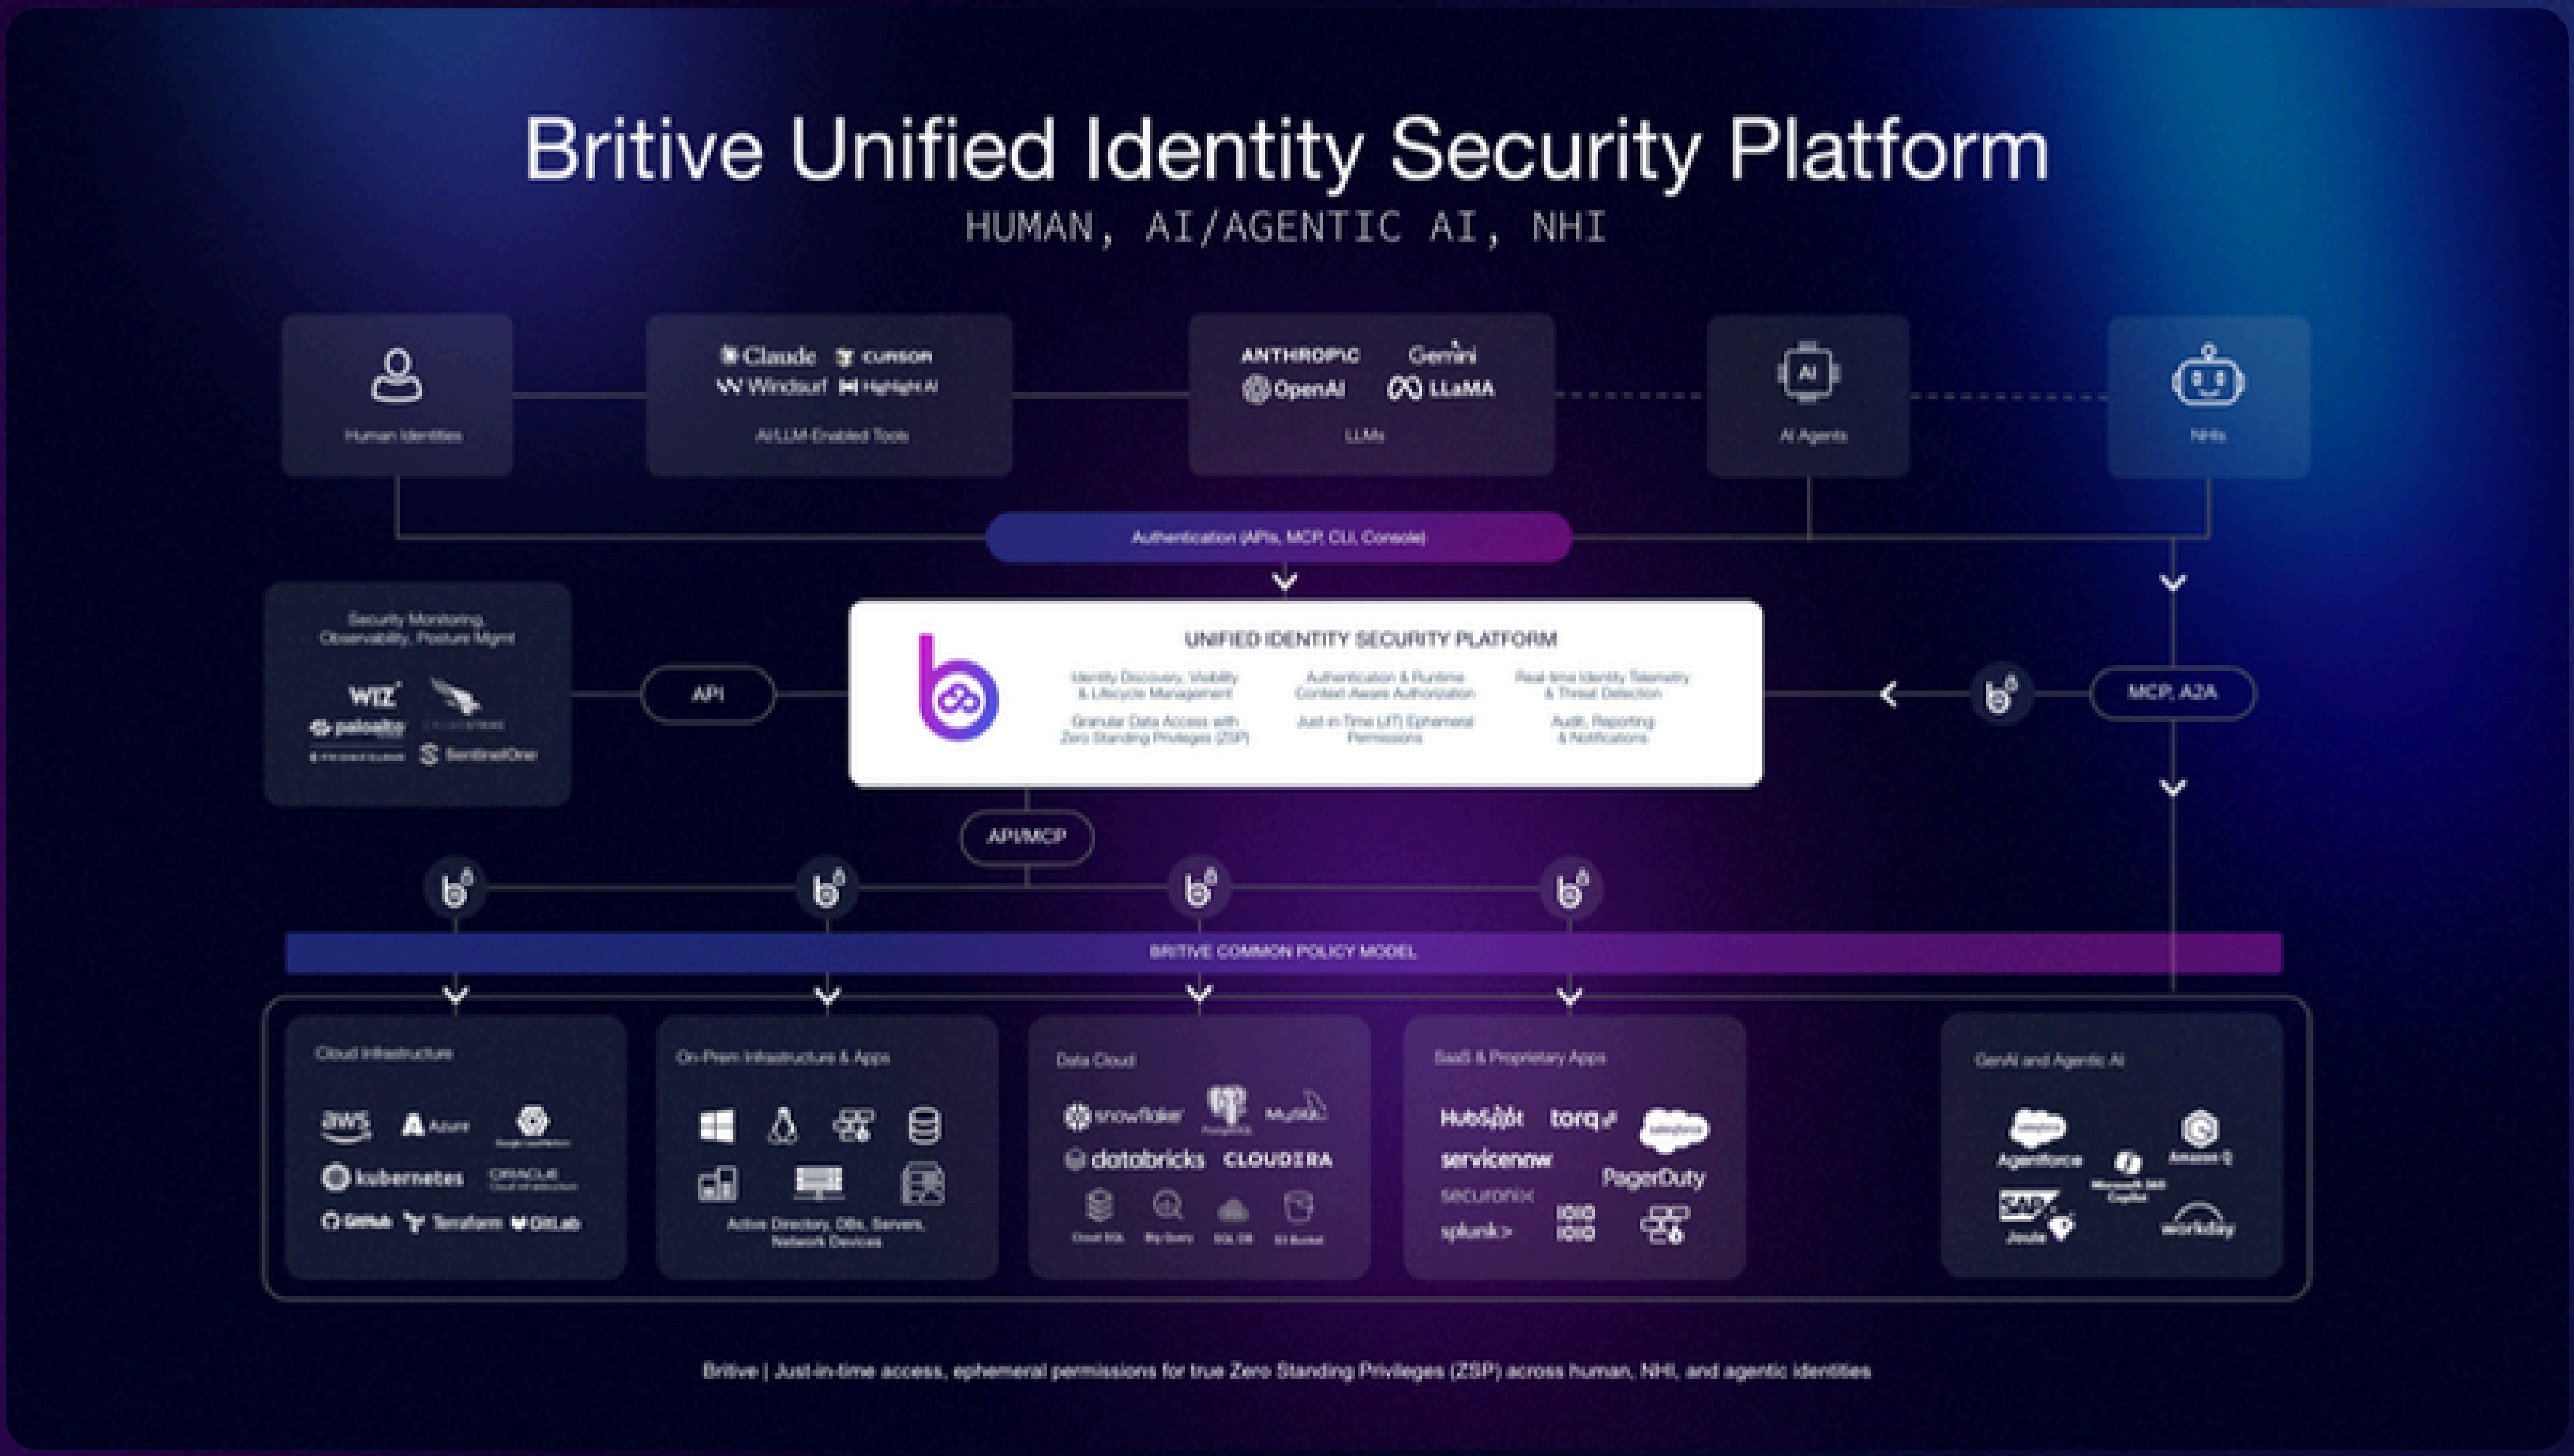

Britive Unified Identity Security Platform

HUMAN, AI/AGENTIC AI, NHI

Claude
CURSOR
Windsurf
Highlight AI
Human Identities
AI/LLM-Enabled Tools
ANTHROPIC
Gemini
OpenAI
LLaMA
LLMs
AI
AI Agents
NHIs
Authentication (APIs, MCP, CLI, Console)
Security Monitoring, Observability, Posture Mgmt
WIZ
paloalto
CROWDSTRIKE
SentinelOne
API
UNIFIED IDENTITY SECURITY PLATFORM
Identity Discovery, Visibility & Lifecycle Management
Authentication & Runtime Context Aware Authorization
Real-time Identity Telemetry & Threat Detection
Granular Data Access with Zero Standing Privileges (ZSP)
Just-in-Time (JIT) Ephemeral Permissions
Audit, Reporting & Notifications
MCP, A2A
API/MCP
BRITIVE COMMON POLICY MODEL
Cloud Infrastructure
aws
Azure
Google Cloud Platform
kubernetes
ORACLE Cloud Infrastructure
GitHub
Terraform
GitLab
On-Prem Infrastructure & Apps
Active Directory, DBs, Servers, Network Devices
Data Cloud
snowflake
PostgreSQL
MySQL
databricks
CLOUDERA
Cloud SQL
Big Query
SQL DB
S3 Bucket
SaaS & Proprietary Apps
HubSpot
torq
salesforce
servicenow
PagerDuty
securonix
splunk>
GenAI and Agentic AI
Agentforce
Amazon Q
Microsoft 365 Copilot
SAP Joule
workday

Britive | Just-in-time access, ephemeral permissions for true Zero Standing Privileges (ZSP) across human, NHI, and agentic identities

Traditional PAM models rely on a linear cost curve, as scaling privileged access security requires substantially more spend on vaults, agent- and proxy-based infrastructure, and operational headcount. However, by leveraging a cloud-native architecture managed via Infrastructure-as-Code (IaC) and Policy-as-Code (PaC), Britive decouples security enforcement from asset expansion.

This allows security teams to dynamically mint permissions at runtime for all identity types, maintaining consistent control across global environments without the compounding operational debt of vault-based systems that rely on static infrastructure.

DETAILED CAPABILITY ANALYSIS

Profiles as Portable Policy Containers

Britive's capability goes beyond just defining a set of permissions for access. Each access profile is a reusable and portable construct that an administrator creates once and deploys across diverse environments, including cloud infrastructure, SaaS applications, and on-premise resources, without the need to rewrite entitlement logic for every system.

These profiles encapsulate the necessary permission set, such as:

- S3 access for a specific account.
- Elevated Kubernetes access restricted to a particular cluster and namespace.
- Snowflake administrative roles limited to a certain warehouse and database/schema.
- Temporary Salesforce admin privileges restricted to specific roles or permission sets.
- Domain admin membership within Active Directory.

Furthermore, each profile includes a policy envelope, which incorporates:

- ✓ Configurable expiration with an automatic forced check-in.
- ✓ Requirements for human-in-the-loop approval.
- ✓ Segmentation based on identity type (e.g., humans, service identities, AI agents).
- ✓ Conditional controls.

This matters for hybrid and multi-cloud environments because, in enterprises spanning AWS, Azure, GCP, or legacy on-premises infrastructure, this model collapses operational complexity. A security team can enforce a single approval workflow for privileged access whether the target is an AWS IAM role, an Azure role assignment, a Google Cloud service account, a Kubernetes cluster, Snowflake, or a database behind the firewall.

The Access Broker technology extends the cloud-native Britive platform into private and on-premises zones, treating them as "resource profiles" governed by the same policy primitives. For example, a GitLab pipeline provisioning EC2, querying on-premises SQL, and writing to cloud-based MySQL illustrates how a single non-human identity can navigate the full stack through one authorization authority. Operationally, this model also benefits from rapid onboarding of target systems, enabling teams to bring cloud and application entitlements into scope and construct a common policy layer without deploying a complex network of connectors and session proxies.

Britive explicitly contrasts this approach with "basic just-in-time tools" that simply add a user to a group for a fixed duration and remove them afterward, with no intermediary policy layer or cross-system governance. Those tools inherit the limitations and inconsistencies of each target system. Britive's abstraction compensates for these gaps while still honoring native constructs such as AWS role assumption or Azure PIM. The result is heterogeneous infrastructure managed through homogeneous policy and a control plane that scales with organizational complexity rather than multiplying it.

Developer-Centric CI/CD Integration

Britive addresses developer friction, a frequent byproduct of legacy PAM, by embedding security controls directly into the DevOps pipeline through its PyBritive CLI and SDK. The platform's dynamic CLI tool is typically downloaded and executed as part of a CI/CD job and exists only for the duration of that job. The only setup required is configuring the pipeline to download PyBritive and execute the desired commands. The tool leverages OIDC and workload federation to authenticate without static API tokens or hardcoded secrets.

This "secretless pipeline" approach ensures access exists only for the duration of a specific task, effectively enforcing Zero Standing Privilege (ZSP) for automated workflows. Furthermore, Britive provides branch-level scoping, allowing security teams to restrict permissions based on specific repositories or environments, preventing accidental or unauthorized production changes during development cycles.

Unified Identity Security for Hybrid Environments

A core architectural pillar of the Britive platform is its Common Policy Model, which establishes a consistent security control layer across heterogeneous cloud, SaaS, and on-premises environments. By abstracting platform-specific entitlements, such as those found in Snowflake, SQL Server, or Active Directory, into a single framework, the platform allows teams to define fine-grained policies once and enforce them universally. This model applies the same Just-in-Time (JIT) and ZSP rigor to non-human identities as it does to human users, an essential capability given that machine identities often outnumber human identities by a factor of 40 to 50. To bridge the gap with legacy systems, Britive utilizes access broker technology to extend modern API-driven privilege controls into traditional on-premises infrastructure.

Agentic AI Identity and Security

Britive has positioned itself as an early entrant in the emerging category of Agentic Identity and Security Platforms (AISP). The platform manages the full lifecycle of autonomous AI agents, treating them as distinct identities requiring intent-aware, context-sensitive authorization. Key governance features include impersonation controls that restrict an agent's ability to act on a user's behalf to specific matching profiles, thereby containing risk even in scenarios involving prompt injection. For high-risk operations, the platform supports human-in-the-loop (HITL) workflows, enabling administrators to intercept and approve actions prior to execution. To ensure accountability in multi-agent systems, Britive establishes agent provenance and agent-to-agent trust using verifiable identity standards such as SPIFFE.

PRACTITIONER ASSESSMENT: WHY BRITIVE MATTERS

Blast Radius Reduction

By enforcing True ZSP, the platform ensures no privileged accounts exist to be compromised when not in use.

Native Role Assumption

Britive leverages native cloud constructs, such as role assumption in AWS, to grant ephemeral credentials. This ensures users do not maintain standing accounts or persistent "keys to the kingdom." By using direct API functionality, Britive can administer native services across AWS, Azure, GCP, and OCI without the operational overhead of deploying middleware or agents for services such as Lambda or compute.

> THE AGENTLESS, PROXYLESS DESIGN ALLOWS DEVOPS AND SRE TEAMS TO MAINTAIN SPEED WITHOUT ALTERING CORE WORKFLOWS.

Separation of AuthN and AuthZ

Britive allows organizations to use existing Identity Providers (e.g., Okta, Ping) for authentication (AuthN), while Britive manages the granular runtime authorization (AuthZ)

Agentic Readiness

Although still in early stages, Britive provides the runtime authorization framework needed to audit and govern agentic tasks.

Compliance & Audit

A single Unified Audit Trail provides a comprehensive record of all human, machine, and AI activities for audit and forensics.

Identity Intelligence and Analytics

Beyond raw logs, Britive provides reporting and analytics designed to turn identity and access data into measurable insight. This includes dashboards and reports to visualize privilege assignment and usage over time, detect potential identity and access risks, and track KPIs such as progress toward Zero Standing Privileges (ZSP), helping security teams produce cleaner audit evidence without manual data stitching.

THINGS TO WATCH

Britive's architecture is optimized for API-driven role assumption and workload federation, such as OIDC. While Britive uses an Access Broker to extend controls to on-premises Windows and Linux servers and some databases, it does not offer the same exhaustive out-of-the-box support for the long tail of legacy enterprise hardware and niche industrial (OT/ICS) protocols found in CyberArk's 350+ integrations. For organizations with a high percentage of non-cloud-native infrastructure that cannot support ephemeral access, Britive's reliance on its Universal Secrets Manager may feel less robust than a specialized legacy vault.

Additionally, Britive's session monitoring is runtime-centric and designed to be lightweight and frictionless for developers. However, it does not focus on endpoint PAM use cases, such as securing local administrator rights across thousands of corporate laptops, nor does it provide the same level of granular, keystroke-level isolation for RDP or SSH sessions that proxy-based PAM platforms like CyberArk deliver for maximum compliance.

CONCLUSION

For IT and identity security practitioners operating cloud- and developer-centric workflows, Britive represents a shift from static security to dynamic authorization. It is best suited for cloud-forward organizations where developer productivity and the rapid scaling of non-human and AI identities make traditional vault-based PAM a bottleneck. Britive should be viewed not just as a tool for securing administrators, but as a foundational control layer for modern identity infrastructure built on abstraction and runtime policy enforcement.

REPORT CONCLUSION

Privileged Access Management is shifting from a credential-centric control to a runtime authorization layer for modern identity security and increasingly, for agentic systems. Recent consolidation (most notably Palo Alto Networks' acquisition of CyberArk) signals that privileged access is now viewed as a foundational control plane, not a standalone compliance product. Platform vendors are prioritizing privilege because it is where identity becomes enforceable and where breach containment is ultimately decided.

This transition is being driven by two structural forces.

- ↳ First, cloud ephemerality has made standing privilege operationally and defensively unsustainable: infrastructure changes continuously, access needs are time-bound, and static roles cannot keep pace without creating excessive risk or friction.
- ↳ Second, the fastest growth in privileged users is no longer human. Non-human identities already operate at machine scale, and agentic AI compounds the risk by introducing autonomous, multi-step execution that can amplify errors or adversarial manipulation at production speed. In this environment, privileged access maturity becomes a prerequisite for safe agentic adoption, not a downstream enhancement.

The vendor landscape highlighted in this report points to a clear direction of travel: modern PAM is moving toward Zero Standing Privilege (ZSP), policy-driven just-in-time access, and deeper integration with identity and infrastructure signals. Apono, Britive, P0 Security, and Teleport represent distinct architectural approaches to the same market requirement: privilege must be provisioned dynamically, scoped narrowly, verified continuously, and audited comprehensively across human, machine, and emerging agentic identities. Legacy vault-based PAM will remain durable in hybrid enterprises, but it is increasingly complemented by authorization-centric layers that better match cloud and automation realities.

Looking ahead, the outlook is constructive. As organizations modernize privileged access, they will not only reduce blast radius and credential sprawl, they will build the governance foundation required to safely unlock automation and agentic workflows. The next phase of identity security will be defined by how effectively enterprises can convert privilege from a static entitlement into a measurable, adaptive, and continuously enforced control system.

To see how this model applies in your environment, request a demo:

Find a time that works for you here.

This summary is drawn directly from Francis's analysis. You can find the full report here: https://softwareanalyst.substack.com/p/the-evolution-of-the-privileged-access