CYBERARK®
The Identity Security Company™

WHITEPAPER

The Backbone of Modern Security: Intelligent Privilege Controls™ for Every Identity

Table of Contents

Introduction

If you could look across your IT environment on any given day, what would you see? An IT administrator maintaining servers or configuring cloud services. A finance worker processing payments. Maybe your chief revenue officer reviewing customer account data. A developer writing code. And all your cloud workloads — what exactly are they up to?

All of these identities — employees, IT admins, developers, machines — have two things in common: 1) they comprise your modern workforce and 2) they need access to applications and services to do their job. Your job is to ensure each identity is authenticated and authorized, ideally by following the principle of least privilege (PoLP). Then, depending on their role and permission level, security controls must be applied in line with the risk of their access, with minimal friction to their user experience.

If only it was so simple.

Vaulting, session isolation and the ability to control freestanding privileged accounts, aka secure standing privileged access management (PAM), remain critical parts of security. But an evolving threat landscape calls for a new approach that can isolate and stop attacks, protect critical assets and grant right-time, right-level access.

In this whitepaper, we'll examine how attack paths have changed, the unique security challenges of each identity group and how intelligent privilege controls provide the backbone of a modern identity security strategy that can protect everything your enterprise is building.

New Identities, New Attack Methods, New Environments

Research tells us that the number of human and machine identities is predicted to more than double this year, expanding all possible attack points. Up to 50% of machine identities have access to sensitive data.[1] While non-human identities are widely viewed as the riskiest type, nearly two-thirds of organizations fail to recognize them as privileged users.[2]

Organizations that only protect traditional human identities or IT admins are not minding this widening gap. Gone are the days when only the most privileged users, typically IT admins, had access to critical systems and sensitive data. Today, all identities — IT admins, developers, workforce and machines — can potentially become privileged or high-risk according to what they can access. Each has the power to be the gatekeeper or a gateway to a breach.

While organizations face proliferating identities and requests for elevated access, they must also navigate the impact of new environments and new AI-driven attack methods. The number of SaaS applications will likely grow. Organizations are eyeing their third- and fourth-party risks with increasing alarm. It's no wonder 93%[3] of organizations reported two or more identity-related breaches in the last year.

Intelligent Privilege Controls: the Baseline of Identity Security

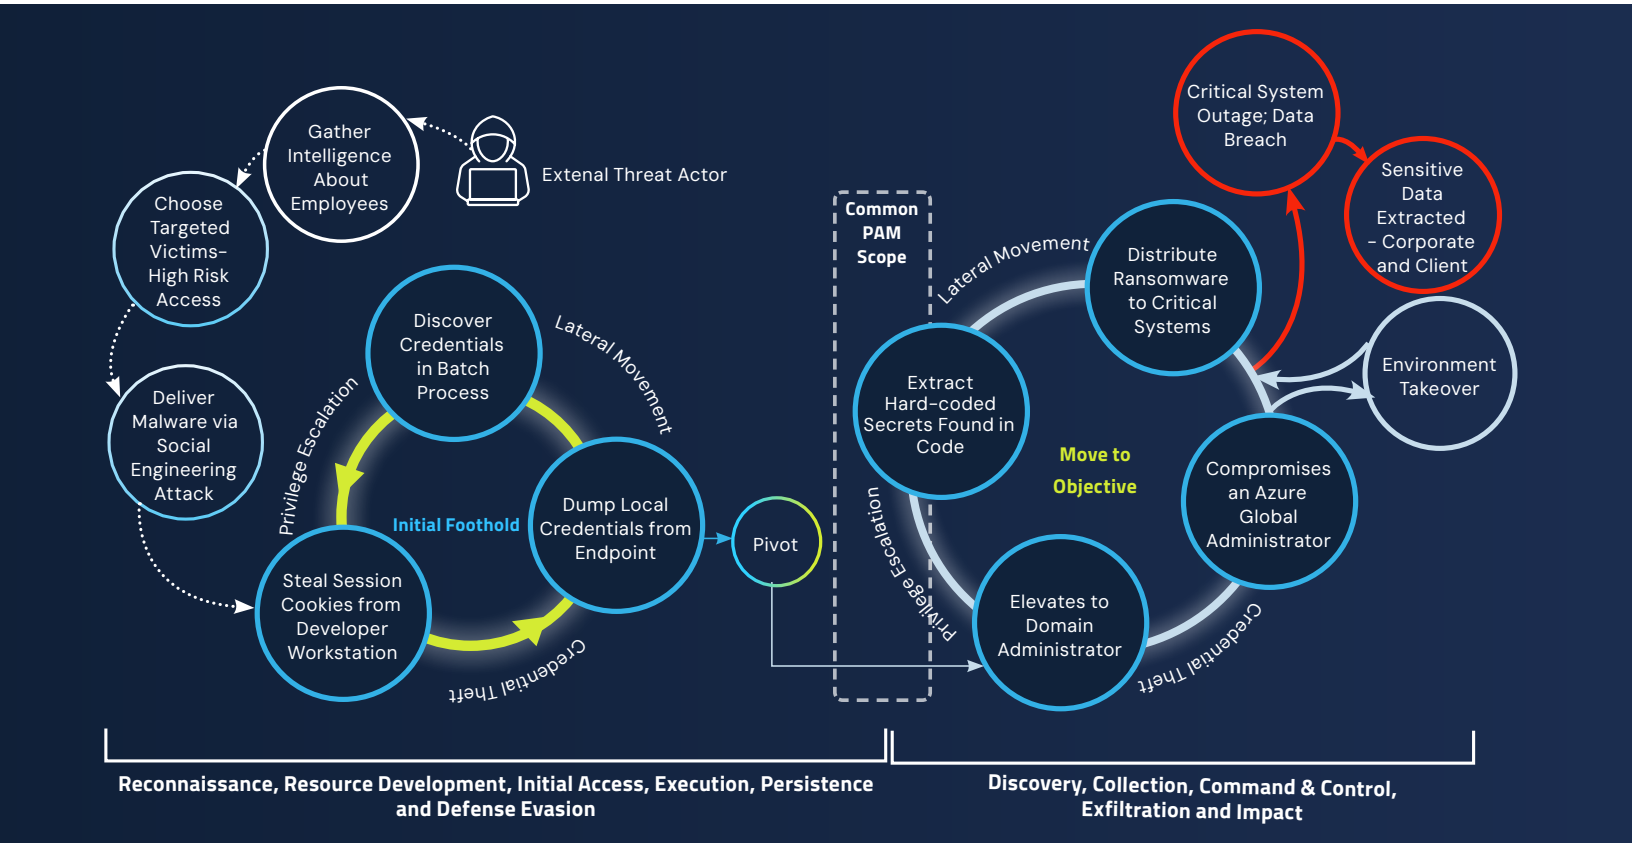

Despite novel methods and new environments, the attack path stays the same.

[1,2,3] CyberArk, " 2024 Identity Security Threat Landscape Report," May 2024.

Every identity, human or machine, needs the right level of privilege controls to be secure. But as you assess the spectrum of identities, it's imperative to consider the complexity of each identity's role, the complexity of the environment and the risk of access.

This brings us to two essential questions: **When are basic security controls enough and when are heightened adaptive security controls required?**

No matter how complex your environment, how sophisticated an attack is, or whose identity is targeted, bad actors follow a predictable plan of attack: compromise identities, move laterally across the environment and escalate and abuse privileges.

Privileged access does not begin or end with the traditional user anymore. Every identity has the potential of becoming a privileged user or accessing some sort of elevated privileged data at some point in their day. An identity security-based strategy helps organizations clearly define and manage who has access to which resources and under what conditions. It keeps them in front of which identity is accessing what data — and what it's doing with it.

> Every identity within your organization needs to be secure. The most effective minimum controls must adapt and correlate to the risk posed. And, no matter how risk fluctuates, user experience must be consistent.

Intelligent privilege controls are security measures that organizations can apply dynamically to protect an end user's access in a high-risk session. They include an important but often missing piece of effective risk assessment: **user expectation.** While every identity within your organization needs to be secure, not every employee needs or expects to jump through multiple hoops to access basic HR functions. Each security practitioner must decide how much risk they're willing to tolerate for each identity type, the type of resource they're accessing — and how much friction that user will tolerate while doing their job. One-size-fits-all constraints frustrate users. We believe the most effective minimum **controls should adapt and correlate to the risk posed.** And, no matter how that risk fluctuates, user experience must be consistent.

In a threat landscape where every identity can potentially become privileged under the right circumstances, organizations need to establish an identity security baseline with intelligent privilege controls that correlate to risk.

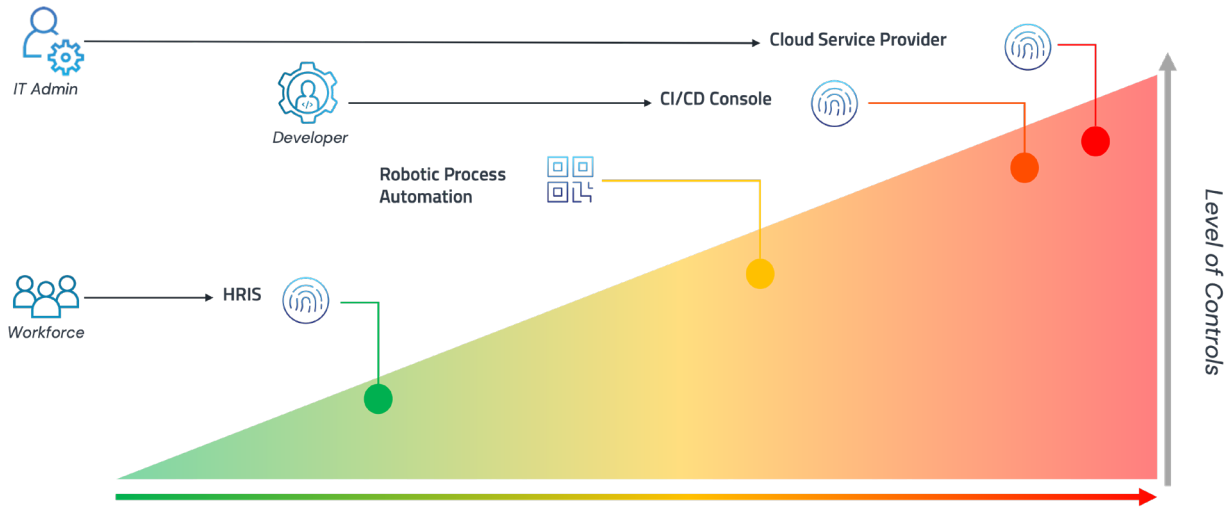

Different roles, privileges and resources require different levels of intelligent privilege controls.

Common Factors That Influence Identity-Related Risk

Risk can be defined in many ways by different organizations. Here, we're referring to the sum of three factors that are common to all organizations and widely exploited by malicious actors:

- **Level of privilege** refers to the type of privileged actions the identity can perform against a given resource.
- **Scope of influence** (also referred to as the blast radius) refers to the number or percentage of systems an identity or account can access, either directly or indirectly. The larger the scope of influence, the higher the level of risk the identity poses.
- **Ease of compromise** refers to how easy or challenging it is for a malicious actor to compromise an identity.

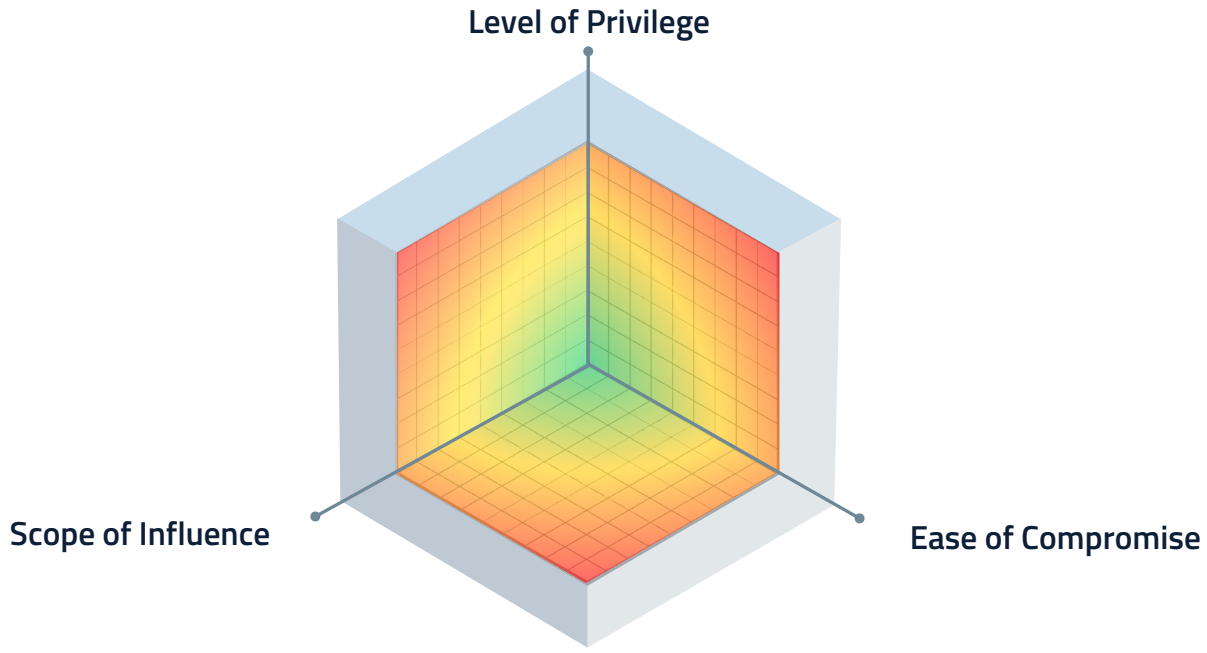

The larger each of these values, the larger the impact. The larger the impact, the larger the risk. The larger the risk, the larger the risk mitigation by implementing intelligent privilege controls.

Holistic security encompasses all the identities that access enterprise resources, whether operational (personal), system (built-in local admins and break-glass accounts) or machine (non-human workload) access. You can't call a specific enterprise resource secured until all identities that access it are secured.

Organizations only have so much time, effort and resources to achieve their identity security goals. Let's look more closely at the emerging security challenges posed by each identity group and how organizations can evolve their identity security strategy and provide the right level of intelligent privilege controls for the right situation.

Workforce: Modern Attacks Require a Fresh Approach

From employees to external vendors and everything in between, your workforce includes a host of human identities working from both managed and unmanaged endpoints. Each one cycles through a range of access to sensitive corporate resources. Organizations rely on multi-factor authentication (MFA), single sign-on (SSO) and identity threat detection and response (ITDR) to secure workforce access.

However, with thousands of new ransomware variants constantly entering the wild, attackers find new ways to turn off or bypass endpoint detection and response (EDR) by abusing privileged credentials.

Even if users don't have administrative rights on an endpoint, they might have web access to sensitive data, line-of-business (LOB) applications, infrastructure management consoles, etc. The attacker can sidestep MFA and target the employee's browser to steal passwords stored there and cookies that allow for web session hijacking. Or, let's say a user attempts to log into an app containing the company's financial records outside working hours, using an unknown device, from a region the employee never travels to. Legacy MFA and SSO can't detect this anomaly.

WORKFORCE: WHO ARE THEY?

Workforce identities include contractors, employees and application admins. As non-IT users, they don't have access to highly sensitive resources by default but can become privileged under certain circumstances.

Target Resources

Endpoints, data and business applications.

The Challenge

Traditionally, a combination of MFA and SSO is used to secure access for workforce identities. But new attacks like MFA bombing and cookie theft can bypass MFA and gain access to your high-risk HR systems. A more dynamic, risk-relevant authentication is required.

The Solution

By layering an AI-powered user behavior analytics (UBA) engine with an adaptive MFA, security teams can authenticate based on insights from the user's login history. UBA can flag suspicious logins, ask for a phishing-proof physical verifier (along with regular credentials) and, if not passed, terminate the session. Continuous session monitoring can ensure that access is minimal and abnormal activities are detected and contained.

NEW ATTACKS ON YOUR WORKFORCE

In September 2023, the Scattered Spider ransomware group used an MGM employee's stolen credentials to launch a major cyberattack. Instead of the tried-and-true malware methods, the attackers used generative AI to voice phish (or 'vish') employees and help desk teams to steal credentials and reset MFA settings.[4]

In October 2023 it happened again: attackers compromised the personal laptop of an employee at Okta.

Even ITDR agents deployed on all endpoints wouldn't have helped. These attacks circumvented the endpoint, relying on social engineering in a medium (phone calls) where detection-focused security tools had no visibility or context.[5]

[4] CyberArk Blog, "The MGM Resorts Attack: Initial Analysis," September 2023.

[5] CyberArk Blog, "Piecing Together the Attack on Okta's Support Unit," October 2023.

IT Runs on Privileged Access. So Do Cyberattacks.

Your IT admins need powerful access rights to do their jobs. Cloud computing has fundamentally transformed — and complicated — the environments they must manage. Their purview not only includes long-lived systems such as data centers, OT environments and lift-and-shift infrastructure but also elastic cloud workloads, virtual machines (VMs) and databases.

IT admins are also charged with configuring cloud provider services powering customer-facing applications, as well as SaaS apps hosting sensitive data like payroll systems. Meanwhile, every endpoint, server and networked device in an organization has a built-in local admin account with privileged access.

IT security teams also rely on SaaS security tools such as SentinelOne, CrowdStrike and Splunk to automate high-volume tasks like scanning, patching and threat detection. Access to these systems is highly privileged — and highly coveted by bad actors since they contain the keys to the kingdom.

As organizations mobilize to protect themselves against these complex threats, 94% report using more than 10 vendors for their identity security initiatives.[6] Worse, new-fangled tech can introduce new risks by failing to secure new or evolving identities and IT environments, requiring specialized skills and labor for integration) and driving up costs.

Regulators, auditors and cyber insurance providers are well aware and are cracking down on privileged access controls, in particular domain admin accounts, service accounts and local admin accounts on employee laptops. Compliance frameworks and regulations (PCI DSS, SOX, NIST, EU NIS 2, NERC-CIP, HIPAA, GDPR, CCPA and SWIFT CSCF) have increased their requirements to control, manage and monitor privileged access — and organizations that don't clear the bar could lose their ability to service key customers or incur steep fines.

IT ADMINS: WHO ARE THEY?

IT admins include traditional IT, third parties and CloudOps. They have widespread access to business-critical applications, cloud servers and sensitive infrastructure — and are the most targeted identities in any enterprise.

Target Resources

Data, apps, workloads, cloud consoles and CLIs.

The Challenge

Since IT admins work across on-premises, cloud and hybrid environments, they often have long-standing access to sensitive enterprise resources to ensure operational ease. Attackers can exploit this to gain access to high-risk resources.

The Solution

Security teams need a holistic set of privilege controls that protect zero standing privileges and standing access across high-risk SaaS, cloud-native services, elastic cloud workloads and long-lived systems. This provides access for just the amount of time required, reducing the attack surface and maintaining tight security without increasing admin workloads.

[6] CyberArk, " 2024 Identity Security Threat Landscape Report," May 2024.

Many organizations use a combination of on-premises, operational technology (OT) and public cloud environments. Most (87%)[7] have chosen a multi-cloud strategy and need to secure access to multiple cloud service providers. As the lines between traditional users and administrators keep blurring, operational and system accounts require advanced controls to prevent the misuse of permissions and to maintain clear lines of responsibility. All cloud services providers recommend federated access — not shared accounts — for operational access to workloads and services. Best practices like providing access with Zero Standing Privileges have emerged, but many organizations have yet to adopt them. As a result, IT admins see inconsistent logging standards that complicate user activity tracking.

Cyber insurance premiums continue to climb: almost 42% of companies[8] paid more in 2023. Specific requirements include use of administrative privileged access, including secure remote access and strong MFA.

NEW ATTACKS ON PRIVILEGED ACCESS

In 2022, attackers gained access to Uber after purchasing workforce user credentials on the dark web. Once inside, they discovered PowerShell scripts on a network drive that contained the "keys" to the organization's PAM service.[9] From there, they pilfered all the organization's secrets, including the administrative accounts to their domain, MFA provider, identity provider and cloud providers.

In 2023, after a slow but stealthy password spray attack, nation-state attacker Cozy Bear accessed a legacy, non-production test tenant account in Microsoft's environment.[10] Overly permissive trust relationships between non-production and production environments allowed the attackers to escalate privileges and move laterally across systems all the way to executive emails. Inadequate auditing and logging meant malicious activities from compromised accounts weren't detected quickly.

Both attacks demonstrate that need to extend intelligent privilege controls not only to IT admins, but also to additional employees with privileged access such as engineers, third-party vendors and developers. These identities are increasingly under attack in emerging breaches.

[7] Flexera Blog, "Cloud Computing Trends: Flexera 2023 State of the Cloud Report," April 2023.

[8] CyberArk, " 2024 Identity Security Threat Landscape Report," May 2024.

[9] CyberArk Blog, "Unpacking the Uber Breach," September 2022.

[10] CyberArk, "Ep 46: Behind the Data Breach - Dissecting Cozy Bear's Microsoft Attack," February 2024.

Let Me Do My Job: Developers vs. Security

An entire new population with 'admin rights' has appeared: software developers. Developer identities live largely in the back end, accessing codes, cloud services and workloads. A staggering 77%[11] of IT security decision-makers agree that developers have way too many privileges, making them highly attractive targets for attackers.

Why are developers overprivileged? For the sake of expediency. They need to innovate; they can't constantly put their work on hold while the right permissions are sorted out. If you must tell your developers, for example, that they need to use an external credential source rather than a cloud-native secret key store (Azure Key Vault, AWS Secrets Manager, etc.) they may revolt. This creates unacceptable delays to engineering timelines and deviates from cloud provider best practices.

But giving developers unfettered administrative rights and excessive permission can lead to attacks like the one Microsoft endured earlier this year. Developers likely had configured applications with elevated privileges or had not periodically audited their permissions.

Developers with privileged access must be protected. This is especially true as organizations accelerate digital transformation projects and develop their own software applications.

DEVELOPERS: WHO ARE THEY?

The engineering, data scientist, site reliability engineers (SREs) and DevSecOps teams are the heart of the enterprise with access to data and workloads in on-premises, cloud and hybrid environments.

Target Resources

Data, apps, workloads, code repositories and cloud services.

The Challenge

Sensitive developer environments are often not fully secured. If this identity is compromised, bad actors can move laterally, escalating privilege.

The Solution

Secure access for hybrid and multi-cloud environments, based on the principle of least privilege and zero standing privileges, allows data scientists and cloud engineers to work natively and efficiently.

Seamless integration with DevOps workflows enables dynamic, just-in-time access that doesn't slow productivity. Secrets management should sync with native tools (e.g., AWS Secrets Manager) for secure retrieval without changing workflows.

Monitoring and analytics should create audit trails and behavioral analytics to identify abnormal usage patterns and provide compliance transparency.

[11] CyberArk, "Identity Security Threat Landscape Report," June 2023.

Machine Identities: Cloud Misconfigurations and Shadow Admins

Bots, applications, IT devices, workloads and even modern code — machine identities live at the front lines of where threat actors are attacking — and are perhaps the most difficult identities to secure.

Ongoing digital transformation and pervasive cloud computing are driving an exponential number of services and applications and all of these need to be discovered, managed, secured and automated to keep their connections and communications safe. Access logic, permission sets, resources, capabilities and risks are much more difficult to manage.

The public cloud is a great enabler for business: cloud architecture is cheaper and it scales. Developers can get full-stack applications up and running within seconds. In fact, organizations report that **84% will leverage more than three cloud service providers (CSPs) in the next 12 months.**[12] But every new service rolled out by CSPs comes with many new roles and entitlements. Often, existing identity controls may not align seamlessly with the new roles, necessitating integration. In the name of speed and convenience, machine identities become over-permissioned beyond what their roles require and often have powerful standing access. Incredible efficiency gains can be eclipsed by massive vulnerabilities as hackers now have thousands of new ways to gain footholds within an organization.

Machine identities use secrets (username/password, SSH keys, API and access keys, tokens, certificates, etc.) to authenticate and authorize themselves when accessing cloud environments, databases, APIs and other systems. But organizations need protections beyond just secrets management. They need an end-to-end machine identity security management solution that can ensure the proliferation of machines are secured in a cloud-first, GenAI, post-quantum world. Traditional on-premises environments had strict but simple roles for Windows, Linux, domain and database administrators. You could always draw a clear line at each infrastructure layer, from the physical data center, physical network cables and physical racks — all the way up the IT stack. This is no longer the case.

MACHINES: WHO ARE THEY?

Devices, data, apps, infrastructure and code including service accounts, serverless functions like AWS Lambdas, etc.

Target Resources

Devices, data, apps, infrastructure and code.

The Challenge

Security teams often lack visibility into how many applications and secrets are deployed across various cloud environments, making it hard to identify which machine identities pose potential risks.

The Solution

An end-to-end machine identity security management solution must protect workload secrets across the full application spectrum, from dynamic and ephemeral to static and monolithic, including homegrown code and commercial software. Centralized secrets management and rotation deliver consistent vaulting, discovery and rotation. Developer-friendly PKI and certificate management functionalities that deliver seamless integration and reporting capabilities.

NEW ATTACKS ON YOUR MACHINES

Attackers were able to infiltrate SolarWinds[13] software development environment and introduce malicious code into the Orion software update. This code then utilized the compromised machine identities to distribute the malware to well over 18,000 SolarWinds customers.

A misconfigured Identity Provider (IdP) was also to blame in the MGM breach in 2023, which cost the resort giant well over $100 million. Attackers used the vulnerability to pivot into MGM's infrastructure in the Okta tenant. The IdP solution granted the attackers highly privileged access to Azure, which ultimately allowed the cloud-originated attack to reach MGM's brick-and-mortar operations.[14]

[12] CyberArk, " 2024 Identity Security Threat Landscape Report," May 2024.

[13] CyberArk Blog, "The Anatomy of the SolarWinds Attack Chain," February 2021.

[14] CyberArk, "Ep. 39: Analyzing the MGM and Okta Breaches – The Identity Connection," November 2023.

How To Put Intelligent Privilege Controls Practice

CyberArk pioneered privileged access management (PAM) for IT, an essential security layer on every CISO's security checklist. Hybrid environments, novel attack methods and new identities with privileged access require a new security paradigm — one that enables organizations to discover, secure and manage identities across the entire lifecycle. The backbone of this approach is intelligent privilege controls that correlate to risk.

Empathy to the End User

Every user expects some sort of constraint on their access. Yet as we raise the level of control, the user's expectation of constraint will push against that. The more controls you put in place, the harder it is for a user to do their job. What is the right amount of constraint that we can put in place without making anyone's job unduly difficult?

Native access for users is essential — leading identity security platforms will apply privilege controls within a user's expected workflow, whether that's a web-based session, a session through an RDP or SSH client or a CLI to AWS, Azure or GCP.

Let's examine five privilege controls that can reduce risk, get out of the user's way and give you the best odds for success.

Credential Vaulting and Rotation

Even in the cloud, business-critical systems come with a built-in administrative account. Admins need these root accounts to stand up services in case an SSO connection is lost or Active Directory is down. It's imperative that businesses protect these keys to the kingdom in a vault with regularly rotated credentials. This is why secure standing privilege controls are — and always will be — essential to identity security.

Zero Standing Privileges (ZSP)

But in today's threat landscape, we have to protect not only the keys but also all the different types of access to your kingdom. Zero standing privilege controls combine just-in-time access and least privilege so organizations can dynamically elevate user rights as needed and then quickly revoke them.

Session Isolation and Monitoring

Malware spread is another privileged access concern. Organizations can provision isolated privileged sessions to workloads to prevent the spread of malware – without sacrificing the native user experience. Meanwhile, credentials can be kept secure and not exposed to end users or their endpoints, reducing the risk of credential theft and lateral movement.

On any given day, an administrator might need operational access to a target system with a vaulted account. But what if a developer or a database admin needs access to that same system?

Rather than defaulting to a standard privilege session management server connection, a SaaS connection can deliver right-time, right-level access without burdensome infrastructure. Organizations can elevate permissions for a vaulted account with a rotating password and shift them to match the right level for a given user's session. Then, once complete, the trail of breadcrumbs gets wiped away.

Step up and Continuous Authentication

With context-aware workflows using machine and user behavioral learning, MFA can be everywhere but only prompt the user when necessary. Intelligent controls can organize the information, measure the level of risk to user experience, correlate it with risk and prompt for more authentication.

Likewise, if a user needs to make an elevation request, triggers are automatically built in, freeing up your security team to focus on more important risks.

Identity Threat Detection and Response (ITDR)

With the tsunami of new identities, our security tools must be able to do more than just search the data — we need to talk to it. ITDR can help security teams establish a baseline of normal activities so they can rapidly identify deviations, such as requests for unusual access levels within your CSPs. Analytics powered by generative AI can auto-recommend appropriate policies for potential threats, including malicious insider activities or compromised identities. A centralized system can also streamline the audit process, enable quick access to data on user activities and simplify compliance with regulatory demands.

How Do I Get There From Here?

Check out the CyberArk Blueprint for Identity Security Success for strategic guidance across people, process and technology domains. Get simple and prescriptive guidance aligned to the latest digital transformation initiatives and risk-based advice to accelerate the development of an effective and mature identity security program.

For more than two decades, CyberArk has been securing the workforce. Today, tools like MFA and SSO have evolved to embed session control and isolation, step-up authentication and deliver an end-to-end passwordless experience from the moment you log into the endpoint to the end of the workday. We've reinvented the entrance point — the browser — where most of the workforce spends their day. And we've unified all of those capabilities in the CyberArk Identity Security Platform.

Our integrated identity security platform is built to protect every identity (human and machine) and every resource they touch, across every one of your cloud or on-premises environments. We believe it will help all of you reimagine the way we approach securing the workforce in this new threat environment.