

Idira Identity Security Blueprint

Achieve More Impactful Identity Security
Outcomes, Faster

Executive Summary

In today's AI-powered, multicloud, SaaS world, identity is the new perimeter. Physical and network barriers have dissolved, and all identities can be an attack path to an organization's most valuable assets. This means that identity security has become the backbone of zero trust. Organizations must strengthen the security of all identities, which is well known among many IT and security teams. However, implementing an effective identity security program that is inclusive of human, machine, and AI identities is a challenge for many organizations, as the identity landscape is large, complex, and continuously evolving.

Idira™ Identity Security Blueprint (Idira Blueprint), by Palo Alto Networks, helps organizations assess and prioritize identity-related vulnerabilities, strengthen security, and reduce risks. This comprehensive resource leverages our vast experience and deep subject-matter expertise, laying out a prescriptive, risk-aligned best practice guidance for establishing and maintaining an effective identity security program.

In This Whitepaper

- Gain insights on common identity security challenges.
- Learn how to improve identity security systems and practices.
- See how to reduce identity security vulnerabilities.
- Explore how to mitigate risk across human, machine, and AI identities.

Get the Most from Idira Identity Security Blueprint

Idira Identity Security Blueprint is a collection of resources you can use in various ways—with a specific purpose tailored to what you want to achieve.

1. Understand the Attack Chain

Leverage Idira Identity Security Blueprint to better understand the identity attack chain. Learn how malicious actors exploit your human, machine, and AI identities to execute their end-game, and how you can prevent those types of attacks. [Learn More](#).

2. Assess Your Security

Security assessments are vital to information security. Learn how to accelerate your security efforts, identify weaknesses and security control gaps, and gain insight about your security posture and how to prioritize next steps. [Learn More](#).

3. Build Your Roadmap

A roadmap serves as the foundational guide for your security initiatives, enabling you to execute on a specific plan and proactively target areas of security deficiency. Accelerate the creation of a prioritized, risk-based plan for success. [Learn More](#).

4. Best Practice Education

Successful programs are about more than good security tools. They require the right people, process, and technology guidance. Take the guesswork out of planning and delivery with practical best practice guidance. [Learn More](#).

Introduction

Identities represent the number-one attack vector for organizations today. Unit 42® found 89% of their investigations over the past year showed that identity weaknesses played a material role in the breach.¹ Further research backs up this claim, finding that 9 out of 10 organizations have faced a successful identity-related breach over the last 12 months.² Attackers find identities attractive because they exist throughout an organization's entire IT spectrum and are integral to authenticating and authorizing access to an enterprise's sensitive data, business processes, and systems.

1. *Unit 42 Global Incident Response Report 2026*, Palo Alto Networks, February 17, 2026.

2. *2026 Identity Security Landscape*, Palo Alto Networks, May 11, 2026.

The increase in the complexity of hybrid and cloud environments and the innovative AI-powered attacks used by bad actors have all led to a dramatic increase in the threat landscape. Every organization has a spectrum of identities across the entire enterprise—humans, machines, and AI identities—and the number of identities is increasing at an exponential rate, for example:

- **Human users seek access** to endpoints, applications, data, cloud services, DevOps, and AI tools—spanning sensitive data to mission-critical access.
- **Machine workloads and automation need high privileges** to access various corporate resources so they can perform their tasks and improve organizational efficiency.
- **AI agents and nondeterministic identities are built to navigate across corporate resources** to execute complex tasks, often operating with broad, delegated authority.

All these identities and their access represent varying degrees of risk to the organization and, therefore, require the appropriate level of intelligent privilege control. IT and security teams can overcome these challenges and minimize the growing risks tied to identities by:

- **Taking a close look at how attackers exploit privileged identities:** Ask: What are the most common identity and privilege attack vectors? And, how does the perpetrator think and behave in each case?
- **Taking a practical, phased approach to identity security:** Identify the most-sensitive identities and their related access, credentials, and secrets. Zero in on identities that could jeopardize mission-critical infrastructure or expose confidential data.
- **Developing a prioritized plan to reduce vulnerabilities and strengthen security:** Consider your priorities: the most important actions, the items that can be achieved quickly and with minimal resources, and the ones that require significant time and effort.
- **Continuously reassessing and improving the identity security plan** to address evolving threats and new technologies.

Idira Identity Security Blueprint Helps Reduce Security, Compliance, and Operational Risks

We have developed a prescriptive blueprint framework to help organizations establish and evolve an effective identity security program. Idira Identity Security Blueprint provides a comprehensive, prioritized security framework that closely aligns identity security initiatives with potential risk reduction, helping organizations address their greatest liabilities as quickly as possible.

Core Identity-Centric Risks

Identity has a broad impact across the enterprise, including three core identity-centric risks.

Security Risk

Failure to secure identities (human, machine, and AI) from unauthorized access, abuse, and misuse or compromise. Examples include identity compromise, lateral and vertical movement, and privilege escalation and abuse. The consequences of poor mitigation include data exfiltration and breaches, ransomware attacks, service disruption, loss of intellectual property, and erosion of customer trust.

Operational Risk

Inefficient identity processes that hinder the agility, reliability, or scalability of IT and security operations. Examples include manual access provisioning, inconsistent access reviews, and orphaned accounts. The consequences of poor mitigation include downtime and service disruption, increased help desk workload, slower onboarding and offboarding, and higher total cost of ownership (TCO).

Compliance Risk

Failure to meet regulatory, audit, and internal governance requirements tied to identity and access. Examples include incomplete audit trails, inadequate segregation of duties, and unattested privileged access. The consequences of poor mitigation include regulatory fines, failed audits, legal exposure, and brand and reputational damage.

When organizations mitigate security, operational, and compliance risks effectively, they build identity resilience—the ability to withstand, adapt, and recover quickly from cyber incidents, process breakdowns, and compliance failures.

Our guidance covers the full spectrum of identities, including human, machine, and AI identities to mitigate these common identity risks. Idira Blueprint assists in accelerating the practical implementation of intelligent privilege controls for these identities and their enterprise resources, an area where traditional cybersecurity frameworks (CSFs) fall short. Use the blueprint as a tool to help guide your own identity security program roadmap development in combination with your current state, internal priorities, and desired business outcomes.

Idira Blueprint consists of a broad ecosystem of self-service educational resources and collateral, including videos, whitepapers, toolkits, and success knowledge articles. Our blueprint is not built on theoretical advice. It's built on the collective experience of Palo Alto Networks and [Unit 42](#) battling threats in the identity security space. Each component provides best practice guidance across the people, process, and technology domains—all designed to help you achieve faster identity security outcomes.

A Recognized Leader in Identity Security

These insights are gathered from lessons learned across our global customer base, postbreach experience, frontline remediators and red-team, and cutting-edge researchers.

As a recognized leader, we are uniquely positioned to deliver a thorough and effective identity security plan:

- Our security solutions are trusted by over 70,000 customers worldwide, including more than 50% of the Fortune 500, across a wide range of industries, such as financial services, insurance, manufacturing, healthcare, and tech.
- Unit 42, our threat research, incident response, and security consultancy is front and center in helping companies recover from some of the largest breaches of the 21st century.
- Our Professional Services and Customer Success organizations have decades of real-world implementation and support experience. They also have a detailed, firsthand understanding of the risks present within human, machine, and AI identities and best practices.
- Leading research and advisory firms recognize CyberArk, a Palo Alto Networks company, as a leader across multiple identity security categories.³ They include access management, privileged access management, secrets management, nonhuman identity management, enterprise password management, identity threat detection and response, and identity fabrics.

Understanding the Identity Attack Chain

While every organization's IT environment is unique, perpetrators can attack virtually any organization by following well-established steps in the identity attack chain. Idira Identity Security Blueprint helps organizations strengthen their security posture by showing them how to think like an attacker. It also explains how to defend against the three common techniques adversaries typically use to access identities, steal data, take down systems, and execute their endgame: identity compromise, lateral and vertical movement, and privilege escalation and abuse.

3. *Gartner® Magic Quadrant™ for Privileged Access Management*, Felix Gaehtgens, James Hoover, Michael Kelley, Brian Guthrie, and Abhyuday Data, Gartner, September 5, 2023; *The Forrester Wave™: Workforce Identity Platforms, Q1 2024*, Forrester® Research, March 20, 2024.

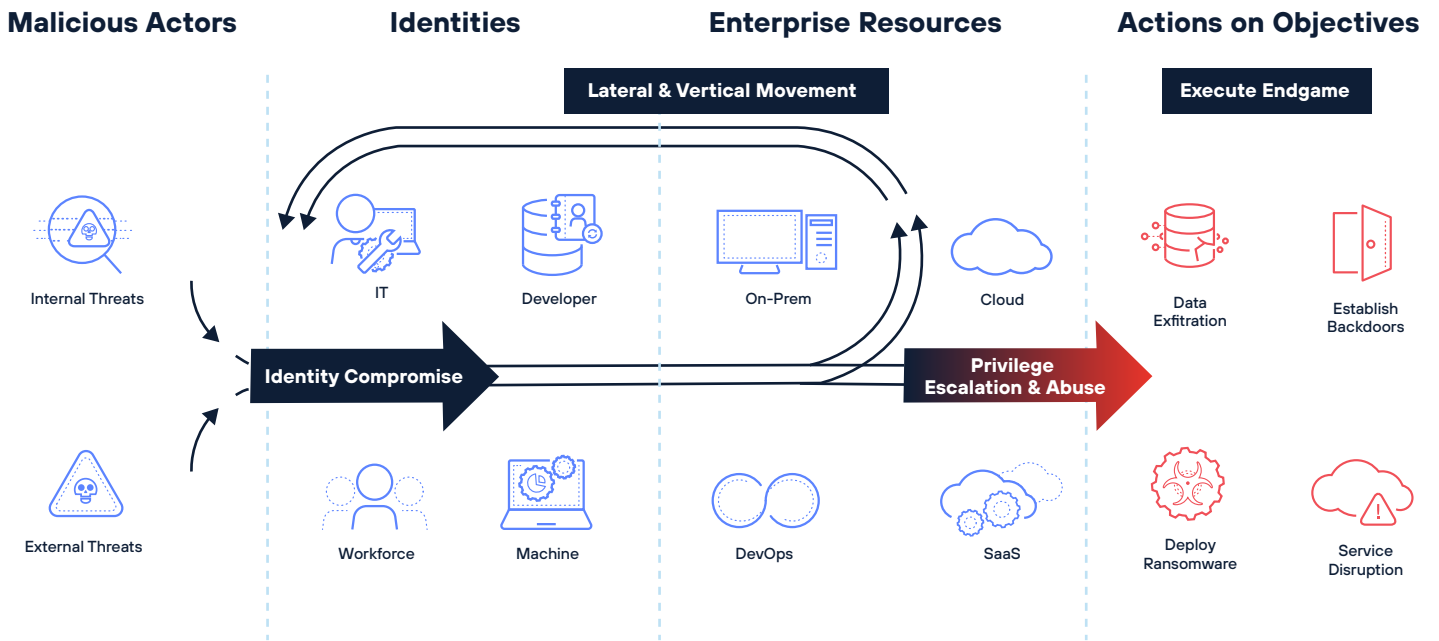


Figure 1. The identity attack chain

For example, a malicious actor might target a developer via a phishing campaign and compromise their standard workforce credentials. Having compromised their workstation, they can also steal any session cookies the developer has to cloud service providers (CSPs). Because that developer was granted standing access with administrative privileges to those CSPs, the malicious actor can bypass controls like MFA. They can immediately begin to abuse those privileges and perform actions like deploying ransomware in the cloud or exfiltrating data.

Idira Blueprint provides security control recommendations that rest on three guiding principles. They are designed to disrupt those common techniques and break the attack chain for all identities, including human, machine, and AI agents:

- **Prevent identity compromise:** Using techniques, such as social engineering, MFA bypass, keystroke logging, and credential repository scraping, to harvest passwords, hashes, SSH Keys, hard-coded credentials, and browser session cookies.
- **Stop lateral and vertical movement:** The movement across an organization’s resources. It can be laterally from within a risk tier (e.g., workstation to workstation or server to server). Or, it can cross vertically into another risk tier or environment (e.g., workstation to cloud service provider or workstation to a DevOps tool).
- **Limit privilege escalation and abuse:** The process of elevating privileges to then carry out malicious actions or behavior against the organization, enabling the action on objectives.

As good security practitioners, your teams want to put identity security controls in place that mitigate the risks associated with this attack chain, which is why the Idira Blueprint uses these three guiding principles. Looking at access in this way provides a valuable rule of thumb to help you make decisions on when, where, and how to engage with security initiatives. These principles are also supported by the controls and processes we use to mitigate operational and compliance risks as well, because failures in these often compound security risks.

Understanding Risk, Access, and Privilege Controls

A key aspect of effectively applying the appropriate level of intelligent privilege control to a specific identity and enterprise resource is first understanding that privilege controls are not a one-size-fits-all model because not all access is equal. Certain identities and their access tend to be higher risk or lower risk, which is fairly consistent across most organizations. We don't want to apply too restrictive controls against users for whom it's not warranted. Inversely, we need to make sure high-risk access is protected properly.

Generally speaking, as the level of risk increases for a particular identity's access to a given resource, so should the level of control and protection through the appropriate intelligent privilege control. This thought process aligns with the concept of minimum effective control. We must put the appropriate control in place in accordance with risk, but not overburden the users or the operations and administrative teams.

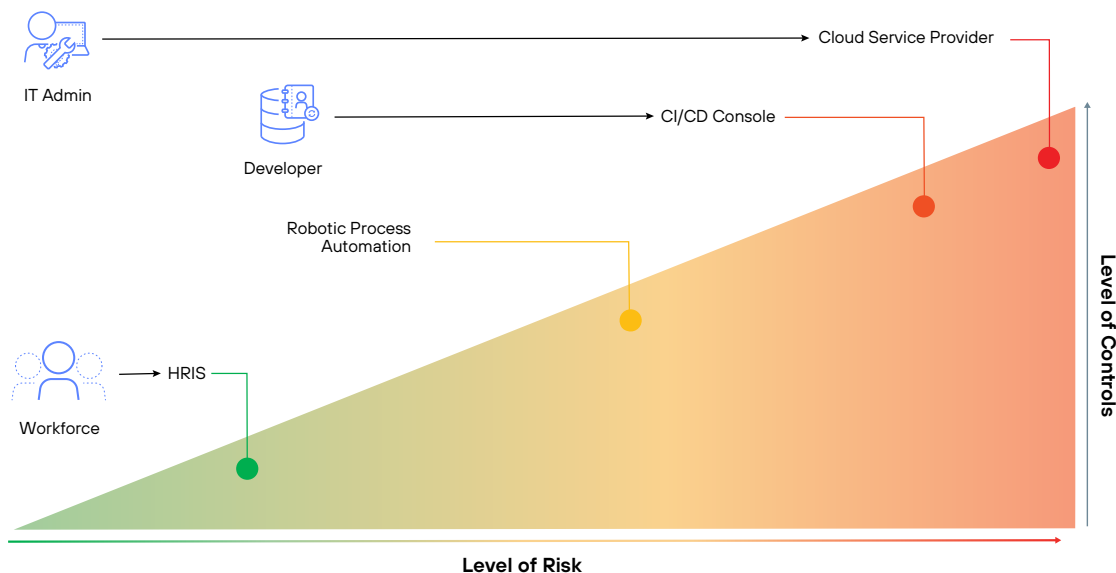


Figure 2. Level of controls aligned to the level of risk

For example, an IT administrator role such as a site reliability engineer commonly has administrative privileges across the organization's CSP accounts. This extremely high-risk access scenario requires applying higher levels of intelligent privilege controls. Inversely, a general member of the organization's workforce and their individual personal access to the company's HR system is not as risky to the organization.

Applying the same level of intelligent privilege control to that identity as the IT admin would not be considered the minimum effective control.

Common Factors That Influence Identity-Related Risk

Another major aspect to implementing effective intelligent privilege controls is understanding what constitutes identity-related risk. Risk can be defined in many ways depending on the context and situation. Many factors influence risk or the methods to calculate risk, such as data classification, data sensitivity, and monetary impact. For the purposes of developing a common language surrounding identity-centric risk and accelerating security outcomes, risk is a combination of three factors: the level of privilege, scope of influence, and ease of compromise.

Level of Privilege

This factor refers to the type of privileged actions the identity can perform against a given resource. Specific privileges and entitlements vary system to system but generally consist of a spectrum of privileges that include read-only, read-write, service-level administrator, modify user permissions, and full administrator permissions.

Read-only access is the least risky permission because the identity can only read data, not modify it. On the contrary, **modify user permissions** allows an identity to escalate their (or others) privileges as a result of being able to modify permissions. Naturally, **full administrator permissions** would be the highest level of risk. The higher the level of privilege, the higher the level of risk the identity poses.

Beyond the technical ability to execute actions, privilege is also defined by the sensitivity and regulatory classification of the data being accessed. From an operational and compliance perspective, an identity is considered privileged if it has access to proprietary corporate IP, strategic roadmaps, or trade secrets that provide a competitive advantage.

Furthermore, access to personally identifiable information (PII), such as Social Security numbers or financial records, carries a high level of privilege due to the legal and regulatory frameworks (like GDPR or HIPAA) that govern its protection. In this context, even a read-only permission can represent a significant security risk if the data being viewed is highly regulated.

Therefore, the risk profile of an identity is a product of both its technical entitlements and the inherent value or sensitivity of the information it is authorized to handle.

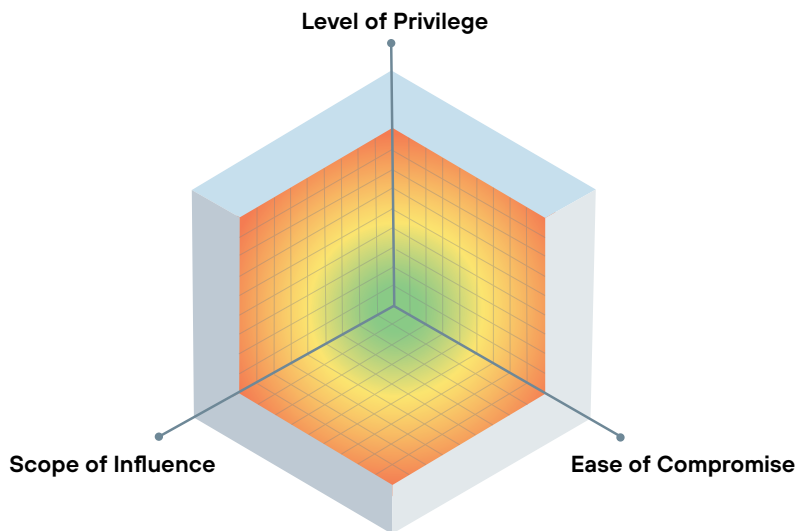


Figure 3. Three elements of identity risk

Scope of Influence

This factor, 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—a singular resource, 10 resources, a percentage of resources, or all resources. The larger the scope of influence is, the higher the level of risk is that the identity poses.

When you think about the scope of influence, keep in mind how interconnected and hierarchical IT and enterprise resources can be. You'll often need to think about access in terms of a specific-resource type (e.g., only access to Windows Servers) or more broadly to multiple types of resources, even all resources (e.g., access to all elastic workloads on the cloud and all cloud-native services). Another way of thinking about this is "downstream" or "inherited" access.

With the privileges an identity has against a given resource, is there subsequent access that the resource itself would then provide to other systems?

In this scenario, the nature of administrative privileges means an identity with a high scope of influence can impact a wide swath of other systems, such as administrative access to an automation service like Terraform, which has downstream administrative permissions to other IT systems for IaC automation.

Ease of Compromise

This factor refers to how easy or challenging it is for a malicious actor to compromise an identity. It includes common attack patterns, technical vulnerabilities related to the identity, access, or resource that can be exploited, poor internal processes that expose the identity, and a lack of minimum effective controls applied.

Examples of Common Factors

Example 1: Domain Admin Login to Workstations

Consider a domain administrator account with administrative privileges to all the servers and workstations in the environment. Such an account would have a high scope of influence, a high level of privilege, and, if used from an unprotected workstation, a moderate-to-high ease of compromise. The workstations of the workforce tend to be the main landing place for malicious actors. Also, as a result of how Windows authentication works, a domain administrator-privileged account that logged into a workstation might leave credential residue in the form of a hash that could be exploited by a bad actor. This scenario would result in a relatively high ease of compromise.

Example 2: SMS MFA for IdP-Based Freestanding Access to CSPs

Organizations that use SMS-based MFA for identity provider-federated access into CSPs would be considered a relatively high ease of compromise. The reason is that SMS-based MFA is considered weak and easily bypassed via TTPs like SIM swap fraud, phishing attacks, and SS7 interception. Using freestanding permissions for these identities means that actors can immediately exploit the permissions. Any identity, including those administrator-type users of the CSPs, would be particularly vulnerable to attacks.

Example 3: Static Service Account Secrets in AI Business Assistants

Consider an AI business assistant that's configured with a long-lived service account secret that's hard-coded into its system prompt to access internal data like CRMs. This identity would have a high scope of influence and privilege over sensitive data. The scenario presents a high ease of compromise because AI agents are vulnerable to prompt injection, where users use natural language to trick the model into disclosing its configuration. Because the secret is static, a malicious actor could easily extract the credential to gain access to internal systems, resulting in a high ease of compromise.

Overlapping Access to Common Resources Drives Lateral and Vertical Movement

One last concept that rounds out the picture of identity risk is overlapping access to common resources. Your need to think about identity security holistically, from the perspective of all these different identities, is partly because of their overlapping access to common enterprise resources. Malicious actors know how to circumnavigate organizations and exploit those areas of overlapping access to their advantage. This overlapping access is what enables malicious actors to move laterally from resource to resource, and eventually move vertically based on the access they find. Protecting IT admins' privileged access to virtual machines (VMs)—but leaving developers privileged access to those same VMs unprotected—means that your VMs are not fully protected. Your IT admins might also be at an increased risk for compromise as a result.

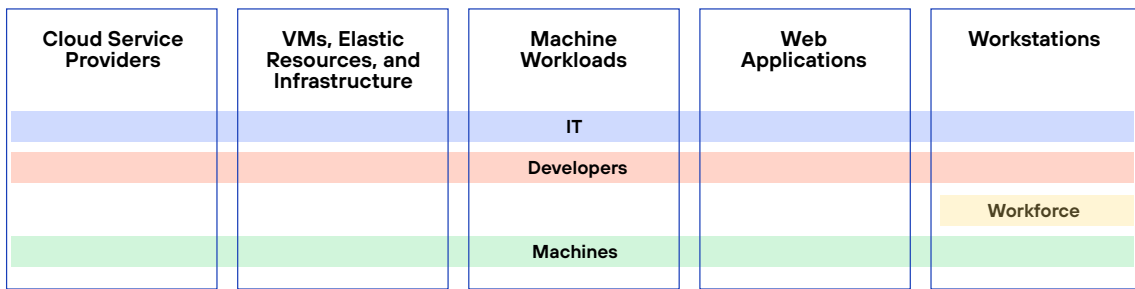


Figure 4. Lateral and vertical movement via overlapping access

Holistic security encompasses all the identities that access enterprise resources, whether it has operational (personal), system (built-in local admins and break-glass accounts), or machine (nonhuman workload) access. You can't call a specific enterprise resource secured until all identities who access it are secured. Many PAM programs traditionally focused on protecting IT and machine access to VMs and infrastructure. But, with all these overlapping rows and columns, malicious actors can navigate their way around and continue to wreak havoc with the overlapping access. Only focusing on one identity's access to a specific resource type leaves additional avenues for malicious actors to compromise the resource and then circumnavigate your organization.

Building a Risk-Based Plan: Aligning Prescriptive Actions with Risk Reduction

Idira Identity Security Blueprint is designed to protect any customer environment, strengthening identity security for cloud-native services, elastic cloud workloads, and high-risk and workforce applications. It lays out a pragmatic, risk-based approach to protecting enterprise resources that splits use cases by security control family and then sorts them by our experience-based, risk-effort index. By using this index, the blueprint can help your organization address the most critical security needs in the short term, while providing a long-term plan to address more advanced security use cases. When combined with your own internal priorities, drivers, and goals, you can build an informed, measurable, and risk-based plan to protect your identities across the enterprise.




Figure 5. Risk-based prioritization quadrant

Risk-Based Prioritization Methodology

As organizations start to see the expansive scope of access for various identities, they quickly begin to realize they need a method for prioritization of securing these identities, access, and enterprise resources. As good security practitioners, you want to focus on reducing risk. We prioritize use cases by using a the risk-effort index, which compares two key factors: the amount of risk reduction and the amount of effort required to mitigate a risk for a use case.

Risk can be defined in many ways to different organizations. In this context, we're referring to the sum of those three factors: level of privilege, scope of influence, and ease of compromise. These factors drive the impact of a compromised identity, are common to all organizations, and are widely exploited by malicious actors. The larger the impact, the larger the risk. The larger the risk, the larger the risk mitigation by implementing intelligent privilege controls.

The second factor, level of effort, refers to the amount of difficulty, including technical, political, or social complexity, and the time it takes to implement a particular identity security use case. Scenarios include where implementation requires little customization or how native integrations are less technically complex. Also, scenarios that include where the identity user group is typically more stand-offish or resistant to change, are more politically or socially complex to implement, and might require a higher level of organizational change management. Not every organization can have a top-down mandated approach to identity security, so advice should default to which use cases are typically easiest, or easier, to implement.

This risk-versus-effort logic enables us to prioritize the identity security use cases that give your organization the best risk protection for the effort you want to invest in it. If your organization has only so much time, effort, and resources to accomplish anything, we want you to get the maximum risk reduction for your effort (i.e., risk reduction ROI or shareholder value). Idira Identity Security Blueprint recommendations are primarily driven from this perspective, starting with the high-priority tasks, or quick wins, that will net your organization the most value for relatively low effort. From there, we conceptually move to either of the medium-priority buckets, and then come around to the low-priority use cases.

Building a Custom Plan for Your Organization

While Idira Identity Security Blueprint leverages a security risk reduction-first strategy, it's not the case for every organization. Other factors, like audit and compliance regulations, cyber insurance requirements, agility or productivity reasons, security incidents or breaches, or executive leadership priorities, all shape an organization's security roadmap. Whatever the program driver, as good security practitioners and trusted advisors, advise security leaders on why Idira Blueprint is prescriptive in its manner, and help them ultimately make the most informed decisions based on our experience.

Idira Blueprint recommendations are built on our combined knowledge and experience with Unit 42 in battling threats in the identity security space. These insights are gathered from lessons learned across our 70,000 customers, postbreach experience, frontline remediators, red team, and innovative researchers. As such, we've developed a prescriptive perspective on what identity security use cases move the needle the most for human, machine, and AI identities.

Idira Blueprint is only one part of the process to create an identity security roadmap. It serves as a guide and tool to help you along the way so you can make the most informed decisions about protecting identities and their access to enterprise resources.

Organizations need to take stock of their current state of security, factor in their internal priorities and drivers, desired business outcomes and metrics, and understand their bandwidth and capacity for security efforts. When all five of these considerations are aligned, you can craft an identity security roadmap that's specific to your organization but is still reflective of industry and security best practices, as well as designed to have the best odds for success.

Security Control	Month 1	Month 2	Month 3	Month 4	Month 5	Month 6
Zero standing privilege for cloud service providers	IAM admins and operators	Developers	Service-level admins	Other privileged and read-only users		
Standing privilege access management for cloud service providers	Root and registration accounts	Freestanding CSP accounts				
Standing privilege access management for AD, VMs, and infrastructure		Active Directory domain admins	Windows server built-in local admins	Workstation built-in local admins	Database built-in local admins	
Zero standing access management for VMs and infrastructure	Windows server admin access	UNIX server admin access	Database admin access			

Figure 6. Example identity security roadmap

Conclusion

Malicious insiders and external attackers can exploit identities to steal confidential data or disrupt critical applications. Idira Identity Security Blueprint helps organizations formulate and maintain an effective risk-based identity security program that takes full advantage of our vast knowledge and expertise. Designed to defend against the three most common attack scenarios, Idira Blueprint provides a prioritized framework that closely aligns prescriptive actions with risk reduction. It helps organizations address the vulnerabilities that pose the greatest potential threat as quickly as possible. Its self-service ecosystem provides on-demand best practice guidance from industry experts.

By following the recommendations and guidelines in Idira Blueprint, organizations can strengthen their security posture, reduce risks across the full spectrum of identities, and make the most of their identity security technology investments.

Develop Your Plan with Idira Identity Security Blueprint

Developing and executing an effective identity security program can be a complex undertaking. We provide the experience, solutions, and security services you need for a successful identity security program. Start developing your roadmap toward identity security success today.

Go further with Idira Blueprint and customize a plan to fit your needs and goals. Along with our partners, we offer a wide range of success services to help you with your identity security program. In addition to the self-service Idira Blueprint resources, these offerings include supplementary guided support through engagements led by us and our partners.

To explore all the ways Idira can secure the identities across your organization, visit www.paloaltonetworks.com/idir.

About Palo Alto Networks

Palo Alto Networks (NASDAQ: PANW), the global AI cybersecurity leader, protects our digital way of life with a comprehensive portfolio of cybersecurity solutions and platforms across Network, Cloud, Security Operations, AI, and Identity. Trusted by more than 70,000 customers and powered by Unit 42 threat intelligence, our AI-driven platforms eliminate complexity, empowering enterprises to modernize with confidence and securing the speed of innovation. Explore the future of security at www.paloaltonetworks.com.

GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and is used herein with permission. All rights reserved.



3000 Tannery Way
Santa Clara, CA 95054
Main: +1.408.753.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087
www.paloaltonetworks.com

© 2026 Palo Alto Networks, Inc. A list of our trademarks in the United States and other jurisdictions can be found at <https://www.paloaltonetworks.com/company/trademarks.html>. All other marks mentioned herein may be trademarks of their respective companies.
idir_wp_idira-blueprint-for-identity-security_050526