

# Best Practices for Securing Cloud Identities

---

## Table of Contents

<b>Executive Summary</b> .....	3
<b>The Cloud Has Changed Security</b> .....	4
Cloud Service Providers and Idira Blueprint .....	4
Challenges with Identity in the Cloud .....	4
<b>The Anatomy of Cloud Identities</b> .....	5
Understanding Cloud Services, Resources, and Identities .....	5
Cloud Identities .....	6
Comparing the Cloud Provider Access Models .....	6
Access Models for Elastic Cloud Workloads .....	10
<b>Identity Security in the Cloud</b> .....	10
Understanding the Identity Attack Chain .....	10
Identity Security Controls .....	11
Key Tenants of Strong and Secure Access to the Cloud .....	13
Alignment to Well-Architected Frameworks .....	14
<b>Prioritization Strategies</b> .....	15
Security Control-Based Prioritization .....	16
Identity- and Persona-Based Prioritization .....	17
<b>Alignment with Cloud Adoption Strategies</b> .....	19
Digital Native Businesses and Digital Native Enterprises .....	19
Lift-and-Shift Organizations .....	20
Operationalizing Cloud Security at Inception with Infrastructure as Code .....	20
<b>Concluding Insights</b> .....	21
<b>Next Steps for Securing Your Cloud Identities</b> .....	21

---

## Summary

The advent of cloud computing has fundamentally transformed organizational IT infrastructure and service deployment, necessitating dynamic security strategies to safeguard digital assets and software development. Idira™ Identity Security Blueprint (Idira Blueprint), by Palo Alto Networks, provides a comprehensive framework to help organizations secure cloud identities by following a holistic approach that includes both human and machine identities.

By adopting Idira Blueprint, organizations can ensure a comprehensive, risk-based approach to cloud security that pursues zero standing privileges (ZSPs) to safeguard their cloud workloads and services. Follow the guidance in this whitepaper to help your organization secure your cloud environments against increasing complexity and threats.

## What You'll Learn



- **Cloud security landscape:** The evolving cloud landscape poses challenges. Grasp the importance of dynamic identity security strategies to address the proliferation of services and identities, coupled with the diminishing traditional security perimeters.



- **Risk-based guidance:** Get a focused view of Idira Blueprint through the lens of cloud service providers (CSPs) and cloud resources. This battle-tested best practices framework is designed to reduce risk with prescriptive guidance for all types of identities.



- **Critical identity security controls:** Identity security is the cornerstone of a successful cybersecurity strategy. Learn about the controls to help prevent identity compromise, stop lateral and vertical movement, and limit privilege escalation and abuse.



- **Minimize the attack surface:** The goal is to minimize the attack surface within cloud environments while advocating for ZSPs, just-in-time (JIT) access, and entitlements provisioning as part of your identity security program.



- **Prioritization strategies:** A one-size-fits-all path doesn't exist for adopting identity security controls in the cloud. Multiple approaches are available, each of which considers the existing capabilities, risk-based prioritization, compliance requirements, and cloud adoption approaches.

---

# The Cloud Has Changed Security

## Cloud Service Providers and Idira Blueprint

The public cloud has revolutionized how organizations function, how information technology runs, and how businesses provide services to their consumers. However, with the rapid ascent into the cloud and the growing number and complexity of services, it's easy for malicious actors to find identity-centric entry points and compromise a public cloud environment. Keeping the cloud secure is vital to an organization's business success.

As CSP offerings increase and the shared responsibility model continues to put ownership on the customer, organizations are required to bring their own safeguards and protections to the cloud. Each cloud provider has its own shared responsibility model with clear guidance that their customers are responsible for securing the configuration within their own environments. This responsibility is extremely challenging at a multicloud scale. Amazon, Microsoft Azure, and Google Cloud offer approximately 1,400 native services with over 40,000 individual access controls. On top of that, the move to the cloud has erased the traditional security perimeter, leaving identity as the key to security in this new environment.

As cloud environments expand, the security risks for identity security programs also increase. Expansion creates a proliferation of human and machine identities, each of which can be configured with a growing number of permissions. Every day more cloud operators, IT administrators, developers, cloud-native services, and cloud workloads gain greater access. Since authentication (AuthN) and authorization (AuthZ) methods vary within each organization and across the cloud providers, centralization and standardization of identity security controls is a critical foundation for cloud initiatives.

Cloud service providers do not use unified terminology in reference to identities, access control methods, or privileges. This paper provides a common language around cloud identities and provides comprehensive guidance across the cloud providers, including their various services, and their elastic resources, infrastructure, and workloads. Much of this common language and comprehensive guidance comes from our Idira Blueprint best practices framework for identity security success. This framework was designed to help organizations measurably reduce risk, based on our lessons learned on the frontlines, providing prescriptive guidance to secure any identity, human, and machine.

## Challenges with Identity in the Cloud

While all cloud service providers offer some level of identity and access management services, the same cybersecurity challenges that exist in all identity and access management (IAM) programs also exist in the cloud. Cloud security uses a [shared responsibility model](#). CSPs, like Amazon, Microsoft, and Google, are responsible for the underlying infrastructure and software that make the services function, such as physical security, system availability, and uptime. However, the shared responsibility model means that your organization, not the CSPs, is responsible for securing your access to their cloud services or your workloads on the cloud.

Your organization is responsible for solving the challenges related to four key areas of identity security:

- **Authentication:** CSPs support a wide array of authentication mechanisms for various types of identities. Different identities use different methods for authentication depending on the use case or scenario. There are valid business justifications for using SAML-based, username and password, or access and API key authentication. Additionally, each CSP offers authentication options in slightly different ways, adding to confusion for end users and security teams.
- **Authorization:** Countless permissions are available within each CSP for use by all these different identities. Identities can be assigned permissions directly through roles or groups or by other means altogether, making the overall process of permission and authorization management increasingly complex. The distinct IAM paradigms of each CSP further add to this complexity for organizations with a multicloud footprint.

- **Access:** Methods of access vary from identity to identity and scenario to scenario. Managing and controlling the various access planes for each unique situation is a new problem for many organizations.
- **Audit:** Security organizations want centralized visibility and control, whether in a single cloud or multicloud architecture. Cloud security might be worthy of its own security initiative, but that doesn't mean organizations want to keep audit processes separate from their other internal security programs.

The combination of these challenges makes cloud identities ripe targets for malicious actors to take advantage of. It also almost certainly ensures they'll find a way to do it. The same benefits that cloud provider accounts grant—such as centralized and simplified deployment, all-in-one data, and application hosting—are the same reasons bad actors target them. They contain high-value information for attackers to exploit through data exfiltration, ransomware, or service disruption.

To help your organizations secure all cloud identities and their access methods, follow the prescriptive guidance in Idira Blueprint. It helps you easily identify the types of privileged entities within an organization's cloud provider accounts and prioritize securing those entities based on risk impact and the effort required.

## The Anatomy of Cloud Identities

### Understanding Cloud Services, Resources, and Identities

To better understand the challenges your organization faces when securing your cloud provider environments, you must first understand the anatomy of cloud identities. All CSP environments, whether Amazon Web Services (AWS), Google Cloud, or Azure, consist of two major parts:

- **Management platform and services:** This part allows various identities to administer and operate different services like identity and access, compute, or secrets vaults. The management console (accessed via the web interface, CLI, or API) is the main access point into the cloud provider.
- **Infrastructure workloads:** These are created by different cloud services like virtual machines, containers, serverless functions, and cloud-native apps and storage. Services are available to administer and create these various workloads, but they also require a separate layer of identity security controls, separate to that of the platform or services.

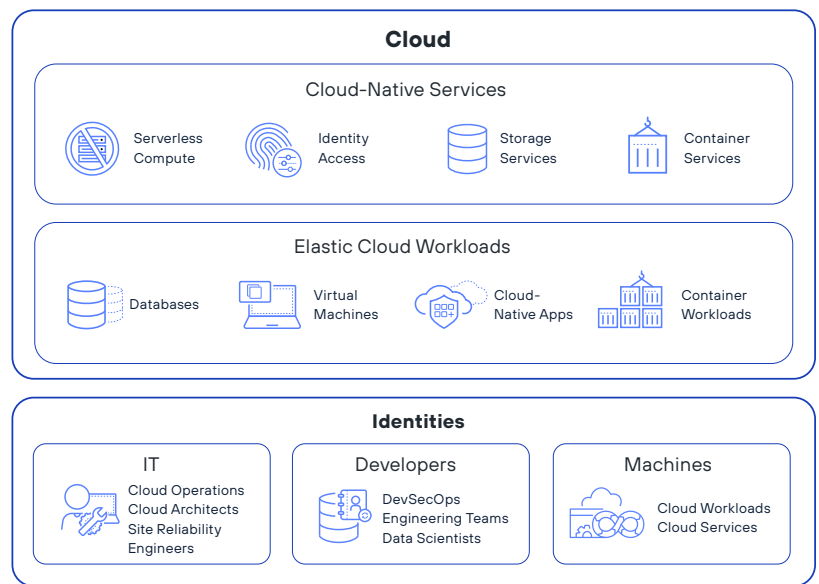


Figure 1. Anatomy of the cloud

## Cloud Identities

Organizations also have identities that access the management console, various services, and resources created:

- **Cloud operations** that have evolved from traditional IT roles, such as infrastructure operations, networking engineers, or database admins, into roles like cloud operators, architects, and site reliability engineers. Cloud operations include:
  - › **Cloud administrators** who have complete administrative access and the ultimate permission to affect every service and resource within the CSP account.
  - › **Service-level administrators**, such as engineers with a specialization in networking or databases, who can administer only a smaller scope of services or resources.
- **Developers** who self-administer various cloud services, create cloud-native applications, push workloads into the cloud, and access supporting resources.
- **Other application and audit teams** with lesser privileges, like read-only access, to various services.
- **Various machine identity workloads**, such as cloud-native applications, services, automation tools, and processes that run your organization.

These identities all authenticate to the cloud using various methods, including standing federated access via an identity provider (IdP); long-lived freestanding local accounts, like user passwords and keys; or in an emergency, using the root or registration credentials.

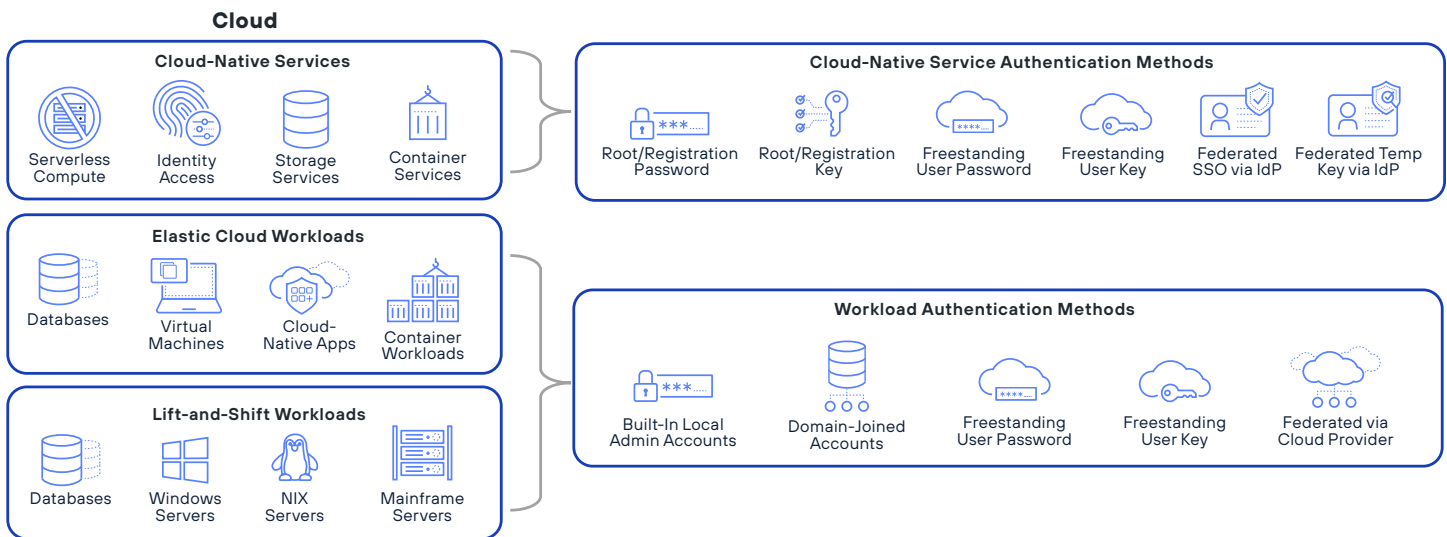


Figure 2. Cloud provider services/workloads and common authentication mechanisms

## Comparing the Cloud Provider Access Models

Three dominant identity security concepts are worth noting the similarities and differences among the various cloud service providers.

### Federated Access

Federated access refers to identities that are created within a centralized directory service bound (federated) to the CSP or its workloads. This access is the standard enterprise approach to granting access to public cloud and SaaS resources for various reasons such as simplicity, administration, and security. Securing federated access is a key underlying element of establishing a secure cloud environment.

Federated access models vary slightly across all the major CSPs. Google Cloud requires the use of one of two top-level Google directory services—Google Workspace or Google Cloud Identity—to establish federation. Both services share the same backend directory structure and act as the Google Cloud directory service for all federated access. You can bring your own identity provider to Google Workspace or Google Cloud Identity. If you’re planning on federating access, you must first join your directory service to Google Workspace or Google Cloud Identity and then grant access to Google Cloud projects. Access and permissions granted to these projects are persistent.

In Azure, federated access uses Microsoft Entra ID and supports bringing your own identity provider, including support for additional IdPs via B2B integration support. Without Microsoft Privileged Identity Management, federated access and the assigned permissions are standing (persistent and long-lived).

When using AWS Organizations to deploy account IDs at scale, use AWS IAM Identity Center (the Amazon IdP service). However, you can still integrate your existing identity provider into AWS IAM Identity Center. Federated access into nonorganizational AWS accounts can use any identity provider directly without requiring integration to Identity Center. Another major distinction for AWS is that it uses role assumption for JIT access, meaning that the accesses are not freestanding, but the permissions to the roles are.

Table 1. Cloud Service Provider Federated Access			
Access Model	AWS	Azure	Google Cloud
Federated Access	<b>AWS Organizations:</b> Any identity provider indirectly via IAM Identity Center	Any identity provider	Any identity provider
	<b>Nonorganization:</b> Any identity provider directly into IAM (or indirectly via IAM Identity Center)	Requires sync/IdP config into Entra ID	Requires sync/IdP config into Google Workspace (Google Cloud Identity)

### Nonfederated Access

Nonfederated access refers to identities, including both human and machine, that are created locally within the CSP’s platform. It commonly applies to emergency or break-glass usage, organizations using CSPs for the first time, newly acquired testing environments, scripts and batch processes, and other applications needing to authenticate to access various APIs. Nonfederated access is similar across CSPs, except for Google.

For human access, AWS enables creation of internal IAM directory users directly within the AWS platform, although we don’t recommend implementing this practice at scale. Azure is similar and also allows you to create nonfederated internal Entra ID directory users. These are considered local accounts because they exist solely within the CSP platform. Google Cloud does not include the capability to create nonfederated local directory users within a project or organization directly, as that requires the use of Google Workspace for all humans. However, IT teams can create local (nondirectory federated) Google Workspace or Google Cloud Identity users and grant them access to Google Cloud projects.

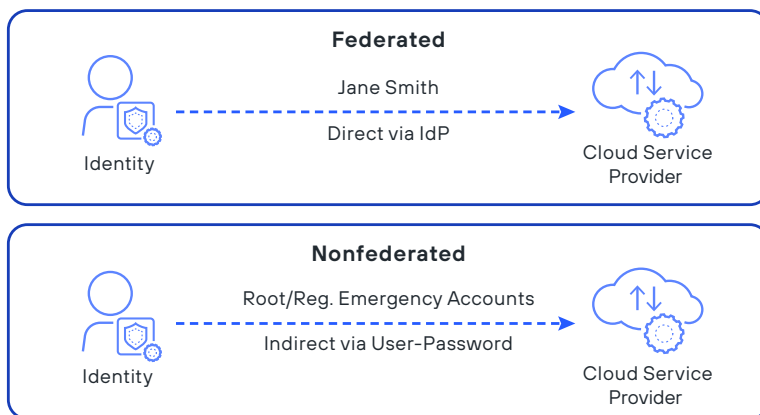


Figure 3. Federated vs. nonfederated access

Machine identity access becomes more complicated because access is provided to workloads, applications, scripts, and services that exist both inside and outside the CSP platform. AWS enables IT teams to use freestanding IAM users with access keys, a method typically used for external workloads, as well as IAM roles for both internal services and external workloads.

Similarly, Azure also enables creation of nonfederated internal Entra ID principals, called **app registrations**, which leverage application keys for authentication and can be used for external application access. For internal access, Azure leverages a **managed identity**, which resides in the Azure subscription and creates a service principal in the Entra ID tenant.

Google Cloud uses internal IAM services within the project, instead of Google Workspace, to create service accounts, which is similar to a traditional directory-based account and credential when compared to the AWS and Azure methodology. The Google Cloud approach is more inline with the external application approach that AWS and Azure use, but it is also used for internal applications.

**Table 2. Cloud Service Provider Freestanding Access**

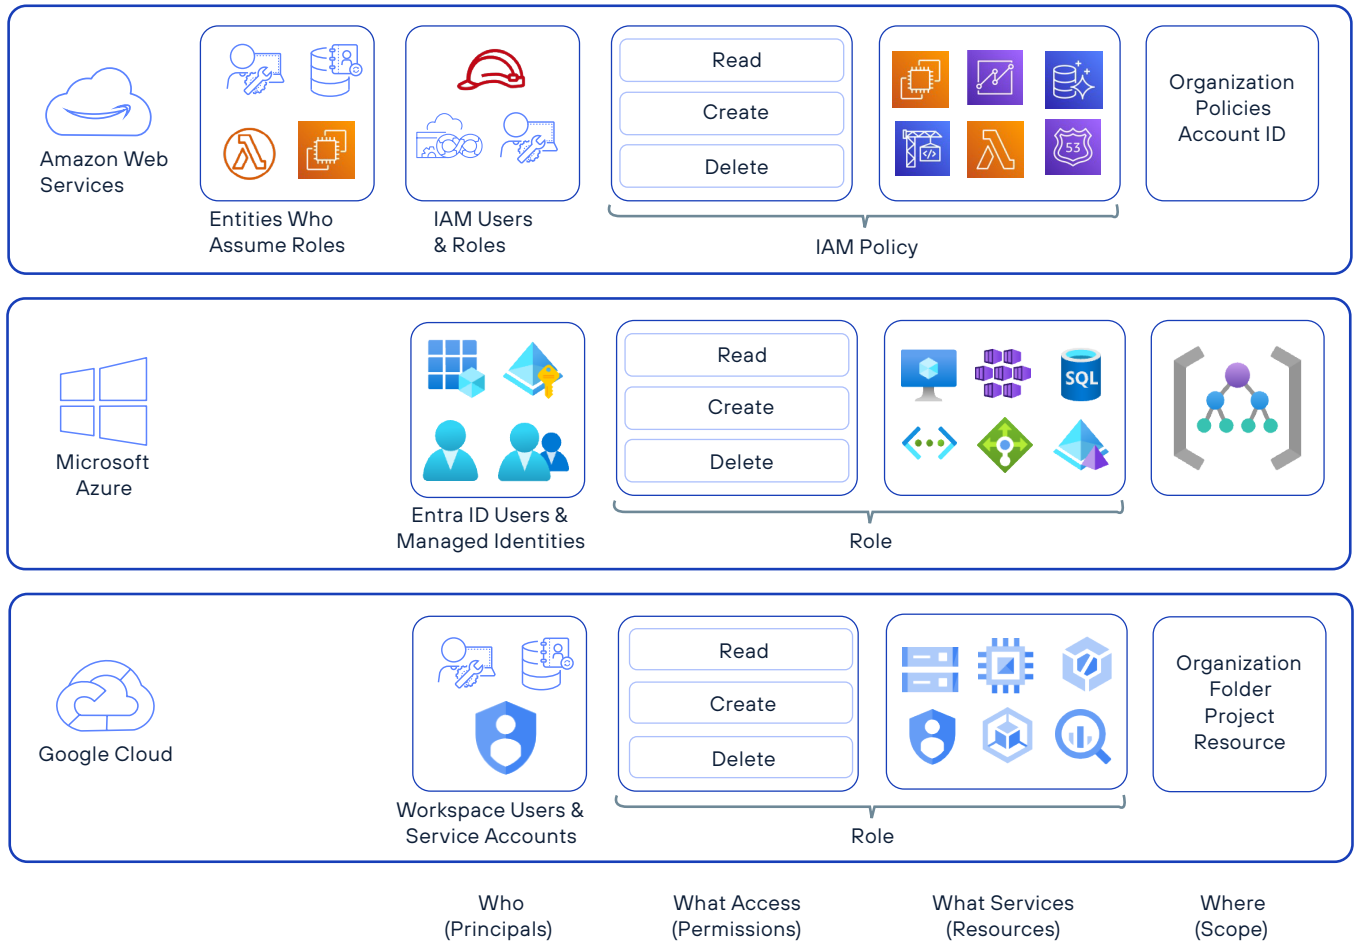
Identity Type	AWS	Azure	Google Cloud
Humans	Root, IAM Users, Access Keys	Entra ID Users	Google Workspace User (External to Google Cloud)
Machines	IAM Users, Access Keys, or IAM Roles	Entra ID Users, Application (Client) ID Keys, or Managed Identity	IAM Service Account

## Permission Assignments and Authorization

Permission assignment and granting permissions are where the biggest differences lie among the CSPs.

- **AWS IAM** uses a policy, either an Identity Policy or Resource-Based Policy, to control an entity's permissions. Policies can be applied via a permission set or applied directly to an IAM user, group, or role.
- **AWS IAM Identity Center** (formerly AWS SSO) uses a permission set. Permission sets are defined by AWS policies and then are mapped to AWS accounts, allowing for any assigned users and groups to access appropriate accounts in your AWS organization.
- **Azure** uses a role. The roles are defined by the configured permissions for them, where entities, like a user or group, get assigned roles.
- **Google Cloud** is more inline with Azure, where it uses the same concept of a role. Think of a permission policy in AWS being the equivalent of a role in Azure and Google Cloud.
- **Both Azure and Google Cloud** have options for both built-in and custom roles, compared to AWS, which encourages the use of custom roles.

For ease of understanding, figure 4 shows how the entities and permissions are structured with the CSP-specific terminology.



**Figure 4.** Cloud provider permission assignment

### Permission Scope and Structure

Permission scope refers to the level in the hierarchical structure in which permissions can be applied to and inherited for. This concept impacts how organizations secure identities within their CSPs, because they might grant permissions at one or more levels.

When using AWS Identity Center, users, groups, and permissions are administered at the management account level, which is where permissions are mapped to member accounts. If AWS Identity Center isn't used, permissions are assigned locally within the respective billing account (AWS account ID).

Azure and Google Cloud can be assigned at multiple levels. Azure can assign permission scope at the subscription, management group, and root management group levels. Google Cloud can assign permission scope at the project, folder, and organization level.

Safeguards include configuring security policies to boundaries, sessions, virtual private clouds (VPCs), and resources that help mitigate some of the risks associated with access. Because CSPs are continuously changing, safeguards undergo significant change. Review your provider's features and documentation to see what's relevant.

All three CSPs are capable of assigning permissions at a service or resource level (the resource group level for Azure), although only Microsoft and Google explicitly call that out when referring to permission scope. Because CSPs use different terminology, it can be confusing to understand how different levels relate to one another. Table 3 highlights the different scopes of the CSPs and how their concepts relate to one another.

<b>Provider/Scope</b>	<b>Overarching Level</b>	<b>Group Level</b>	<b>Account Level</b>	<b>Resource Group Level</b>	<b>Resource Level</b>
<b>AWS</b>	Organization	OU	Account	N/A	Resource
<b>Azure</b>	Root Management Group	Management Group	Subscription	Resource Group	Resource
<b>Google Cloud</b>	Organization	Folder	Project	N/A	Resource

## Access Models for Elastic Cloud Workloads

Access models to the elastic cloud workloads have more commonalities across the providers.

Federated access to elastic cloud workloads involves the use of a centralized identity provider to manage access to resources like virtual machines, databases, and other workloads hosted by the CSPs. Federated access to resources is typically integrated in one of two ways: using traditional methods joined to AD (similar to a self-hosted infrastructure) or through cloud-native services like AWS Systems Manager Session Manager, Azure AD, or Google Compute Engine Authenticate.

Nonfederated access to elastic cloud resources is typically reserved for built-in local administrative accounts only, the Windows administrator SID-500 account, NIX Root UID0 user, or RDS primary user type-scenarios. Unlike within self-hosted infrastructure, nonfederated access is not a common access model outside of those built-in accounts due to the dynamic nature of cloud resources.

Permission assignments and authorization for elastic cloud workloads depend on your access method. Permissions are granted and abstracted via the CSP's native services, managed via AD groups, or locally managed on each resource or instance.

## Identity Security in the Cloud

### Understanding the Identity Attack Chain

To best understand which security controls to apply in which circumstance, organizations also need to understand how bad actors attack the cloud to begin with. Idira Blueprint outlines the common attack path malicious actors (internal and external) take to compromise identities and execute their endgame.

These actors compromise identities by using several techniques, such as social engineering, MFA bypass, credential theft, or cookie hijacking. From there, they use the identity's access to move around laterally, looking for more access to give them power, as well as vertically when they can, eventually escalating and abusing the privileges they've obtained. Malicious actors often target cloud operators, site reliability engineers, developers, and cloud engineers, because they have high privilege levels. They also target the workloads and services in the cloud because their privileges are often overprovisioned.

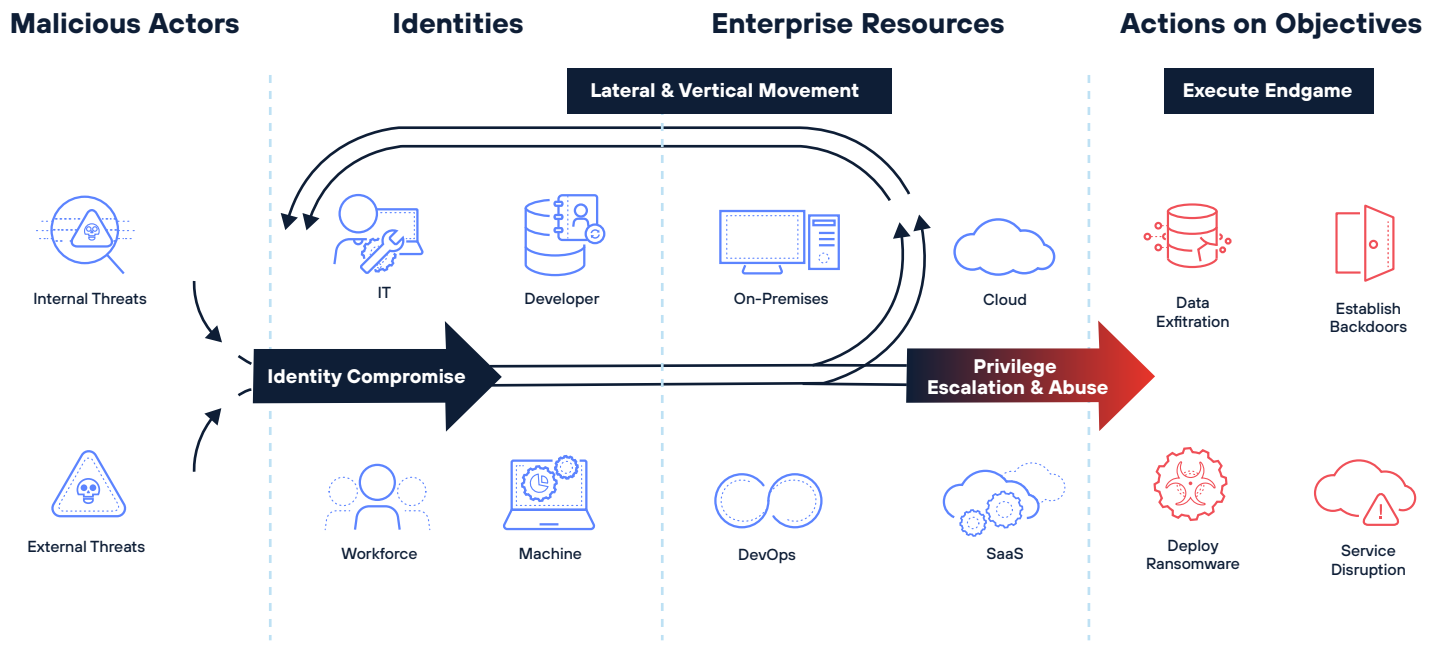


Figure 5. Identity attack chain

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 CSPs. Since that developer has been granted standing access with administrative privileges, that malicious actor can immediately begin to abuse those privileges and perform actions like deploying ransomware in the cloud or exfiltrating data.

## Identity Security Controls

To holistically address these challenges, we need a combination of effective intelligent privilege controls and prioritization. Think of it like securing a building with many doors. If you only lock the front door and keep all the side doors unlocked, you've left yourself vulnerable. By adhering to our recommendations, you'll gain control and visibility into the human-interactive and machine identities with access to your CSP environments and workloads.

### Mitigate Risks Associated with the Identity Attack Chain

Organizations should seek to implement security controls that mitigate the risks associated with the identity attack chain. This includes preventing credential theft, stopping lateral and vertical movement, and limiting privilege escalation and abuse. These risks can be mitigated using a combination of zero standing access and secure standing privileged access, secrets management, least privilege, and identity governance (lifecycle management and compliance) controls.

### Federated Access to the CSP and Its Workloads

For federated access to the CSP and its workloads, organizations should strive to achieve ZSPs. This can be accomplished using controls such as JIT elevation of access, JIT assignment of entitlements, and limited time-bound durations for the access. By following this approach, organizations reduce the risk of both identity compromise and lateral movement from a compromised identity.

Additional defense-in-depth layers, like session protection, recording, and audit, help further deter bad actors. This fundamental shift moves away from traditional freestanding federated access via SSO and standing entitlements. Organizations should strive for all interactive access to be with a ZSP approach.

## Nonfederated Access

For nonfederated access, organizations should strive for the objective of achieving secure standing access. They can use controls such as credential vaulting, password, and key management and rotation, complex password policy, multifactor authentication (MFA), session isolation, session monitoring, and audit and threat detection and response. Organizations should apply these controls to root and registration accounts and to any remaining nonfederated freestanding access directory user passwords and access or application keys, such as shared or emergency access accounts. Overall, organizations should strive to minimize the number of freestanding credentials to reduce the attack surface.

Furthermore, workloads and resources hosted within the cloud should have their built-in local administrative accounts (e.g., SID-500 Admin, UIDO Root, and Master DBA) protected with these same controls.

## Secrets Management

Secrets management controls include functions like secrets vaulting, secrets rotation, complex secret value policy, removal of hard-coded secrets from workloads and applications, and JIT secret delivery and dynamic secrets to those workloads. These controls build upon privileged access management (PAM) controls, extending credential management capabilities to machine workloads. Any machine identity, like cloud-native secrets, dynamic applications, scripts, or other services, should leverage secrets management controls to mitigate the risk of identity compromise.

## Least Privilege

Least privilege controls include two key concepts: the move to on-demand or as-needed privilege and limiting an identity's permissions to only those needed to perform its function or responsibilities. Moving privilege to an on-demand or as-needed basis, instead of granting always-on access, is a key mechanism to minimize the risk of standing privileges. Privileges are not authorized on the target resources at any point until they are needed. Reducing standing privileges provides immediate risk reduction while organizations refine and limit the permissions and privileges required through privilege analysis and policy creation or modification.

## Identity Governance

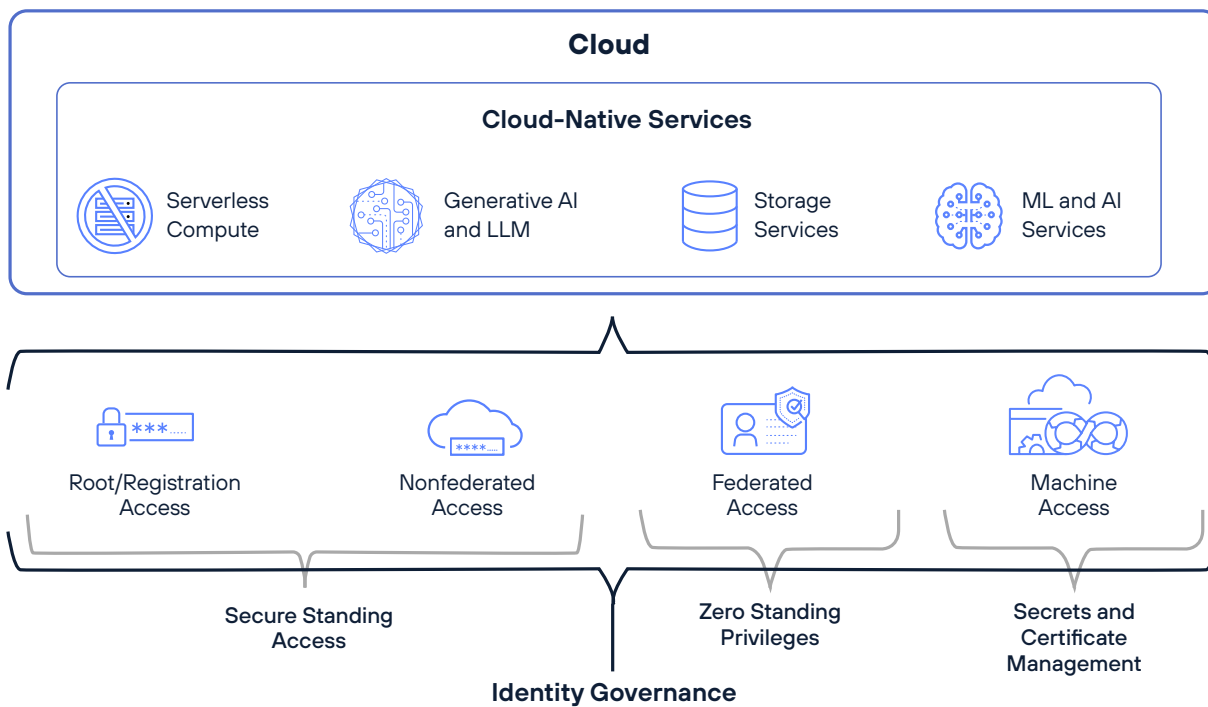
Identity governance controls include:

- **Lifecycle management** enforces the process of granting authorized identities access to the resources they need (via the appropriate control plane) at the time of hire or inception and revoking access when those identities no longer require them—your traditional Joiner-Mover-Leaver process.
- **Compliance controls** enforce the periodic certification and attestation that identities still require access to the things they currently can access and revoke that access if no longer required.

## Food for Thought

Holistic cloud security encompasses the objectives of achieving ZSP for federated access and secure standing access for nonfederated access. While the emphasis is on ZSP for humans, all CSPs and workloads have built-in local admin-type credentials that also require protection. Also, machine identities with admin access to your CSPs require secrets management controls.

Your cloud isn't secure until all objectives are accomplished.



**Figure 6.** Alignment of access methods and identity security controls

## All Identities

All identities, regardless of access method or whether they're human or machine, should be assigned only the minimum necessary permissions for their job function—inline with the least privilege methodology. This greatly reduces the risk of lateral movement and privilege escalation and abuse. Similarly, all identities should be subject to identity lifecycle governance to ensure that identities are granted correct access, at the right time, and have it revoked when it is no longer required.

When all these controls are put together properly, you can develop an effective access model with JIT access and ZSPs at the center.

## Key Tenants of Strong and Secure Access to the Cloud

The recommended identity security controls fit into five key tenants that should guide your security journey for cloud access.

### Implement Zero Standing Privileges for Federated Access

When possible, all federated access to cloud service providers and services should be provisioned with ZSPs. If necessary, move from freestanding access to JIT, and refine permissions over time to achieve ZSP.

### Minimize Accounts with Freestanding Access

Reduce the number of cloud service provider internal directory accounts and users leveraging passwords and keys to the absolute minimum. Freestanding accounts pose higher risk due to their long-lived nature and static permissions.

### Manage Defense-in-Depth PAM Controls on Remaining Freestanding Access

Any remaining freestanding accounts should have their passwords and keys managed with PAM controls. Mitigate the risk of identity compromise, lateral and vertical movement, and privilege escalation with vaulting, MFA, rotation, isolation, and audit.

---

## Protect Root and Registration Accounts with Extreme Care

Apply these same critical PAM controls to the accounts and emails used to register for the CSP account or subscription. Remember to also protect access to the inbox of the email addresses used for registration.

### Remember Machine Identities

Don't wear blinders just for human access, because there are often many more machine identities than human ones. Focus your efforts on protecting any machine workloads that have administrative permissions into your CSPs, such as your IaC tools and pipeline.

## Alignment to Well-Architected Frameworks

The well-architected frameworks from major cloud service providers outline the guidelines that help organizations build secure, high-performing, and resilient cloud environments by focusing on effective design principles and practices, including those related to identity and access management.

We've identified the following common principles that Amazon, Google, and Microsoft recommend to be compliant with their frameworks. Organizations must secure their cloud environments by aligning their identity security posture and security controls with the following principles:

- **Principle of least privilege:** Assign the minimum necessary permissions to users, processes, and systems to perform their tasks, reducing the risk of unauthorized access. Even when providing access with ZSPs, no user should have unnecessary permissions for the job at hand.
- **Authentication and authorization:** Implement strong authentication mechanisms, like MFA, and ensure proper authorization controls to manage access to cloud resources effectively.
- **Centralized identity management:** Use a centralized identity management system for user authentication and authorization, facilitating better control and monitoring of access across your cloud environment.
- **Credential management:** Regularly rotate and manage credentials securely, including for machine identities like service accounts. Avoid hard-coding credentials and use identity and access management roles whenever possible.
- **Audit trails and monitoring:** Implement comprehensive logging and monitoring for all identity-related events, enabling timely detection and response to security incidents. Incorporate cloud log solutions and cloud monitoring capabilities.
- **Automated compliance checks:** Employ automated tools and processes to regularly assess and ensure compliance with security best practices, IAM policies, and configurations.
- **Secure DevOps practices:** Integrate security measures into the DevOps pipeline, ensuring that identity and access controls are considered and tested throughout the development lifecycle.


These common principles are woven into the fabric of identity security and the Idira Blueprint recommended controls. An integrated identity security strategy that is aligned with these well-architected guidelines is a critical element of defense against attacks in today's threat landscape. Keep these principles in mind as you develop your own prioritization approach and cloud security strategy.

## Prioritization Strategies

Securing cloud service providers is a critical priority for many organizations. While all identities with access to the cloud are considered privileged, everything from read-only permissions to IAM administrator, you have to prioritize them little by little. Keep in mind the many identity and security controls. Organizations need to have a method to deploy security controls effectively and efficiently.

A major factor in our prioritization logic is balancing risk and effort. But how do we define risk? Each organization may include other factors in their definition of risk, such as data classification or sensitivity. We define risk through a common lens as being a combination of three factors:

- **Level of privilege** refers to the type of privilege granted to the identity, ranging from read-only access to the ability to modify other identity's permissions and access to full administrative control.
- **Scope of influence** (also referred to as the blast radius) refers to the amount or percentage of systems and resources an identity can access. It can range from access to a single cloud-native service to multiple services with access to elastic workloads to full access to every resource and service.
- **Ease of compromise** refers to how easy or challenging it is for a malicious actor to compromise the access, including the technical vulnerabilities that exist and the level of controls applied to protect the identity.



**Figure 7.** Three elements of identity risk

Idira Blueprint takes you through two recommended approaches based on real-world experience. Both are risk-based and maximize the impact for the level of effort. One focuses on prioritizing security controls and the services you might need to secure the cloud as a whole. The other strategy looks at prioritizing the identity or role that needs to be secured based on their privileges, population, and level of risk.

---

## Security Control-Based Prioritization

In this prioritization method, organizations focus on a single security control family at a time. Not every organization has solutions or services to implement all the security controls for each identity. In this approach, organizations prioritize by the type of security control they want to apply, which typically correlates directly to a service or solution, and the risk impact and effort level required to mitigate the risk.

Above all else, always secure your root and registration accounts first. Even a simple, temporary one-time passcode (TOTP) MFA application, like what's available in your typical mobile authenticator application, will be valuable here. We don't need to implement full PAM controls to begin the process of mitigating the risk of these sign-up accounts.

### Controls That Support Progress Toward ZSPs

The first type of controls to implement are those supporting progress toward ZSPs. This includes functionality like role-based federated JIT access, ZSP, MFA, session protection, session recording, and audit. Grant privileges on demand or as needed to start, and over time, and refine them with least-privileged enforcement. Controls that enable ZSP have the highest priority because they cover the largest swath of human access. Implement these controls early to avoid the sprawl of freestanding access that can quickly accumulate as cloud footprints grow. Roll out these controls to identities with a risk-based mindset: IT admins, developers, other service administrators, like networking or DBA roles, and those with read-only access.

### Standing Privileged Access Controls

The second control family to implement is standing privileged access controls. These include functions like credential vaulting, password, and key management and rotation, complex password policy, MFA, session isolation, session monitoring, and audit. Here, first return to the root and registration account passwords and access keys, and then move into the freestanding access (passwords and keys) for those same identities listed before. Strive to minimize freestanding access at all costs, but when required or necessary, controls like credential management are critical. Work toward refining the privileges for nonbreak-glass emergency accounts over time, focusing on well-known roles first, enforcing least privilege.

### Secrets Management Controls

Secrets management includes controls like secrets vaulting, secrets rotation, complex secret value policy, removal of hard-coded secrets from applications, dynamic secrets, and JIT secret delivery to those apps. These controls build on the foundation of PAM and require an additional discovery and prioritization effort, which is why they follow PAM controls. Refine the privileges for machine workloads over time to enforce least privilege. Apply these controls to any machine passwords and keys that are consumed by workloads, scripts, or services to mitigate the risk of identity compromise and privilege abuse.

### Identity Governance Controls

Identity governance controls consist of lifecycle management and compliance mechanisms. Lifecycle management is the process of granting authorized users access to the resources they need (via the appropriate control plane) while hiring and revoking their access when they no longer require it—a traditional Joiner-Mover-Leaver process.

## Controls Rollout

Roll out these controls to the explicitly defined IT admin roles first, followed by the developers, and then the other privileged roles. Lifecycle management requires the explicit definition of roles to be effective, so organizations should focus on the well-known roles first. Identity compliance is about periodically certifying and attesting that users still require access to the things they currently can access and if not, revoking that access. Roll out identity compliance controls simultaneously across all human users. This set comes last in the list of controls because users must first have access to resources via control planes to conduct compliance campaigns against them.

Zero Standing Privilege	Secure Standing Privilege	Secrets Management	Identity Governance
<ul style="list-style-type: none"> <li>Least privilege role-based access</li> <li>Zero standing privileges</li> <li>Federated just-in-time access</li> <li>Adaptive multifactor authentication</li> <li>Session protection and audit</li> </ul>	<ul style="list-style-type: none"> <li>Credential vaulting</li> <li>Rotation and isolation</li> <li>Multifactor authentication</li> <li>Session isolation, monitoring, and audit</li> <li>Protection for shared and breakglass emergency accounts</li> </ul>	<ul style="list-style-type: none"> <li>Secrets vaulting</li> <li>Rotation and complex secret value policy</li> <li>Removal of hard-coded secrets</li> <li>Dynamic secrets and just-in-time secret delivery</li> </ul>	<ul style="list-style-type: none"> <li>Joiner, Mover, Leaver processes</li> <li>Grant authorized identity access at the time of hire</li> <li>Revoke access at time of change or departure</li> <li>Certification and attestation processes for privileged access</li> </ul>

## Identity- and Persona-Based Prioritization

In this prioritization method, organizations can apply multiple security control families simultaneously to a persona or type of access. Not every organization has that capability, but the benefit is that this approach is exclusively risk-based. Even if your organization does not have that capability, keep this information in mind as you build your cloud security strategy.

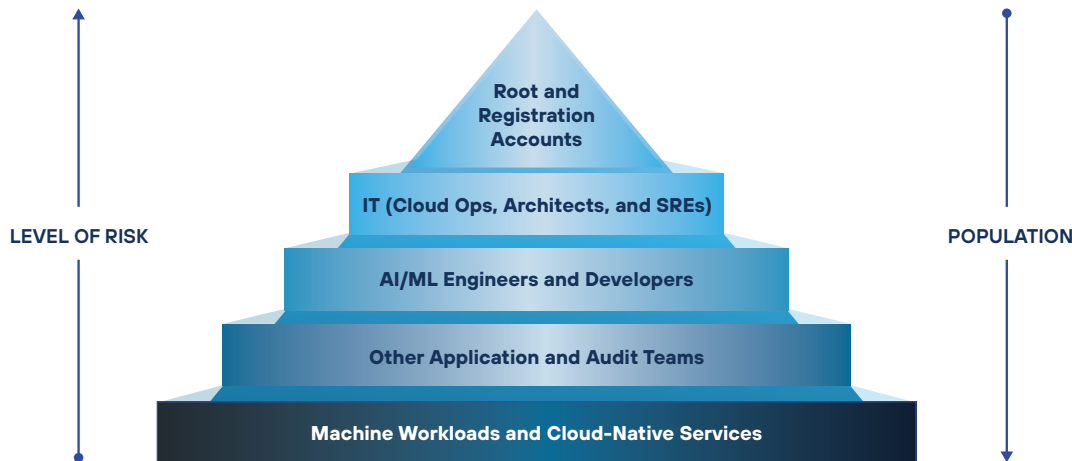


Figure 8. Identity and persona risk prioritization pyramid

This prioritization method assumes the organization is applying all security controls to each type of identity at the same time, leaving no gap or hole in any of these types of access. This theoretical ideal scenario helps when you have the capability to deliver comprehensive identity security controls simultaneously.

---

## Root and Registration Accounts

Always start with the root and registration accounts. Each one represents the email address used to register for the cloud service provider account. Only one account exists for each cloud service provider account you have, so the population is small. These accounts can do anything on your cloud service and should have a minimum of vaulting and MFA applied, but ideally have additional secure standing privilege controls as well. In AWS, this is your root, organizational root, or member root user. In Azure, this is the global administrator you signed up with. In Google Cloud, this is the owner you signed up with. For Azure and Google, you can likely identify this account by looking for a shared corporate email address (e.g., LoBCloudProvider@example.com).

## IT Administrator Personas

These personas are typically internal roles such as cloud operations, cloud architects, and site reliability engineers. They are usually granted full-fledged administrative access across CSP accounts. These identities have the highest risk because they have explicit access to control all aspects of IAM and permissions, meaning they can create new identities, modify permissions, and manipulate all aspects of the CSP and its resources. They have the ultimate privileged access to affect every service and resource within the CSP account. These users should have ZSP access to the CSPs and should be protected in addition with least privilege and identity governance controls.

## Developers and Service Administrators

Focus on securing your developers and service administrators who have privileged access to various services and resources within the CSPs, like serverless compute, database, or secret vaults. Unlike the IT Admins, these users typically don't have power across the CSP itself, but they can be powerful for the services they access or administer. Since CSPs have hundreds of services, many identities with this type of access need to be secure. These users should have a combination of secure standing and operational ZSP access with least privilege, identity governance carefully considered.

## Other Application and Audit Teams

Following developers, apply the same controls to other application and audit teams and users with fewer privileges, like read-only access, to various services.

## Machine Workloads, Cloud-Native Services, and Other Application Identities

Lastly, prioritize machine workloads, cloud-native services, and other application identities. Many machine identities are likely to be in your cloud service provider accounts, but many of them are not likely to have cloud admin, service admin, or resource admin permissions. Your automation and orchestration workloads are most likely to have higher-risk or more sensitive permissions. These machine identities should have a combination of secrets management and least privilege controls.

## Prioritize What Makes Sense for Your Organization

Every organization has unique prioritization needs. Looking at this list, you might see areas that you want to implement first before the others. In that case, choose the order that makes the most sense for your organization. Whatever is driving your initiatives and is important to your organization's goals should be your priority. Consider this prioritization guidance and rationale so you understand the tradeoffs and make the most informed decision for your organization.

**Table 5. Security Control Prioritization**

Identity/Controls	Zero Standing Privilege	Secure Standing Privilege	Identity and Persona Prioritization and Control	Identity Governance
Root & Registration Accounts		X		X
IT (Cloud Ops, Cloud Architects & SREs)	X	X		X
Developers and Service Administrators	X	X		X
Other Application and Audit Teams	X	X		X
Machine Workloads and Cloud-Native Services			X	X

## Alignment with Cloud Adoption Strategies

Most organizations fall into one of two cloud adoption patterns, which will influence their overall prioritization and security strategy. Recently established organizations tend to favor the more recent advancements in cloud services. They build their business applications and processes on native services in the cloud, which is why we refer to their adoption path as **digital native business (DNB)** and **digital native enterprises (DNE)**.

Organizations that have existed longer and previously used traditional IT services must undergo a much larger IT modernization strategy to leverage the value of the cloud, migrating their once self-hosted resources and applications to the cloud. These types of organizations have an adoption path called **lift and shift**.

As these two adoption paths vary significantly, their approaches to cloud security vary as well. Each of these cloud adoption patterns influences an organization’s identity security strategy.

### Digital Native Businesses and Digital Native Enterprises

DNBs and DNEs use cloud-native services to build the applications that run their digital businesses. The personas that need to be protected in these organizations are cloud architects, developers, and product owners in site reliability engineering, cloud engineering, and architecture teams under the CIO and CDO organizations.

In this scenario, DNBs and DNEs are likely to prioritize securing access with ZSPs. Key controls to implement include role-based access control, federated JIT access, MFA, session recording, and auditing. These controls provide immediate risk reduction to the organization while simultaneously delivering efficient access methods to enable these teams to build the applications and services their businesses require.

Extend their security posture to machine identity workloads with secrets management controls. Focus on workloads like cloud-native services (e.g., serverless compute functions and container services), DevOps pipelines, and source code repositories (Chef, Puppet, GitHub, and GitLab), as well as any custom code that uses embedded secrets. DNBs and DNEs tend to favor federated access whenever possible, so freestanding accounts and keys are typically kept to a minimum. However, do not overlook these freestanding accounts. A holistic strategy incorporates PAM controls to protect credentials that grant the necessary standing access as well.

## Lift-and-Shift Organizations

The lift-and-shift cloud adoption strategy is prevalent among organizations with a substantial investment in existing on-premises infrastructure and applications. These organizations, often bound by historical IT decisions, opt to migrate their current systems to the cloud with minimal changes. This approach enables them to capitalize on the cloud's scalability, flexibility, and cost-efficiency without the need for immediate, extensive redevelopment of their applications. The personas that typically need to be protected include IT administrators, application support teams, and developers who are responsible for the maintenance, operation, and security of migrated systems. These teams' reporting structures tend to be much more dispersed than digital native organizations.

In this scenario, lift-and-shift organizations are likely to have implemented some form of existing PAM controls onto their self-hosted, on-premises resources, which they will likely be migrating along with the resources as they move to the cloud. Ideally, simultaneously, lift-and-shift organizations should seek to further protect themselves with the same ZSPs, federated access controls to the cloud providers themselves that their digital native counterparts are implementing to achieve immediate risk reduction. From there, they can continue to evaluate whether to implement more operationally efficient security controls for their VM and infrastructure access or expand to protect machine identity workloads.

## Operationalizing Cloud Security at Inception with Infrastructure as Code

The cloud security concept of long-term operationalization relates to how you ensure all newly created identities, cloud provider accounts, and resources will inherit the same security controls as deployed for the existing identities. It also addresses how you ensure no gaps exist in your security posture for new tools and solutions. In either of the two cloud adoption strategies, security teams must solve this operationalization problem.

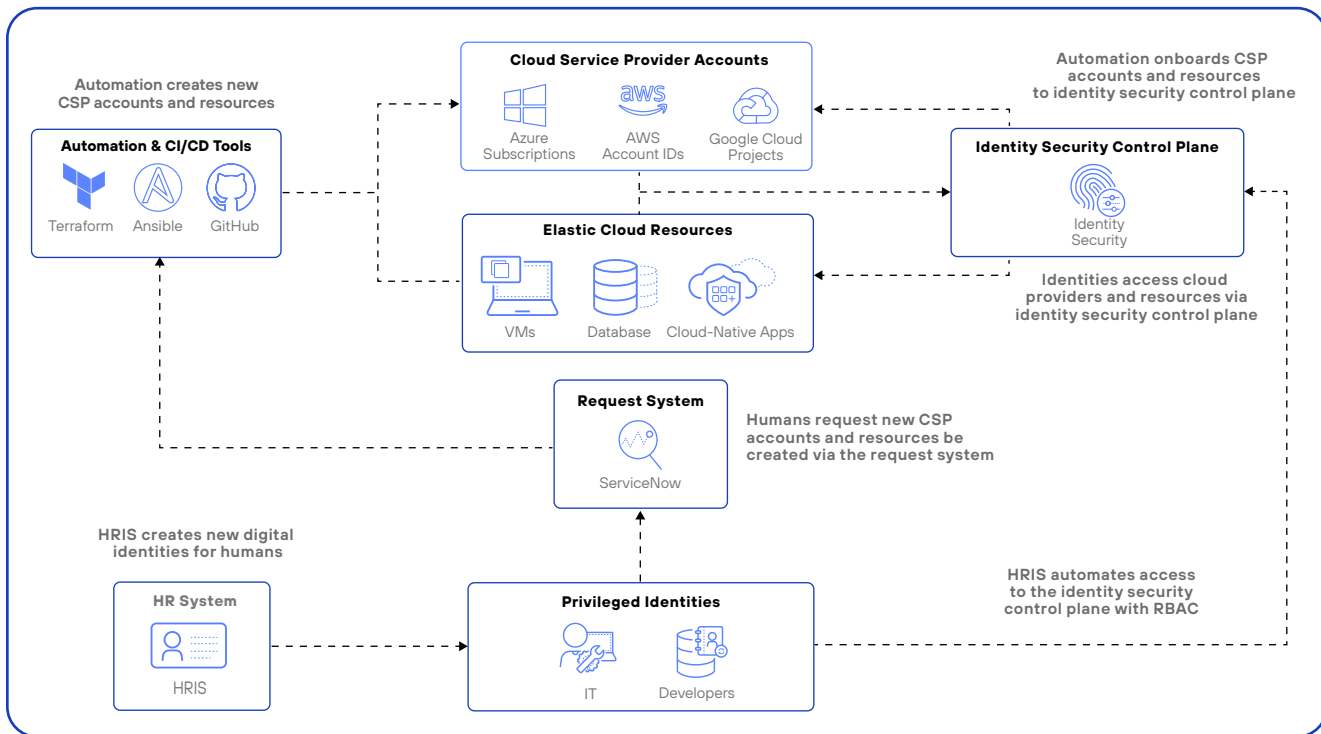


Figure 9. Security at inception in the cloud

---

The term **operationalizing cloud security at inception** highlights the critical importance of securing access to resources from the moment they are provisioned or created, using the power of automation. As organizations continue to automate their IT infrastructure provisioning, including the creation of new cloud service provider accounts, virtual machines, databases, and other critical components, they must secure access at inception. This approach both streamlines the provisioning process and significantly reduces the risk of manual errors and security vulnerabilities that malicious actors can exploit.

Integrating your existing automation services, infrastructure as code (IaC) tools (e.g., Terraform or Ansible), and HRIS systems (e.g., SuccessFactors or BambooHR) with your identity security control plan enables the automatic application of intelligent privilege controls at the onset. This security-first approach applies to the protection of CSP accounts, the resources they create, and any built-in local accounts related to them. It effectively protects access for any identity, from developers to operations teams, and other stakeholders interacting with cloud resources. By automating the provisioning of intelligent privilege controls, organizations can ensure a consistent application of security policies across all cloud resources, enhancing the overall security posture.

Organizations can tap into their existing automation processes that use services, like AWS Control Tower Account Factory or Google Cloud Deployment Manager, adding a step to onboard the relevant resource to the appropriate identity security service. Operationalizing cloud security this way presents a proactive and automated approach to access. By embedding security controls and access management from the initial stages of resource provisioning, organizations can achieve a more secure, efficient, and compliant cloud environment. This strategy enables organizations to maintain agility and innovation while ensuring that their cloud infrastructures are protected against the evolving landscape of cyberthreats.

## Concluding Insights

- **Inclusive security strategies:** Cloud security encompasses developers, cloud operations, IT administrators, service administrators, and machine identities that have high-risk access. This mix of identities necessitates strategies that address the diverse needs and access patterns of these personas and roles.
- **Risk-based approach:** Consider a pragmatic strategy to secure cloud identities, focusing efforts on areas with the most significant impact based on risk.
- **Efficiency with Idira Identity Security Blueprint:** Idira Blueprint makes implementing cloud security measures more efficient and effective, helping organizations quickly secure privileged entities within their cloud environments.
- **Alignment with architectural and compliance frameworks:** Our guidance aligns with the well-architected frameworks and major regulatory frameworks of the major CSPs. It ensures applicability and compliance within organizations' existing cloud architectures.

## Next Steps for Securing Your Cloud Identities

No organization can secure their cloud overnight. The complexity of CSPs and their access methodologies is too great. However, by following the prioritized guidance of Idira Blueprint, your organization can develop an informed, risk-based plan with the best odds of success.

See how Idira Blueprint can help you build the right framework to secure your cloud identities.

[Request a demo.](#)



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](http://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.

idira\_wp\_best-practices-for-securing-cloud-identities\_050826