CYBERARK®
The Identity Security Company™

WHITEPAPER

Best Practices for Securing Cloud Identities

A CyberArk Blueprint Whitepaper

Table of Contents

Executive 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. This whitepaper, leveraging the CyberArk Blueprint, presents a comprehensive framework for securing cloud identities using a holistic approach that includes both human and machine identities and advocates for a practical, risk-based strategy to enhance cloud security postures.

What You Will Learn

- **Cloud Security Landscape:** The evolving cloud landscape poses challenges and this whitepaper stresses on the importance of dynamic identity security strategies due to the proliferation of services and identities, coupled with the diminishing traditional security perimeters.
- **Risk-Based Guidance:** The CyberArk Blueprint, a battle-tested best practices framework designed to reduce risk with prescriptive guidance for all types of identities, will give a focused view through the lens of cloud service providers and cloud resources.
- **Critical Identity Security Controls:** Identity security is the cornerstone of a successful cybersecurity strategy with controls aimed at preventing identity compromise, stopping lateral and vertical movement and limiting privilege escalation and abuse.
- **Minimize the Attack Surface:** The goal is to minimize attack surface within cloud environments while advocating for zero standing privileges (ZSP), just-in-time access and entitlements provisioning as part of your identity security program.
- **Prioritization Strategies:** There isn't a one-size-fits-all path to adopting identity security controls in the cloud. There are multiple approaches, taking into consideration existing capabilities, risk-based prioritization, compliance requirements and cloud adoption approaches.

Key Takeaways

- **Inclusive Security Strategies:** Cloud security encompasses developers, cloud operations, IT administrators, service administrators and machine identities with high-risk access. This necessitates strategies that address the diverse needs and access patterns of these personas and roles.
- **Risk-based Approach:** A pragmatic strategy for securing cloud identities should be considered, focusing efforts on areas with the most significant impact based on risk.
- **Efficiency with CyberArk:** The CyberArk Blueprint can make 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 both the major cloud service providers (CSP) well-architected frameworks and major regulatory frameworks, ensuring applicability and compliance within existing cloud architectures.

In summary, this whitepaper is a vital resource for enterprises aiming to secure their cloud environments against increasing complexity and threats. By adopting the CyberArk Blueprint, organizations can ensure a comprehensive, risk-based approach to cloud security, pursuing zero standing privileges to effectively safeguard their cloud workloads and services.

Cloud Has Changed Security

Cloud Service Providers and the CyberArk 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 ever-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 the offerings of cloud service providers increase and the shared responsibility model continues to put ownership on the customer, organizations will be required to bring their own safeguards and protections to the cloud. Each cloud provider has its own shared responsibility model with clear guidance that customers are responsible for secure configuration within their own environments.

This responsibility is extremely challenging at multi-cloud scale; Amazon, Azure and GCP 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, so do the security risks for identity security programs. Expansion creates a proliferation of human and machine identities, each of which can be configured with an ever-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. In this whitepaper, we attempt to develop a common language around "cloud identities" and provide comprehensive guidance across the cloud providers (inclusive of their various services) and their elastic resources, infrastructure and workloads. Much of this common language and comprehensive guidance comes from our best practices framework for identity security success, the CyberArk Blueprint. This framework was designed to help organizations measurably reduce risk, based on our lessons learned in battle, providing prescriptive guidance to secure any identity, human and machine.

In subsequent sections, you will see a comprehensive set of guidance that's relevant to securing your cloud identities from the CyberArk Blueprint, and blend that with other important cloud security best practices, enabling you to develop a holistic, dynamic and prioritized cloud security strategy – and do it fast.

Challenges with Identity in the Cloud

While all the 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 exist in the cloud as well. Cloud security uses a shared responsibility model. Cloud service providers 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

CSPs are not responsible for securing your access to their cloud services or your workloads on the cloud. It's up to your organization to secure both of those things.

That means you are responsible for solving the challenges related to four key areas of identity security: authentication, authorization, access and audit.

Authentication: Cloud service providers 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/password or access/API key authentication. Additionally, each cloud service provider offers authentication options in slightly different ways, adding to confusion for end users and security teams.

Authorization: Countless permissions are available within each cloud service provider for use by all these different identities. Identities may 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 multi-cloud 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 desire centralized visibility and control, whether that's in a single cloud or multi-cloud architecture. Cloud security may 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 not only makes cloud identities ripe targets for malicious actors to take advantage of, but also almost certainly ensures they'll find a way to do it. The same benefits cloud provider accounts grant (such as centralized and simplified deployment, all-in-one data and application hosting) have become the same reasons bad actors target them. There's lot of juicy information there for attackers to exploit through data exfiltration, ransomware or service disruption.

This is why it's so important to secure all of your cloud identities. But how do you secure them and their access methods? And how can you be pragmatic about where to focus your efforts? This is where the CyberArk Blueprint's prescriptive guidance comes into play. It helps you easily identify the types of privileged entities that exist 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

Management Platform, Services and Workloads

In order to better understand the challenges organizations face when securing their cloud provider environments, we need to first understand the anatomy of cloud identities. Whether you're an AWS, Google Cloud or Azure shop, your environment consists of two major parts:

- First, the **management platform and services** allows various identities to administer and operate different services like identity and access, compute or secrets vaults. The management console (whether accessed via the web UI, CLI or API) is the main access point into the cloud provider.
- Second, there are the various **infrastructure workloads** that are created by different cloud services, like virtual machines, containers, serverless functions, cloud-native apps and storage. There are services to administer and create these various workloads, but what's unique about these resources is that they too require a separate layer of identity security controls, separate to that of the platform or services.

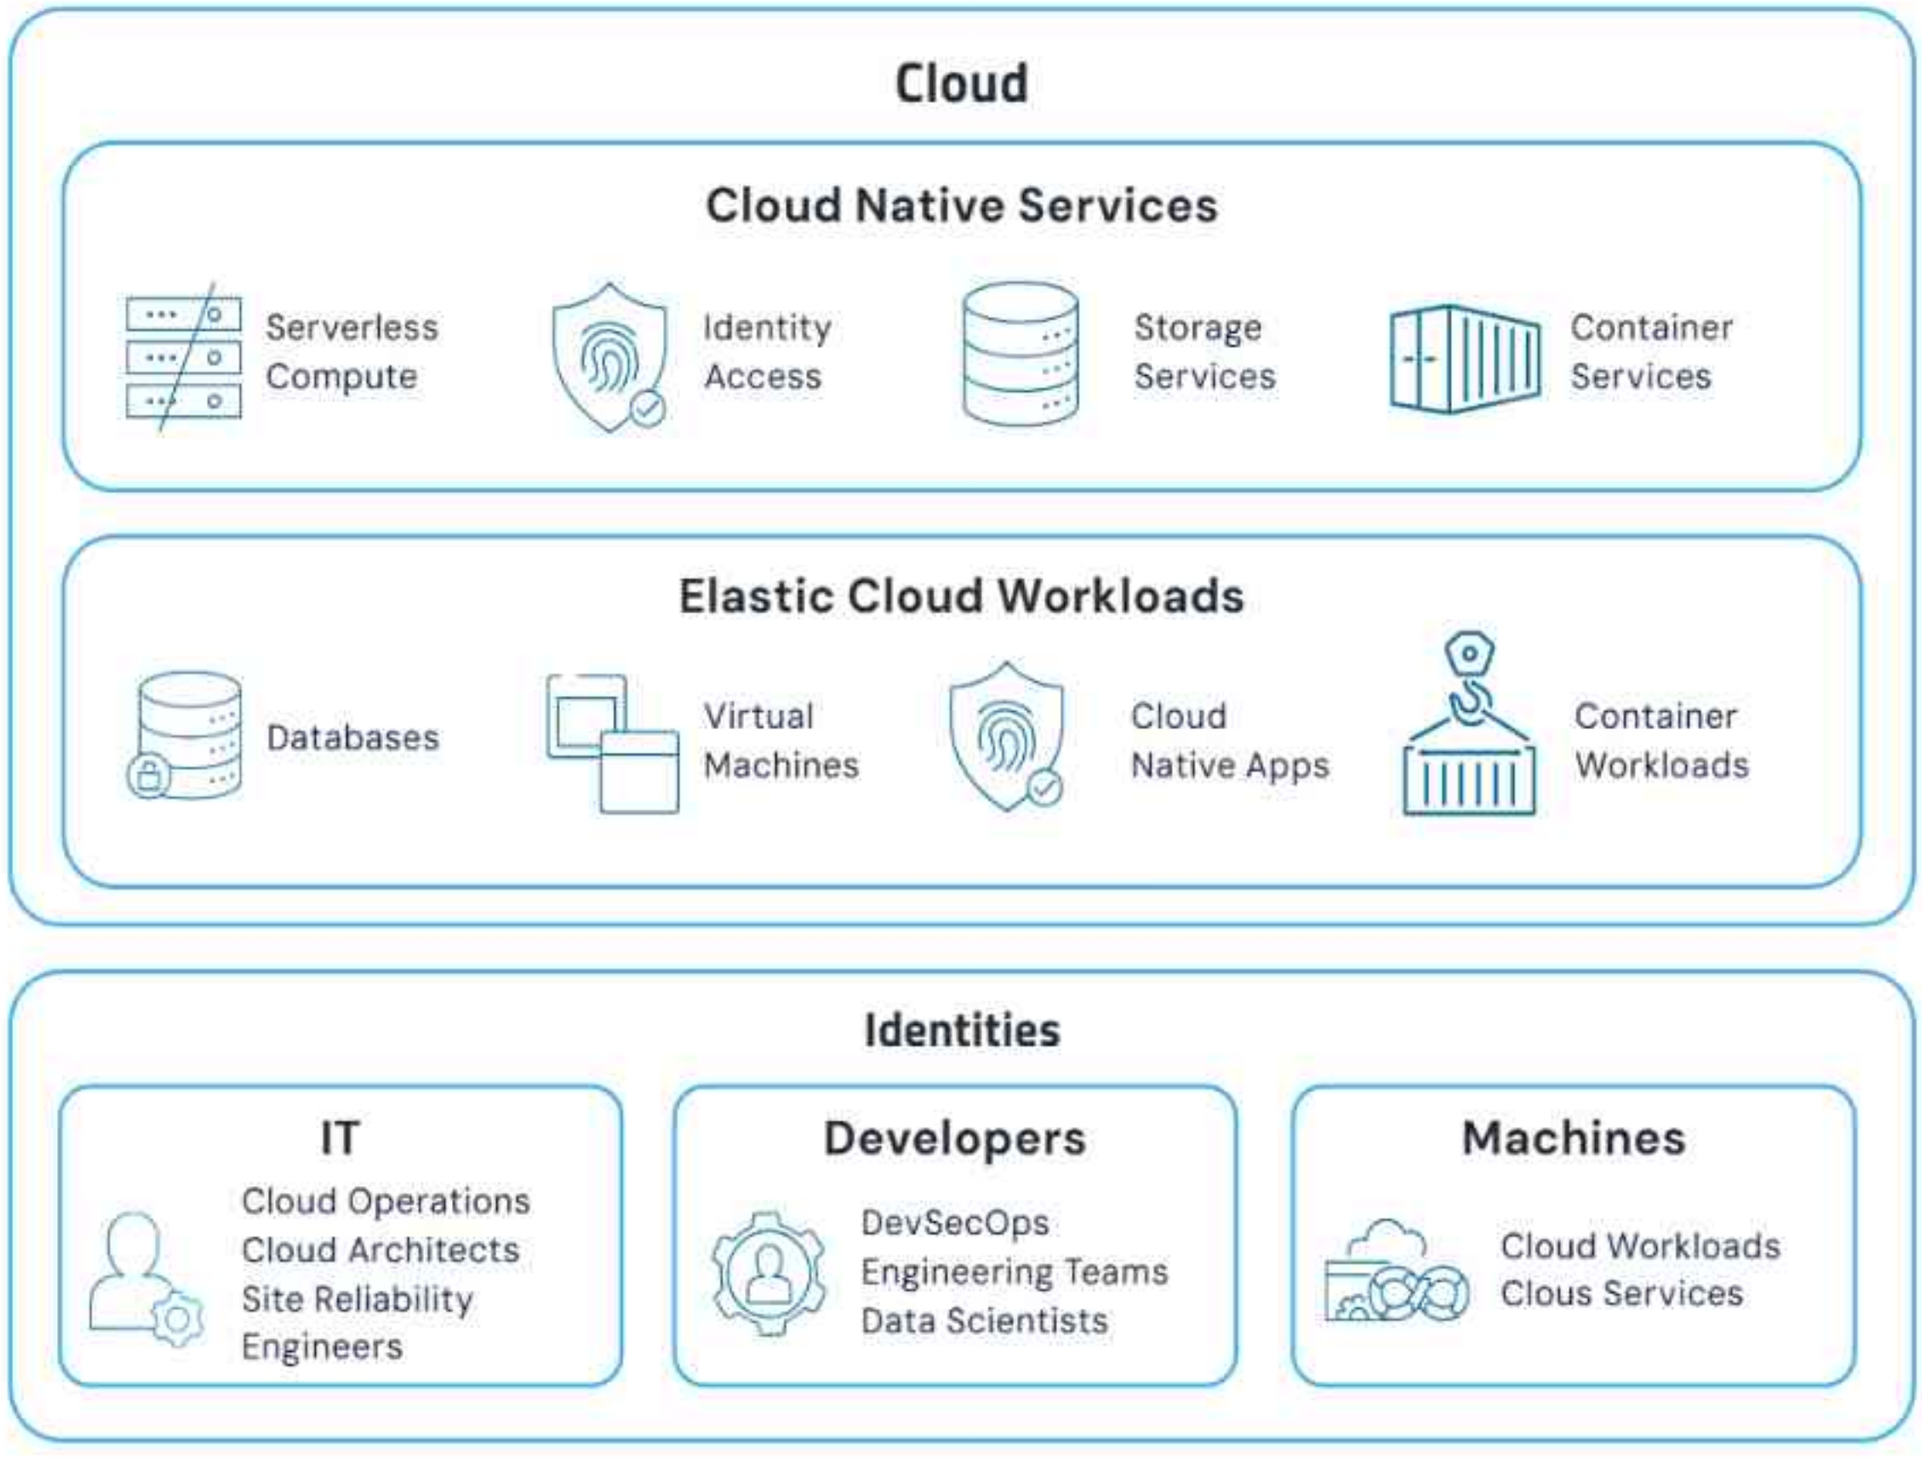

Figure 1: The Anatomy of the Cloud

Cloud Identities

We also have identities that access the management console, various services and resources created. That typically consists of:

- **Cloud operations** that has evolved from traditional IT roles such as infrastructure operations, networking engineers or database admins into roles like cloud operators, architects and site reliability engineers.
    - Within cloud operations, there are **full cloud administrators**, those with complete administrative access and who have the ultimate permission to affect every service and resource within the CSP account.
    - Additionally, there are also **service-level admin roles**, such as engineers with a specialization on networking or databases, that can only administer a smaller scope of service(s) and/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.
- The various **machine identity workloads**, such as cloud-native applications, services, automation tools and processes that run your business.

All these identities authenticate to the cloud using a variety of methods, including standing federated access via an identity provider (IdP), long-lived freestanding local accounts like user passwords and keys or (in case of emergency) using the root or registration credentials.

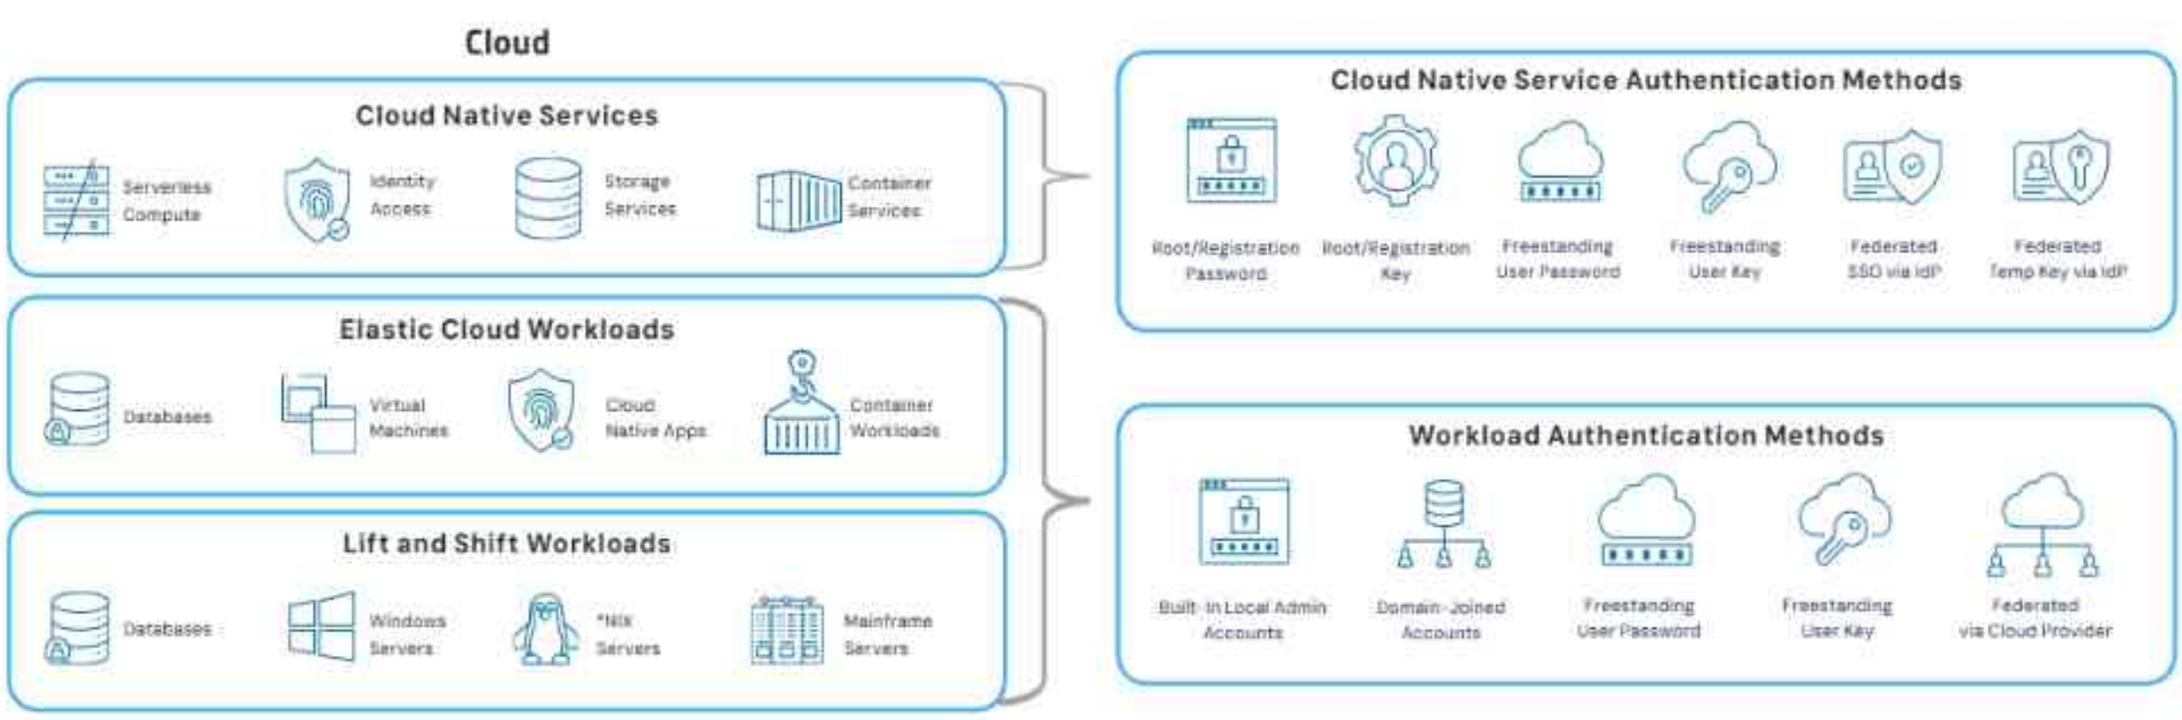

Figure 2: Cloud Provider Services/Workloads and Common Authentication Mechanisms

Similarities and Differences Among the Cloud Provider Access Models

There are three dominant identity security concepts that are worth noting the similarities and differences among the various cloud service providers: federated access, non-federated access and permission models.

Federated Access

When we refer to federated access, we're referring to identities that are created within a centralized directory service bound (federated) to the cloud service provider or its workloads. Federated access is the standard enterprise approach to granting access to public cloud and SaaS resources for a variety of reasons (simplicity, administration, security, etc.). This makes securing federated access a key underlying element of establishing a secure cloud environment.

Federated access models vary slightly across all the major cloud service providers. Google Cloud requires the use of one of two top-level Google directory services to establish federation: Google Workspace or Google Cloud Identity. Both of these services share the same backend directory structure and act as GCP's 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 you can grant access to GCP projects. Access and permissions granted to these projects are persistent.

Federated access in Azure uses Entra ID, but it also supports bringing your own identity provider, including support for additional IdPs via B2B integration support. Without the use of Azure PIM, federated access and the assigned permissions are standing (persistent and long-lived).

When using Amazon Web Services (AWS) Organizations to deploy Account IDs at scale, it is considered a best practice to leverage AWS IAM Identity Center (Amazon's IdP service). However, you can still integrate your existing identity provider into AWS IAM Identity Center. Federated access into non-organizational AWS accounts can leverage any identity provider directly without requiring integration to Identity Center. Another major distinction for AWS is that it leverages role assumption for just-in-time access, meaning that the accesses are not freestanding, but the permissions to the roles are.

Access Model	AWS	Azure	GCP
Federated Access	**AWS Organizations:** Any identity provider indirectly via IAM Identity Center **Non-Organization:** Any identity provider directly into IAM (or indirectly via IAM Identity Center)	Any identity provider Requires sync/IdP config into EntraID	Any identity provider Requires sync/IdP config into Google Workspace (Google Cloud Identity)

Figure 3: Cloud Service Provider Federated Access Table

Non-Federated Access

When we refer to non-federated access, we're referring to identities that are created locally within the cloud service provider's platform. This includes both human and machine access scenarios. Non-federated access is commonly used for 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. Non-federated access is mostly similar across cloud service providers, with the exception of Google.

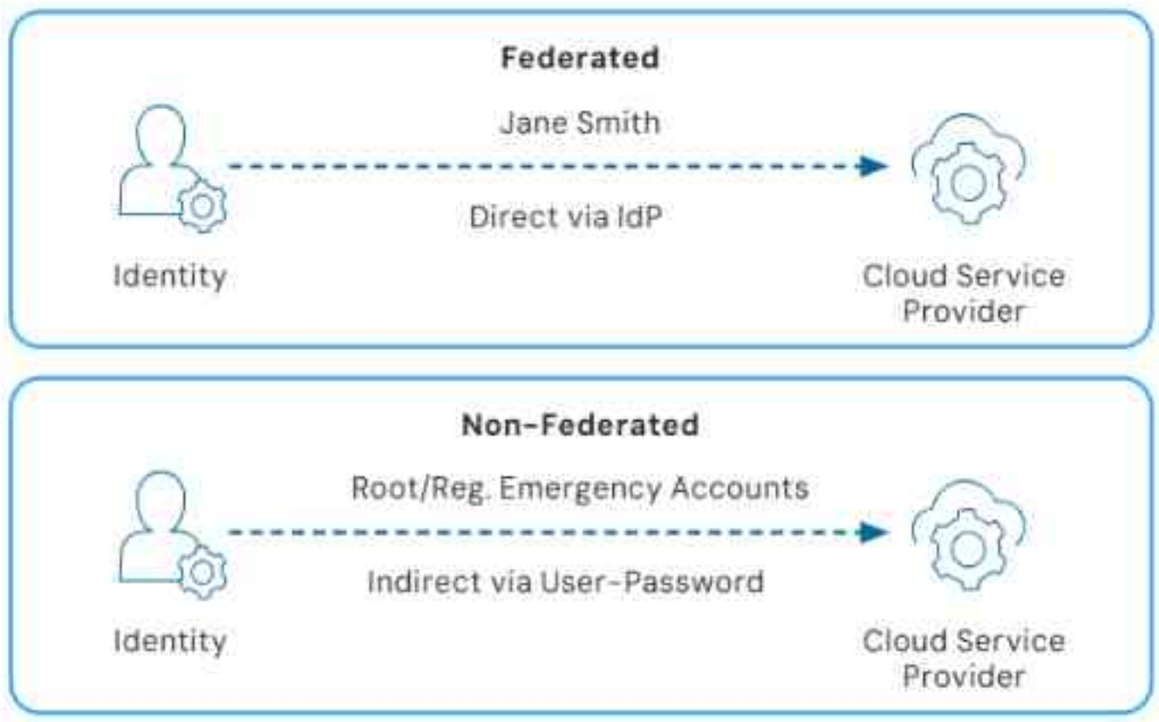

Figure 4: Federated vs. Non-Federated Access

For human access, AWS allows you to create internal IAM directory users directly within the AWS platform, although this is not a recommended best practice to implement at scale. Azure is similar and also allows you to create non-federated internal Entra ID directory users. These are considered "local" accounts because they exist solely within the cloud service provider platform. GCP does not include the capability to create non-federated local directory users within a project or organization directly, as that requires the

use of Google Workspace for all humans. However, you can create local (non-directory federated) Google Workspace or Google Cloud Identity users and grant them access to GCP projects.

For machine identity access, it gets a bit more complicated because we're dealing with access for workloads, applications, scripts and services that exist both inside and outside the cloud service provider platform. AWS allows you to leverage 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 allows you to create non-federated 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 concept called Managed Identity, which resides in the Azure subscription and creates a service principal in the Entra ID tenant.

GCP uses internal IAM services within the project (instead of Google Workspace) for creating service accounts, which is similar to a traditional directory-based account and credential when compared to AWS and Azure's methodology. This is more in line with the external application approach that AWS and Azure use but is also used for internal applications.

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

Figure 5: Cloud Service Provider Freestanding Access Table

Permission Assignments and Authorization

Permission assignment and the granting of permissions are where the biggest differences lie among the cloud service providers.

- AWS IAM leverages the concept of a policy, either an Identity Policy or Resource-Based Policy, to control an entity's permissions. Policies can be applied directly to an IAM user, group or role, or applied via a permission set.
- When leveraging AWS IAM Identity Center (formerly AWS SSO) the concept of a permission set is used. Permission sets are defined by AWS policies, and then are mapped to AWS accounts. This allows for any assigned users and groups to access appropriate accounts in your AWS organization.
- Azure, on the other hand, uses a different concept of "role." In Azure, roles are what permissions get configured in, and then entities, like a user or group, get assigned roles.
- GCP is more in line with Azure here, too, leveraging the same concept of a role. Think of a permission policy in AWS being the equivalent of a role in Azure and GCP.
- In both Azure and Google Cloud, there are options for both built-in and custom roles, compared to AWS which encourages the use of custom roles.

For ease of understanding, we've included the diagram below as a reference for how the entities and permissions are structured with the CSP-specific terminology.

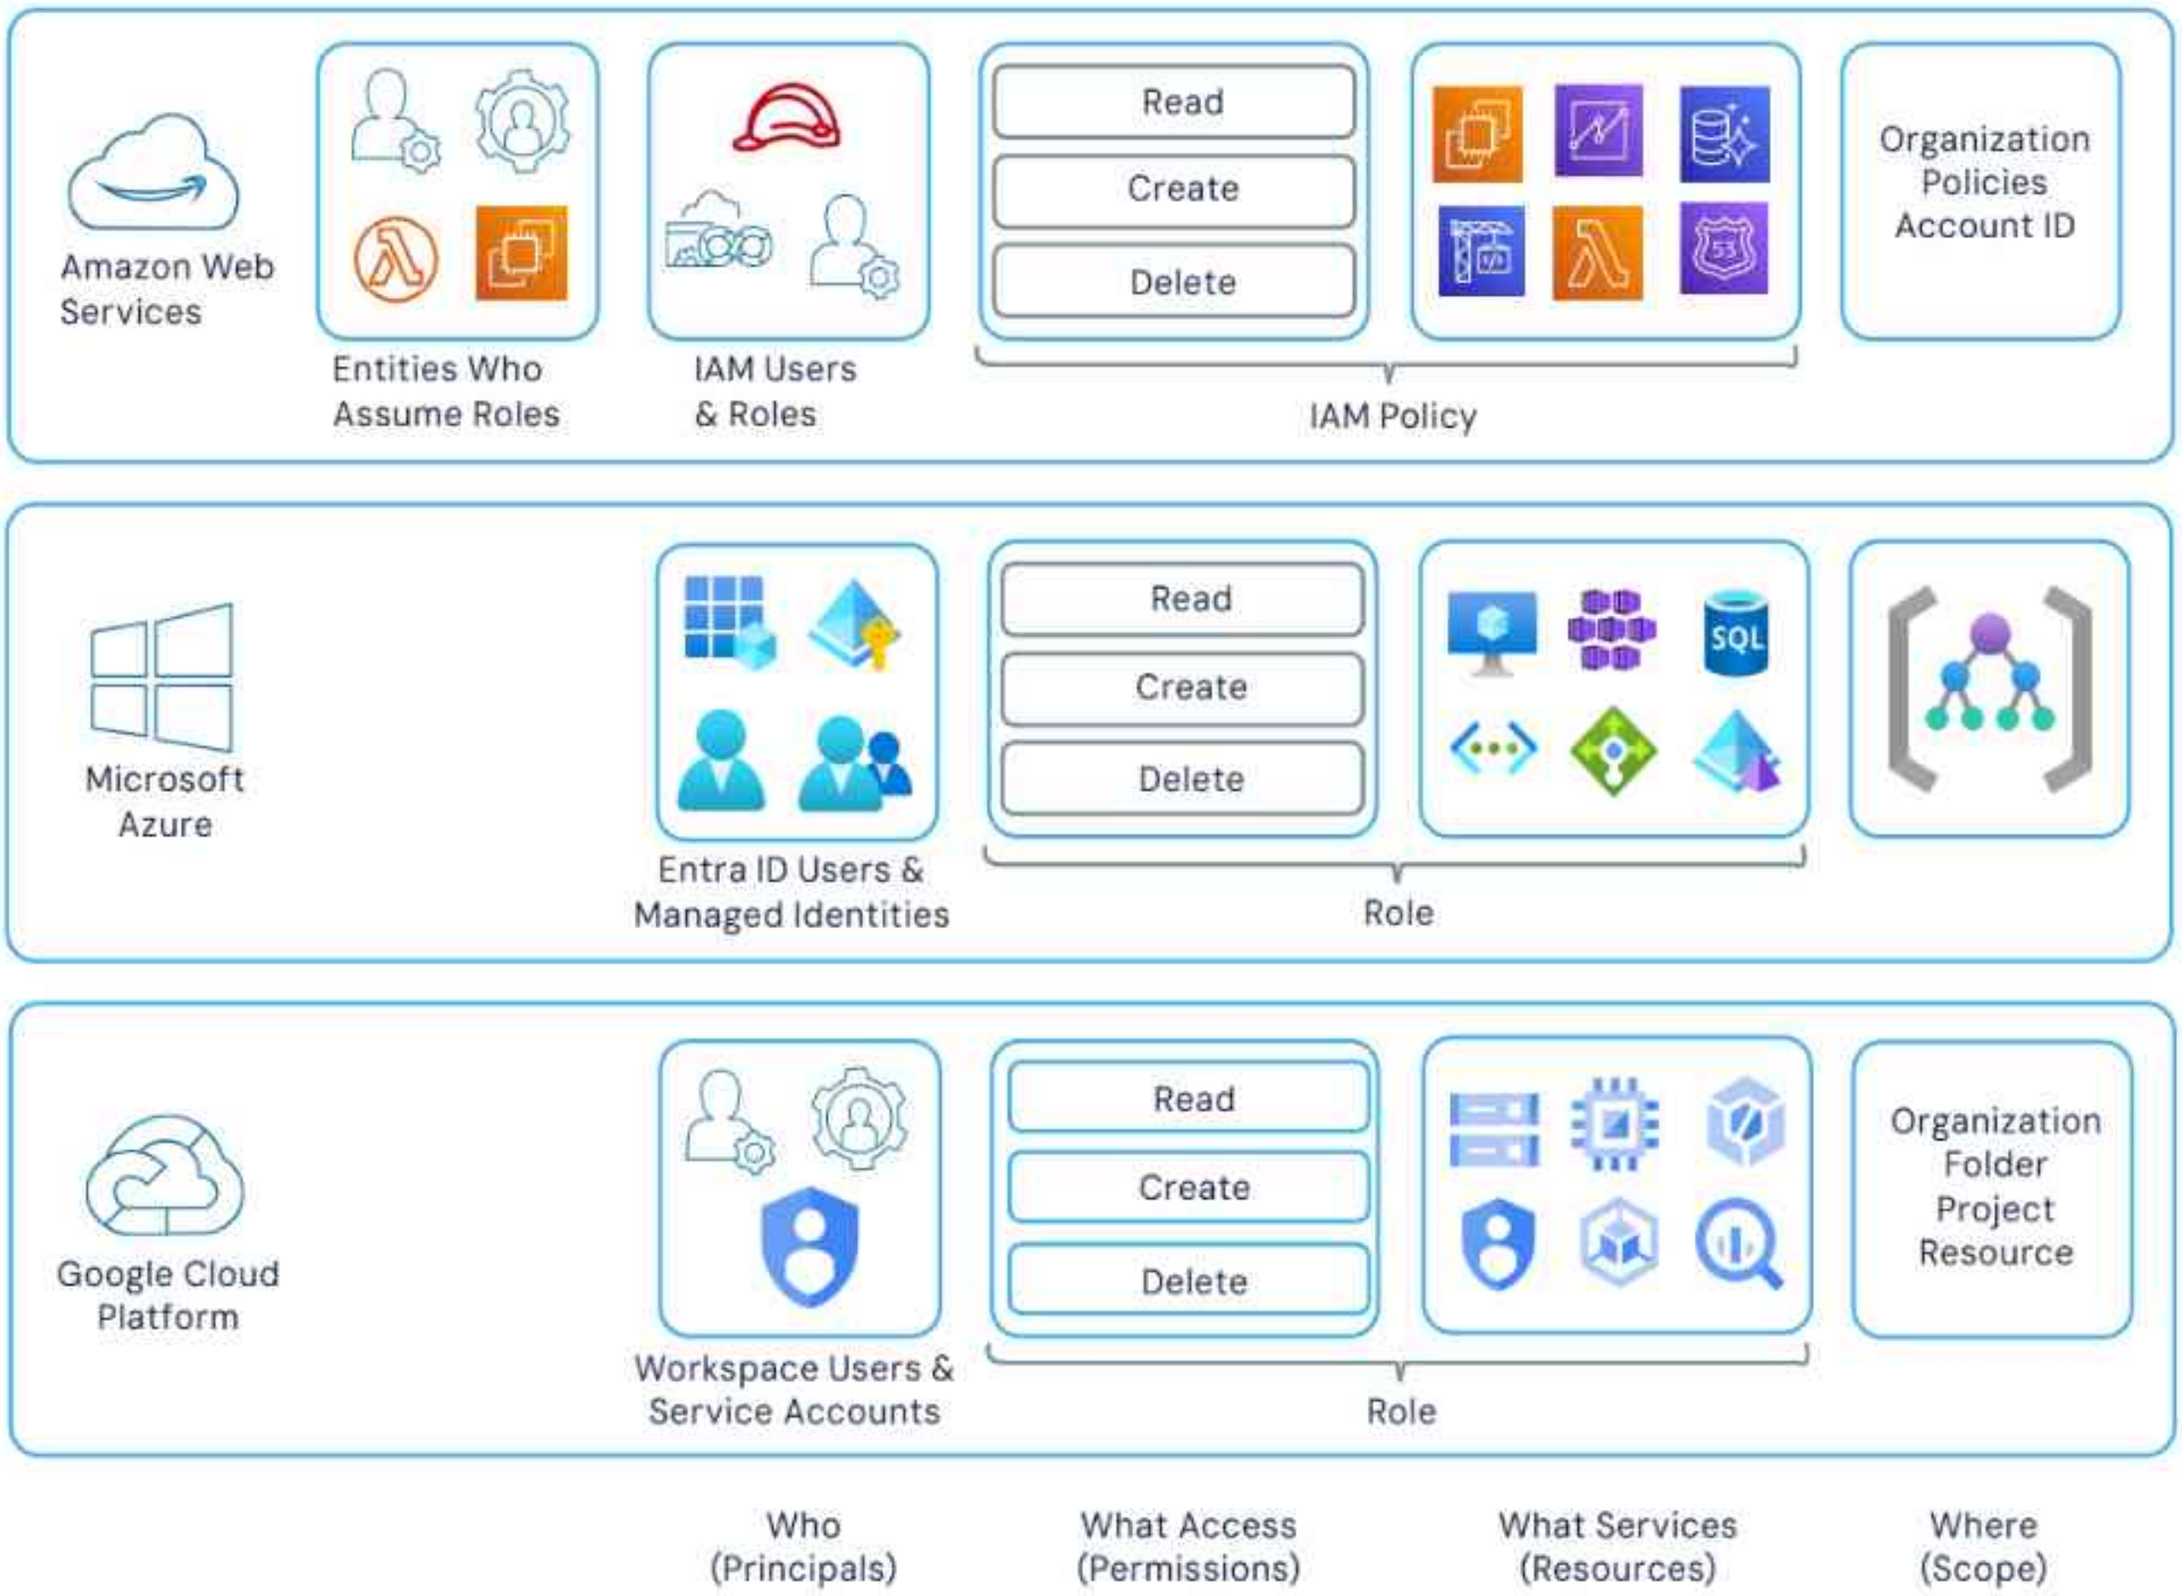

Figure 6: Cloud Provider Permission Assignment Diagram

Permission Scope and Structure

The last major category of differences worth diving into is that of permission scope: at what level of hierarchical structure can permissions be applied to and inherited for. This is an important concept that impacts how we secure identities within our cloud service providers, as we may be granting permissions at one or more levels.

When utilizing AWS Identity Center, users, groups, and permissions are administered at the management account level. This is where permissions are mapped to member accounts. If not using Identity Center, permissions are assigned locally within the respective billing account (AWS Account ID). Azure and GCP, on the other hand, can be assigned at multiple levels. Azure can assign permission scope at the subscription, management group and root management group levels. GCP can assign permission scope at the project, folder and organization level.

Also worth noting is the concept of "safeguards". Examples of safeguards include configuring security policies to boundaries, sessions, virtual private clouds (VPCs) and resources that help mitigate some of the risks associated with access. Cloud service providers are constantly changing, and this is one area where there is a lot of change. We recommend reviewing your provider's features and documentation to see what's relevant.

All three cloud service providers are capable of assigning permissions at a service or resource level (in Azure's case, also the resource group level), although only Microsoft and Google explicitly call that out when referring to permission scope. Since cloud service providers all use different terminology, it can be confusing to understand how different levels relate to one another. We've included the table below to showcase the different scopes of the cloud service providers and how their concepts relate to one another.

Provider/ Scope	Overarching Level	Group Level	Account Level	Resource Group Level	Resource Level
AWS	Organization	OU	Account	N/A	Resource
Azure	Root Management Group	Management Group	Subscription	Resource Group	Resource
GCP	Organization	Folder	Project	N/A	Resource

Figure 7: Cloud Service Provider Permission Scope Table

Access Models for Elastic Cloud Workloads

Thus far, we're primarily covered access models for the cloud service providers themselves, which all have their own unique intricacies. Access models to the elastic cloud workloads, on the other hand, have more commonality across the providers, and as such, will only be discussed briefly.

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 cloud service providers (CSPs). Federated access to resources is typically integrated in one of two ways: using traditional AD-joined methods (similar to self-hosted infrastructure) or through cloud-native services like AWS Systems Manager Session Manager, Azure Active Directory or Google Compute Engine Authenticate.

Non-federated access to elastic cloud resources is typically reserved for built-in local administrative accounts only, the Windows "Administrator" SID-500 account, *NIX Root UIDO user or RDS master user type-scenarios. Unlike within self-hosted infrastructure, non-federated 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 is dependent on your access method. Permissions are either granted and abstracted via the cloud service provider'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, we also need to understand how bad actors attack the cloud to begin with. The CyberArk Blueprint outlines the common attack path malicious actors (internal and external) take to compromise identities and execute their endgame. First, identities are compromised via a variety of techniques, such as social engineering, MFA bypass, credential theft, or cookie hijacking. From there, bad actors leverage the identity's access to move around laterally looking for more access of power, and move vertically when they can, eventually escalating and abusing the privileges they've obtained. Malicious actors will often target cloud operators, site reliability engineers, developers and cloud engineers, as these users have high levels of privilege. They will also target the workloads and services in the cloud as their privileges are often over-provisioned.

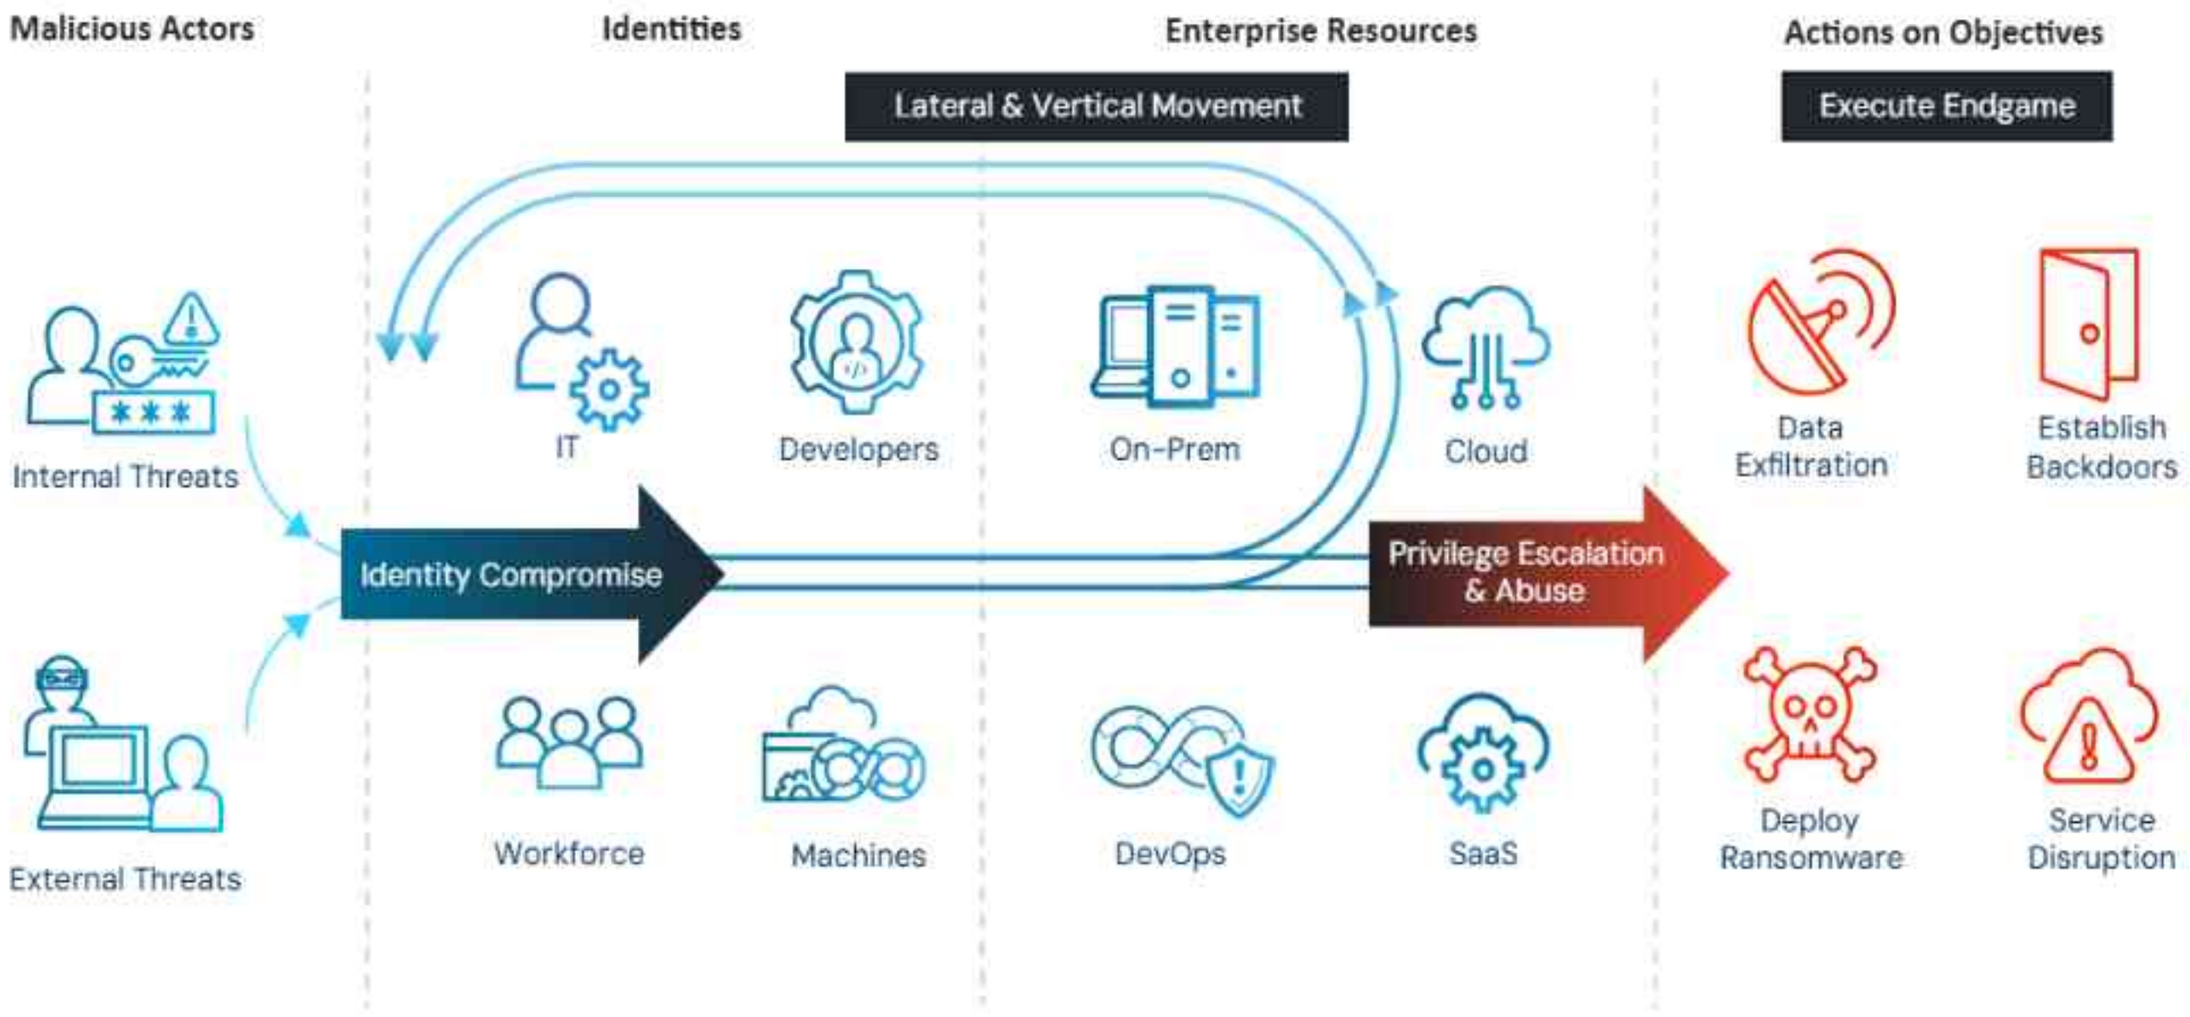

Figure 8: The Identity Attack Chain

For example, a malicious actor may 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. 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 the recommendations below, you'll gain control and visibility into the human-interactive and machine identities with access to your cloud service provider environments and workloads.

To learn more about each specific identity security control and how they mitigate risk, check out our Success Blog article Understanding the Identity Attack Chain with the CyberArk Blueprint!

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.

For federated access to the cloud service provider and its workloads, organizations should strive to achieve zero standing privileges. This can be accomplished using controls such as just-in-time elevation of access, just-in-time assignment of entitlements and limited time-bound durations for the access. This reduces 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 is a fundamental shift away from traditional freestanding federated access via SSO and standing entitlements. Organizations should strive for all interactive access to be with a zero standing privilege approach.

For non-federated access, organizations should strive for the objective of achieving secure standing access. This can be accomplished using controls such as credential vaulting, password and key management and rotation, complex password policy, multifactor authentication, 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 non-federated freestanding access directory user passwords and access/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.

Food for Thought: Holistic cloud security encompasses both the objectives of achieving zero standing privilege for federated access and secure standing access for non-federated access.

While the emphasis is on ZSP for humans, all CSPs and workloads have built-in local admin-type credentials which also require protection.

Furthermore, machine identities with admin access to your CSPs require secrets management controls.

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

Furthermore, workloads and resources hosted within the cloud should have their built-in local administrative accounts (e.g., SID-500 Admin, UID0 Root, Master DBA) protected with these same controls.

Secrets management controls include functions like secrets vaulting, secrets rotation, complex secret value policy, removal of hard-coded secrets from workloads and applications and just-in-time secret delivery and dynamic secrets to those workloads. Secrets management controls build upon 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 controls include two key concepts: the move to privilege on-demand or as-needed 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. This means that 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/modification.

Identity governance controls include both lifecycle management and compliance functions. Lifecycle management enforce 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 their access when those identities no

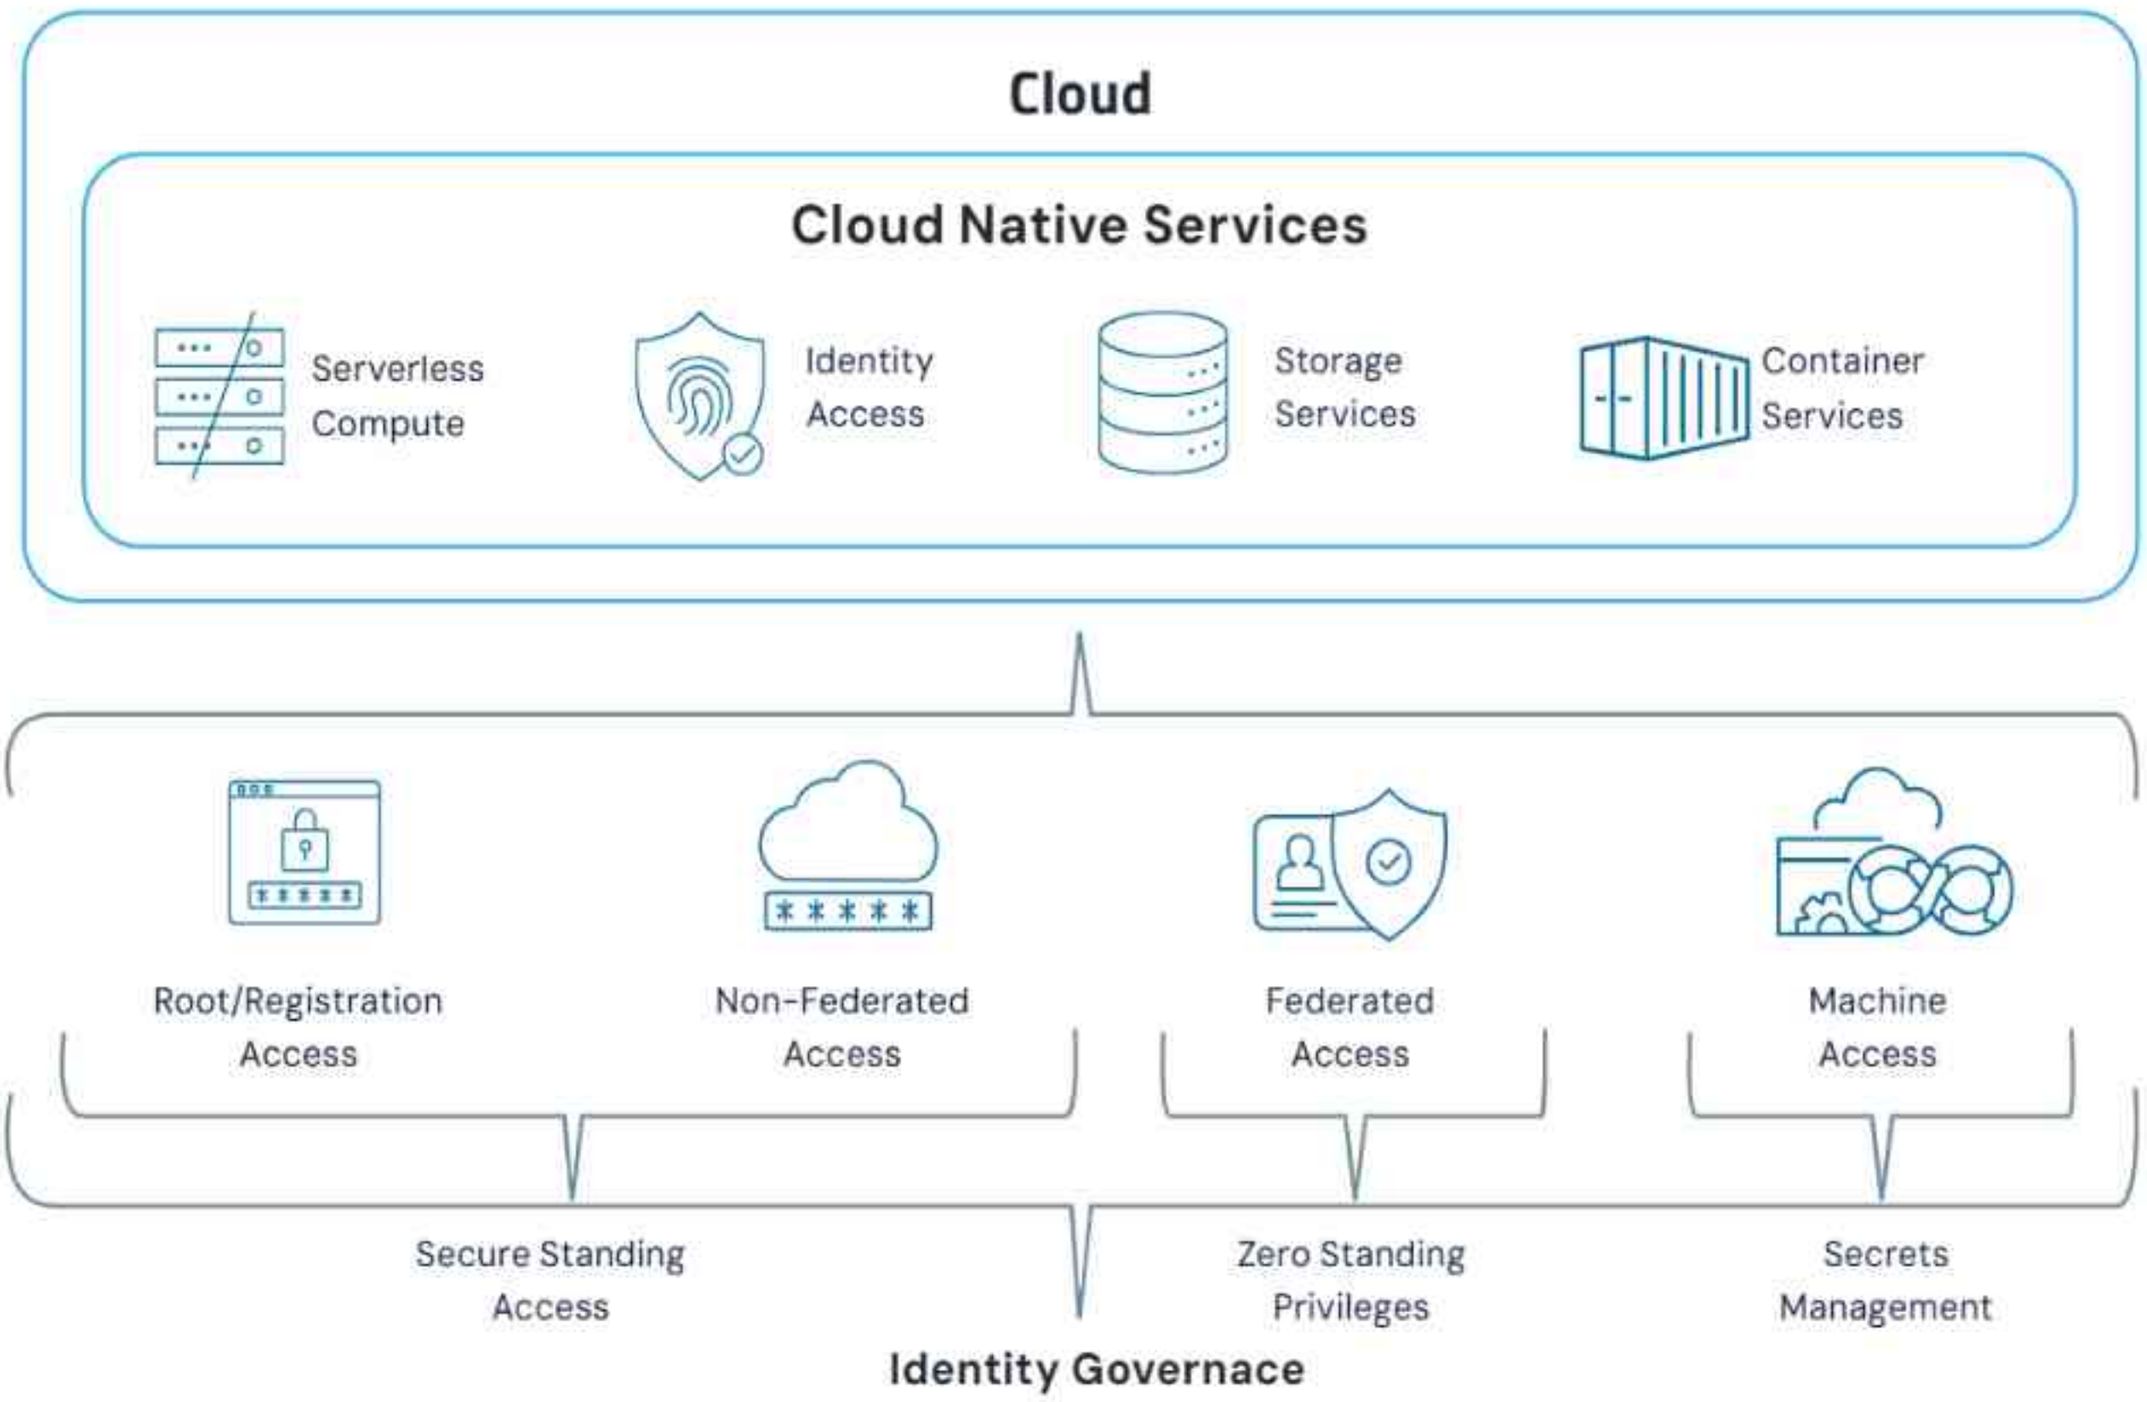

Figure 9: Alignment of Access Methods and Identity Security Controls

longer require them. This is 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. Together these controls form identity governance.

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 – in line 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 just-in-time access and zero standing privileges at the center.

Key Tenants of Strong and Secure Access to the Cloud

The identity security controls recommended above can be summarized 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 Zero Standing Privileges. If necessary, move from freestanding access to just-in-time, and refine permissions over time to achieve ZSP.

Minimize Accounts with Freestanding Access

Reduce the number of CSP 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 Cloud Service Provider account or subscription. Don't forget to protect access to the inbox of the email addresses used for registration too.

Remember Machine Identities

Don't wear blinders just for the human access, as there are often many more machine identities than human ones. Focus your efforts on protecting any machine workloads that have administrative permissions into your cloud service providers, such as your laC 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.

In this section, CyberArk has consolidated and highlighted the seven common principles that Amazon, Google and Microsoft recommend to be compliant with their frameworks. Organizations should secure their cloud environments by aligning their identity security posture and security controls with the following principles:

1. **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 Zero Standing Privileges, no user should have permissions unnecessary for the job at hand.

2. **Authentication and Authorization**

   Implement strong authentication mechanisms, like multi-factor authentication (MFA), and ensure proper authorization controls to manage access to cloud resources effectively.

3. **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.

4. **Credential Management**

   Regularly rotate and manage credentials securely, including for machine identities like service accounts. Avoid hard-coding credentials and utilize identity and access management roles whenever possible.

5. **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.

6. **Automated Compliance Checks**

   Employ automated tools and processes to regularly assess and ensure compliance with security best practices, IAM policies, and configurations.

7. **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 CyberArk Blueprint's recommended controls below. 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 can't boil the ocean all at once. As you can see from the section above, there are many identities and security controls. Organizations need to have a method to the madness to deploy security controls effectively and efficiently.

As good security practitioners, we want to provide organizations with risk-based insight. A major factor in CyberArk's 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 can define risk through a common lens as being a combination of three factors: level of privilege, scope of influence and ease of compromise.

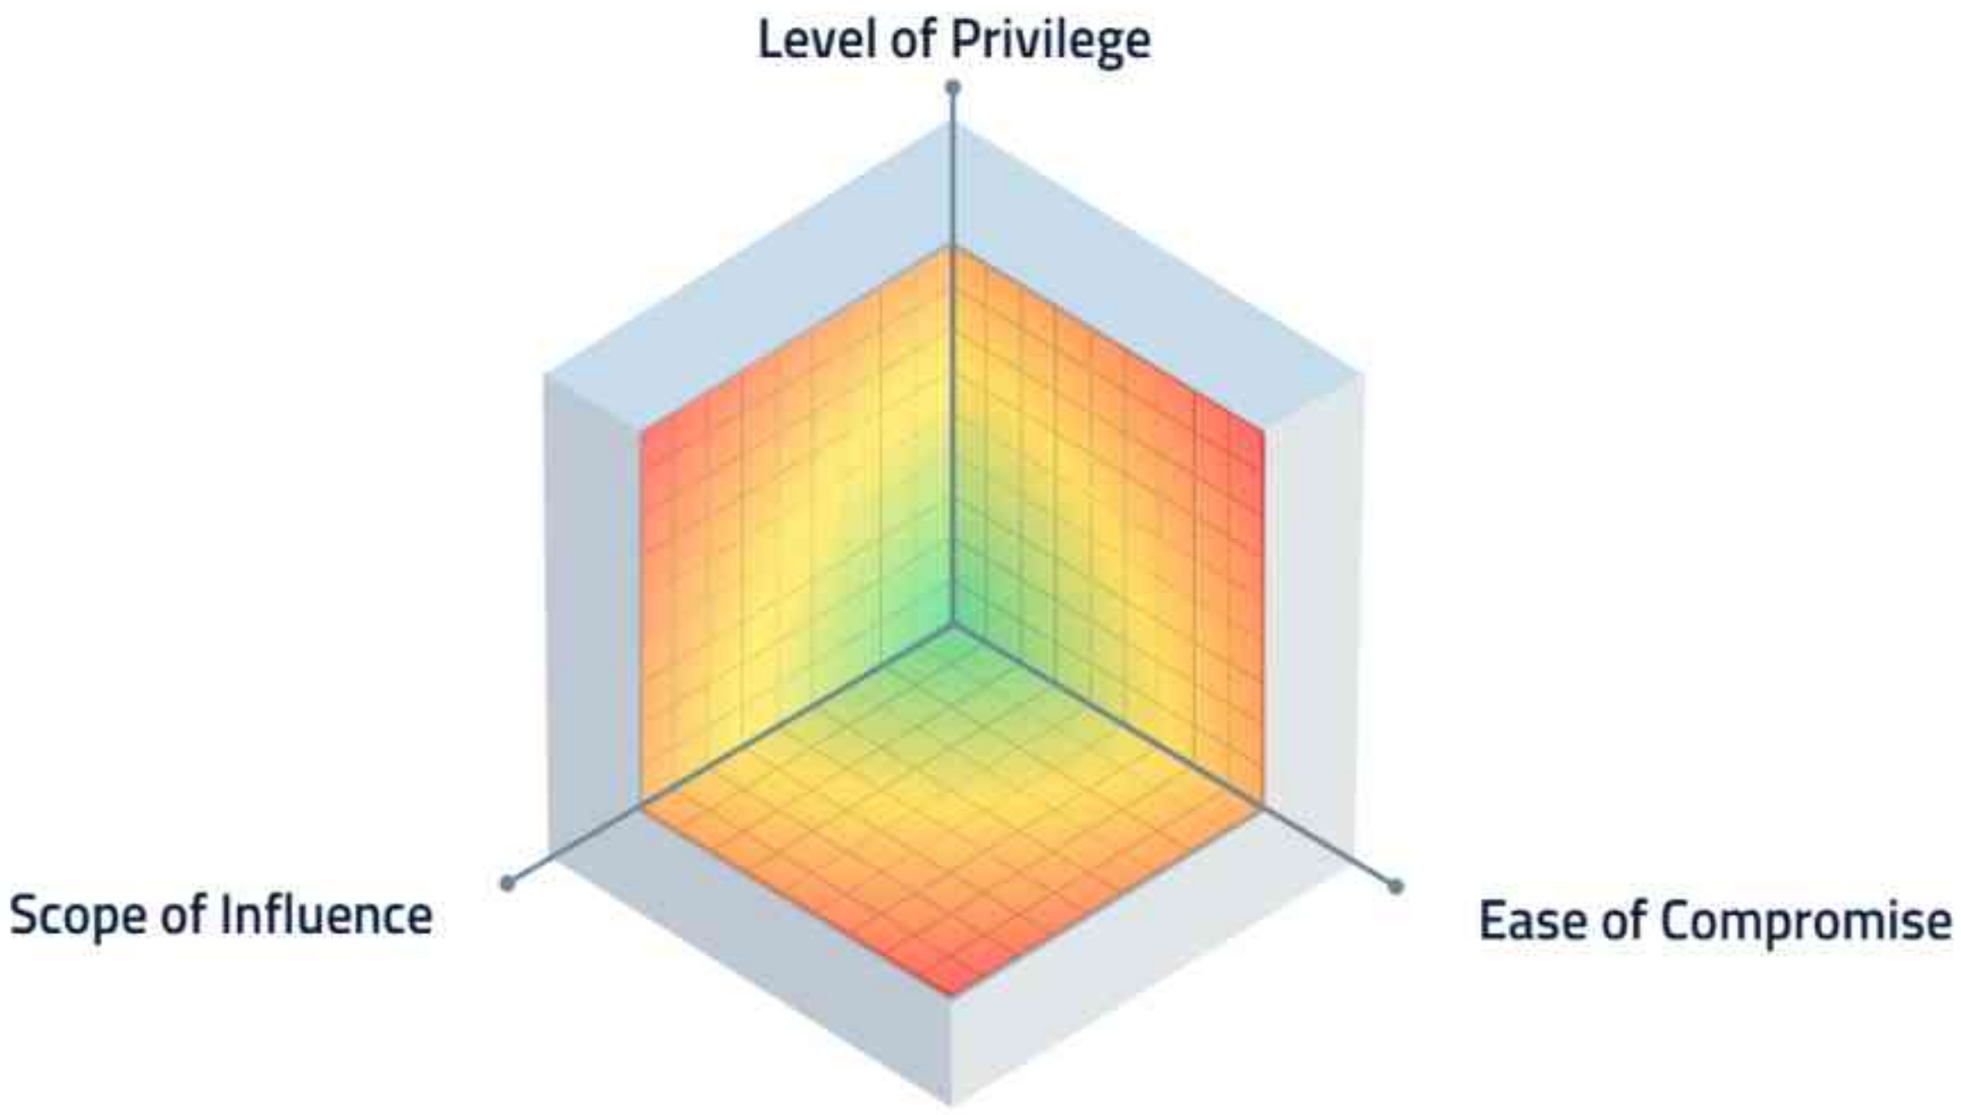

Figure 9: The Three Elements of Identity Risk

Level of privilege refers to the type of privilege that's been 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 the amount or percentage of systems and resources an identity can access, which 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.

The CyberArk Blueprint takes you through two recommended approaches based on real-world experience. Both are risk-based, using the concept above, and maximizing impact for the level of effort. One focuses on prioritizing security controls and services you may need to secure the cloud as a whole. While the other strategy looks at prioritizing the identity or role needing 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. So, in this approach, organizations prioritize by a combination of the type of security control they want to apply (which typically correlates directly to a service or solution) while also considering the risk impact and level of effort 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.

1. The first type of controls to implement are those supporting progress towards **zero standing privileges.** This includes functionality like role-based federated just-in-time access, zero standing privilege, multifactor authentication, session protection, session recording and audit. Privileges should be granted on-demand or as-needed to start, and over time, should be refined with least privilege enforcement. Controls enabling ZSP are considered the highest priority because they cover the largest swath of human access. It's important to implement these controls early in order to avoid the sprawl of freestanding access that can quickly accumulate as cloud footprints grow. These controls should be rolled out to identities with a risk-based mindset: IT Admins, Developers, other Service Administrators (like Networking or DBA roles) and finally those with read-only access.
2. 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, multifactor authentication, session isolation, session monitoring and audit. Here, organizations should circle back to the root and registration account passwords and access keys first, then move into the freestanding access (passwords and keys) for those same identities listed before. Organizations should strive to minimize freestanding access at all costs, but when required or necessary, controls like credential management are critical. Organizations should strive to refine the privileges for non-breakglass emergency accounts over time, focusing on well-known roles first, enforcing least privilege.
3. The third control family is **secrets management.** This includes controls like secrets vaulting, secrets rotation, complex secret value policy, removal of hard-coded secrets from applications, dynamic secrets and just-in-time secret delivery to those apps. Secrets management controls build on the foundation of PAM and require an additional discovery and prioritization effort, which is why they follow PAM controls. Organizations should strive to 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.
4. The fourth control family is **identity governance controls.** Identity governance consists 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. This is your traditional "joiner, mover, leaver" process.

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

ZERO STANDING PRIVILEGE	SECURE STANDING PRIVILEGE	SECRETS MANAGEMENT	IDENTITY GOVERNANCE
• Least privilege role-based access • Zero standing privileges • Federated just-in-time access • Adaptive Multi-Factor Authentication • Session Protection and Audit	• Credential vaulting • Rotation and Isolation • Multi-Factor Authentication • Session Isolation, Monitoring and Audit • Protection for shared and breakglass emergency accounts	• Secrets vaulting • Rotation and complex secret value policy • Removal of hard-coded secrets • Dynamic secrets and just-in-time secret delivery	• Joiner, mover, lever processes • Grant authorized identity access at the time of hire • Revoke access at time of change or departure • Certification and attestation processes for privileged access

Figure 11: Security Control Prioritization Diagram

Identity/Persona-Based Prioritization

In this prioritization method, the assumption is made that the organization is capable of applying 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, it's great information to take into consideration as you build your cloud security strategy.

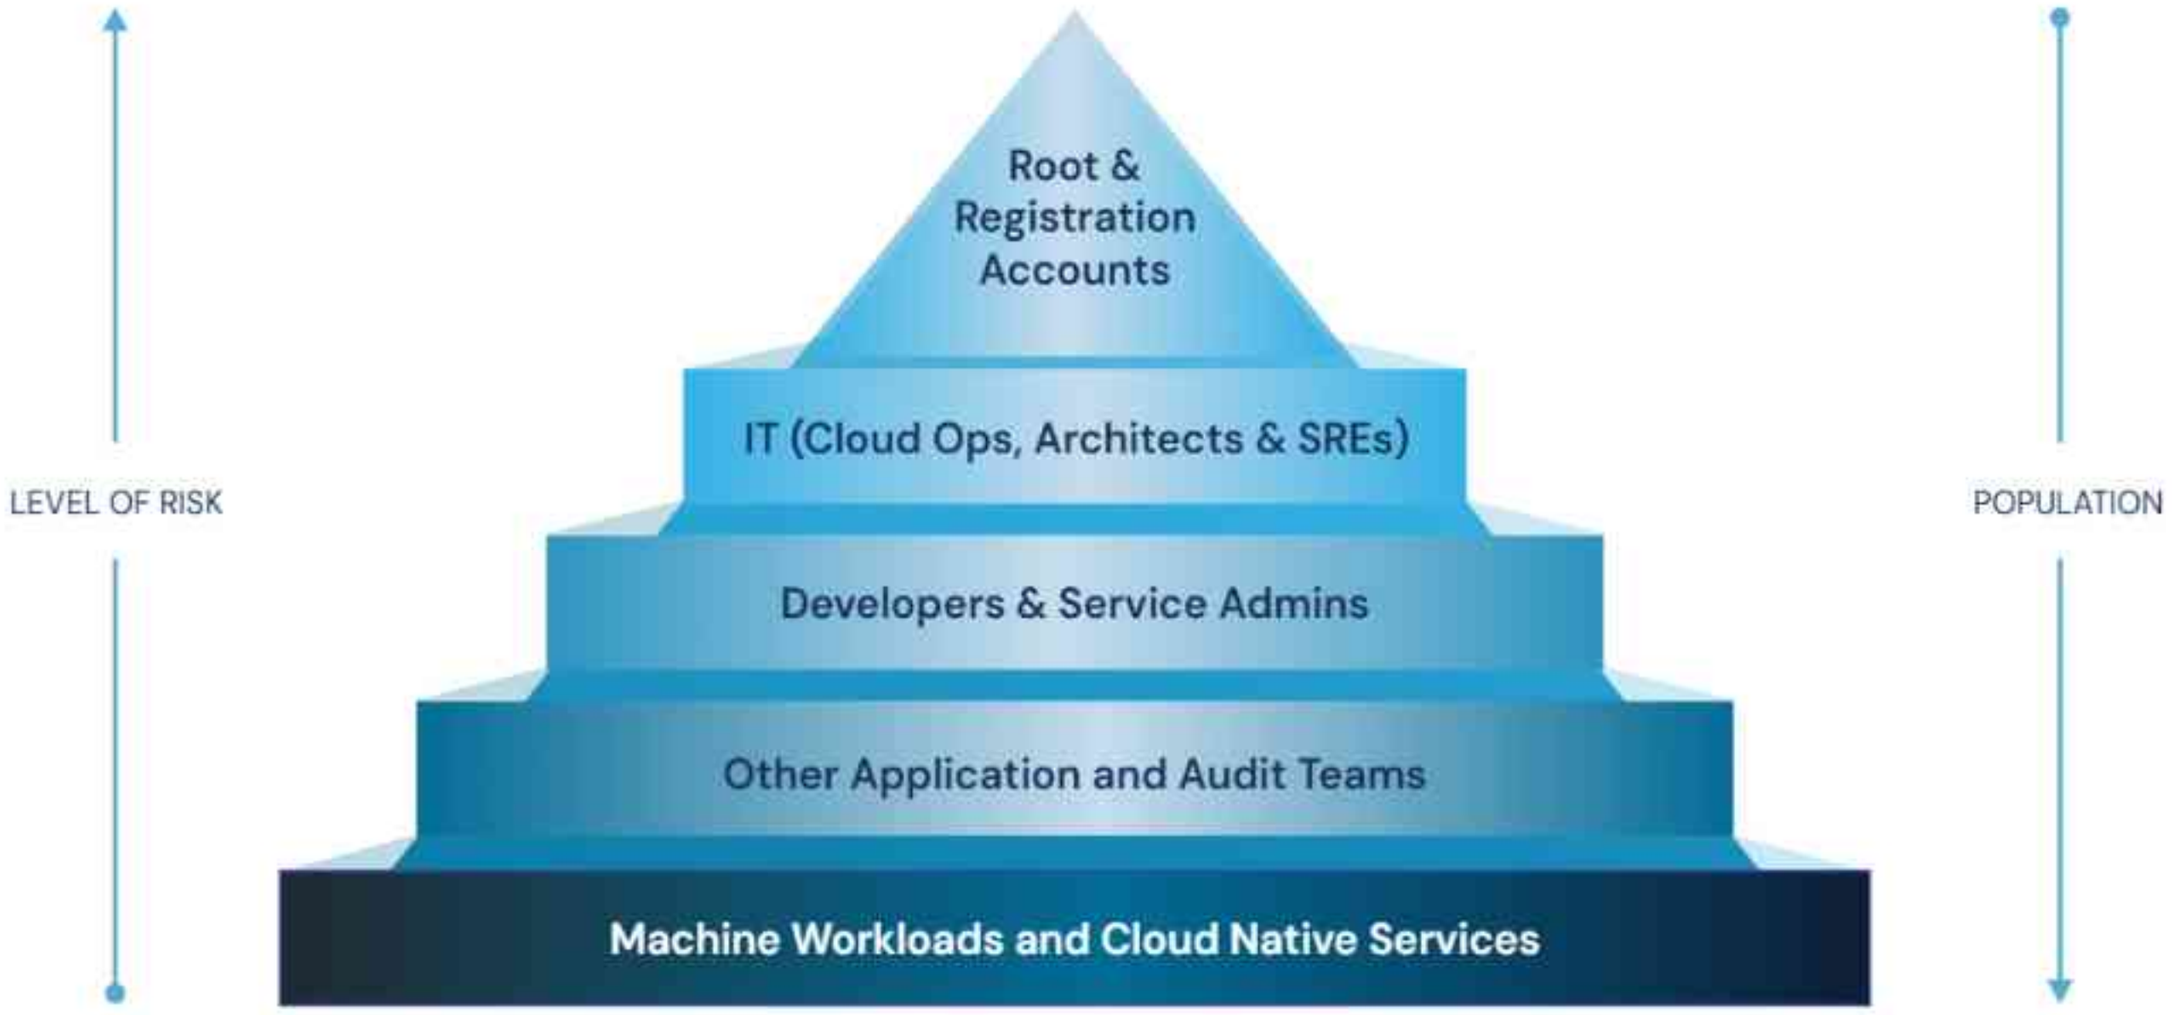

Figure 12: Identity/Persona Risk Prioritization Pyramid Diagram

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 is the theoretical ideal scenario in the event that you have the capability to deliver comprehensive Identity Security controls simultaneously.

1. Always start with the **root and registration** accounts first. This is the account that represents the email address used to register for the cloud service provider account. There will only be one of these for each cloud service provider account you have, so the population is small. These accounts can literally do anything on your cloud service provider 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 and/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. In the case of Azure and Google, you can likely identify this account by looking for a shared corporate email address (e.g., LoBCloudProvider@acme.com).
2. Next up are the **IT administrator** personas. This is typically internal roles such as Cloud Operations, Cloud Architects and Site Reliability Engineers. These roles are typically granted fully-fledged administrative access across the Cloud Service Provider accounts. These identities are considered 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 zero standing privileged access to the CSPs and should be protected in addition with least privilege and identity governance controls.

3. Following that, organizations should focus on securing their **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 are not typically all-powerful across the cloud service provider itself, but they can be all powerful for the services they access or administer. Since cloud service providers have hundreds of services, there can be many identities with this type of access to secure. These users should have a combination of secure standing and operational zero standing privileged access with least privilege, identity governance carefully considered.
4. Following our developers, apply the same controls to **other application and audit teams**, users with lesser-privileges, like read-only access, to various services.
5. Last but certainly not least, are **machine workloads, cloud-native services and other application identities.** There are likely to be many of these machine identities within 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.

However, every organization has unique prioritization needs. Looking at this list, you may see areas that you would want to implement first before the others. If that's the case, go with 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. This prioritization guidance and rationale is something to take into consideration, so you understand the tradeoffs and make the most informed decision for your business.

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

Figure 13: Identity/Persona Prioritization and Control Diagram

Alignment to Cloud Adoption Strategies

Most organizations fall into one of two cloud adoption patterns, which will influence their overall prioritization and security strategy. Recently established businesses tend to favor the more recent advancements in cloud services and build their business applications and processes on native services in the cloud, as such we refer to their adoption path as Digital Native Business and Enterprises (DNB/DNE).

Separately, organizations which have existed longer and have previously leveraged traditional IT services at one point 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. In this section, we'll talk about how each of these cloud adoption patterns influences an organization's identity security strategy.

Digital Native Business and Enterprises (DNB & DNE)

Digital native businesses (DNB) and enterprises (DNE) utilize 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/CDO organizations.

In this scenario, DNBs and DNEs are likely to prioritize securing access with Zero Standing Privileges. Key controls to implement include role-based access control, federated just-in-time access, multifactor authentication, 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. These organizations should seek to extend their security posture to machine identity workloads with secrets management controls, focusing on workloads like cloud native services (e.g. serverless compute functions, container services), DevOps pipelines and source code repositories (Chef, Puppet, GitHub, GitLab) and any custom code that's using embedded secrets. DNB/Es tend to favor federated access whenever possible, so freestanding accounts and keys are typically kept to a minimum. However, these freestanding accounts should not be overlooked. A holistic strategy would incorporate PAM controls to protect credentials granting 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 zero standing privilege, 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 (IaC)

To this point, we've talked about the various approaches for prioritization of securing existing cloud service providers and workloads and the journeys that organizations take to adopt intelligent privilege controls in the cloud. However, what hasn't been covered yet is the concept of long-term operationalization of cloud security. How do you ensure that when new identities, cloud provider accounts and resources are created, they will all inherit the same security controls as deployed for the existing identities? How do you ensure there are no gaps in your security posture for new things? In either of the two cloud adoption strategies, this operationalization problem must be solved.

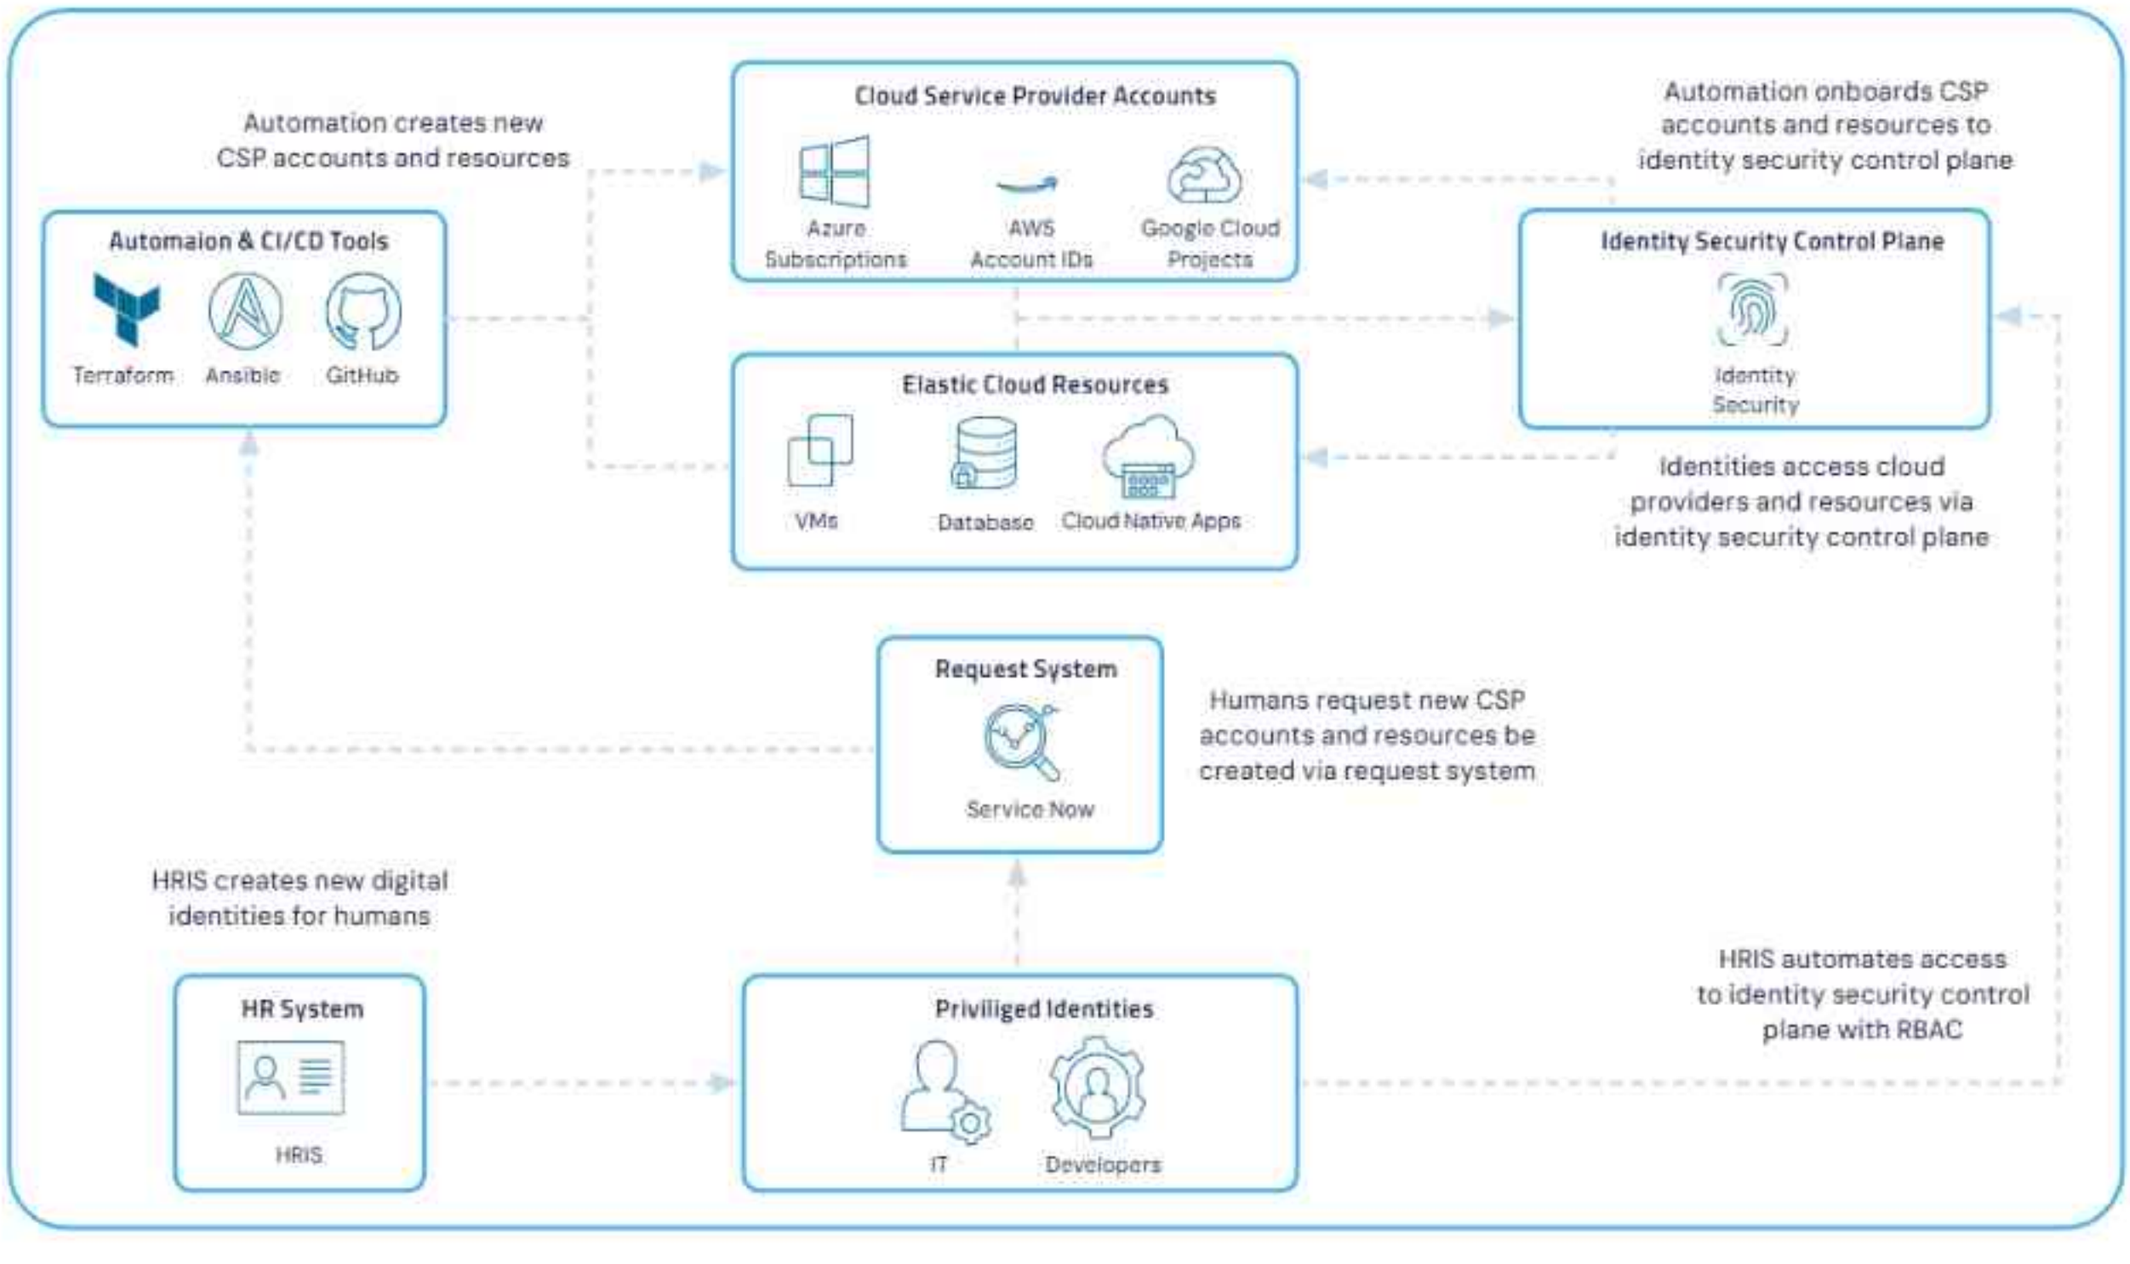

Figure 14: Security at Inception in the Cloud

This is where the term "operationalizing cloud security at inception" comes in. This concept highlights the critical importance of securing access to resources from the moment they are provisioned or created, leveraging 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, ensuring access is secured at inception is paramount. This approach not only streamlines the provisioning process but also significantly reduces the risk of manual errors and security vulnerabilities that can be exploited by malicious actors.

Integrating your existing automation services, infrastructure as code (IaC) tools like Terraform or Ansible, and HRIS systems like Success Factors or Bamboo HR with your identity security control plan enables the automatic application of intelligent privilege controls at the very onset. This security-first approach applies to the protection of cloud service provider accounts, the resources they create and any of the built-in local accounts related to them, effectively protecting 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, thereby enhancing the overall security posture.

Organizations can tap into their existing automation processes that utilize 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 in 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 businesses to maintain agility and innovation while ensuring that their cloud infrastructures are protected against the evolving landscape of cyber threats.

Next Steps for Securing Your Cloud Identities

No organization can secure their cloud overnight. The complexity of the cloud service providers and access methodologies is simply too great. However, with the CyberArk Blueprint's prioritized guidance, your organization can develop an informed, risk-based plan with the best odds of success.

To learn more about the CyberArk Blueprint, check out www.cyberark.com/blueprint, download our Blueprint Toolkit or check out the CyberArk Success Blog.

Connect with our Cloud Security Architecture Team here at CyberArk to discuss your Identity security strategy in the cloud and the solutions required to protect your digital business and move fearlessly forward.

Interested in a free 30- day tiral of CyberArk Secure Cloud Access? Give it a try here.

Request a meeting

About CyberArk

CyberArk is the global leader in identity security. Centered on intelligent privilege controls, CyberArk provides the most comprehensive security offering for any identity — human or machine — across business applications, distributed workforces, hybrid cloud workloads and throughout the DevOps lifecycle. The world's leading organizations trust CyberArk to help secure their most critical assets.

CYBERARK®
The Identity Security Company™