Privileged access to the cloud—why PAM fails you

Atul Tulshibagwale, CTO, SGNL

Introduction

That privileged access to the cloud is one of the most exploited methods to breach enterprises is clear evidence that conventional privileged access management has failed to secure cloud access. However, the shared responsibility model of cloud security places the burden of managing access entirely on cloud customers. Organizations today depend on a mix of network-level access control technologies such as Zero Trust Network Access (ZTNA), which is a generalization of Virtual Private Network (VPN), Cloud Access Service Broker (CASB), and Privileged Access Management (PAM) products to meet the unique security challenges of privileged access to the cloud. These strategies, which haven't changed in decades, have a number of issues:

- They are unable to take advantage of the rich features of cloud services;
- Their reliance on manual access granting processes creates an unacceptably large attack surface; and
- Their dependence on a shared credential model of access creates difficulties in auditing while increasing risk.

PAM solutions are notoriously difficult to implement and maintain: A Gartner report says: "Implementing privileged access management (PAM) tools is a difficult process under the best circumstances."

In this white paper, I will demonstrate how a modern privileged identity management model can be implemented to secure access to the cloud effectively. This model does not define privileged access narrowly to just administrator access but also includes business user access - as many cloud breaches have occurred due to business user compromise. This model is far easier to implement and maintain, more reliable, and is responsive to data and policy changes. Customers have seen several orders of magnitude reduction in the number of roles, entitlements, and policies required to specify access.

Gartner

"Implementing privileged access management (PAM) tools is a difficult process under the best circumstances"[1]

June 2024

[1] "5 Interlocking Strategies for a Successful PAM Implementation" - Michael Kelley, Felix Gaehtgens, James Hoover, Abhyuday Data, June 28, 2024

How the cloud is different

While PAM technology has worked (with varying degrees of success) in the context of on-premise systems, the nature of cloud services being different is one of the main reasons why the way privileged access works needs to change. To determine its efficacy, here's how the cloud is different:

- **Variety:** As described in the background section, there are a wide variety of cloud services, including SaaS, PaaS, and IaaS. An organization typically needs many instances of each in order to modernize its business.
- **Scale:** The number of cloud services that require privileged access is far greater than on-premise systems that require privileged access; each cloud service may be specialized in a different way, and innovation in the cloud has led to a proliferation of services available and in use by organizations.
- **Architecture:** Cloud systems are, by definition, accessible from outside the boundaries of your traditional network perimeter. As a result, they are naturally more exposed to attackers. Attempting to create an artificial network perimeter around cloud systems has typically been challenging and hasn't removed all of the risk from the equation.
- **Intricacy:** Digging into a cloud service will reveal many different types of privileged access that any one cloud service may require. For example:
  - Should an AWS administrator have access to all your resources in the IaaS system? (Answer: almost certainly not). Attacks like that on MGM Resorts® occurred because an IaaS administrator had access to everything.
  - Under what circumstances should a customer service agent have access to a specific customer's data in a SaaS system like Salesforce? (While this may seem benign, some of the most devastating attacks (e.g., Okta, Snowflake/AT&T) have occurred because of sensitive customer data having leaked from such services)
  - Who should have access to data in a PaaS system like Snowflake and under what circumstances? (Another cause of some of the worst data leaks)

  Applying traditional access control strategies in such situations makes it really difficult to manage access, as you will see on the next page.

While PAM technology has worked (with varying degrees of success) in the context of on-premise systems, the nature of cloud services being different is one of the main reasons why the way privileged access works needs to change.

Traditional cloud access technologies

There are three aspects to traditional privileged cloud access strategies:

- **Privileged Access Management (PAM)**: Administrator access is secured by vaulting shared passwords, and a proxy is required to access target systems. They provide workflows for requesting and granting access, time-bound access, credential rotation, and session recording as unique security features.
- **Security Service Edge (SSE)**: Typically consists of ZTNA/VPN and CASBs, configured using static roles that are leveraged to assure endpoint posture and high-level access control (e.g., user X has access to the entire app Y).
- **Single Sign-On (SSO)**: Enables organizations to ensure only a single credential (preferably using phishing-resistant MFA and/or passkeys) is used by users to access their systems. Sometimes, the PAM product will leverage SSO to let administrators log in to the PAM product, but the credential for the cloud is still vaulted and shared.

These aspects are represented in the diagram below:

Figure 01. Conventional privileged cloud access model

Access decisions in traditional technologies

Regardless of the technologies used (SSO, PAM or SSE), the decision-making process for controlling access to the cloud is similar.

Admin time decisions

Managers are periodically asked to determine whether a specific user belongs in a specific category of access. This category might be expressed as a role, attribute, or entitlement. Regardless of the terminology and mechanism, the result of this admin-time decision is that the user is ultimately added or removed from a static list that is checked at the time of access, typically by the target application or system.

Login / “run” time decisions

Single sign-on (SSO) systems will often make additional checks when establishing a login session. If a user requests to access a cloud application such as Salesforce, the single sign-on system (typically, an identity provider) can be configured to verify other factors, such as the user’s IP address is within a pre-configured range and the user’s device posture is acceptable. PAM systems also offer log in time decisions when their users (administrators of the target systems) either login directly to the PAM system or use single sign-on to log in to the PAM system.

Managers are periodically asked to determine whether a specific user belongs in a specific category of access. This category might be expressed as a role, attribute, or entitlement. Regardless of the terminology and mechanism, the result of this admin-time decision is that the user is ultimately added or removed from a static list that is checked at the time of access, typically by the target application or system.

Limitations of traditional strategies

The traditional technology offerings and the decision-making processes they necessitate together constitute the traditional strategies for controlling privileged access. When applied to cloud services, these strategies have the following limitations:

- **Coarse-grained:** Cloud services often have a rich set of resources and data. The inability to control which of those resources a user may get access to greatly increases your organization's risk, because an attacker who assumes a privileged user's identity now has access to all of those resources and data. Traditional privileged access protections don't provide fine-grained access control beyond the "enable access to the service," and thereby exposing organizations to unnecessary risk.
- **Manual processes:** Even though they may be simplified using graphical workflow UIs, manual processes to grant and revoke access are error-prone and susceptible to social engineering attacks. Traditional strategies almost always involve manual workflows to make and adjudicate access requests. Manual access decisions made in the past also need to be periodically reviewed. The review process is expensive and susceptible to rubber-stamping because the reviewer does not have the context of whether a specific user belongs in a specific category or not. Moreover, given the complexity of the cloud, manual processes can increase the attack surface over time by creating more pathways for a user to gain privileged access to the target systems.
- **No direct single sign-on to the cloud:** Cloud services are able to leverage the investments your organization already makes in strong authentication (e.g., passkeys or phishing-resistant MFA solutions), but many PAM products have their own way of authenticating users, thereby making your organization unable to leverage those investments for cloud access.
- **Non-native IAM:** Cloud services almost always provide sophisticated native IAM capabilities, which can dynamically map users in your organization to specific roles that enable access to specific collections of resources or classes of resources that share specific attributes. However, because of their shared credentials model, PAM systems cannot leverage these capabilities.
- **No single logout:** The SSO system issues a token (e.g., a SAML token) in order for the user to access the target system or application. Typically, this system will issue its own token (sometimes in the form of a session cookie) that enables the user to access it. Even if the SSO system later detects that the user should no longer have access to the target system, there is no way to terminate the logged-in user session in the cloud system.
- **Data proliferation:** Another often overlooked factor is that sometimes sensitive data is required to determine access. For example, a user's citizenship or certifications may determine their access to specific data within cloud systems. To enforce such access rules, specific categorizations are created in access control systems, and users are added to these categories based on the sensitive data. This categorization is then available to all applications, exposing organizations to liability from such data leaking out.

Continuous identity-first security

A different (and newer) way of thinking about cloud access is to continuously evaluate access at every cloud system. Cloud services are often secured using the zero-trust architecture. A core principle of the zero-trust architecture is to be able to evaluate access independently for each request. However, since each cloud service is dependent on a login provided by an independent PAM, SSO, or SSE system at the time the user logs in, the cloud service is blissfully unaware of any changes to the user's privileges or any issues with the user's device or behavior that may have taken place after the login. Organizations often rely on smaller token lifetimes (e.g., one hour), forcing the user to go back to their authentication system to log in again to the cloud system. This results in a particularly poor user experience and can disrupt critical administrative operations, which can cause a multitude of operational problems.

CAEP

An open standard, the Continuous Access Evaluation Protocol (CAEP), was **introduced by Google** in February 2019. It proposed an asynchronous publish and subscribe framework for conveying events that modify session properties. It has since merged with the OpenID Foundation to create the "Shared Signals Framework" (SSF), and CAEP is now a profile of SSF (so the acronym now stands for the Continuous Access Evaluation Profile). CAEP provides events like "session revoked", "credential changed", "assurance level changed", etc., to continuously convey any changes to session properties. CAEP, however, does not dictate what a receiver must or should do upon receiving such events. This is because the level of trust between the transmitter and the receiver may vary, and each receiver should have the flexibility to determine what it should do in response to receiving specific events. CAEP now sees increasing adoption from large and small companies like Okta, Apple, SGNL, and others.

Principles of continuous identity-first security

By following the principles behind CAEP, one can achieve the same effect through proprietary integrations. The principles of continuous identity security are:

- Determine which data sources to trust for which information.
- Continuously obtain data required for making access decisions from these data sources. If an asynchronous notification mechanism (e.g., CAEP) is not available, continuously poll for updated data from the data sources.
- Update internal state to incorporate new data
- Determine which policies need to be re-evaluated with the new data.
- Internally enforce access decisions and updated policies.
- Convey any decisions or updated data to other parties that rely on your updated data.

While CAEP can deliver continuous identity security and is gaining traction with major industry players, it is not yet widely implemented. In the meantime, continuous identity-first security systems can use proprietary integration methods while support for CAEP develops among products.

Zero standing access

Continuous identity-first security enforces zero standing access, meaning no default privileged access to cloud services. If credentials are compromised, the impact is contained, making it a Security Operations Center (SOC) issue rather than a C-Suite concern. However, users should have seamless access to necessary cloud resources when justified. Justification depends on contextual factors—who, what, when, where, and why—requiring data from various systems to make informed access decisions. As an example:

- In your HR system
- In your on-call tracking system
- Your case or ticket management system

For organizations that are adopting a continuous identity security model, all this data is continuously updated and available right where it is needed for enabling such privileged access. And, as the data changes over time (say the case is closed, or the user goes off-duty), their access is automatically removed.

Modern Privileged Identity Management

Zero standing access is implemented by Modern Privileged Identity Management (MPIM) systems. They fill the gaps that traditional PAM leaves behind. Continuous identity security is achieved by Modern Privileged Identity Management (MPIM) as follows:

Continuous data ingestion

Source systems (user directories, HR, ITSM, CRM, etc.) may lack the availability and speed needed for just-in-time decisions, making direct reliance impractical for enforcing access. An MPIM system continuously updates and organizes this data, ensuring fast responses to access queries without disrupting business operations.

Single sign-on integration

MPIM uses single sign-on for direct user login with individual credentials secured by MFA or passkeys. Similarly, API access goes to target systems, with the API gateway verifying each request with MPIM.

High-level policies

To enforce fine-grained policy with dynamic data, policies should be general, not user- or system-specific. Instead of "user X has entitlement Y for system Z," they should be expressed in a more general way, e.g., "SRE users that are not on leave and on duty may access cloud resources related to an open case that is escalated and approved for emergency access."

Removing access

MPIM removes cloud access when granted conditions change. While CAEP will standardize this soon, MPIM integrates with proprietary APIs for non-CAEP systems, dynamically revoking access when a case closes or a device becomes non-compliant per the device management service.

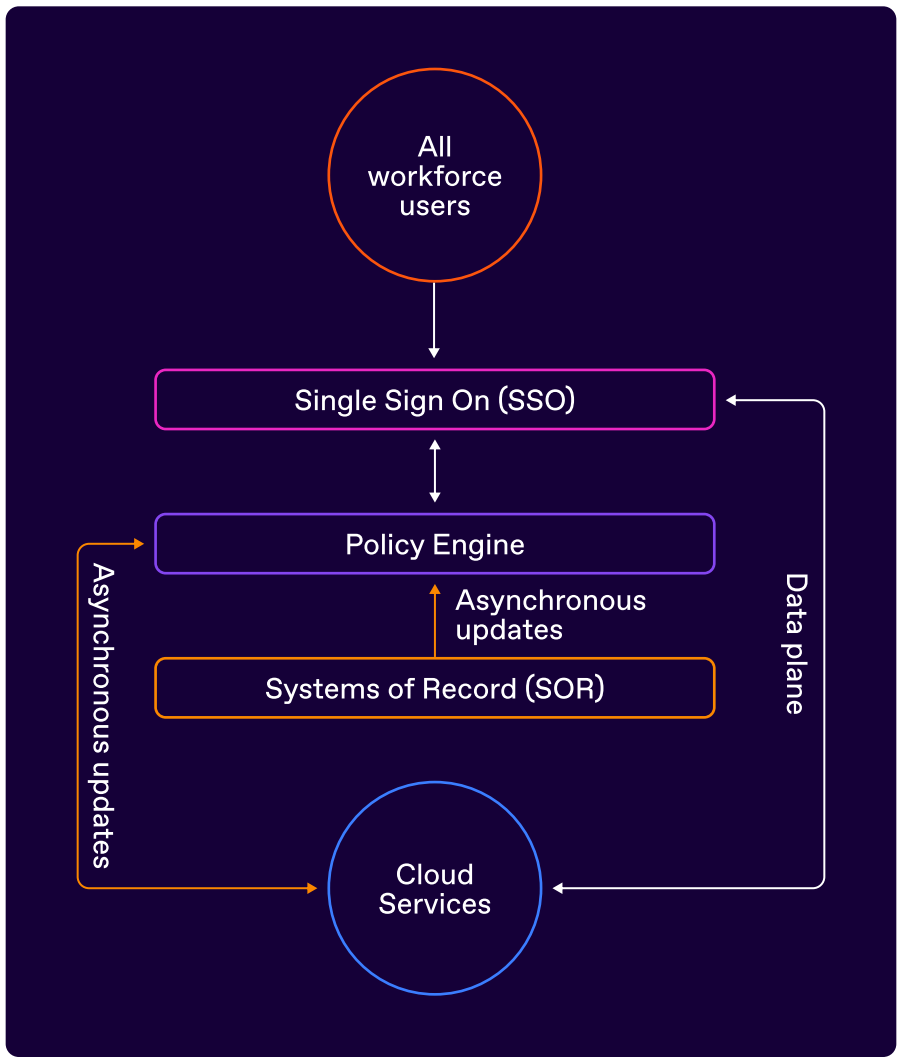

Figure 02. Modern privileged identity management architecture

Building out your Identity Fabric

One of the reasons security becomes weak is that individual access control systems are siloed and do not act in concert with each other. MPIM breaks these silos and strengthens security by effectively building an identity fabric. But this does not mean you need to rip out everything you have.

Augmenting existing identity solutions

Organizations already have a wide range of identity technologies, such as SSO / identity providers, Identity Governance and Administration (IGA), PAM, and SSE systems. These systems require constant maintenance to be effective, such as access reviews or updating entitlements and workflows to meet changing business policies. A good way to introduce MPIM in this mix is to identify where you need zero standing access the most.

If you are manually doing role assignments in the target cloud system (either directly or through an IGA), then you can deploy MPIM to eliminate the manual role assignments there. This will also provide the added benefit of eliminating periodic access reviews for those role assignments. Doing this will dramatically reduce your exposure to credential-based attacks on those target systems and reduce your maintenance costs by eliminating access reviews or manual access granting processes.

In Summary

As you can glean from all of the above, MPIM can offer zero standing privileges by overcoming all of the limitations of traditional strategies of privileged access. It drastically reduces an organization's exposure to credential-related attacks and even API-based attacks:

- **Fine-grained:** By leveraging data from the source systems, MPIM is able to determine which specific cloud resources within the target system the user needs access to, and only grant access to those resources.
- **Automated:** Since MPIM obtains data from your systems of record, which you already use to run your business, no identity-specific manual processes are required to grant or revoke access.
- **Native single sign-on:** In MPIM, there are no shared credentials, so users use their own identities to log in directly to target systems using the single sign-on provider your organization already uses
- **Native Cloud IAM:** The target applications' permissions are directly updated by MPIM, either using standards such as CAEP or using proprietary integrations
- **Single logout:** When the conditions that afforded users access change, MPIM is able to reassess user access and appropriately logout the user from target systems, either using CAEP or proprietary integrations.
- **Data confidentiality:** Since policies are dynamically enforced, sensitive data used to determine policy is not propagated to any target system.

Security is best implemented not at admin time, not at login time, not even at event time. It is best implemented continuously.

ABOUT SGNL

In today's era of persistent identity attacks, high-risk standing access is a serious threat to critical enterprise systems. Traditional IGA, RBAC, and PAM approaches fall short because they simply weren't designed for today's identity-centric security perimeter. SGNL's dynamic approach to access management achieves Zero Standing Privilege across your cloud applications like Azure, AWS, GitHub, and Salesforce, as well as on-prem systems.

It's why global enterprises and fast-growing mid-market companies alike are turning to SGNL to reduce their identity attack surfaces, and why SGNL is backed by top security technology investors including Cisco and Microsoft.

SGNL IS BACKED BY LEADING TECH INVESTORS

CISCO investments

M12 MICROSOFT'S VENTURE FUND

REQUEST A DEMO