Defining Shadow Access: The Emerging IAM Security Challenge

The permanent and official location for Identity and Access Management Working Group is https://cloudsecurityalliance.org/research/working-groups/identity-and-access-management/

Acknowledgments

Lead Authors

Sasi Murthy
Venkat Raghavan
Steven Schoenfeld

Contributors

Philip Griffiths
Shruti Kulkarni
Michael Roza
Dhaval Shah
Heinrich Smit

Reviewers

Senthilkumar Chandrasekaran
Ivan Djordjevic
Rajat Dubey
Ahmed Harris
Shraddha Patil
Alberto Radice
Osama Salah

CSA Analysts

Ryan Gifford

Editor

Larry Hughes

CSA Global Staff

Claire Lehnert

Table of Contents

Shadow Access

Shadow Access is the unintended and/or undesired access to resources, such as applications, networks and data. This new problem has arisen with the growth of cloud computing, DevOps velocity, cloud native architectures and data sharing.

Shadow Access is increasingly a cloud issue, resulting from increased use of access and entitlements that connect cloud services together, coupled with automated infrastructure and software development, resulting in incorrectly or unexpectedly permissioned accounts and resources. Organizations from small to large often find out the hard way that what was once a secure starting point has silently evolved into an unsecure one. The above problem, combined with the common practice of cloning accounts and permissions (typically during onboarding or account creation) have multiplied the shadow access problem by providing access that is not truly not required.

The consequences of Shadow Access are potentially catastrophic, and threaten to impact any organization that has an evolving cloud. This short document intends to summarize the background, causes, impact and path forward to regain the benefits of a dynamic and secure cloud environment.

Background

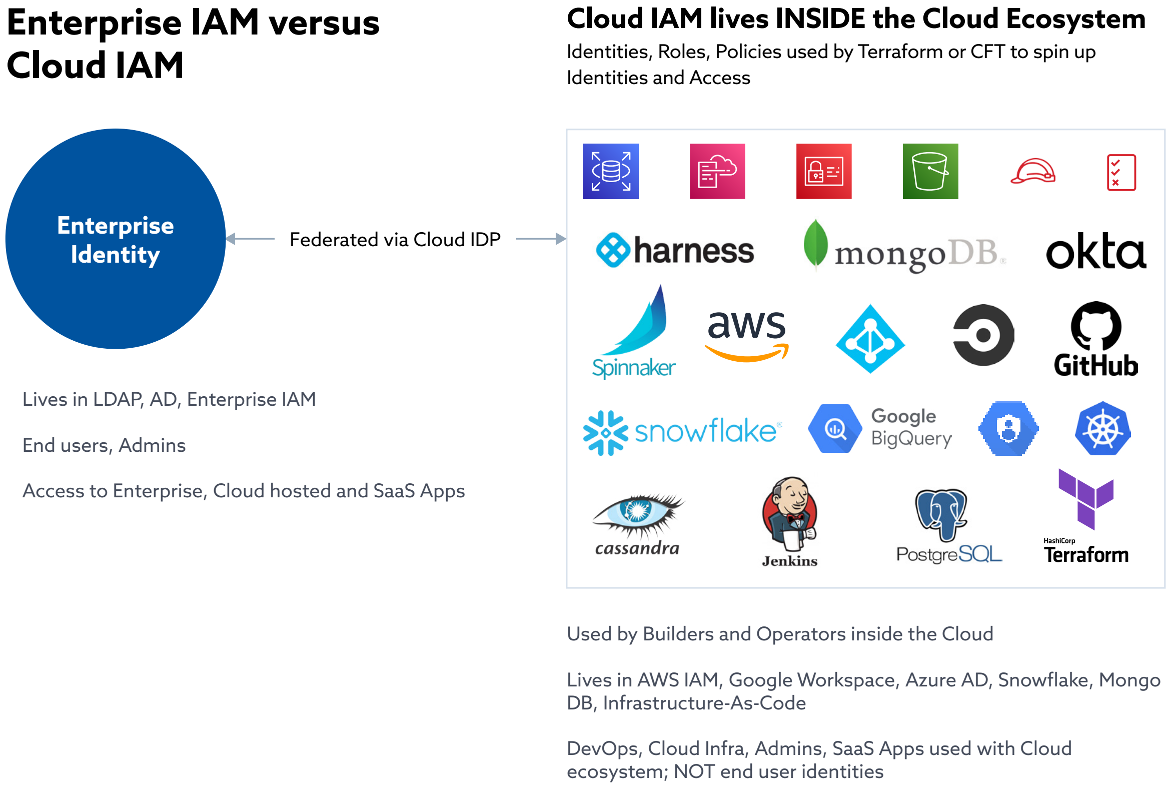

Enterprise IAM versus Cloud IAM

Cloud IAM lives INSIDE the Cloud Ecosystem

Identities, Roles, Policies used by Terraform or CFT to spin up Identities and Access

Lives in LDAP, AD, Enterprise IAM

End users, Admins

Access to Enterprise, Cloud hosted and SaaS Apps

Used by Builders and Operators inside the Cloud

Lives in AWS IAM, Google Workspace, Azure AD, Snowflake, Mongo DB, Infrastructure-As-Code

DevOps, Cloud Infra, Admins, SaaS Apps used with Cloud ecosystem; NOT end user identities

Figure 1: Enterprise IAM versus Cloud IAM

Traditional Enterprise IAM (Identity and Access Management) systems have been deployed for decades, and are often built on popular services or protocols like LDAP and Active Directory. Enterprise IAM systems provision entitlements and credentials to identities, and typically rely on an enterprise HR system as an authoritative source of truth. Established policies and processes surround the invocation of an action and the resulting access by end-users to resources and applications are typically hosted "inside" the corporate enterprise firewall accessed by employees and contractors "outside" of the firewall through a secure VPN connection.

As cloud applications grew in prominence, Cloud IDP (Identity Provider) systems began to emerge. Cloud applications by definition are not hosted inside the enterprise. Therefore many enterprises today operate Enterprise IAM (for on-premises hosted applications) in conjunction or in tandem with a popular Cloud IDP like Okta, Azure AD, or Ping Identity.

The advent of cloud computing has introduced a new concept called Cloud IAM. Cloud IAM is used to provision and control access and entitlements to resources, applications and data, residing inside public cloud ecosystems like AWS, Google Cloud, and Azure Cloud, as well as private clouds powered by Kubernetes.

While similar at first glance, in reality, the concept is substantially different, and for that reason, we need to separately classify and examine these Cloud Identities.

So why a new concept called Cloud Identities and how are they different?

> - Everything you spin up in the cloud has an identity with access to a critical cloud service, supply chain element, or data. The Cloud Service Provider (for example, AWS, GCP or Azure) systems control the provisioning of all identities and their access via a lynchpin service like Cloud IAM.
> - Inside the cloud, every access requested is authenticated and authorized before access is granted.
> - Cloud identities can be human or non-human identities. Human identities are mostly end-users, developers, DevOps, and cloud administrators. Non-human identities are the majority of the rest, composed of identities attached to cloud services, APIs, microservices, software supply chains, cloud data platforms, etc.
> - One power of the cloud is its "programmability." This power is unleashed via the developers by programmatically combining cloud services, APIs, and data to create applications. This difference is not to be taken lightly. Modern cloud applications are really an assembly of many distributed services driven from APIs, across providers and their ecosystems. As developers combine cloud services, they create automated identities having access pathways to data.
> - Another power of the cloud is automation. Cloud teams use the power of automation using infrastructure-as-code to easily spin up and define their cloud resources, cloud identities, and their access. Automation first, governance second.

Cloud computing has created an identity-centric world and the differences surrounding them lead to the root causes of Shadow Access.

Causes

The root causes of Shadow Access stem not just from having Cloud Identities, but also from the fundamental complexity and processes driven by the cloud.

Complexity

The "power of the cloud" referred to above is primarily released through developers and automation, and is significantly more complex than previous environments. Some notable differences include:

> - Data is no longer stored in a single data store. There is a proliferation of cloud data stores and data-sharing applications across cloud and SaaS environments.
> - Data stores are constantly evolving, expanding or contracting, as new types appear with new or updated use by applications.
> - An application is not monolithic, but a dizzying combination of interconnected identity systems, cloud services and data.
> - Vast increases in the use of SaaS applications that connect with cloud ecosystems.
> - Each cloud service has associated permissions and entitlements that provide authorization to sensitive data and operations.
> - The scale of permissions and entitlements are wildly more expansive and orders of magnitude more complex compared to conventional on-prem environments.
> - Organizations use multi-cloud and a combination of public/private cloud environments.

To illustrate the complexity, in AWS alone there are 12,800 cloud services with 13,800 permissions attached, creating an enormous set of permutations and combinations of cloud access.

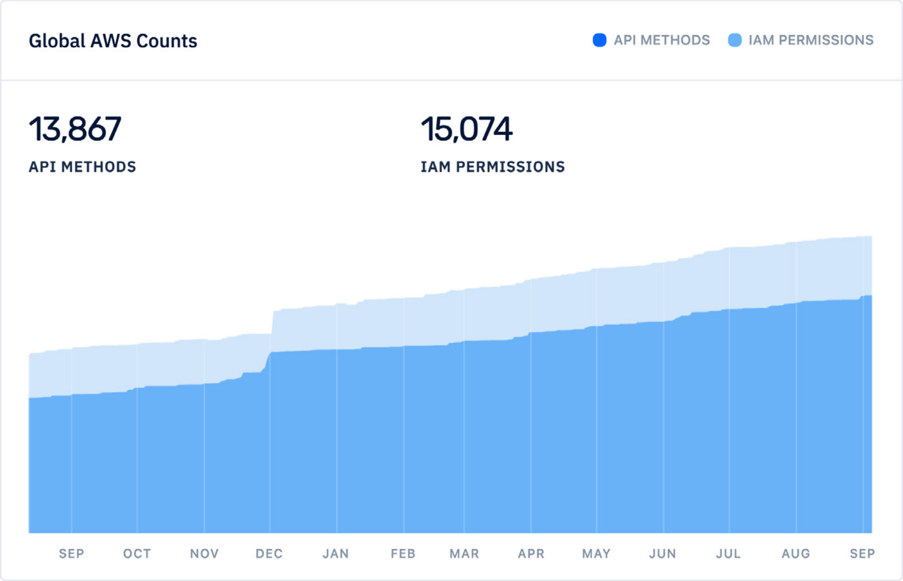

Figure 2: from https://aws.permissions.cloud/

Process Changes

In previous environments strict policies and processes were usually in place before an identity would be created and given access. For organizations that had established controls, this governance process typically regarded not just the creation, but also a consistent review and approval process. Here again, the cloud is very different:

- New identities and access are created centrally, often by the developers using infrastructure-as-code.
- The profile of a new identity is usually copied from a template that hopefully has a central review procedure for organizational standards.
- New identities and accesses are automatically created with little, if any, governance.
- Applications that the identities access are constantly changing, without full access reviews.
- Application components are often re-used, copied, or used for multiple applications in order to facilitate speed.
- Increased use of SaaS and third-party applications have no formal security reviews.
- The data stores the applications access are constantly changing.

Absent from the creation is the continuous monitoring, review, and rightsizing that needs to be as automated as the original identity and access creation. Since cloud applications are distributed and ever-evolving, a change in one element could have unintended consequences for the overall exposure.

It is this highly complex and evolving nature of the applications, and disruption of the processes surrounding the creation and ongoing review of cloud identities, that leads to Shadow Access and a series of potentially massive exposures for an organization.

Impact

As stated earlier, Shadow Access is the unintended and/or undesired access to resources, such as applications, networks and data.

To illustrate the impact, the Verizon Data Breach Investigations Report (DBIR) report highlights that 80% of breaches were related to identity and access. Zettabytes of data are being stored in cloud platforms which is driving a massive demand for access.

The impacts of Shadow Access include:

- Existing tools are blind to the multitude of cloud identities and access pathways.
- Governance and visibility gaps make it very difficult to implement IAM guard rails.
- Unrecognized access pathways allow vulnerabilities to be exploited to breach cloud data.
- Threat actors can weaponize programmable access to cause harm far beyond breach of data.
- Third-party and SaaS Applications that connect to cloud ecosystems introduce lateral movement risks.
- The existence of Shadow Acces creates data security, audit and compliance exposures and creates policy and governance gaps.

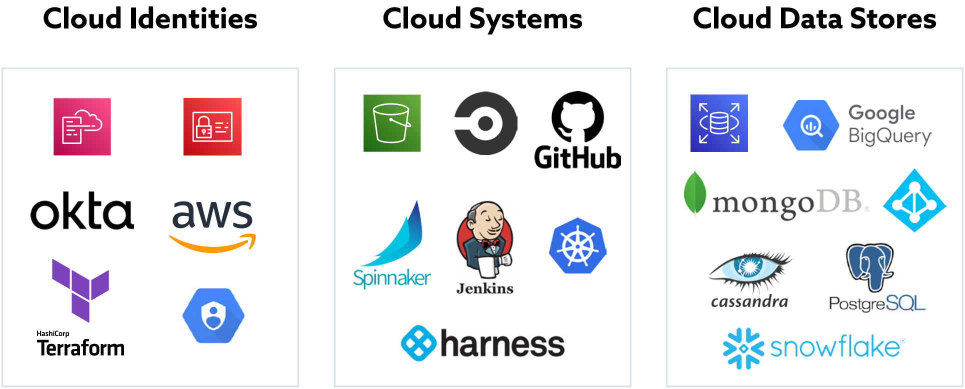

Figure 3: Shadow Access exists across multiple systems in public cloud ecosystems (e.g., AWS)

In essence, the true security state of an environment is never known, and the mechanisms and processes to derive that information are typically outdated before the analysis is complete. The result is an environment that is vulnerable, and owners of the environment have no way to truly assess the risk.

> "The existence of hundreds (or sometimes thousands) of identities - both **human and programmatic** - across the CI/CD ecosystem, paired with a lack of strong identity and access management practices and common usage of overly permissive accounts, leads to a state where compromising nearly **any** user account on any system, could grant powerful capabilities to the environment, and could serve as a segue into the production environment."
>
> Verbatim from OWASP Top 10 CI CD SEC-2
> https://owasp.org/www-project-top-10-ci-cd-security-risks/

Conclusion

Shadow Access, as a new phenomenon, impacts many areas in cloud computing. A new generation of tools and processes needs to be established and instrumented to address the issue, re-establish the intended state of access and data security, and allow the full benefits of the cloud to be achieved.

The work on understanding Shadow Access is just beginning. The broader trends of automation, AI, and data creates a fertile environment for Shadow Access to thrive. It impacts just not access, but also broader aspects, such as Zero Trust. The relationship between Shadow Access and Zero Trust and many other areas will be discussed in more detail in future documents.