bpi
BITS

Employee Identity and Access Management: A BITS Primer

MARCH 3, 2022

Table of Contents

Member Contributors

Trish Slavin, Bank of America

Gina Amos, Citizens

Alfred Bonilla, Mastercard

Anthony Morello, PNC

Jennifer Emerson, Regions

Eric Levita, State Farm

Joseph Frantz, Synovus

Ed Goff, Truist

Andrew Strear, Truist

Aaron Wilson, Truist

Noel Latsha, USAA

Mark Baran, U.S. Bank

BITS Staff

Andrew Kennedy

Executive Summary

In recent years several security incidents have been widely reported whereby bad actors have infiltrated a financial institution (FI) and gained access to confidential information or impacted business operations. Many FIs have been implementing point solutions to stop bad actors at the perimeter or address individual weaknesses. However, how effective can point solutions be as bad actors evolve their techniques?

IAM FRAMEWORK OBJECTIVES

- Reduction of security risk. The IAM Lifecycle Framework will better ensure that persons will be granted authorization to access the right resources at the right time for the right purpose.
- Creation of stronger identities through a process that begins when onboarding a subject. Rather than starting with a technological process that cannot be tied to a person, each subject's identity begins with a government ID.
- Position FIs for evolving to a Zero Trust, identity-centered paradigm that allows attributes, behavioral, and other factors to be considered before granting access to a subject.
- Reduce the complexity of meeting regulatory requirements by using closed loop provisioning and access review processes.
- Identify policy violations when a privileged user goes outside of the authorization process.

THE PROBLEM

For several years, the annual Verizon Data Breach Investigations Reports have concluded that individuals are weak links in cybersecurity – we fall for malware / phishing attacks and can be socially engineered to provide information and even personal credentials. Further, networks have been designed to facilitate access to applications and data. Because of these conditions, once bad actors get into an FI's network, they are able to move largely undetected from enterprise resource to resource. This puts business and client data, as well as business operations, at risk.

THE SOLUTION

Our proposal consists of a robust Identity and Access Management (IAM) Lifecycle program. Such a program will provide a framework that utilizes robust technology capabilities, end-to-end provisioning processes, fine grain entitlement control, privileged access management and includes robust monitoring and measurement mechanisms. By facilitating successful establishment of Identity, a strong IAM program has the potential to reduce the likelihood of security risks that can destroy corporate reputation, damage our clients, and impair corporate profits.

The IAM framework described covers several topics:

- Positions FIs to evolve into Zero Trust. The major capabilities and processes discussed in the paper are requisites for an FI to adopt Zero Trust – a direction that the federal government and our regulators are beginning to require.
- Describes the development of Identity Governance standards.
- Identifies the Identity Lifecycle from joiners through leavers, as well as events during a subject's tenure including request, approval workflows, and reviews/attestations.
- The capability of providing fine-grained access control through Entitlement and Role Management. No longer do FIs need to allow users to have broad access to a wide swath of data – that is unnecessary to perform a subject's business duties.
- Discusses the need to differentiate between primary and privileged or secondary accounts. Discuss the need for nonhuman (e.g., privileged accounts used by applications/services and APIs).

- The protection of privileged accounts – what technical capabilities are needed to secure, or vault, these accounts until the moment they are needed – and the processes necessary to monitor their access and use.
- The monitoring, alerting and use of accounts.
- Integration into cyber threat processes including incident response functions such as detection, containment, remediation, and recovery.

Implementing any framework involves risk. However, this paper is different. The capabilities and processes described within the framework have been implemented at many other financial institutions – and the paper represents our learnings.

- Most technology frameworks are written at an architecture level and leave it to individual FIs to interpret and apply them. This paper goes a level deeper and describes the characteristics of the capabilities and the processes that are to be engineered.
- The IAM framework discussed in this paper is based upon the NIST framework as well as successful implementations at several FIs.
- While most framework documents focus on capabilities and process, this paper goes deeper and provides guidelines to mature processes and capabilities; suggests which capabilities are foundational versus those that are built in later phases; and provides specific examples that are practical to implement.
- Gives guidance to allow an FI to evolve into the IAM framework – as we know it is impractical to believe that an FI can pivot on a dime from a low level of maturity to a high level of maturity across each area.

Why Identity Lifecycle Management is Key

Many financial institutions rely on either manual Identity Lifecycle processes or have Identity capabilities and/or standards inconsistently applied across the organization. A robust Identify Lifecycle Management program is critical to securing an organization's data and ensuring availability. Consider a robust lifecycle management program:

- Creates operational efficiencies within the organization by reducing provisioning effort and duration; reduces administrative errors.
- Provides insights to understand the access level each individual has across the organization and if there are separation of duty conflicts; as well as insights to understand who has and has utilized access to enterprise resources.
- Serves as the basis of a strong identity and entitlement program where access is tailored to an individual's organizational role.
- Serves as the basis of a strong Authentication (AuthN) and Authorization (AuthZ) program so that access can be provisioned at a finer level than in the past -- thus minimizing the risk footprint of an FI.
- Serves as the basis of an evolving Zero Trust security model that aims to restrict access to only those who have authorization to access the data, in a manner that the individual is anticipated to access the data.

The Identity Lifecycle Management process and capabilities will put an FI on the road to a more effective and efficient risk management program. To achieve success, enterprise strategies must consider the entire access management lifecycle, and be flexible and robust enough to accommodate a variety of use cases including offshore, cloud, and privileged access, all while relying on a trusted source of truth for identity information.

For the readers' convenience, a glossary of key Identity Lifecycle Management terms is included in table Apx-1.

Identity Creation and Assurance

Establishing a user's identity in enterprise systems starts before a username and secret credential are generated. An enterprise must first ensure a person is who they claim to be before creating credentials and granting access. Similar to a chain of custody, an enterprise should ensure a process is created and risks understood from identity proofing through credential delivery.

EMPLOYEE ONBOARDING

With this mind, a department responsible for credential generation and delivery must understand the business partners and stakeholders in the process. Typically, these tend to be Human Resources and an Onboarding Team.

Human Resources tends to be a common stakeholder given it's normally responsible to verify a person is who they claim. Human Resources can use a variety of methods to perform identity verification. Some organizations utilize and trust state-issued identification material. Others may rely on third-party identity verifiers who have their own methods. Either way, a cybersecurity organization normally entrusts this initial step to another organization within the business. It is important to understand who that is and to understand their responsibilities and accountability[1].

Following this, a process ensuring there is correlation between who Human Resources on-boards and how that person obtains network credentials must be established. The process should consider identity proofing practices within the context of an organization's risk tolerance.

For example, if the organization is handling highly sensitive data (i.e., mass amounts of personal identifiable information), that organization may require in-person identity verification prior to the hand-off of network credentials. The onboarding team may then rely on visual identification, where Human Resources is responsible for maintaining a photograph of the individual in a system.

If the organization wishes to mitigate risk further, they may put other processes in place. Some organizations may utilize Personal Identification Verification (PIV) cards containing biometric data along with an identifying photograph. Human Resources or some other party would be responsible for getting that data from the individual. Communicating credentials to that individual would require the presence of a PIV card.

Alternatively, an organization may wish to accept more risk. They may have processes where an individual is provided a secret passphrase that is verified with a helpdesk associate. If the passphrase is correct, network credentials are communicated to that individual. Whatever the method, the organization must adopt a process sufficient for its risk appetite.

An organization should also utilize similar considerations for secondary accounts, whether privileged accounts or accounts for other systems. It is up to the organization whether the use of an existing identity (like through a ticketing system) or some other form of identity verification is required prior to the creation and delivery of a secondary account.

CONTRACTOR ACCOUNTS

The same analysis for contractors must be completed. Here, however, a department may have different business partners and stakeholders to work with when onboarding contractors. A vendor management team and/or contracts team may be involved. Regardless of who, in each case, a business regularly relies on the supplier to

[1] https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-63-3.pdf

perform the initial identity verification steps. A process needs to be developed to connect that vendor's assertion to delivery of network credentials to an individual.

Identity is the New Perimeter - Heading Towards Zero Trust

Zero Trust architecture should be incorporated as organizations formulate and mature their Identity and Access Management programs. The focal point of Zero Trust is Identity – specifically the Authentication and Authorization aspects (see Figure 1). Organizations will need to implement Federation, Multifactor Authentication, Machine Identity, and verbose logging together with effective monitoring. These combine to create a strong identity that acts as the foundation of Zero Trust. The absence of a well-thought-out program that includes a strong Identity will hinder an organization's ability to adopt cloud solutions and ultimately a Zero Trust architecture. Please review the NIST special publication on this topic.[2]

"A zero-trust architecture is designed and deployed with adherence to the following zero trust basic tenets:
*1. **All data sources and computing services are considered resources***
*2. **All communication is secured regardless of network location.***
*3. **Access to individual enterprise resources is granted on a per-session basis.***
*4. **Access to resources is determined by dynamic policy***
*5. **The enterprise monitors and measures the integrity and security posture of all owned and associated assets.***
*6. **All resource authentication and authorization are dynamic and strictly enforced before access is allowed.***
"This includes the use of multifactor authentication (MFA) for access to some or all enterprise resources."
*7. **The enterprise collects as much information as possible about the current state of assets"***

Figure 1

CHARACTERISTICS OF A DIGITAL IDENTITY

Incorporating a Zero Trust approach to Identity requires that an organization define the characteristics of that identity. Several data points make up an individual's digital identity. Like physical attributes, there are digital attributes we can use to get higher confidence that an individual is who they claim to be.

Examples of Digital Identity Characteristics

- Credential a person is using (Single Factor vs. MFA)
- The device a person is using including who owns it. There are generally two classes of devices. The first is corporate owned devices (e.g., laptop or tablet) and managed by Mobile Device Management (MDM) software. The second is where an employee uses a personal device, whether in the office or when connecting to the corporate network or SaaS service. These devices are generally referred to as Bring Your Own Device (BYOD)
- The device types a person is using (Mobile OS vs. traditional OS)
- Geolocation of a person and person's device
- Geolocation of asset being leveraged

[2] https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf

Organizations can allow authentication to resources or data based on a combination of these characteristics. Risk tolerance and regulatory or legal requirements can determine how a user asserts their identity via a sign-on method or combination of sign-on methods.

SIGN-ON METHODS

A non-exhaustive list of sign-on methods are described below.

Memorized Secret

This is the basic form of authentication and the most susceptible and attacked by malicious actors. Enterprises must protect these secrets by enforcing complexity, length and the use of special characters. Organizations must evaluate these requirements to find a balance between security and usability -- for example, a very complex password will be difficult to remember, resulting in reduced productivity and increased Service Desk calls. Other compensating controls can be implemented to protect passwords, such as preventing password reuse and enforcing expiration/rotation. These controls help to promote the use of new passwords. A good compensating control is to restrict the use of vetted words such as username, enterprise name, and common passwords (Password123!). Additionally, it is ideal to compare passwords with leaked passwords to help mitigate dictionary attacks. Implementing a lockout threshold for wrong password attempts could mitigate brute force attacks. Password vaulting solutions could help users to centralize credentials and avoid passwords written in clear text on files or on paper. Corporate policies are also very important to protect passwords (e.g. noSecrets, noSecrets sharing, no Secrets hard coded in programming code, no plain text files with Secrets, password phishing awareness, etc.). These policies and practices must be promoted by enforcing security training at least once a year. It is critical to ensure regular testing of these controls.

Certificate-Based Authentication

A strong authentication method that is based on Public Key Infrastructure (PKI), where possession of a private key is considered a strong proof of identity. Enterprises must protect certificate private keys by making them non-exportable or secured by a vaulted secret.

This sign-on method + a PIN is the most common way to replace basic authentication by eliminating traditional passwords (passwordless). This is a recommended authentication method that potentially reduces the risk of credential leaks and the user friction when complex password requirements are enforced.

Single Sign-On or (SSO)

Reducing the number of times users are prompted for credentials is crucial for usability and security. If users are prompted too often for credentials, they will be more prone to enter their credentials during a phishing attack. One consideration for this method is that some critical or sensitive applications might require user authentication each time they log in - this is a common use case for Administrator logins.

Biometrics

Human physical recognition in a variety of forms is recommended as a second factor for authentication, where the first factor is often Certificated-based or Secret-based. Protecting biometric data is critical because once leaked it cannot be changed. Enterprises must consider that users will not be always able to complete biometric challenges so a fallback mechanism must be implemented - consider using a device PIN or One Time Passwords (OTP).

One Time Passwords (OTP)
OTPs are used as a second factor of authentication. The OTP delivery method must be selected based on the risk appetite. For example, using SMS or phone calls to deliver OTPs can be spoofed easier than a hard-token or soft-token authenticator. The access to the OTP codes should be protected by a PIN whenever possible.

Risk tolerance and regulatory or legal requirements can determine how a user asserts their identity via a sign-on method or combination of sign-on methods (commonly known as multifactor authentication, or MFA).

CONSIDERATIONS FOR SINGLE FACTOR VS MULTIFACTOR AUTHENTICATION

The decision to leverage single factor vs. multifactor authentication will be based on various options, some at the business's discretion and others dictated by external regulations. Companies will need to identify and prioritize use cases to mitigate these challenges appropriately.

Regulatory and Legal Requirements
The decision to require multi-factor authentication (MFA) will be based upon several considerations. In making the determination to use MFA, it is critical to look at regulatory requirements. As an example, the New York Department of Financial Services (NYDFS) has levied very large fines against corporations that failed to implement MFA consistent with their regulations[3][4]. The SEC has also fined financial services companies for failing to protect email accounts with sufficient controls, thereby allowing breaches even after the organizations were initially warned of the vulnerabilities[5].

Similarly, Europe's General Data Protection Regulation (GDPR) requires that data considered a high risk must be protected by MFA[6]. In these cases, the organization may face a grey area regarding applicability of the requirements, as there can be ambiguity in the language. Bearing that in mind, MFA should be considered for use cases that provide access to sensitive data, and for which increased friction is acceptable to prevent risk of data loss/exposure. Be aware that some regulations require specific controls when customer data is involved, including strong authentication. This must always be a primary consideration when planning when and where to enforce MFA in a system.

Beyond Legal and Regulatory Requirements
A userid and password alone are no longer adequate controls for protecting an organization from malicious actors. Both recent and past cyber-attacks demonstrate the disruption caused by compromised credentials in the form of ransomware and data breaches. The implications go beyond the attacks themselves and have a ripple effect in brand damage and the subsequent fines levied for those said breaches.

In addition to regulatory requirements, there are several factors to consider when deciding to employ MFA:

1. **Is a user traversing network zones?**

[3] https://www.dfs.ny.gov/reports_and_publications/press_releases/pr202105131
[4] https://www.jdsupra.com/legalnews/nydfs-imposes-heavy-fine-and-adopts-7165817/
[5] https://therecord.media/sec-fines-three-companies-over-hacked-employee-email-accounts/
[6] https://www.irisclasson.com/2017/10/06/not-so-stupid-question-303does-the-gdpr-require-two-factor-authentication-from-companies-processing-personal-data/

MFA must be used when traversing zones. If your users are remoting into your network from outside, that user must be required to assert their identity and MFA. Also, if a user is communicating, connecting to, or utilizing information resources from one network zone to the next within the network, that user should have to assert their identity and MFA.

In cases where users are connecting to a remote cloud solution, they will also need to leverage MFA (see Figure 3). The New York Department of Financial Services has set the precedent for the cloud-based MFA requirement with its decision to fine First Unum Life Insurance Company of America ("First Unum") and Paul Revere Life Insurance Company ("Paul Revere") 1.8 million dollars combined.

"The investigation uncovered, among other things, that First Unum and Paul Revere violated the DFS Cybersecurity Regulation by failing to implement Multi-Factor Authentication ("MFA") without implementing reasonably equivalent or more secure access controls approved in writing by the Company's Chief Information Security Officer."[7]

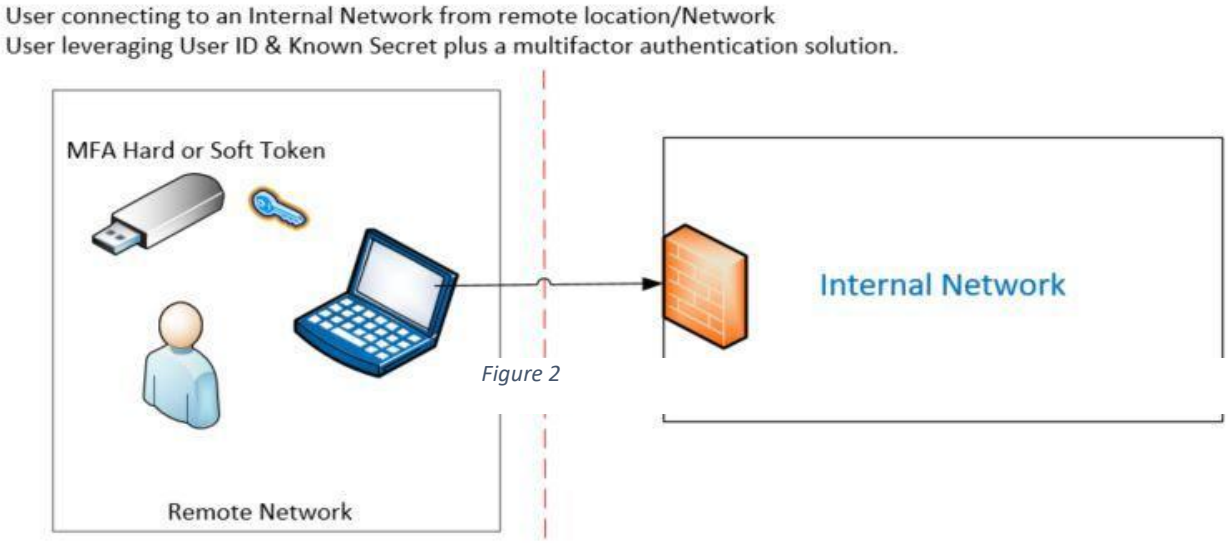

Figure 2

Figure 2

2. Is the access privileged?

In cases where the access is privileged, the identity asserting that privilege should be proven via MFA. It is up to the organization how to define what privileged access means. For most, it could be defined as an administrator or individuals who can change configurations that could expose sensitive data. However, privileged access can also mean the type of data the identity is authorized to access. If a particular application holds highly sensitive data, an organization may require a user to use MFA to authenticate to the app. Contrast this to an application holding non-sensitive data that is accessible through either Single Sign On or single factor authentication.

[7] https://www.dfs.ny.gov/reports_and_publications/press_releases/pr202105131

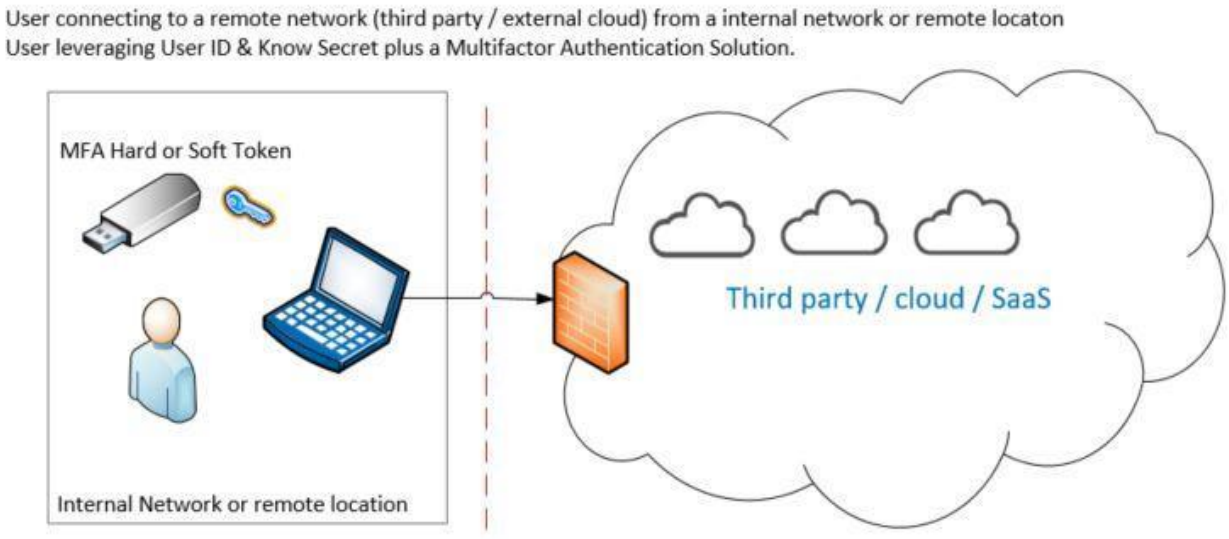

Figure 3

3. **How they will meet increasing internal "user experience (ux)" expectations.**

As the digital disruption wave continues to impact our personal lives, employees are aware how frictionless the customer experience can be. For example, in the past, we laboriously filled out a screen of information to checkout from a website. Today, we have one click checkout. In many stores near field (NFC) technology allows us to tap and go. The exposure to modern digital experiences in employees' personal lives have pushed for better user experience in their business life.

While employees are enjoying less friction in their personal consumer lives, FIs are increasingly using new technology to reduce risk (increase security) and fight fraud. Additionally, as FI environments become hybrid – spanning on-premise and cloud environments -- Identity practitioners need to find more ways to secure client and business data.

Modern, strong authentication requires a balance between user experience and security to stay ahead of cybersecurity threats. Business efficiency will suffer if security controls are too onerous. FI employees will quickly lose patience if the security protocols are too onerous to use. At minimum, difficult to use controls can drive an increase in calls to the help desk. Worse, poorly implemented controls will drive employees to circumvent them. Multifactor Authentication (MFA) solutions should balance the level of protection and usability.

For example, users are familiar with the use of biometrics to unlock their mobile devices. They may not realize that this represents a multi-factor authentication (something you have + something you are). The familiarity with this type of MFA is likely to ease introduction and acceptance within corporate environments. Conversely, a higher friction second factor will not be well received. For example, a hard-token used to generate an OTP may be perceived as a much higher friction MFA implementation, and would not be suitable for general users.

When considering the use cases to employ MFA, FIs must take a proportional approach to balance the risk of lateral movement in the network; privilege escalation; and employee user experience. If the data value or risk are deemed to be high, it is reasonable to add friction to the user experience. With Identity as the new perimeter, the sign-on methods discussion is critical for the FI to establish confidence that the subject is who they claim. Equally important to establishing strong identity is for an FI to be able to uniquely identify a subject as they remain connected to the FI's network; and throughout their lifetime.

Identity Governance

Identity Governance is the set of standards and processes a firm uses to orchestrate its Identity life cycle events (following section) – from onboarding to off-boarding. It also encompasses the set of standards and processes that support access reviews / attestations / audits. This section explores a few aspects of Identity Governance

NAMING CONVENTION

Each organization's lines of business, products and organization structure is unique, and as such, will more than likely have a different naming convention for identities from one organization to another. However, it is important for naming conventions to quickly identify characteristics about the resource. For example, whether a userid supports privileged or standard access; whether it supports a production or lower environment workload; or some other useful information. Consider that different information will be relevant for different enterprise resource types (e.g., userid versus a device name).

Identities represent an "entity" which can be a person, organization, or device. Each identity will have a standard set of attributes assigned, which provide additional information on the identity to other applications and users and may determine what access is appropriate.

Naming conventions for identities should be leveraged wherever technically feasible and maintain a consistent pattern within an organization. Well-defined naming conventions provide an organization with the ability to administer identities, reduce time spent on governance and provide consistent data that is translatable to audit. These conventions should be clearly documented within organizational policy, outlined in operational guidance, and enforced through audit. An effective naming convention must meet the following requirements:

- Is unique across the organization
- Meets specifically defined criteria such as minimum length/only number/only characters etc.
- Works with all current and future target systems
- Works for users with multiple affiliations
  Works with systems between multiple departments/locations

Example naming conventions include:

- First initial + Last Name, where numerics would be added to the end for common names (ie, JSmith, JSmith2)
- First Name + Last Name, where numerics would be added to the end for common names (ie, JohnSmith, JohnSmith2)
- Email address (ie, john.smith@org.com, john.smith2@org.com)
- Randomly generated Unique ID of fixed length (ie, D0PM4, 03596)
- Randomly generated Unique ID of fixed length with designation for non-employees (ie, X03596 for contractor, T03596 for Temp)

IDENTITY REPOSITORIES AND THE EMERGENCE OF IDENTITY GOVERNANCE AND ADMINISTRATION PLATFORMS

An Identity Repository stores all Identities for an organization. Identity repositories serve functions such as:

- Identity repositories enforce unique identities. In the previous section, Name standards were discussed as a way to enable uniqueness. To enforce uniqueness, it is important for a central identity repository to exist. For example, two independent database servers not using a central identity store can each create a user "MyDBadmin." However, the owner of each of the local accounts can be two different people. Further, using identical names makes it difficult to distinguish between the two (though other factors can assist, it would be laborious).
- Identity repositories simplify Identity Lifecycle Management processes by creating a single source of truth; by providing comprehensive views of what users can access or all those who have access to an application.
- Identity Repositories typically have connectors that allow various applications to consume Identity services (e.g., entitlements). In many cases, the central repositories can consume identities from applications that locally administer AuthZ - so that the organization can have the views described above.

Identity Governance is more than the repositories. It also includes core functionality such as:

- The ability to discover local accounts and entitlement through connector applications that consume this data from applications
- The ability to support role engineering so that AuthZ can be done at a group level rather than to specific individuals.
- The ability to discover segregation of duties issues based upon an organization's rules
- Access Review and Certification so that organization can manage risk footprint by keeping entitlements as granular and tightly administered as practical.
- Automated Provisioning based upon agreed upon workflows

The systems that perform Identity Administration and Identity Governance have come together over the last several years. These Identity Governance and Administration (IGA) platforms (e.g., SailPoint, Saviynt, RSA IGA, etc.) can be leveraged to define authoritative sources for identities (such as the HR system, Vendor Management System, other key Identity stores) and aggregate account information from other platforms (Active Directory, UNIX, mainframe, cloud, etc.) and applications. With the appropriate rules and workflows these systems can fully automate the identity management processes, provide robust governance and provide a holistic view of the identity across the organization.

Identity Life Cycle

Identities have defined life cycle events, typically broken down into several main categories and driven by changes in HR attributes. While many of the account management processes can be handled through various scripts and/or manual processing, IGA systems are specifically designed to automate and audit these processes.

JOINER

Also known as "onboarding" or "new hire" events, this represents the initial creation of the identity and, at a minimum, a primary account. Depending on the program, certain "birthright" access may be granted, such as access to the company intranet home page, the Microsoft Office suite, and a blank home drive to store information. Other applications or software may also be provisioned automatically, depending on the defined roles.

Once the initial accounts are created, job-specific access requests can be placed for applications required to perform daily functions.

Below are two renditions that demonstrate the creation of an Identity in a joiner process. The process always begins in the system of record in order to ensure a novel identity is being created, leveraging established roles and naming conventions, and the appropriate approval process.

The second process is similar to the first as it represents a joiner process. In addition to the process steps represented in the below diagram, this diagram shows interaction between the system of record (Identity Repository) and the Identity Directory (e.g., Active Directory) that consumes these identities.

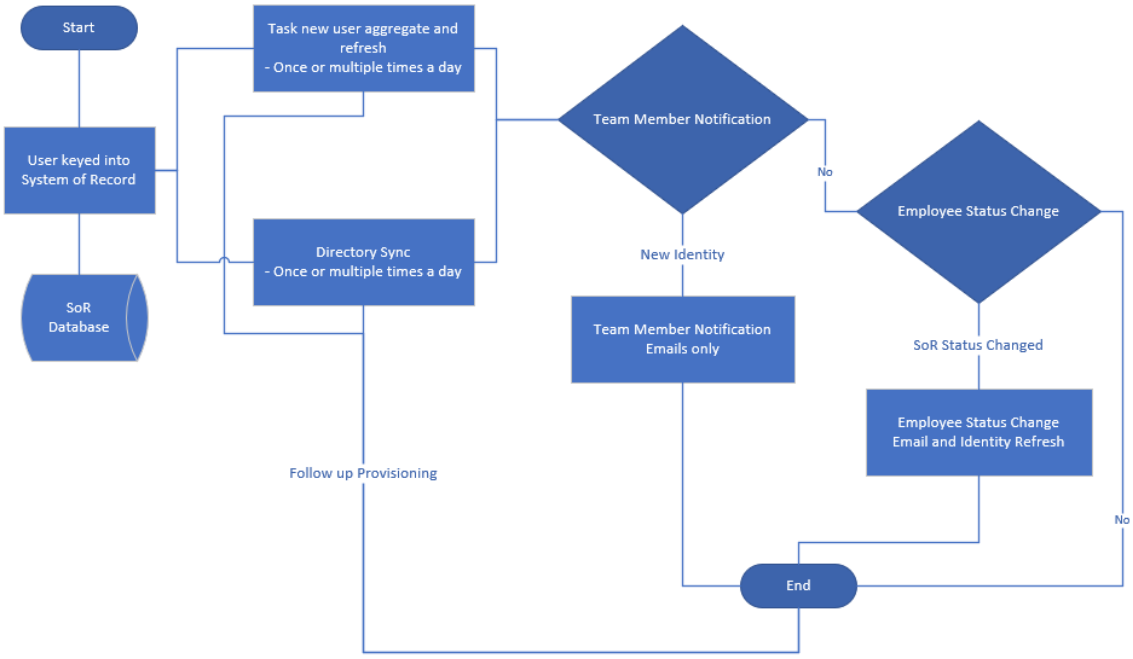

Figure 4

MOVER

When a person switches functions, best practice dictates that access be reviewed or reduced to baseline or birthright access, depending on the function change and the rules within the organization. These changes are usually indicated by a change in HR attribute, such as a job code, cost center, manager or a combination of multiple attributes.

For switches between front line and other lines of business, access should be reduced to new hire status and job-specific access requested and approved by the new manager. For moves between a single line of business, a re-certification of access by the new manager should be performed.

Additionally, during job function changes within the organization, ownership rights assigned must be evaluated, and either recertified or removed to ensure proper ownership is maintained. Ownership to be evaluated includes, but is not limited to, ownership of a Role, Entitlement, Non-human and Secondary Accounts, Resources, and Applications.

LOA/RLOA

For temporary leaves of absence (LOA), disabling logical and physical access is likely good enough to ensure security. Access can be resumed quickly upon their return (RLOA). These events are typically triggered by a change in HR status.

LEAVER

A person leaves an organization for a few reasons: voluntary resignation, termination due to a reduction in force ("lay off"), or termination for cause. People can also be placed in a suspended status pending an investigation, if, for example, if the person's account is suspected of being used by a third party. The deprovisioning process will likely be different depending on the use case.

For voluntary resignations, logical and physical access can be revoked at the end of the notice period, with full revocations following. Access for the underlying entitlements can occur through normal BAU processes in line with the service level agreements and pre-defined metrics. These events are typically triggered by a change in HR status.

In cases where a person is terminated with cause, physical access and logical account access should be revoked immediately. Underlying permissions can then be revoked as soon as possible, but the immediate threat has been removed.

In cases of immediate suspension, it is prudent to terminate active sessions after disabling logical account access. Once suspicion is lifted, the person's access can be retained with minimal disruption.

If a person's employment is terminated, ownership of resources assigned to that person should be reassigned. Typically, that person's manager receives the reassigned ownership. If no manager is available, the items should be assigned to an appropriate alternate team member or system administrator until further review. Ownership to be evaluated includes, but is not limited to, ownership of a Role, Entitlement, Non-human Accounts, Resources, and Applications. Also consider, if the owner of a secondary account within the cloud is terminated, whether the account must be disabled or deleted.

Identity Accounts

Each identity within the organization can have one or more accounts. As with the identity, each account should have a standardized naming convention across a set of attributes that determine the account type as well as provide a mechanism to tie it back to the identity. This can be done either through standardized naming convention or through attribute values. Where technical constraints may cause one system to use a naming convention, this should be documented, enforced, and audited to maintain consistency.

In general, users will have two types of accounts:

PRIMARY ACCOUNT

A primary account is the main account leveraged for day-to-day activities and is typically lower privileged. This account will allow the user to log into the network(s), check email and perform routine tasks. Ideally the naming convention of the primary account will match the naming convention of the identity.

SECONDARY ACCOUNT

Secondary accounts provide access to non-standard activities. These may include privileged accounts, test accounts, internal application accounts or accounts in external/cloud systems. Ideally the naming convention of a secondary account will match the identity. In instances where the secondary account provides specialized access, such as administrative or testing privileges, the account name should include the identity as well as a mechanism to determine the account type. For example:

- Privileged accounts may have an "ADM" prefix to indicate Administrative access
- Test accounts may have a "T" or "-test" appended to the account ID

These accounts require additional levels of approval, monitoring and alerting.

Entitlement and Role Management

ENTITLEMENTS

Entitlements are what a user requests to gain access to specific resources. These may include Active Directory groups, distribution lists, internal application access, UNIX groups/netgroups, database access, cloud access, etc. The larger the organization and the more applications that are supported the larger the list of entitlements will be.

As with user accounts there are certain attributes for Entitlements that are key for governance functions, such as requests, approval workflows, access reviews, etc. Whether using an IGA solution or managing access manually, the following are key considerations when implementing an Access Governance program for Entitlements.

Managing Entitlement Scope And Nesting

Each entitlement's scope should be as granular as possible to ensure what it grants access to is clear and defined. It also provides a solid foundation for building roles, which will be covered later.

Nesting is the practice of adding one group to another group. Administrators have done this to simplify administration. And while this is an effective practice in a small, new environment. Nesting quickly becomes an unwieldy and risky practice. For example:

- Over time, a nested group may evolve its purpose or expand its membership resulting in unintended consequences such as unknowingly providing Privileged Access to additional individuals.
- Access Reviews and Attestations are cumbersome because of the need to recursively iterate to find all userids.
- Inheritance rules might create unexpected outcomes.
- Larger, unwieldy groups may cause segregation of duties issues.

Naming Conventions

While entitlement attributes are very useful for designating purpose, a standardized naming convention can provide users with an additional context of the entitlement's purpose. As an example, if your SQL Server databases leverage AD group membership to authenticate, the naming convention may include "SQL", the application code and the type of privilege: SQL_MYAPP_R (read only access to the database for MYAPP) or SQL_MYAPP_RW (read/write access to the database for MYAPP).

While the example above was simple, naming conventions can be as granular as the organization and standards permit. As another example, the naming convention may be [Directory]-[GroupType]-[ResourceName]-[Environment]-[Permission]-[Appendix as Needed],

For an example of this breakdown, refer to table APX-2 in the appendix.

Detailed Description
Having robust business descriptions for entitlements is essential for ongoing management. When pulled into a request catalog, the descriptions assist users in determining the correct level of access to request. They also help the approvers determine if the access being requested is appropriate for that user in both the request and review process. Finally, they provide additional information to the risk/audit partners to assist with any documentation requirements without direct input from the IAM team.

In general, the description should include what level of access is provided by the entitlement as well as what group(s) the access may be appropriate for. Any additional justifications for the entitlement can also be included.

Defined Owner
Ownership of entitlements is critical to future management. Each entitlement should have a defined owner (or group) who is responsible for any access approvals and ongoing access reviews. Owners should be full-time employees within the organization and should function at a high enough management or responsibility level to understand the responsibilities of ownership. Contingent workers (contractors or temporary employees) should not generally be owners of an entitlement.

Entitlement owners are also responsible for communicating any broad usage changes to the IAM team and would field any questions from Risk/Audit as to the purpose of the entitlement.
Given these responsibilities, ownership should be reviewed on a periodic basis to ensure it is appropriate and authorized.

Linked to Application
Where possible each entitlement should be tied to an application. This provides an avenue for application-level access reviews as well as a mechanism for ensuring associated entitlements are removed when an application is retired.

Other Information
Depending on the system, there may be several other attribute fields that may be leveraged. Some additional information that may be helpful are:

- Ticket number – To correlate the information back to the request
- Data classification – May be used for review purposes
- Regulatory types or jurisdictions as these may impose additional or specific controls
- Requester – This may be different than the owner
- Privilege flag – Is this entitlement considered privileged?
- Requestable – Is this entitlement available for users to request? If it is tied to an application being decommissioned, has become stale or has been included in a role it may no longer be eligible for independent requests.

ROLES

Roles are groups of entitlements bundled together to provide a standardized set of access and ease provisioning requirements. Roles ease provisioning by: operating on collections of users at one time which help enforce standards and uniformity; creating granular collections of userids with human digestible group names that make their purpose obvious; and more efficient for OS applications to process. Roles may consist of entitlements within a single application, single platform or, in more mature models, may span across multiple applications/platforms.

Roles can simplify the access review and request process by consolidating numerous smaller items, sometimes with many different owners and approvers, into a single item with a single approver. Many roles are requested based on specific job function or business need. In some cases, roles can be tied to user attributes and be provisioned automatically. All roles should have a defined purpose and a defined population of applicable users. In many instances users can be assigned multiple roles or a combination of roles and attributes.

Like entitlements, roles have attributes that are critical to governance, such as purpose, naming conventions, descriptions, and defined owners. However, since there may be numerous entitlements included in a single role, each with different descriptions and owners, additional due diligence and oversight is usually required.

Purpose
The purpose of a role is typically related to a specific function, whether that function is within a specific application, job role or platform. Some items to consider when determining the purpose are:

- Consider the number of users that would be impacted by the creation of a group and the administrative ease/burden – If there is a large group of users with similar job function and access needs, spending the time on entitlement analysis to create a role would be beneficial. However, if the population is very small or the job function is very narrowly defined building a role may not make sense. In general, users who are granted a specific role should require at least 80% of the access defined to complete their daily job duties.

- Balance competing needs of HR and Technology – many times there is a tug of war between HR and the IAM team that can drastically impact roles. HR's goal is to consolidate as many different positions as possible into job titles/families that span multiple departments across the organization. Developers, analysts, managers, auditors, etc. – these are all generalized "buckets" that personnel may fit into.

However, when it comes to access, these users may have significantly different access needs, depending on the department, project or other criteria. This makes it very difficult to have a single role that will fit all the unique needs.

- Consider how dynamic role group membership may be – building adequate roles can be very time consuming and may require additional documentation/signoff to ensure all audit requirements are met. Access required only for short durations or access needs that may change frequently as part of a job function are typically not suited for roles.

- Leverage identity attributes when deciding on creating a role – this could be job code, cost center, department, full-time/part-time status, manager designation, contractor, etc. These attributes can be leveraged alone or in combination with other attributes (i.e., job code + cost center) and can provide additional opportunities for automated access management.

Naming Conventions
Role naming conventions are typically tied to the purpose and should provide information on how the role is used and to whom it applies. Below are some example conventions, based on usage:

Role Type	Description	Format and Examples
Application Role	A set of permissions specific to a particular Application, such as Salesforce or Workday.	[ApplicationName]-[RoleName] Salesforce-CRMAdmin
Business Role	A set of permissions for a particular job title	[Dept]-[RoleName] IT-SecOpsAdmin
Platform Role	A set of permissions specific to a particular cloud provider	[Directory]-[CloudServiceType]-[PlatformService]-[RoleName]-[Environment] Azure-Networking-DNS-DNSAdmin-Prod

Detailed Descriptions
Since roles include multiple entitlements, the descriptions often contain a generalized access summary as well as purpose and applicable population. As with single entitlements, the goal is to provide enough information for requesters and approvers to make informed decisions for request/reviews and for auditors to understand how the roles should be applied.

Example:
The Salesforce-Marketing Assistant-Member Role provides all Marketing assistants access to the Salesforce Application

Role Ownership
The same rules that apply to entitlement owners also apply to role owners – they can be one or more people, should be full time employees and be at a high enough level to understand the responsibilities of ownership. Additionally, role owners should be responsible for reviewing and evaluating the role composition on a periodic basis (at least annually) to communicate any changes that may be required.

Other Attributes
Additional attributes can vary, depending on the organization. Some that may be beneficial include:

- Status: Is the role currently Active or Inactive?
- Requestable: Can the role be requested manually?
- Criteria: Are there any HR attributes that drive provisioning of the role?
- Sensitivity: Does the role provide access to sensitive information?

Provisioning & Deprovisioning

The previous section discussed roles and groups. These two constructs can be used to facilitate the provisioning and deprovisioning of access to related Identities.

Access to applications and other information resources is granted through the process of provisioning and revoked through deprovisioning. An approval step is added to the process to ensure access being requested is appropriate for the person's role in the organization. This approval is usually performed by the person's manager. In many cases manager approval may be sufficient. However, access to higher risk applications, confidential information of clients or the enterprise or privileged/administration access typically require additional approval.

Provisioning and deprovisioning processes may be manual or automated. The work can be centralized in a security function, a specialized provisioning and deprovisioning group or managed by the individual application teams.

Having a well-defined approval framework and consistent provisioning and deprovisioning processes are key components to access management operations. While the approval workflows help prevent inappropriate access from being granted, timely revocations ensure access that is no longer required is removed as quickly as possible. SLAs for both approvals and request fulfillment as well as ongoing monitoring of acceptable thresholds are typically used as Key Risk Indicators (KRIs).

In addition, access events should be captured on at least a daily basis and stored for auditability per your organization's policies and standards. In addition to event logging, access logs can provide key insights to system access, system use and reasonable usage patterns. If anomalous behavior is detected, the logs can be used to verify activity and usage patterns.

ACCESS CONTROL MODELS

Access control establishes an individual completing a particular task, authenticating them by their credentials and granting them only what they need access to, nothing more. Access is at the heart of every organization, and it is imperative that it is appropriately controlled to avoid misuse or abuse of privileged access.

The primary objectives of an Access Model include:

- Enforcement of least privileged access
- Enforcement of separation of duties
- Efficient and timely access provisioning
- Proper approvals and access reviews pursuant to the company's risk appetite (user's manager and where necessary resource owner)

  Note: Resource owners should be considered leadership for dual authorization[8] *purposes as defined by NIST and outlined in* NIST 800-53 AC-3. [9] *Dual authorization puts the onus on the responsible party of a given asset to ensure the access being provided to the underlying asset is warranted.*

Access Control Models have four primary types:

- Mandatory Access Control (MAC)
- Role-Based Access Control (RBAC)
- Discretionary Access Control (DAC)
- Attribute Access Control (RBAC or RB-RBAC)

Mandatory Access Control (MAC)

MAC is a network-based access control model that utilizes policies, passwords, and settings to govern access centrally. System administrators with elevated privileges maintain access from a central controller (computer). The system administrators have full control to an individual's access. Administrators can prescribe customized access patterns per individual, making it highly customizable but resulting in additional overhead for administrators. MAC could be a good option for privileged access management models as the system administrators have complete control. However, there are many instances where users not in a system administrator role will need privileged access. The system administrator would need to create a customized pattern for that individual and subsequently must maintain it.

[8] dual authorization - Glossary | CSRC https://csrc.nist.gov/glossary/term/dual_authorization
[9] NIST Risk Management Framework | CSRC https://csrc.nist.gov/Projects/risk-management/sp800-53-controls/release-search#!/control?version=5.1&number=AC-3

Pros

- Only system administrator can control the access, deeming it more secure
- Security error reduction

Cons

- Network configuration drives policy decisions

Role-Based Access Control (RBAC)

RBAC is a model where every individual is assigned a specific role based on their job duties. Those roles are associated with permissions or access that will be granted to the members of the role. This requires establishing appropriate team roles, using access objects (such as Active Directory groups) to represent the role, and partnering with division leaders and stakeholders to associate team members to corresponding role. RBAC is a popular option for enterprises when it comes to access management models. Once roles are established, privileged entitlements can be defined and maintain in a central place. This allows for consistency in distributing, management, and enforcing privileged access.

Pros

- Centralized and covers the entire organization
- Less work and more streamlined for users of roles

Cons

- High difficulty to implement and maintain for administrators of roles

Discretionary Access Control (decentralized)

DAC is a model where the custodian (IT admin) determines access permissions for a given individual, ultimately giving full control to the end-user on any object owned by them. DAC allows flexibility from a privileged access management perspective; however, it can easily become too much flexibility. The decentralized model may need to have centralized guiding standards to ensure consistency from a PAM perspective to make it truly viable on an enterprise level.

Pros

- Users may transfer ownership of an object
- Users may determine another user's access type
- Authorization failures restrict user access.

Cons

- Inherent vulnerabilities
- Less secure than other models
- Limited central control

Attribute-Based Access Control (ABAC)

ABAC is model based on attributes (also known as **policy-based access control**). These attributes could be related to a particular environment, resource or individual and the model supports "IF THEN" Boolean logic. For example: If the requester is the technical manager of the application, then allow access. ABAC is commonly used within privileged access models as policies can be defined, implemented, and maintained at a very granular level. Be aware that the granularity poses challenges, as minor changes in the organization may require repeatedly adjusting policies. ABAC is a good model to use on conjunction with another model like RBAC.

Pros

- Highly customizable
- Access can be defined at a granular level

Cons

- Attributes must be defined up front which causing initial set up to be more time consuming

Hybrid Approach
Given the multiple access model approaches, many organizations use a hybrid approach based on the resource type. For example, a role-based approach may be used for application access while an attribute model is applied for accessing data.

Practical Implementation Example
Analysis can be conducted on a server access model to determine improvement opportunities, which may allow for automated access provisioning at build time while maintaining least privilege using the RBAC (role-based) model.

Prior to moving forward with implementation, identify all stakeholders and work with each individually to review the suggested improvements. Collaborating with stakeholders is critical to creating an access model that will meet enterprise needs. Flexibility is also key when working through differences of opinions amongst interested parties -- this facilitates reaching a consensus on roles and entitlements per environment.

Many models fail the implementation phase when a model is developed but there is insufficient follow-through. The final model should be communicated to individual teams via a top-down approach and it's critical to include reference guides and documentation.

With a uniform role-based model implemented, it is much easier to establish accurate preventative controls aligned with the model. With less nuances to account for, the PAM team can enforce the enterprise access model and flag any diversions for quick remediation.

Example of general roles and entitlements for an access control model related to server access:

Aqua indicates privileged access

Role	Development		Non-Production		Production	
	Non-operational Servers	Operational Servers	Non-operational Servers	Operational Servers	Non-operational Servers	Operational Servers
Developer	Administrator	Administrator	Ability to read log files	Ability to read log files	No access (Logs via Splunk or other tool)	No access (Logs via Splunk or other tool)
Integration (Build) Engineer	Administrator	Ability to read log files	Administrator	Ability to read log files	Administrator	No access (Logs via Splunk or other tool)
Decentralized Operations / Production Support	Ability to read log files	Administrator	Ability to read log files	Administrator	Ability to read log files	Administrator
Centralized Operations / Production Support	Ability to read log files	Administrator	Ability to read log files	Administrator	Administrator	Administrator
Application Support	Ability to read log files	Ability to read log files	Ability to read log files	Ability to read log files	Ability to read log files (If Splunk or another tool is not sufficient)	Ability to read log files (If Splunk or another tool is not sufficient)

Periodic Access Reviews

USER ACCESS REVIEWS

User access reviews should be performed on access to resources, applications, and platforms to ensure access is restricted to authorized individuals and is appropriate based on job responsibilities. These reviews may be performed monthly, quarterly, semi-annually or annually, depending on organizational need. Frequency may be driven by asset classification, type of access (i.e., regular or privileged), calculated risk score, special designation (i.e., SOX, PII) or a combination of requirements. This frequency needs to be planned and implemented with all stakeholders represented. They may be performed at the application, resource, or platform level or as part of an identity-level review.

User access reviews should include the following:

- Entitlements
- Roles
- Privileged Access
- Dormant Accounts
- Orphan Accounts
- Non-human and Secondary Accounts

Access reviews may be performed at the manager level, application owner level, or a delegate. Reviews can be performed by a single user or group of users. For applications with a larger footprint, tiered approaches can be taken where the manager does the initial review and the application owner does a secondary review to ensure any remaining access is appropriate. Birthright access that is automatically provisioned or deprovisioned based on HR attributes may be excluded from periodic reviews.

Each access review campaign should have a designated Service Level Agreement (SLA) for review completion, and monitoring must be in place to ensure the reviews are completed within the appropriate time frame. For any reviews not completed within the allocated timeframe, all access should be considered for revocation.

OWNERSHIP REVIEWS

Entitlement and role ownership, including privileged role ownership, must be reviewed on a periodic basis to ensure ownership is authorized and appropriate. Ad-hoc reviews may be required after a departmental reorganization.

ROLE COMPOSITION REVIEWS

All roles should be certified annually, at a minimum, by the role owner. Documentation should be prepared that lists the role attributes (name, owner, description, requirements, etc.) as well as each individual entitlement (including the description and entitlement owner) within the role. The role owner will review, document any required modifications and sign off on the configuration.

Since there may be instances where underlying systems and/or entitlements have changed and are no longer appropriate for the established role, a secondary sign-off from each of the entitlement owners is also recommended. This ensures that all parties agree that the access being granted is appropriate.

ADDITIONAL CONSIDERATIONS

- Keep your target audience in mind when designing your program. Generating a single but very large review or a smaller but very frequent review can lead to review fatigue and encourage rubber stamping. Ideally, systems should be grouped together by risk rating, regulatory designation, asset classification, department or some other mechanism in order to keep the line items being reviewed to an acceptable level.
- Metrics are key. How long, on average, does it take for reviewers to complete the review? Once the access is marked for removal, how long does it take for the access to be removed? How many line items are being reviewed, and what is the revocation rate for a given period? If there is a spike in a revocation rate, does it indicate there was previously excessive access that is being addressed, a major change in functionality, a reorg? Leveraging artificial intelligence (AI) in the IGA tool can assist with identification of such metrics.
- Whenever possible, only exception access should be reviewed for role or model-based access. For example, if all members of Department H should have access to resource H, then only exceptions to the rule need to be reviewed.

- Evidence of the completed access review, including the users that were reviewed, the date and timestamp of when the review was completed, and the results of the review must be retained to support future audits. In addition, a copy of the review results should be provided to the application owner and contacts with the following expectations:
    - To retain a copy as evidence for any business line controls
    - To review the results to ensure all entitlements were included in the review and the decisions match what they would expect
    - To validate that any access marked for removal was natively removed from the system
- Where possible, provide audit with direct access to any documentation such as onboarding templates, change requests, access review summaries, etc., allowing them real-time access to the evidence they need without engaging IAM resources and management.

Privileged Access Management

WHAT IS PRIVILEGED ACCESS?

Privileged access is a general term to represent different access levels, depending on the group. From an infrastructure perspective, it could be defined as domain administrators, server administrators, workstation administrators, software deployment access, etc. From an application perspective, it could be defined as administrative rights within the application or specific access levels that are deemed critical. From a database perspective, it could be limited to database administrator (DBA) privileges or anything outside of read only access. The various interpretations of privileged access can lead to inconsistent controls and monitoring, resulting in additional risk and audit findings. If defined and tracked effectively, the IAM team should be able to quickly produce lists of any elevated access for any audit and/or risk requests.

Defining Privileged Access

Providing objective criteria on what is considered elevated and how that access can be identified and managed is key to reducing risk within the program. Some of the questions that can be asked to help identify the various categories are:

- Who has the “keys to the kingdom”? This represents broad rights across the infrastructure, such as OS admins (Active Directory, UNIX, Mainframe) or storage admins.
- Who has the ability to deploy code or perform other functions that could have a broad impact to the organization? This could include software deployment, middleware teams, firewall teams, file transfer admins, DBAs, etc.
- Who has administrative rights on production servers?
- Who are the workstation administrators?
- Are there any additional defined groups or exceptions?
- Do members of any of the above groups require secondary accounts, additional authentication measures, etc.?

DETECTING PRIVILEGED ACCESS

Once you know the various types of access you need to manage, the next step is determining which entitlements fit into each category. This may include a cursory review of entitlements and privileges, such as AD groups with elevated rights, UNIX groups and netgroups that provide the ability to sudo to root, database roles that grant DBA privileges, etc., or it could leverage discovery tools that can look for specific configurations within the environment.

The detection process can also help identify any privileged groups/entitlements that have become stale or are no longer required. In these cases, the entitlements should be removed through the proper change control to close any potential exposure gaps.

After considering the types of privileged accounts in the environment, an organization needs to have proper process and controls to know where they exist. In order to control risk, an organization's privileged account monitoring must be comprehensive. Most mid- and large-sized IT organizations have robust lifecycle management processes in place which impose controls on provisioning userids. However, it is necessary to simultaneously wear both "belt and suspenders" with privileged accounts. Why? Imagine if an administrator used their privileged account to create another privileged account – how would an organization know? There are several sources to consider.

- **Active Directory Groups** contain the userids that have rights granted in active directory and allow the userid to perform actions on a subset or all objects within the domain. Inventory scripts as well as many commercial products allow an organization to collect those accounts from a convenient central source.
- **Local machine privileged accounts** are accounts defined on the enterprise resource and require a tool to visit each machine and gather that set of accounts.
- **IaaS privileged accounts** may be defined in an organization's identity store, such as Active Directory, and are federated with the IaaS application. Federated applications are more amenable for this type of inventory. From time to time, IaaS applications are not federated with an organization's identity store, thereby requiring local accounts. This is a more complicated situation that requires a custom approach.
- **Cloud Roles** are similar to IaaS in that on premise accounts can be associated with cloud roles and privileged userids can be easier to inventory. However, many of the cloud providers have sophisticated IAM and PAM capabilities that should be utilized. Additionally, some cloud automation services (e.g., Amazon's kubernetes services, EKS) create a container-specific instance of OIDC to support their operations.

Although this paper addresses privileged access accounts**, don't ignore IoT and Compute and Storage identities** like containers, EKS, lambda functions, S3 buckets, etc. While not included in the definition of most organizations use of "privileged accounts", these IoT items are identities and, like service accounts, can access sensitive data and perform critical operations (e.g., provision and de-provision of objects) in an environment.

To prevent teams from overlooking important alerts, a risk-based filter should be placed on viewing logged alerts. Organizations should be more concerned with anomalous privileged authentication events than typical privileged authentication events. In summary, alerts should be actual events. Non-actionable alerts should be kept in the log to flesh out context or provide additional data points.

This is done by prioritizing high risk issues and combinations of events that point to a likely issue. Typically, SIEM platforms allow for the combination of alerts from multiple contributing applications as well as real time analytics of events. Event Analytics should consider events individually, as part of a combination of events within the environment, and over a configurable period of time.

The types of alerts that may be of interest to a financial institution fall into three categories: those that indicate use of privileged access; those that indicate misuse of privileged access; those that indicate changes to elements of privileged access.

- **Use of privileged access** is not interesting in and by itself. Anomalous use, however, is very interesting. Anomalous use can be based upon attributes (privileged authentications originating from unusual place, such as off-network or time of day); or failures coming from the same source but across multiple privileged userids. For the most critical applications, such as those with money transfer or DBAs accessing a PII data store, any use of privileged access might be something to investigate. Create alerts to meet the risk appetite of your institution.

- **Misuse of privileged access** is interesting. Misuse can arise from a privileged user perform an operation that is not aligned with a business need (e.g., using the account to unmask data and download); or performing an embargoed operation (e.g., creating another privileged account; or turning off logging of events associated with the use of a privileged userid). Those events should be raised to the appropriate personnel for follow-up.
- **Changes to elements of privileged access**. A fully automated Identity Provisioning application adding a userid to an Active Directory group – following all organization standards – will not give rise to an actionable alert (although it is a logged function). However, privileged users making direct changes to Active Directory groups that provide entitlements is an actionable alert. Similarly, changes to environments, roles, and logging should raise actionable alerts.

MANAGING PRIVILEGED ACCESS

Once the access has been defined, the next question is how do you track and manage it? Privileged access can be managed using the same framework for other entitlements. Ideally there will be a set of attributes that could indicate a risk/ranking level and can drive other processes. For example, within Active Directory, an attribute could be defined to track if it's related to a specific type of elevated access, and, depending on the value, drive any requirements for additional authentication, approvals, review frequency, etc. If this attribute is then pulled into the IAM platform as part of the defined schema, additional rules/workflows can be defined to better manage access (i.e., limit access to specific departments, only allow certain types of accounts to be added, dynamically create access groups, based on value, etc.). Where attributes are not available, naming conventions can also be leveraged.

JUST-IN-TIME PRIVILEGE ACCESS

Why Just-in-Time?

Drastically reduce risk, with a goal toward eliminating persistent privileged access by using Just-in-Time approaches. This will ensure that privileged access is only granted when a valid reason for the users exists in a limited time window, rather than available at will all the time.

Approach to implement Just-in-Time

The first goal of a Just-in-Time PAM approach is better IT security by removing standing privileges, which reduces the attack surface for privilege abuse. Only the right access for only the right person to only the right system is granted for only a limited time.

Just-in-Time Group Membership approach under Active Directory.

A defined access model to perform administrative tasks on endpoints using Active Directory security groups would need to be in place (i.e. AD groups that would allow members to be an administrator on endpoint). When users need privileged access to perform tasks, they would request access through a workflow system for the target endpoint, and an approver (someone who has the appropriate decision rights to determine the necessity of access) receives the request and approves or denies the access for a defined period of time. Access is granted to the requestor's account, meaning they are granted privileged access to perform administrative tasks. The elevation happens by adding the account to specific privileged security groups. Once the task is complete (or the time limit for access has expired), the requester's account will be removed from the privileged security groups, and the requester will no longer have privileged access until another request is submitted.

Access Control – Privileged Access Management

Access control and an established model are critical to a mature privileged access management program. General access models provide the overall standard for controlling access across the enterprise and need to harmonize with model used to control privileged access. There are several commonly accepted access models, and each has its advantages and disadvantages.

Privileged Access Considerations for Non-Human Accounts

Don't ignore IoT and Compute and Storage identities like containers, EKS, lambda functions, S3 buckets, etc.

INTERNET OF THINGS (IOTS)

The Internet of Things (IoTs) present an interesting challenge to cyber security. There are numerous controls financial institutions utilize to secure them (e.g., isolated networks; frequent firmware upgrades; continual scans to detect new / unauthorized devices) that are the subject of a different discussion. There are a few aspects that might come within the purview of privileged access management.

- **Passwords of the Devices themselves** – consider leveraging the rotation capabilities of a PAM to frequently change the password used to access each device. While the refrigerator in the breakroom does not seem like a particularly threatening device, a bad actor accessing that device can turn it into an actor on a DDOS attack. A more harmful attack would be accessing a network enabled copier/printer in the HR department to forward all copies/print jobs to another device for later inspection and exfiltration. Similarly, consider the password vaulting capabilities that will enable longer and complex passwords.

- **Treat all Robotic Process Automation (RPA) and Internet of Things (IOT) accounts as privileged.** While the previous examples considered IoTs that were devices, an IoT can be a Robotic Process Automation bot (RPA). Here it is important to leverage similar controls for other privileged accounts. RPAs are designed to operate in the background, unsupervised, and perform countless repetitive transactions. It's important that the underlying userids are not used for some other purpose.

- **Machine Identities of IoTs are critical**. Be certain to keep track of the authorized IoTs (the machines they run from if they are an RPA). Similar to the accounts themselves, it is important to be able to detect, inventory, and alert on the discovery of a new (or unexpected) IoT device. Machine Identity management is not merely keeping track of the serial number. Consider other techniques such as certificates (and they need to be managed).

- **Consider IoT behavior**. IoTs, like human accounts, likely have a particular usage pattern (amount of data, what it connects to, what connects to it, time of day, etc). Create alerts for anomalous behavior.

APPLICATION PROGRAMMING INTERFACES (APIS)

Application Programming Interfaces (APIs) enable access to on premise and cloud based (IaaS and SaaS) enterprises resources and need to be properly secured. This section considers privileged manager aspects of APIs. APIs are becoming more plentiful in financial institution (FI) environments as FIs automate infrastructure deployment and code deployment. Other papers will cover security controls such as protecting data at rest, data in transit, and IP/Geo-location. With APIs being the conduit to access both data and configurations, it is imperative that we protect them as we protect other privileged accounts.

- Scripts that embed credentials (even those that start APIs) need to be protected. Consider leveraging PAM tools that allow a program to authenticate to the PAM vault and **programmatically retrieve credentials** that are then passed to the called object. These credentials can be secured in the PAM platform (rotation, length, complexity) similar to other privileged accounts.

- APIs may leverage an authentication server that provide **access tokens** (grants authorization) that are passed to the client application. Rather than sending a set of static credentials to the enterprise resource, a request is sent to a service with some type of credentials. A successful authentication results in a bit more back and forth between the API and the identity provider (IdP) – with an access token eventually being returned for signing each subsequent API request to the enterprise resource.

Consider **API gateways** as a mechanism to implement behavior controls such as throttling, geo-location, inspecting data for malicious payloads, and generally looking for malicious behavior.

Privileged Access Monitoring and Alerting

Privileged Access is essential to the operation of every organization. These credentials provide access to the most sensitive and vulnerable (operational, data leakage, security and alerting systems, etc.) in the technology ecosystem. As a financial institution, regulators scrutinize privileged use – and the potential for misuse by insiders and bad actors.

Bad Actors and Insiders (malicious and accidental) leverage privileged user credentials when accessing enterprise resources such as systems and data. Monitoring and alerting -- two sides of the same coin -- are critical activities for an organization to undertake in order to protect enterprise resources. While an effective Access Review process is a critical piece of Privileged Access Monitoring and Alerting (PAMA), this section will focus on real time activities – those activities that enable an organization to detect, protect and provide data for respond to privileged credential misuse will be the focus.

Monitoring informs an organization on how privileged accounts are being used; *alerting* informs an organization when privileged accounts are being misused. In the NIST framework, alerting is related to Detection as an organization enables the timely discovery of an event; and monitoring is aligned with Identity as it assists in creating an organization understanding of risks within the environment.

Monitoring

Typically, in technology, we monitor people, process and technology in order to understand performance. With people, it is generally about productivity. With process and technology, it is adherence to some type of standards – architectural, operational, security, etc. However, with privileged access, organizations monitor to understand risk footprint and to detect misuse of privileges.

Monitoring - Understanding the Size Risk Footprint

The privileged access risk footprint can be measured across multiple vectors. For example, the number of users with elevated access to systems; how tightly specific roles/access levels are assigned to userids; and how broadly an individual's privileged access reach across an organization.

RISK FOOTPRINT - TYPES OF ACCOUNTS IN YOUR ENVIRONMENT THAT NEED TO BE MONITORED

Organizations provision several types of privileged user accounts for both human and non-human actors. With regard to human actors, the size of the risk footprint is related to the scope (breadth or depth) of access.

For example, there are privileged accounts that provide:

- **Administrative** access to individual workstations or allow the manipulation of an individual application's entitlements to other users
- **Domain** level Administrative access that can allow wholesale changes to managed systems or the structure of the domain in and by itself.
- **Local** administrative access that exists on a machine and is more laborious for central IT to track

Regarding non-humans (or Internet of Things (IoTs)), there are several types of administrative accounts that affect the size of the risk footprint. For example:

- **Service** accounts used by an application to interact with other services (e.g., file transfer) or the operating system (e.g., application or component startup).
- **Accounts embedded in application scripts** (with their credentials) that are used for intra-application processes such as data loads. For these accounts, attention should be given to how accessible the credentials are to others with access to the server / application.

Lastly, an organization should consider the definition of privileged access. Initially, most organizations create an IT-centric definition – those userids that can manipulate the configuration of an enterprise resource, change logging levels, and add/delete objects. An organization should consider expanding the definition to consider business processes such as:

- Those userids that can initiate or approve **high risk actions** such as money movement or affect the credentials of client or employee access
- Those userids that can access **sensitive data**

An additional control to consider for privileged accounts is whether they can be used from another enterprise resource or must be leverage directly from the enterprise resource. This topic is out of scope.

Monitoring - What to Monitor

Standards of monitoring should be proportionate to the amount of risk as measured across the depth and breadth of access an account provides the human and non-human actor. Consider monitoring at the following points of control:

What to Monitor – the Inventory and Provisioning Process

For all accounts consider periodically validating the basics regarding **inventory**.

- Is the **inventory still comprehensive** – has the inventory tool been deployed to all enterprise resources in your CMDB?
- **Detect unauthorized accounts**. Perform a periodic reconciliation report to compare changes in provisioned accounts between two periods and compare with provisioning reports. Does the delta between the reports match the provisioning report?
- **Validate who/what has access to each resource**: Review privileged accounts that have access, by account type, by enterprise resource? Does the resource owner agree that this set of accounts should have privileged access to the enterprise resource?

- Are **Provisioning steps being** followed? Are all necessary approvals in place? Is anomaly reporting and follow-up in place?
- Are de-provisioning steps occurring in a timely manner?
- Are there robust, and followed, joiners / movers / leavers processes in place for privileged accounts?
- Is there a sufficient privileged access management education program the covers, among other things, acceptable use? Do we monitor that each employee takes that course and periodic refreshers?
- Is there a way to leverage security champions to augment the security training program?

What to Monitor – Account Usage

Effective monitoring of privileged account use assists organizations to manage risk and protect enterprise resources and data.

Organizations should create a privileged access review process, or merely reporting, that measures:

- How many people have access to which enterprise resources?
- How many people have privileged accounts of each type discussed above?
- When was each privileged account last used and should dormant accounts be de-provisioned?
- Which accounts have been recently used?
- Are privileged accounts being used during expected times of day?
- Is anomaly reporting and follow-up in place?

Organizations should create reports and review how privileged accounts are being used:

- Have recently used userids been used for a legitimate business purpose?
- Could a standard account have been used instead?
- Which privileges/rights is an account using and can they be trimmed to better comport with the least privilege principle?

What to Monitor – Controls related to how a privileged account was accessed

The risk footprint is affected by the process for obtaining the credential. There may be different use cases or controls that we apply to these accounts such as:

- Organizations that **conflate primary and privileged accounts** must rely on detective controls. It is bad practice for a **user's primary business account** to have privileged access too. Conflating the two accounts increases the chances of negative impact from an inadvertent action or from a bad actor obtaining the user's credentials. From a least privilege perspective, a user's standard account should be used for everyday tasks; and the privileged account must be used only when required by the tasks at hand. Please refer to the Alerting section of this paper.
- Persistent Access with **no vaulting** increases the risk footprint as there is no preventative control and misuse requires a weaker, detective control. Please refer to the Alerting section of this paper.

- For all accounts, establish account **usage reporting and monitoring**. Consider looking for these items (with example reports to detect):
    - **Validate Activity:** Daily usage reports that are reviewed by the employee's manager or change control organization.
    - Review privileged logins, by account type, **by enterprise resource**? Does the resource owner agree that this set of accounts should have privileged access to the enterprise resource?
    - Usage over time showing that could indicate a **change in patterns** or frequency of use

- Look for anomalous or changing behavior
- Look for changes in entitlements
- Look for **inactive accounts** and deactivate
- While not specifically a privileged account report, work with the infrastructure team to validate that all environments that use terminal servers or **jump boxes** as a control have them in place for each enterprise that should.
- Are privileged accounts being access from on or off-network?

- **Vaulted Accounts** in privileged access management (PAM) tools. Additionally, the PAM tools can impose the need for a higher level of identity validation by imposing a 2FA control. Further, the PAM tools can impose controls that impose periodic (or mandatory after each use) password rotation which lessens the risk that an administrator will bypass the vault and go directly to the enterprise resource. Consider the usage reports identified above, and consider adding:
  - **Failures to Retrieve Credentials** – report on user accounts that have been unsuccessfully accessed (including over time intervals) because of an invalid password or invalid 2FA. Are specific accounts under attack? Are specific addresses or machines referenced in repeated failures across a wide range of target accounts?
  - **Compliance with Configuration standards** – report on whether any account fails to comport with an enterprise standard such as mandatory password rotation (or duration), requirement for 2FA, and session recording enabled?
  - Usage over time that could indicate a **change in patterns** or frequency of use

- **Persistent or non-persistent** availability to a user. Persistent accounts allow an administrator to utilize credentials, even through a vault, whenever they see fit. Contrast that to a non-persistent account that requires an intervening approval or mouse click from a manager or network operations center. For vaults that support non-persistent access:
  - Are non-persistent accounts **authorized for periods (durations)** within the organization's standards? If the standard allows privileged accounts to be checked out for up to 48 hours, have any exceeded that period? Is the same account being checked out in successive periods, thereby flaunting the duration control?
  - Is the **appropriate approval** chain in place?

- **Break Glass** accounts that are used in the most extreme events and typically have chain of custody around their use. Consider the usage and vaulting reports identified above, and consider adding:
  - Are all **controls** (e.g., sealed envelope with password still sealed) **in place**?
  - Are the additional **controls practical**? For example, if a 2FA item is stored in a locked safe, is it within the reach of a 24x7 staff or does it require employees to first arrive on site? Are conditional controls in place that would frustrate the use of the account in an emergency? Consider the problem of having IP-based controls on a break-glass cloud access account if your primary network has an outage and you cannot originate the authentication from within your IP space.
  - Have we **tested** timely and controlled access to the account (or 2FA)?
  - Reports on **who has access** to these credentials – and is the list sufficiently small?

What to Monitor – Account Compliance with Standards

Each organization should have a set of Identity/Privileged Account standards. Each standard should have a method to measure compliance, and periodic report should be in place to measure compliance rates. Additionally, there are likely password requirements such as complexity, change frequency, and length, as well as remote versus on-prem resource access standards.

Compliance with Configuration Standards – report on whether any account fails to comport with an enterprise standard such as mandatory password rotation (or duration), requirement for 2FA, and whether session recording is enabled.

Alerting

Alerting is a detective control that helps organizations respond to potential events. Preventative controls help organization reduce risk. As with monitoring, it is important to ensure that all the alerts are considered by: capturing alerts from relevant sources; capturing relevant events from those sources; and understanding how each event is raised and must be dispositioned. The concern regarding whether the PAM process and capability is alerting as it should requires that all alerts are being captured from each source.

Alerting – Capturing the Alerts: Sources

Before considering the type of events that should be captured, the organization should consider that all sources (by type) are covered, as well as the entire fleet of devices. Typically, all security-related systems will send alerts to a Security Information and Event Management (SIEM) platform. The systems that will contribute alerts include:

- Identity stores such as Active Directory
- Applications that support federation such as SSO
- Privileged Access Management Vaults
- Segmentation policy engines
- Applications performing AuthN or AuthZ functions
- VPN and VDI access points
- 2FA applications
- Active Directory - Change Auditor applications
- Perimeter security tools such as IDS, IPS, firewall, etc.

The SIEM will provide central storage (separate from the contributing enterprise resource) for all security events. The SIEM will have immutable storage to reduce the likelihood that logs are manipulated by a bad actor. Additional storage must be sufficient to allow for analytics to be performed over days or even a few months.

As with Monitoring, the organization should have proper processes in place to make certain each of the above types of systems is configured to send alerts to the SIEM, and that there is a process in place to verify that this continues to be the case over time.

Alerting – Capturing the Alerts: Types

In general, organizations need to capture all Authentication and Authorization events in a central data store. For example, capture the following alerts:

- Account lockouts
- Users added to privileged groups
- Security group modifications
- Successful and failed AuthN and AuthZ events
- Creation or modification of privileged user accounts.
- Privilege Elevation

The SIEM should maintain a sufficient history of events to allow for analysis to be performed over a sufficiently long window – this is especially useful for anomaly detection analysis.

Alerting – Capturing the Alerts: Filtering and Summarization

To prevent teams from overlooking important alerts, a risk-based filter should be placed on viewing logged alerts. Organizations should be more concerned with anomalous privileged authentication events than typical privileged authentication events. In summary, alerts should be actual events. Non-actionable alerts should be kept in the log to flesh out context or provide additional data points.

Alerting – How are Alerts actioned?

Alerts that are raised by a SIEM can be screened by the Security Operations Center to determine which require further action. Over time, this screening process needs to be fine-tuned with non-actionable alerts being better identified and not being brought to the attention of a SOC analyst. Events that don't require human attention can leverage the organization's Security Orchestration, Automation and Response (SOAR) platform. The SOAR platform can gather relevant information and take automated action based on programmed responses or ML/AI learning.

Once an alert is deemed actionable the SIEM or SOAR should open a ticket in the ITSM application to ensure that the alert, action, and resolution can be logged and tracked.

Cloud Considerations

The principles of privileged access management transcend the deployment model – on premise, IaaS, and cloud. However, the implementation in a cloud provider will likely impact the PAM tools selection. Consider the following:

- Most tools **cannot merely lift and shift** into the cloud. While some tools, such as account discovery tools on a local machine (Windows or Linux) can shift to the cloud quite easily, it becomes impossible if it's merely a container such as a Lambda function or a Java instance such as TC server (e.g., Cloud Foundry container).
- **Passwordless Authentication** - Many privileged access models in the cloud utilize referrer tokens instead of authentication where authentication starts on premise and ends in a cloud facility. Consider whether keys and certificates are part of privileged access management.
- **Dynamic Environments** – where containers and machines are subject to auto-scaling and ephemeral execution environments. This complicates mapping privileged account use to a specific instance.
- **Account discovery** is difficult in a highly scalable and ephemeral environment. Rather than periodic or episodic account discovery, the cloud requires a continuous approach to PAM lifecycle management.
- Difficulty identifying all types of privileged access in the cloud.

The good news is that the major cloud providers are creating robust IAM and PAM tools as well as auditable IAM processes to address these difficulties. There seems to be a convergence of PAM and Identity Governance services (i.e., lifecycle, provisioning). This ensures PAM processes include risk awareness and governance requirements

Measuring Program Success

How do you know your program is successful?
Measuring security is hard[10]. Measuring program success is simpler if measures of success can be identified in advance. Here are some questions to generate ideas for measuring the success of your IAM program:

- Comparison to industry baselines or maturity assessments [NIST, ISACA COBIT]
- Metrics that indicate a level of control, like percentage of timely revocations or percentage of reviews completed on time
- Adoption of strategic platforms which guarantee a level of control. For example, privileged accounts that use a password checkout vault or keystroke logging
- Compliance with laws, rules, regulations and guidelines: how many apply to your business? Do you have any internally identified gaps? Do you have any regulator defined gaps?
- Are your processes operating within a level of tolerance that is acceptable? Examples are measurement of deprovisioning process. Is the percentage of revocations completed on time within an acceptable risk tolerance?
- Are internal policies and standards relevant to your industry and level of maturity? Too aspirational or too undemanding?
- How are policy and standard adherence measured? Is compliance also considered as a metric with tolerance thresholds? How sever are the consequences? Are they uniformly and systematically applied, no exceptions?
- Is there a prevalent security culture, identified through shared values by managers and employees; organizational norms; environmental queues; and established routines?
- Are users well-indoctrinated to security issues and risks? Is the training periodic, continuous, and through multiple modes of education?
- How expensive is your IAM program in relation to your business and the risks inherent in it? ROI or cost effectiveness should be considered but is not a primary objective.
- How many IAM incidents are reported each year? How many near misses? This could include social engineering events, phish-clicking, or insider threats.

Innovation

The complexity of Identity and Access Management requires constant readjustment and forward ways of thinking to evolve access governance processes. Innovation, geared to reduce elements of manual IAM governance oversight to combat increasing requirements placed on business lines, is necessary to assist in meeting regulatory demands with efficiency and efficacy.

Thorough analysis for implementation of new tools and analytics should be considered against return on investment (ROI). Risk reduction, cost optimization for manual resource spend, increased visibility into access clusters to introduce or enhance role-based access controls (RBAC), enhanced metrics, as examples, can all propose value to the organization.

Current vendor offerings in the market offer a variety of innovative solutions for standard and next-generation IGA products and suites. Selection should prioritize must-haves versus nice-to-haves for which IGA capabilities are required.

Artificial intelligence (AI) and machine learning (ML) are now included within some vendor products allowing for a more robust and rigorous application of governance to be applied to programs. AI/ML capabilities assist with capturing a wide range of information for increased visibility into identity, access and review behaviors. Organizations can view current and historical data for users in their systems and query identity data to validate that specific controls are operating as expected. The complexity of reviewing extensive identity data is reduced as AI/ML highlights areas where attention is required. Recommended actions can assist IGA programs in taking proactive action to remain aligned with regulatory and policy requirements. Key features of IGA AI/ML may include, but are not limited to:

- Peer group and access cluster algorithms
  - IAM policy recommender for ABAC and RBAC
  - Could be used for role-mining and attribute recommendations
- Access history
  - Granular analysis at the identity level
- Data explorer
  - Ad hoc data interrogation from historical data
  - Could be used to generate quality-based metrics
- Access review recommender
- SOD policy enhancement
- Risk scoring
  - Risk-based workflows
  - Risk-based access review programs

Conclusion

Recent experience with threat actors has provided a more complete understanding of the controls that a financial institution needs to implement to limit unauthorized access to data or applications. Guidance from NIST, as well as experienced cybersecurity practitioners in the financial services industry, suggest that better control of identities, technical capabilities in the identity platform, and closed loop processes in identity lifecycle (provisioning, deprovisioning, and access reviews) play a critical role in reducing bad actor risk. Further, the suggestions made in this document positions the institution to evolve into a Zero Trust / Identity Centric model. This is valuable as the federal government is adopting this type of model (and regulators may encourage FIs follow suit).

More complete control of identity lifecycle events, as well as finer-grained access to enterprise resources, will reduce risk to data, protect applications, and ensure the availability or integrity of services. In short, strengthening IAM will reduce reputational and financial risk to the institution and U.S. banking system – and better protect our clients' private data.

[10] Pfleeger, S., & Cunningham, R. (2010). Why measuring security is hard. IEEE Security Privacy, 8(4), 46–54. https://doi.org/10.1109/MSP.2010.60

Appendix

Table Apx-1

Application	A software program that performs a particular function for users. Applications may be hosted on-prem or within a Cloud solution.
Birthright Access	Access provisioned automatically based on user attributes
Directory Service	A centralized repository for storing and managing identities and security objects (used to provide access to a wide range of Platforms, Applications, or services. e.g., Cloud Identity Providers, Lightweight Directory Access Protocol (LDAP), Active Directory (AD), etc.)
Dormant Account	An active account that has no recent login activity
Dynamic Group	A Group where membership changes depending on defined membership rules
Emergency (break-glass) account	A highly Privileged account that is not assigned to an individual user within the organization and limited to emergency or "break glass"' scenarios where normal administrative accounts cannot be used
Entitlements	Also referred to as privileges, authorizations, policies, or permissions, entitlements may be granted individually or combined using Roles or Groups
Entitlement Management	An Identity Governance and Administration feature that enables organizations to manage the identity and access lifecycle at scale by automating access request workflows, access assignments, reviews, and access/account expirations.
Environment Variable	A dynamic-named value often used in scripting and automation that could change values used in Identity metadata
Environment	Organized tiers within the cloud service that represent the stages of software deployment (e.g., Dev, Test Prod)
Generic Account	Also referred to as a shared account. This is an account used for system access by multiple individuals who share a common area of responsibility.
Guest & External Account	Users outside of the organization who require some level of access to technology delivered via cloud only technology (e.g., external sharing of content in O365, such as SharePoint, Teams, etc.)
Identity	The digital representation of a person or service that is recognized by a system
Identity Governance and Administration	An enterprise framework in support of automating the creation, management, and certification of user accounts, roles, and access rights for users across an organization
Internal User Account	A standard or Primary User Account that identifies a user within the organization. This account is typically sourced from a Human Resources Management System, such as Workday or PeopleSoft.

Just-in-time (JIT) Access	A model in which users receive temporary permissions to perform Privileged tasks, which prevents malicious or unauthorized users from gaining access after the permissions have expired. Access is granted only when users need it.
Non-human Account	A type of account which is utilized solely by non-human processes/services for the purposes of establishing an authenticated communication channel. Also referred to as service or bot accounts. These accounts are frequently more Privileged than human User Accounts and thus must be treated with a high degree of caution
Orphan Account	An active account that is not attached to a user profile or owner.
Platform	A suite of cloud-computing services (e.g., Azure, AWS, GCP)
Primary User Account	An account used for regular daily interactive login to perform standard job activities. Primary accounts typically include Internal User Accounts as well as External/Guest User Accounts
Principle of Least Privilege Access	Users are granted the minimum privileges needed to accomplish the tasks they are authorized to perform
Privileged	Special access rights that allow a user elevated access over a system or Application giving them the ability to modify security settings, add/remove users, modify system configurations, access sensitive data, or gain full control of a system or Application
Resource	An entity managed by the cloud Platform, such as a web Application, virtual machine, database, etc.
Role	A collection of Entitlements applied to an Identity, Application, or Platform such as Azure Access Packages or AWS IAM Roles and IAM Policies.
Secondary User Account	An account in addition to a user's Primary Account. This includes test accounts, Privileged accounts, External cloud only accounts (AWS, GCP, Azure Developer) or defined lower Environment accounts (development or non-production Environments)
Security Group	A directory object that contains user account(s) and/or group objects that are used to assign permissions related to granular functions within an Application, Resource or Platform. This includes AWS IAM Groups and Google Groups
Self-Approved	Any approval to a Group or Role that does not require a manager or owner approval to support access
User Attribute	Information which determines the properties of a user

Table APX-2: Naming Convention example

Identifier	Description	Examples
Directory	The centralized repository used for storing and managing identities within a cloud Platform.	AD – Active Directory on Prem Azure – Azure Cloud Directory AWS – Amazon Cloud Directory GCP – Google Cloud Directory
GroupType	The type of group being created	SECURITY – Security Group DYN – Dynamic Group
ResourceName	The entity managed by the cloud Platform	WebApp – Web Application VM - Virtual machine DB – Database
Environment	The deployment level the Resource is in	Dev – Development Environment Test – Test Environment Prod – Production Environment
Permission	The type of permission being granted	*READ – Read-only access* *CONTRIBUTOR – Read-write access* *ADM – Administrator access*
Appendix	Used as needed; to allow flexibility in the naming standard and provides additional clarification where needed. If a group is considered privileged it can be noted here.	