

RECOMMENDED BEST PRACTICES FOR ADMINISTRATORS

IDENTITY AND ACCESS MANAGEMENT



DISCLAIMER

DISCLAIMER OF ENDORSEMENT

This document was written for general informational purposes only. It is intended to apply to a variety of factual circumstances and industry stakeholders, and the information provided herein is advisory in nature. The guidance in this document is provided “as is”: Once published, the information within may not constitute the most up-to-date guidance or technical information. Accordingly, the document does not, and is not intended to, constitute compliance or legal advice. Readers should confer with their respective advisors and subject matter experts to obtain advice based on their individual circumstances. In no event shall the United States Government be liable for any damages arising in any way out of the use of or reliance on this guidance.

Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement, recommendation, or favoring by the United States Government, and this guidance shall not be used for advertising or product endorsement purposes. All trademarks are the property of their respective owners.

PURPOSE

The National Security Agency (NSA) and the Cybersecurity Infrastructure Security Agency (CISA) developed this document in furtherance of their respective cybersecurity missions, including their responsibilities to develop and issue cybersecurity recommendations and mitigations. This information may be shared broadly to reach all appropriate stakeholders.

CONTACT

Client Requirements/Inquiries: Enduring Security Framework nsaesf@cyber.nsa.gov.

Media Inquiries / Press Desk:

- NSA Media Relations, 443-634-0721, MediaRelations@nsa.gov
- CISA Media Relations, 703-235-2010, CISAMedia@cisa.dhs.gov

Table of Contents

- Introduction..... 1
- Scope 2
- The Threat Landscape..... 2
- IAM Threat Mitigation Techniques 4
- Identity Governance..... 4
 - What it Does..... 4
 - Why It Matters 5
- Environmental Hardening 6
 - What it Does..... 6
 - Why it Matters 7
 - Setting the Stage for Implementation 7
 - Implementing Best Practice 7
 - Actions to Take Now 9
 - Summary 10
- Identity Federation and Single Sign-On..... 10
 - What it Does..... 10
 - Why it Matters 10
 - Factors to consider when selecting an SSO solution 11
 - Implementing Best Practices..... 13
 - Actions to Take Now 13
 - Summary 13
- Multi-Factor Authentication..... 13
 - What It Does 15
 - Why MFA Matters..... 17
 - Preparation for Implementing MFA..... 18
 - Catalog User Populations, Device Types, and Use Cases 18
 - Evaluate Assurance Requirements 19
 - Evaluate Privacy and Operational Considerations 19
 - Implementing MFA..... 20
 - Actions to Take Now 21
 - Summary 21

IAM Auditing and Monitoring..... 22

 What it Does..... 22

 Why it Matters 22

 Preparation for Implementing Best Practice 23

 Actions to Take Now 24

 Summary 25

Conclusion..... 25

Appendix I: Actions to Take Now Checklist..... 26

Introduction

Identity and access management (IAM) is a framework of business processes, policies, and technologies that facilitate the management of digital identities to ensure that users only gain access to data when they have the appropriate credentials. Beyond the physical users, service and system accounts are also in scope for IAM and critical for IAM administrators to manage within their organizations. Inventorying, auditing, and tracking all of these identities and their access is imperative to ensure that proper IAM, including permissions and active status, is executed on a regular basis. Managing the growing complexities of digital identities can be daunting especially with industry's push toward cloud and hybrid computing environments; however, the need for IAM is more important today than ever. In recent years, we have seen various nation state-led cyber operations successfully access protected data by targeting the trust established within networks or by exploiting vulnerabilities in IAM products and/or IAM implementations. Specifically, the critical infrastructure within the U.S. is an attractive target for the adversaries. In fact, according to the 2022 Verizon Data Breach Investigation Report, 80% of web applications attacks leveraged stolen credentials, a technique used by both basic cyber criminals and nation-state bad actors. Additionally, excluding breaches based on user error and insider misuse, 40% of breaches involved stolen credentials and nearly 20% involved phishing. Recent and notable attacks include:

- In 2021, compromised credentials were used to attack and shut down the Colonial national gas pipeline in the U.S.¹
- In another 2021 cyberattack, an unknown attacker manipulated computer systems in a Florida water treatment plant to increase the concentration of sodium hydroxide in the water supply by a factor of 100.²
- In 2022, another attack targeted a water treatment plant in South Staffordshire, U.K.³

As such, the critical infrastructure organizations have a particular responsibility to implement, maintain, and monitor secure IAM solutions and processes to protect not only their own business functions and information but also the organizations and individuals with whom they interact. It is important to keep in mind that IAM systems implement credential management, authentication, and authorization functions that are foundational to security and also very complex and subject to vulnerabilities if not implemented correctly. Like any kind of software, IAM solutions are subject to software vulnerabilities and must be patched, updated, and managed. A vulnerable IAM solutions can facilitate access to multiple systems and data across the organization. Therefore, securing IAM infrastructure is critical. Ultimately, the goal is that organizations proactively take the

¹ <https://www.bloomberg.com/news/articles/2021-06-04/hackers-breached-colonial-pipeline-using-compromised-password>.

² <https://arstechnica.com/information-technology/2021/02/breached-water-plant-employees-used-the-same-teamviewer-password-and-no-firewall/>.

³ <https://www.zdnet.com/article/confused-cyber-criminals-have-hacked-a-water-company-in-a-bizarre-case-of-mistaken-identity/>.

appropriate action to protect against an attack rather than be in the position of deploying fundamental IAM capabilities far too late.

To address the risk to a wide range of critical public and private sector networks, the Enduring Security Framework (ESF) hosted a working panel staffed by government and industry subject matter experts tasked with assessing the challenges and threats to IAM and identifying recommendations on how to mitigate these risks. While the working group recognizes the need for a broad, layered approach to network defense, this guidance is focused on the aspects of IAM identified as critical in addressing the threats laid out in this paper.

Scope

This paper sets forth the IAM best practices for administrators to implement to address threats that are highly likely, highly impactful, or both. Furthermore, it identifies mitigation areas most effective in reducing the impacts of these threats to IAM.

This paper focuses on identifying mitigations for the following techniques frequently used by bad actors:

- Creating new accounts to maintain persistence.
- Assuming control of accounts of former employees which were not suspended upon employee termination.
- Exploiting vulnerabilities to forge authentication assertions (e.g. Kerberos tickets, Security Assertion Markup Language (SAML) assertions, OAuth2).
- Utilizing or creating alternative access points to systems.
- Exploiting or utilizing users with legitimate access.
- Compromising passwords through a variety of tactics (e.g. phishing, multi-factor authentication (MFA) bypass, credential stuffing, password spraying, social engineering, brute force).
- Gaining system access and exploiting stored credentials.
- Exploiting default passwords in built-in or system accounts, exploiting active attacks to downgrade, and exploiting deprecated encryption, or plain-text protocols to access credentials.

The Threat Landscape

Organizations are subject to attacks from a broad range of threat sources including nation-states, terrorist groups, organized crime, hacktivists, and individuals looking to harm or embarrass an organization. Additionally, organizations are subject to attacks where a trusted user is the source of the compromise (e.g., insider threat). The spectrum of threat sources varies wildly in capabilities, motivations, and methods. For example, nation-state actors have significant resources, and can establish long-term plans to gain access to critical resources. They can also use indirect methods such as exploiting the supply chain.

Exploiting known IAM vulnerabilities could allow a bad actor the same access to resources as legitimate users by mimicking legitimate activity which complicates detection of the bad actor. This provides the bad actor more time to gain access to resources and elevate privileges to gain persistent access.

For example, a recent CISA Alert (AA21-321A)⁴ showed that Iranian government-sponsored advanced persistent threat (APT) actors are actively targeting a broad range of victims across multiple U.S. critical infrastructure sectors by exploiting IAM vulnerabilities to compromise credentials, escalate privileges, and establish new user accounts on domain controllers, servers, workstations, and in directories responsible for authenticating and authorizing users and devices. These actors could leverage this access for follow-on operations, such as data exfiltration or encryption, ransomware, and extortion.

Additionally, exploitation of Single Sign-On (SSO) technology (a component of IAM) is becoming a more prevalent attack vector. Bad actors attempt to exploit the SSO functions with hopes of easily gaining access to protected resources throughout the system and/or organization. Several examples that show the impact of SSO compromise include:

- In September 2021, Palo Alto Networks revealed bad actors exploiting a vulnerability in Zoho's ManageEngine ADSelfService Plus SSO solution. The bad actors were observed deploying backdoor and credential stealing tools to maintain access to the victim's networks including critical infrastructure entities.⁵
- The SolarWinds compromise highlighted the risk of SSO exploitation. The NSA and others characterized the "Golden SAML," Active Directory Federation Services bypass technique, as shown in Figure 1, which gave bad actors access to all of the enterprise's Active Directory authentication.⁶

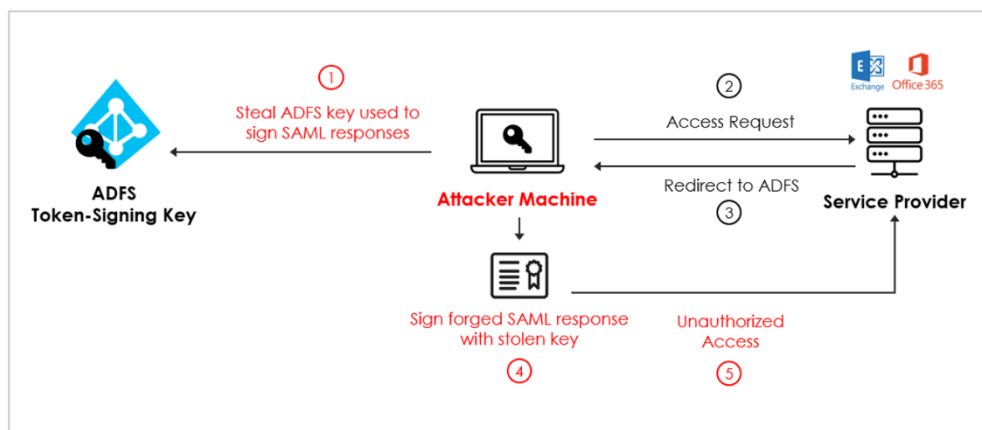


Figure 1 Depiction of "Golden SAML" Attack Process.⁷

⁴ <https://www.cisa.gov/uscert/ncas/alerts/aa21-321a>.

⁵ <https://unit42.paloaltonetworks.com/manageengine-godzilla-nglite-kdc.sponge/>.

⁶ <https://www.darkreading.com/attacks-breaches/solarwinds-campaign-focuses-attention-on-golden-saml-attack-vector>.

⁷ <https://blog.sygnia.co/detection-and-hunting-of-golden-saml-attack>.

Defending against this broad spectrum of attacks requires a comprehensive IAM solution, with operational awareness of the environment to detect anomalies and attribute anomalous activity to adversary exploits.

IAM Threat Mitigation Techniques

The best practices and mitigations discussed in this paper provide tactics that help to counter threats to IAM through deterrence, prevention, detection, damage limitation, and response. Specifically, this paper identifies best practices relating to:

- **Identity Governance** - policy-based centralized orchestration of user identity management and access control and helps support enterprise IT security and regulatory compliance;
- **Environmental Hardening** - makes it harder for a bad actor to be successful in an attack;
- **Identity Federation and Single Sign-On** – Identity federation across organizations addresses interoperability and partnership needs centrally. SSO allows centralized management of authentication and access thereby enabling better threat detection and response options;
- **Multi-Factor Authentication** - uses more than one factor in the authentication process which makes it harder for a bad actor to gain access;
- **IAM Monitoring and Auditing** - defines acceptable and expected behavior and then generates, collects, and analyzes logs to provide the best means to detect suspicious activity.

Identity Governance

Identity governance is the process by which an organization centralizes orchestration of its user and service accounts management in accordance with their policies. Identity governance provides organizations with better visibility to identities and access privileges, along with better controls to detect and prevent inappropriate access. It is comprised of a set of processes and policies that cover the segregation of duties, role management, logging, access review, analytics, and reporting.

What it Does

Identity governance solutions can manage the entire identity and access lifecycle for an organization's workforce. The most critical lifecycle events are often referred to as "Join, Move, and Leave" (JML) events:

- **Join** – when a new employee or contractor joins the organization, the identity governance solution can collect biographical, position-related, and credential data (such as professional certifications or clearances) from recruiting, human capital management, and personnel security systems to build out an identity record for the individual. Identity governance systems can use this data to automatically create

accounts in directories and applications with entitlements based on the collected data.

- **Move** – when an individual’s role in the organization changes, an identity governance system can automate the granting of additional entitlements needed for their new role as well as the removal of entitlements that are no longer needed. Without adequate management of Move events, long-term users tend to accumulate privileges as their roles change, increasing the potential impact of insider abuse or account takeover.
- **Leave** – when users separate from an organization through retirement, termination, or contract expiration, their accounts and privileges must be promptly terminated. Identity governance systems can automate the disablement and removal of accounts in response to separation actions in human capital management systems or other personnel systems. Identity governance systems also provide a record of accounts and privileges associated with the individual, ensuring that access is completely removed.

Why It Matters

Identity governance solutions implement governance policies using orchestration tools that are designed to link people, applications, data and devices, and allow customers to determine who has access to what, what kind of risk that represents, and take action in situations where policy violations are identified. They provide a comprehensive view of an organization’s identity management practices and identify gaps in the identity management lifecycle. This centralized control and visibility helps to mitigate the risk that identities and privileges will be mismanaged, as well as the risk that attackers can exploit different systems within the organization without being detected.

Additionally, identity governance systems maintain an inventory of active accounts and privileges that currently exist in systems and applications, enabling monitoring and analysis. Account creation and modification events can be reviewed and correlated with approved access requests. Policy rules can be created for segregation of duties requirements, enabling administrators to identify and remove non-compliant combinations of privileges assigned to individuals. Automated risk analysis can identify high-risk individuals so that appropriate mitigations can be taken, such as re-assigning privileges or elevated monitoring of those users’ accounts. The access inventory also enables application and data owners to periodically review and reconcile accounts and privileges. Together, these processes support the principle of Least Privilege, ensuring that users have only the privileges required for their job functions.

Further, managing system and application accounts is also critical. Identity governance systems can monitor and manage the creation, modification, and removal of these accounts to ensure they are only created and granted privileges in response to approved, documented change requests. The entitlements policies, monitoring, risk analysis, and access reconciliation processes applied to user accounts as described above can also ensure that system accounts are managed in accordance with least privilege.

Effective identity governance can mitigate the impacts of many prevalent IAM threats:

- **Phishing, spear phishing, or social engineering:** Identity governance cannot directly prevent these attacks, but can reduce the potential impact of user account compromise using these techniques. A compromised account with excessive privileges can do more damage than one whose privileges are contained. In addition, Segregation of Duty controls enforced through identity governance can ensure that compromising a single account does not provide access to key business processes and data.
- **Insider threats:** As with phishing and other account compromise threats, identity governance cannot prevent insiders from abusing their privileges, but it can reduce the impact when these events happen if they do not have excessive privileges.
- **Creating accounts to maintain persistence:** Attackers who compromise privileged accounts may attempt to create additional user accounts to maintain access to a system even if the initially compromised accounts are revoked or disabled. Identity governance systems monitor account creations and can help an organization identify unauthorized account creation.

Privileged accounts require additional monitoring and control and should be separately managed using a Privileged Access Management (PAM) solution with strong identity governance. Modern PAM solutions include advanced capabilities such as just-in-time provisioning, in which users are temporarily granted privileged access in order to complete a specific task or resolve an issue. This further supports the principle of least privilege and reduces the number of privileged accounts that an attacker could target.

Environmental Hardening

Hardening the enterprise environment includes making sure the foundations and implementations of IAM are sufficiently secured, assured, and trusted. The degree of hardening will vary depending on what is being protected. For example, credential issuing systems for cryptographic digital certificates or stores of passwords are more critical since they secure authentication for entire organizations. Implementation of cryptographic mechanisms must also be sufficient to provide the level of security assumed and needed by the system.

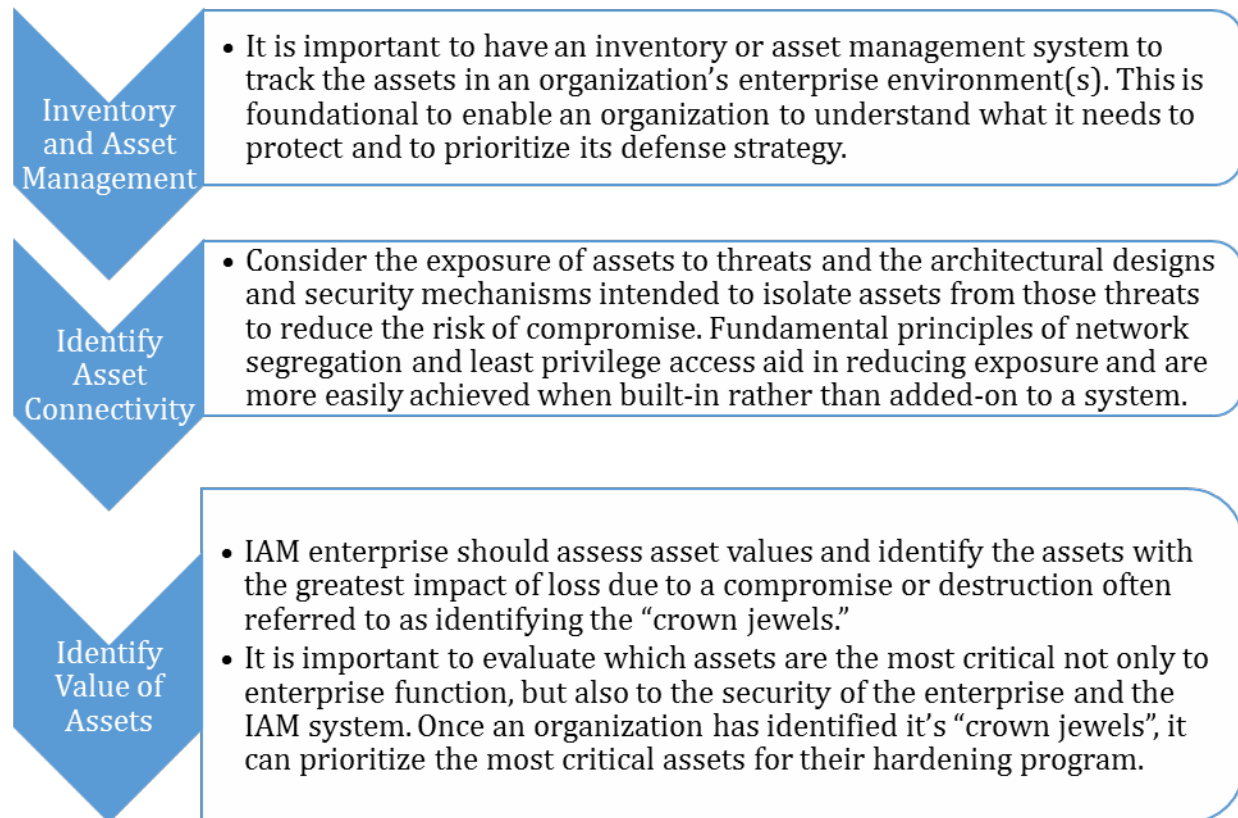
What it Does

Environmental hardening secures the hardware components and software in the enterprise environment around the IAM solution. A defense is only as good as its weakest component. Therefore, it is important when implementing an IAM solution to include securing the other services that are involved. Combining environmental hardening (e.g., patching, asset management, and network segmentation) best practices with sound IAM foundations and implementations reduces the likelihood of a compromise and limits potential damage.

Why it Matters

Environmental hardening generally makes it harder for a bad actor to exploit IAM components and software. Bad actors target IAM solutions because they can provide access to a significant amount of sensitive data, enables persistence, and be used for future malicious cyber operations. IAM solution components must be hardened to prevent footholds for attackers to pivot to more critical systems.

Setting the Stage for Implementation



Implementing Best Practice

<p>Physical and Environmental Hardening</p>	<ul style="list-style-type: none"> • Ensure assets are protected from interruption or data loss due to unauthorized access to a specific physical environment. This can be done by limiting physical access to the data center hosting the IAM assets and the systems controlling logical access to the IAM assets. It is also imperative to use best practices to provide the appropriate resilience of these systems from other physical threats. IAM functions and capabilities should be purposely implemented with system geo-redundancy, if possible, to survive and withstand a physical and/or destructive cyber event at one physical location. • For IAM systems hosted on-site in the organization’s work offices, ensure the server room is located behind a locked door
--	--

	<p>with access granted only to those who have a purpose in that room. A cipher lock or badge access can add MFA capabilities to access the room itself.</p> <ul style="list-style-type: none"> • Ensure that any doors and rooms that provide access to sensitive or critical IAM infrastructure are monitored with cameras that can trigger an alarm if there is unauthorized physical access to the facility (e.g., data center) and room (e.g., on-premises server room). • For IAM systems managed offsite or through a cloud provider, environment hardening needs to ensure remote access is limited by using strong phishing-resistant MFA and limiting access based on other factors (e.g. role-based, normal work hours, location, device, position). It is also key to only engage reputable cloud service providers when choosing to implement the IAM systems offsite. • Ensure disposal of used assets properly by thoroughly wiping or completely destroying the asset depending on the sensitivity of the data.
Network Hardening	<ul style="list-style-type: none"> • When software patches are published for IAM components and/or software, perform a security risk assessment on the patch to assist with installation prioritization. If you have the capacity, consider executing a comprehensive security test plan on all software patches in a non-production environment to ensure compatibility. Proceed to patch and update all impacted devices and/or software as soon as possible. • Ensure an intrusion detection system is in place to alert security operations teams of any suspicious IAM activity. • Develop and set a network baseline so that anomalous network traffic and/or behaviors can be identified and flagged for security analysis to determine if it is a result of malicious or unauthorized activity.
Backups	<ul style="list-style-type: none"> • Follow the “3-2-1 principles” in the event of a disk failure or other disaster: maintain three copies of the data, in at least two mediums, with one being offsite. • Build resiliency in the IAM system in order to prevent access loss due to failure. This resiliency can also have the added benefit of providing better performance through maintaining a lower baseload. Geodiversity should be considered in the resiliency plan for the IAM system.
Least Privileged	<ul style="list-style-type: none"> • Limit user account permissions to those that are necessary to perform their job. IAM solutions can help handle this through locking down privileged accounts, protecting user credentials, and making it easier to assign users to groups with specific permissions.

	<ul style="list-style-type: none"> • Develop policies where normal users, system administrators, and other privileged (e.g., operation and management, application/process, alias, backup, etc.) accounts are separated to ensure that all accesses are using least privilege permissions. • Audit all assets regularly in the organization to identify local identities. Remove unnecessary local identities and investigate to identify who or what process created the local identity. Monitor remaining local identities for anomalous behavior.
Network Segmentation	<ul style="list-style-type: none"> • Carefully design and implement network segmentation with security in mind to limit the spread of an intrusion and to disrupt attempts to escalate privilege. • Isolate IAM systems in a dedicated network segment with layers of security controls between the IAM systems and other systems inside and/or outside the organization.
Network Security Assessment	<ul style="list-style-type: none"> • Perform regular security penetration testing and asset vulnerability security scanning to understand attack surfaces from both outside and inside the organizational boundaries. • Prioritize security hardening efforts on externally exposed assets. • Assess the access allowed internally and the current vulnerabilities that could be exploited by an internal and/or external threat actor. Implement least privilege and access monitoring to reduce risk.
Protect and Manage Critical IAM Assets	<ul style="list-style-type: none"> • Identify your credential/trust stores, control access paths, and provide enterprise-wide management. • Protect keys and certificates at appropriate assurance levels – consider hardware-based security modules for critical items such as signing keys. • Understand tradeoffs between on-premises and cloud based IAM services and ensure visibility into the security of cloud services used. • Recognize and mitigate risks of using 3rd party applications for IAM functions.

Actions to Take Now

- Take an inventory of all assets within the organization. If there is something missing, or if there are additional assets that are unknown, determine the cause of the discrepancy.
- Identify all the local identities on the assets in order to know who has access to which assets.
- Understand what security controls are in the enterprise environment now and what security gaps persist in an organization's enterprise environment.
- Develop a network traffic baseline that can be used to detect security anomalies in the network. Any compromise to any component in a network has the potential to threaten more critical enterprise systems, including IAM.

Summary

IAM solutions are only one part of a wider enterprise environment, where compromises in one area can eventually lead to compromises in another. Hardening the enterprise environment, including the IAM systems as critical resources, helps to limit the potential for a compromise and keep the IAM system safe and accessible.

Identity Federation and Single Sign-On

Identity federation using SSO within and/or between organizations, including the utilization of identity providers, mitigates risks by centrally managing differences in policies and risk levels between the organizations and eliminates wide implementation and dependence on local identities. Without formally defining the policies and levels of trust and assurance between organizations or between multiple identity providers within an organization, the organization is susceptible to attacks based on weaknesses in each federated IAM. SSO provides a risk mitigation capability by centralizing the management and control of authentication and access across multiple systems and from multiple identity providers. Implemented properly, it can also raise the authentication assurance level required for initial sign on and can control and secure the authentication and authorization information passed between systems.

What it Does

Identity Federation and SSO simplifies identity management internally within an enterprise and with trusted external partners by reducing the need for users to maintain multiple identities in both internal and external directories, applications, and other platforms, eliminating the need for local identities at each asset. It allows for seamless integration with other security controls such as privileged access management for step-up authentication and increases confidence that only active users are allowed access. Additionally, it reduces the labor costs associated with managing multiple identities for each user on the various on-premises and/or cloud-based applications.

Why it Matters

Passwords are a vulnerability due to the complexity of requiring a user to remember multi-character passwords that almost every application requires today. SSO nominally reduces the user burden to remembering one solid, complex, and hard-to-guess passphrase, and facilitates the migration to strong MFA, potentially eliminating passwords altogether. Implementing both Identity Federation and SSO supporting strong MFA allows for improved security without compromising the user experience.

Locally provisioned accounts (e.g., user, system, process, admin) on individual assets creates an unmanageable environment and is a lucrative target by bad actors. For example:

- Locally provisioned accounts may or may not allow for security policy enforcement.

- Massive volumes of locally provisioned accounts on individual systems across the enterprise cannot be maintained. These accounts can include shared accounts, vendor default accounts, and unknown accounts (e.g., ex-employee, ex-vendor).
- Security event monitoring is ineffective on locally provisioned accounts. For instance, the ability to monitor and detect shared accounts, stolen credentials, and cracked credentials (e.g., password spraying) is considerably more difficult given the volumes of assets, accounts, and individual asset configurations.
- Adversaries, both internal and external threat actors, can exploit the security policy and/or security event monitoring gaps in one system to compromise the assets it manages and use their access as a foothold to launch exploits against other systems.

Identity Federation and SSO drastically reduce the need for locally provisioned accounts and enables IAM administrators to have more centralized visibility and control over accounts. It also enables more effective management of default and/or shared accounts that are required on an individual asset. For example, most default and shared accounts can be disabled and those that cannot be disabled can have passwords changed to highly random values protected in a password vault.

Factors to consider when selecting an SSO solution

SSO services may use different protocols, such as SAML or Open ID Connect (OIDC). When selecting an SSO service, it is important to keep in mind the following factors:

- What protocol is being used?
- How has the service provider secured the protocol and the service?

SAML

SAML is used for exchanging authentication and authorization data between identity providers and service providers. One of the most common use cases for SAML is facilitating browser-based SSO. Up until the past few years, SAML was considered the industry standard and proven workhorse for passing an authenticated user into applications while allowing these applications to defer authentication to a centralized identity solution.

If the services use SAML, specific implementation and hardening measures are a must to be a secure SSO option as it is prone to exploits if it is not implemented correctly. Every year brings new issues with SAML – in the form of newly discovered exploits – which gives it a reputation of not being the most secure option.

OIDC was created to address some of the flaws in SAML. However, SAML is still considered a relevant option for SSO and there are still requirements for developers to support it in modern environments.

OpenID Connect

OAuth 2.0 is designed only for authorization for granting access to data and features from one application to another. OIDC is a thin layer that sits on top of OAuth 2.0 that adds login and profile information about the person who is logged in. OIDC enables scenarios where

one login can be used across multiple applications (i.e., SSO). An application could support SSO using social networking services (i.e., Facebook or Twitter) so that users can choose to leverage a login they already have. Authorization code flow enables the applications to first get authorization codes instead of getting tokens directly from the authorization callback request. It then uses these codes in a request to another endpoint on the authorization server to exchange them for the tokens they need. The most significant advantage that this flow has in relation to the implicit flow is its security. There are two characteristics of the authorization code flow that make it a better choice than SAML when it comes to security. An example of the authorization code flow is depicted in Figure 2 below.

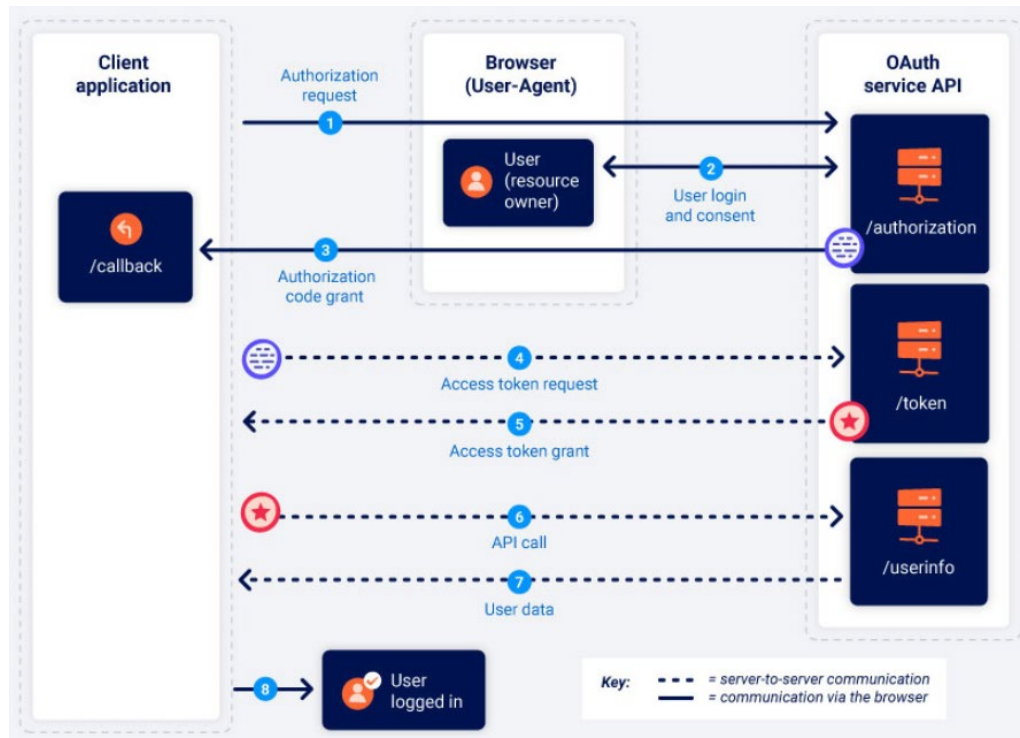


Figure 2 Diagram of Authorization Code Flow⁸

First, the process to exchange codes for tokens happens on back channels. Instead of having tokens traveling through users' devices, the application opens a back channel connection to the authorization server, eliminating the need to pass credentials and other information through the users' devices (like browsers). By establishing a direction connection to each other, the application and authentication server reduce the chances that certain credentials will be exposed. When registering a Web App, the call back configuration is important from a security point of view because it restricts what URLs the OIDC provider is allowed to call after a successful authentication process.

The second characteristic is that, before issuing tokens, authorization servers require applications to authenticate themselves. This authentication process usually happens by applications using credentials that authorization servers assign to them.

⁸ <https://portswigger.net/web-security/oauth/grant-types>.

In summary, OIDC is a more secure and reliable protocol because it uses a direct channel between the applications and the authentication server, protecting identity tokens.

Implementing Best Practices

Organizations should consider the following when assessing their SSO capability and making improvements to counter their organization's top threats and plan for periodic reassessments to ensure updates are made as needs change.

- Define and understand how assets are audited for any local accounts and/or identities configured and active.
- Define and understand how the engagement with trusted partners to audit for any local accounts and/or identities configured and active.
- For any required and authorized local accounts/identities, define a password policy, and auditing to ensure compliance.
- Define a policy that disallows local accounts on any platform.
- Implement a configuration management solution which supports the identification, tracking, and reporting of any local accounts.
- Identify and track all exceptions for systems, platforms, and/or applications that require local accounts. Disable those that are not necessary and establish and enforce password policies for those that are. Review these periodically with the application teams and/or vendors in an effort to drive them to SSO support.
- Ensure SSO availability. If SSO fails, access to all related systems is lost. Therefore, it is key to have a solid high availability design and plan implementation which includes both local and regional geographic redundancy and the appropriate security hardening guidelines.

Actions to Take Now

- Assess your organization's internal on-premises applications/devices/platforms and your cloud providers ability to connect using SSO.
- Determine if your SSO integration can collect user context during SSO logins including location, device, and behavior.

Summary

Organizations should develop and deploy SSO friendly applications and platforms to eliminate all local accounts and/or identities. Doing so will improve the user experience while also significantly reducing the risk associated with local accounts which are difficult to manage and monitor. Local accounts that use shared passwords (e.g., root) create legal and forensic issues for the organization when attempting to identify the attacker's identity.

Multi-Factor Authentication

Since the introduction of multi-user computer systems, user authentication has primarily relied on usernames and passwords. MFA is an approach to strengthen the authentication process by requiring the user to present multiple elements in different categories, or

“factors”, as part of an authentication attempt. These factors are as shown in Figure 3 are something you have, something you know, and something you are.



Figure 3 Multi-Factor Authentication Factors

MFA incorporates more than one of the above factors as part of a login flow. Examples include:

- Typing a password and responding to a push notification sent to a registered smartphone.
- Typing a password and providing a one-time code from a hardware authentication device.
- Using a biometric facial scan and/or passphrase to unlock a cryptographic credential stored on a registered device (i.e. phone, hardware token).

Authentication systems are the front doors to enterprise networks, applications, and data. As such, attackers are highly focused on finding and exploiting authentication vulnerabilities. Authentication systems are also high-volume user interfaces and frequently seen as friction points between users and their ability to perform their business functions. This combination of characteristics poses a challenge for systems engineers and implementers since they must be seamless and user-friendly yet also strongly resistant to attacks.

MFA authenticators may take the form of software that runs on a smartphone or other device or dedicated hardware tokens. Some MFA solutions are designed to augment passwords with an additional factor, whereas, “passwordless” solutions can eliminate the need for passwords altogether. Passwordless MFA solutions typically involve the use of two factors together, such as a cryptographic credential stored on a hardware token that is unlocked using a memorized PIN. Table X below lists some common forms of MFA.

MFA Type	Examples	Relevant Standards
One-time Passwords (OTP)	OTP delivered out of band by simple messaging service (SMS) or email Hardware OTP token Mobile OTP app	HMAC-based OTP (HOTP) – RFC 4226 ⁹ Time-based OTP (TOTP) – RFC 6238 ¹⁰
Out-of-band Push Notification App	Mobile app that presents options to approve or reject a login event from another device	N/A
Cryptographic authenticator	Fast Identity Online (FIDO) hardware token FIDO software token (e.g, Passkey) Smartcard Software Public Key Infrastructure (PKI) credential unlocked with biometric	CTAP2 ¹¹ Web Authentication ¹² NIST SP 800-74 ¹³ NIST SP 800-157 ¹⁴

It is important to note that not all MFA solutions provide equal protection against authentication attacks, and there are critical implementation details that can impact the security and usability of an MFA deployment. The following subsections provide guidance for selecting and implementing an MFA solution. Further guidance is also available in the National Institute of Standards and Technology (NIST) Special Publication (SP) 800-63¹⁵, NSA's publication, *Selecting Secure Multi-factor Authentication Solutions*¹⁶, and the Cybersecurity Infrastructure Security Agency's guidance on MFA.¹⁷

What It Does

MFA was created to address the shortcomings of passwords including the fact that:

- Passwords can be shared with unauthorized users;
- Users can be tricked into giving their passwords to attackers through phishing; and
- Users tend to use the same or closely related passwords across multiple websites, services, and computer systems, meaning a breach of one system allows an attacker to obtain usernames and passwords that can be used in other systems using techniques such as credential stuffing.

⁹ <https://datatracker.ietf.org/doc/html/rfc4226>.

¹⁰ <https://datatracker.ietf.org/doc/html/rfc6238>.

¹¹ <https://fidoalliance.org/specs/fido-v2.1-ps-20210615/fido-client-to-authenticator-protocol-v2.1-ps-20210615.html>.

¹² <https://www.w3.org/TR/webauthn-2/>.

¹³ <https://nvlpubs.nist.gov/nistpubs/specialpublications/nist.sp.800-73-4.pdf>.

¹⁴ <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-157.pdf>.

¹⁵ NIST Special Publication, 800-63AB.

¹⁶ [CSI MULTIFACTOR AUTHENTICATION SOLUTIONS U0017091520.PDF \(defense.gov\)](https://www.defense.gov/CSA/MULTIFACTOR_AUTHENTICATION_SOLUTIONS_U0017091520.PDF).

¹⁷ <https://www.cisa.gov/mfa>.

MFA mitigates common attacks against passwords such as brute force guessing and credential stuffing as well as common misuse practices such as password sharing by requiring the presentation of another factor in addition to the password. Unless an attacker can defeat the MFA authentication mechanism, knowing the password by itself does not enable impersonation of the user. In the case of passwordless authentication systems, passwords are eliminated altogether as an attack vector.

Some, but not all, MFA solutions also mitigate phishing attacks. Given the prevalence of phishing as an attack vector, phishing resistance should be a key consideration in choosing an MFA solution. Figure 4 represents different types of MFA ranging from weakest to strongest.

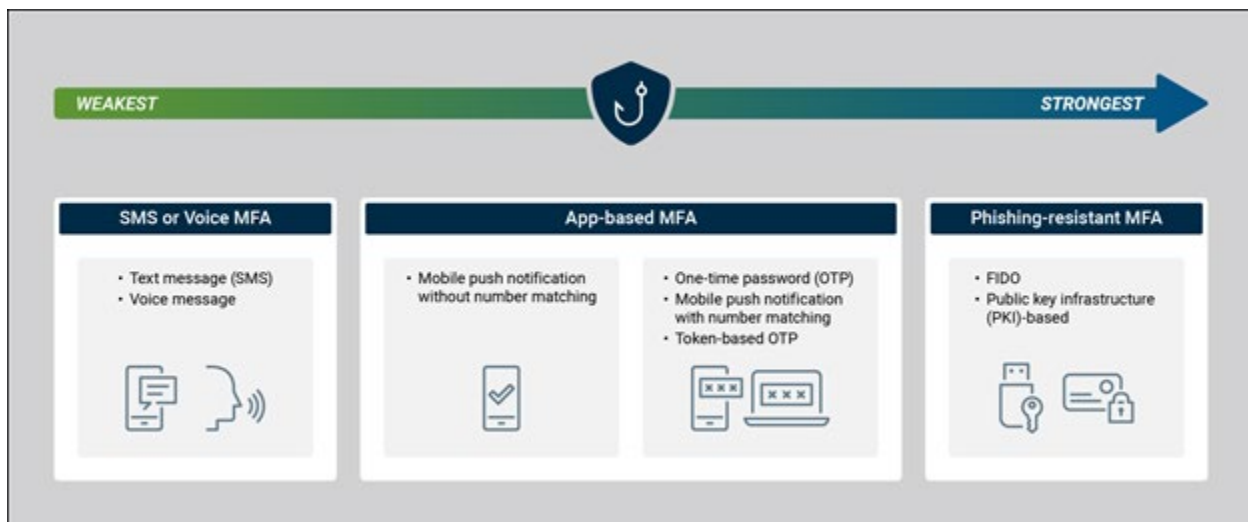


Figure 4 Weakest to Strongest Types of MFA¹⁸

The following are some general guidelines around MFA and phishing:

- One-time passwords, whether generated in an app or hardware token or delivered through SMS, e-mail, or some other out-of-band method, do not protect against phishing unless they are combined with some other phishing-resistant technology such as mutual TLS authentication. Because the user enters the one-time password into a login form, a phishing site can capture an OTP just as easily as a password and replay it to the legitimate site or application in real-time.
- Push notification-based authenticator apps that prompt the user to approve login attempts also generally do not protect against phishing. A phishing site can trigger a login attempt that will send a push notification to the user's registered device and the user may have no way of determining whether the notification is legitimate. Some attackers have had success rates simply triggering push notifications to users who are not even attempting to log in at the time.¹⁹ Some push notification-based MFA solutions provide additional context about the authentication attempt, such as

¹⁸ <https://www.cisa.gov/mfa>.

¹⁹ [What Are Push Attacks? | HYPR](#).

the location from which it originated, to aid the user in determining whether it is legitimate. If the login request came from a phishing site, the detected location of the login attempt should not match the user's current location. However, location can be spoofed, and the basic issue remains that the push notification is not strongly bound to a legitimate authentication attempt and the service to which the user is authenticating.

Further guidance is also available in CISA's publication *Implementing Phishing-Resistant MFA*.²⁰

Phishing-resistant forms of MFA include:

- FIDO authenticators – a wide range of interoperable authenticators, both built into commonly-used operating systems (Windows, MacOS, iOS Android) and available in hardware tokens, based on industry standards maintained by the FIDO Alliance and the World Wide Consortium (W3C).
- PKI credentials, in the form of software crypto modules, smartcards, and other hardware tokens.

Why MFA Matters

MFA solutions can mitigate many of the most common attacks against authentication systems:

- Credential stuffing is the attempt to use known username/password credentials obtained from one system (typically through compromise and cracking of the password database) to access other systems. Credential stuffing takes advantage of the tendency for users to reuse the same credentials for multiple sites and services. With MFA, these stolen credentials are not sufficient to gain access to a user's account because the attacker cannot bypass second-factor authentication.
- Password spraying is a similar attack where the attacker tries a relatively short list of the most commonly-used passwords against a list of known usernames. Typically, the attacker tries a small number of passwords for each user account to avoid triggering the account lockout threshold to reduce the risk of detection. If a system locks users after 10 failed attempts, the attacker may try 9 passwords for each username. Again, MFA can prevent account takeover even if the attacker discovers valid username/password credentials by requiring an additional authentication factor.
- Phishing is an attempt to trick users into logging into an attacker-controlled system and capture their credentials. As described above, some MFA authentication systems prevent phishing by using protocols designed such that a phishing site cannot simply "replay" the authentication protocol messages against the legitimate site.
- Brute-force attacks are the simplest form of password attacks, where an attacker simply tries different passwords in the hopes of finding valid credentials. Account

²⁰ [Implementing Phishing-Resistant MFA \(cisa.gov\)](#).

lockouts make these attacks much more time-consuming, but strong MFA can completely mitigate them.

Preparation for Implementing MFA

Before deploying MFA, it is important to understand the full scope of use cases and scenarios the MFA solution needs to address. An ad-hoc approach can lead to incomplete coverage, multiple systems, and users needing to enroll multiple MFA mechanisms to access all the applications they need. Up-front planning and strategy definitions can help ensure a smooth, coherent implementation. This section details several aspects to consider and further information can be found in the NSA publication “*Transition to Multi-Factor Authentication*”²¹ and in CISA’s publication “*Capacity Implementation Guide: Implementing Strong Authentication*.”²²

Catalog User Populations, Device Types, and Use Cases

Consider the needs of different user groups to best determine how to handle MFA enrollment. Questions to consider include:

- What types of authenticators are suitable for each based on assurance level, usability, supportability, and cost?
- What are the various device platforms your MFA solution needs to accommodate? Desktops, laptops, smartphones, and tablets are common requirements.
- Is MFA also needed for networking devices or other equipment?
- What are the potential device compatibility issues for both software and hardware MFA solutions? It is important to consider up-to-date operating systems and browsers, available USB ports (with their A/C/Micro variations), or support for Bluetooth or Near-Field Communications (NFC). It’s also important to consider the security profiles of the devices and the difference between devices under enterprise management and monitoring, devices managed by a different organization, and personal, unmanaged devices, especially regarding software solutions.

Also consider the different support needed for operating environments. In operating environments with shared workstations, portable authenticators would probably be most appropriate versus software authenticators tied to a specific device. Additionally, in operating environments that have users with managed mobile devices, the iOS and Android platforms both provide built-in authentication capabilities using the FIDO standards and numerous other vendors offer authenticator applications that could meet your MFA needs without buying additional hardware. However, if there are high-security environments such as research facilities where electronic devices are not admitted, smartphone-based authenticators would not be appropriate.

²¹ <https://media.defense.gov/2019/Sep/09/2002180346/-1/-1/0/Transition%20to%20Multi-factor%20Authentication%20-%20Copy.pdf>.

²² https://www.cisa.gov/sites/default/files/publications/CISA_CEG_Implementing_Strong_Authentication_508_1.pdf.

It is important to note that organizations may find that a single MFA solution cannot accommodate all their needs, especially if managing access for external users. Deploying different MFA solutions for different groups of users may be required. This is a situation where ID Federation/SSO will be important.

Evaluate Assurance Requirements

Some use cases, applications, or data types may require higher-assurance authentication than others. For example, privileged users with operating system or database administration rights should have strong, phishing-resistant authentication. The use of separate user-level and administrative accounts and credentials for individuals with privileged access and Privileged Access Management (PAM) systems that provide auditing of privileged access use are additional best practices for managing privileged access that can be deployed in conjunction with MFA. PAM may also provide work-flow management and be a credential proxy for systems that don't support the selected MFA. Other high-risk roles or functions may also require special protections if they involve management of high-value assets or critically sensitive information.

For reference, NIST SP 800-63-3 provides guidelines for performing a risk assessment to guide selection and implementation of identity and authentication systems, including MFA. Also consider any regulatory or compliance mandates applicable to your organization, which may include requirements that are relevant to MFA solutions, such as the use of Federal Information Protection Standards (FIPS) 140-3 validated cryptography or FIPS 201 (PIV).

Evaluate Privacy and Operational Considerations

Many MFA solutions incorporate biometric authentication of the user, which can raise concerns over privacy. The biometric authentication solutions in most widespread use today, such as the facial recognition and fingerprint unlock mechanisms built into smartphones, keep biometric templates in hardware-protected storage and are designed to prevent the removal of biometric data from the device. When these systems are used to authenticate to systems and services, the biometric matching occurs locally on the mobile device, and successful authentication unlocks a private key that is then used in the actual authentication protocol carried out over the network. These types of protections are requirements for FIDO-certified devices. Using solutions that bind biometrics templates to a single device, instead of storing them in a central database, may help alleviate privacy concerns.

Equity across demographic groups is another potential issue with biometrics; some biometric solutions perform differently for individuals of different ages, genders, and/or ethnicities. Pilot testing with a representative cross-section of your user base can help identify any potential issues. Aspects of your users' operating environment may also impact the suitability of specific biometric modalities; the use of gloves or masks, for example, may preclude facial or fingerprint authentication.

Implementing MFA

The following are some best practices and considerations when embarking on an MFA implementation.

Implement MFA as part of an enterprise SSO solution. Integrating MFA with all of an organization's applications can be a daunting prospect; it's also not the best way to go about an MFA implementation. MFA integration is complex, and small mistakes can lead to issues like the ability for attackers to bypass MFA. This is a job for experienced IAM practitioners and vendors, not an additional-duty-as-assigned for application developers. Also, allowing individual applications and projects to choose their own MFA solutions leads to a complex environment where users need to manage multiple authenticators to access all the applications they need. Having multiple MFA infrastructures also expands the attack surface and complicates maintenance.

Instead, as discussed in the previous section, MFA should be integrated into an enterprise authentication and SSO service that uses industry-standard, tested and proven protocols, like SAML, or OpenID Connect and OAuth 2.0, to connect with your applications. A single, centralized authentication service is simpler to test, secure, and maintain than several independent application-level implementations. In addition, a centralized SSO system can enable enterprise risk-based authentication policies to selectively require MFA. When a user has an active session with the SSO service, policies can determine whether they need to authenticate again when accessing additional applications. Policies can trigger the need to re-authenticate or perform step-up authentication (i.e., requiring higher-assurance authentication than was used to initially establish a user's session) when users access sensitive applications or perform high-risk activities. This provides the flexibility to require high assurance when needed without frustrating users engaged in routine, low-risk tasks with repeated MFA prompts. It also provides an integration point for Zero Trust Architecture (ZTA) policies such as requiring re-authentication or step-up based on risk signals from threat defense systems.

Consider the total account and authenticator lifecycle, and exception processes.

Procuring an MFA system and enrolling users is only the beginning of the process. It's important to consider all the needed workflows for authenticator lifecycle management and how edge cases and failure scenarios will be handled. Initial MFA enrollment (or issuance, in the case of hardware authenticators) process must provide adequate assurance that the authorized user is enrolled in the MFA system. Consider the use of multiple communication channels to provide additional assurance. For example, if a hardware token is physically mailed to a user, require additional authentication (e.g., with their password or a one-time secret provided out-of-band) as part of the enrollment process.

Maintain an inventory of the authenticators deployed in your environment.

Vulnerabilities may be discovered in both software and hardware authenticators, so it's critical to be able to identify authenticators in need of replacement or upgrade. Pay attention to vendor announcements and support lifecycles, and plan well in advance for any end-of-life authenticator solutions in need of replacement. For mobile authenticator

apps, consider your device refresh period and how users will enroll a new device. Also have a response plan for lost or stolen authenticators or devices to rapidly disable the lost authenticator and enable the user to enroll a new one. This can be one of the most challenging aspects to manage – if a user can enroll a new MFA authenticator using their password alone, this severely undermines the security of your MFA solution. A best practice, particularly in passwordless environments, is to issue multiple strong authenticators to each user, perhaps with one kept in reserve in a secure location to allow access and enrollment of a new authenticator in case the primary authenticator is lost. A simpler solution is using backup one-time codes, kept in secure storage by the user.

Routinely test and rapidly patch your MFA infrastructure. This is good advice for any system or application, but it is especially critical for MFA and other authentication infrastructure. Promptly test and install any vendor security patches. Routinely test your registration and authentication flows, especially when changes are made to your infrastructure.

Realize that MFA is not the only solution required for securing identities and access. MFA is a critical security control, but it is only one component of securing access to your systems and applications. MFA (and SSO) enable users to establish a session with an application, but the application must implement secure session management with timeouts for inactivity and maximum session lifetimes. Applications and client devices must protect cookies and tokens that can allow impersonation of the user if stolen. MFA cannot prevent malware on client devices from capturing users' credentials or application data. It's important to understand that while MFA addresses some of the most common threats, MFA should be part of a holistic cybersecurity architecture.

Actions to Take Now

- Determine the MFA solution best suited in your organization's operating environment.
- Implement MFA as part of an enterprise SSO solution.
- Maintain a robust inventory of the MFA authenticators deployed in your organization's operating environment.
- Routinely test and patch your organization's MFA infrastructure.

Summary

MFA can provide strong protection against many of the most prevalent attacks against authentication systems. Careful planning will help ensure that your MFA implementation meets your organization's needs and provides both security and usability. As with any enduring capability, it's important to consider the full lifecycle management of MFA authenticators and infrastructure. Integrating MFA with an enterprise SSO system is essential to facilitate application adoption and enable a coherent enterprise authentication policy.

IAM Auditing and Monitoring

IAM auditing and monitoring should not only check for compliance, but also monitor for threat indicators and anomalous activities. This encompasses the generation, collection, and analysis of logs, events, and other information to provide the best means of detecting compliance related infractions and suspicious activities. Attacks such as use of stolen credentials and misuse of privileged access by insiders would not be detected in a timely manner, if at all, without an effective IAM auditing and monitoring program. These auditing and monitoring capabilities can be integrated with automated tools that orchestrate response actions to counter these IAM attacks. Effective reporting from auditing and monitoring also provide situational awareness of the security posture of an organization's IAM.

What it Does

IAM auditing and monitoring:

- Provides deterrent to users especially privileged users who know their actions are being tracked;
- Provides awareness of how system is being used and attempted to be misused;
- Detects problems and potential problems through indicators of attack/compromise and changes in behavior; and
- Collects forensic evidence which also supports evaluation of effectiveness leading to improvements in capabilities.

Why it Matters

There are many types of threats that IAM auditing and monitoring can counter but they tend to fall into one of two buckets; insider threat and unauthorized access. Insider threats range from authorized using their privileges to perform inappropriate actions (e.g. downloading a list of current customers) to administrators seeking to cause harm to the organization, to former employees whose access was not turned off. For example, in September 2022, an individual working as a cybersecurity professional in a Hawaiian-based financial company, pled guilty and admitted that, after severing ties with the company, he utilized the credentials of his former employer to gain access to the company's website configuration settings and purposefully misdirected web and email traffic to computers unaffiliated with the company incapacitating the company's website and email.²³ IAM auditing and monitoring could have potentially prevented this by allowing the system to remove the user's access upon separation from the company.

Unauthorized access can occur when external systems or users with lower assurance (i.e. weaker authentication) inappropriately gains access to an organization's system and data. Further, exploitation of vulnerabilities in security protocols, cryptographic algorithms, and/or third-party programs could also lead to unauthorized access. Additionally,

²³ <https://www.justice.gov/usao-hi/pr/honolulu-man-pleads-guilty-sabotaging-former-employer-s-computer-network>.

unauthorized access can occur with the theft or hijacking of a legitimate user's credentials to attack an organization's system with the stolen or hijacked credentials. In this instance, the impostor's behavior and actions will likely be different from the normal behavior of the legitimate user and can lead to detection of the identity theft. The legitimate user may also receive notifications of log in failures or other activity that they did not perform and can provide out of band information to help detect the impostor.

Preparation for Implementing Best Practice

Below are key considerations for assessing an organization's auditing and monitoring capability to determine which improvements are necessary to counter top threats. It is important to note that this is not a one-time assessment. Assessments should be made periodically, and capabilities updated in order to meet changing needs and be better postured to counter new threats.

- Organization defines and understands what is considered normal/acceptable behavior, suspect behavior, and misbehavior.
- Organization uses defined and de-facto policy rules, requirements/models of systems, and baselines of current activity to identify monitoring and analysis parameters.
- Organization identifies users with access to critical assets (e.g., crown jewels) and focuses enhanced monitoring on critical assets (proprietary information, systems mission critical); Identify, prioritize assets).
- Collect data including standard logs/audit records, and security events as well as other data about the users, systems, applications, and network behaviors. Use the collected data for real-time detection and alerting, storage for forensic use, baselining of current behavior, analysis to detect trends, and indications of anomalous behavior.
- Behavioral analytics will require an initial period (and ongoing updates) of collection and analysis to establish baselines and thresholds. This should address normal day, busy day, and emergency situation baselines.
- Avoid collection and analysis that does not provide useful information such a large number of unprioritized alerts that require human analysis since this is a waste of systems and human resources and will not achieve better cybersecurity. Collecting and analyzing data that provides actionable information to your staff and management to raise security awareness and can support the business case for additional funding to improving your IAM auditing and monitoring capabilities.
- Determine the appropriate tools and capabilities to effectively derive information from the collected data. Consider what data formats and content can be processed, configurability, scalability, growth capability to provide or interface with other systems and capabilities. For example, a SIEM tool that can accommodate SOAR capability or one that can work with advanced analytic tools including machine learning.
- The tools and capabilities should match and best augment your staff skills and availability. Manual review of logs or of overly detailed or too frequent tool outputs will not be effective. If your current tools are at the basic SIEM level, focus on

configuring them to alert on your most critical events and provide the most pertinent info to staff. Organizations with more sophisticated capabilities should start looking for anomalous behavior and developing procedures on how to deal with potential insider threats. For example, when to shut them down immediately versus when to steer them to honeypots and collect more forensics evidence.

Initiatives such as the Defense Advanced Research Projects Agency's (DARPA) Anomaly Detection at Multiple Scales (ADAMS) Project²⁴ provide valuable information for organizations to use as a starting point when attempting to identify and remediate insider threats. The project developed an Anomaly Detection Engine for Networks (ADEN) to detect malicious users and characterize anomalous behavior typical of malicious users, to support improved prediction-based actionable intelligence and response. While only a small percentage of anomalous behavior was associated with malicious users, the project did highlight several key findings associated with behavior of malicious users, including:

- Malicious users were more active and chose to “do nothing” significantly less times than benign users.
- Malicious users fetched significantly more sensitive information than benign users.
- Though malicious users appeared to save more data to removable devices than benign users, these differences were not found to be statistically significant in our study.
- Malicious users edited the data slightly less compared to benign players users. However, these differences also were not found to be statistically significant.
- Malicious users sent significantly more information out of the organization than benign users.
- Malicious users fetched significantly less un-sensitive data in contrast to the benign players.

Actions to Take Now

- Establish baseline expectations of activity levels and policy and monitor privileged user behavior for both acceptable and suspicious activity. Avoid automatic response actions to suspicious behavior that could be important and legitimate (e.g. system administrator that flags as unusual activity due to logging in from a remote location on a weekend however could be responding to an emergency network problem). Include manual procedures to confirm the legitimacy of these actions before determining how to respond. For example, if the activity includes setting up new accounts or changing privileges a first step would be to determine if there are indications that this may be a malicious insider attack versus preparing for the startup of a new program.
- Monitor general user behaviors in both good and bad terms such as how many successful access attempts versus unsuccessful, what hours typically worked, whether remote access allowed, what systems accessed and amounts of data downloaded.

²⁴ <https://www.darpa.mil/program/anomaly-detection-at-multiple-scales>.

- Monitor activity between applications and systems and associated network traffic for changes in connectivity, level of activity, and types of data. If an attacker is attempting to move laterally within your network, this may include accesses and traffic that are unusual.
- Monitor external traffic that may include new interactions with previously unknown sites or different types and levels of interactions. Remember that data exfiltration attacks may be 'low and slow' so a change may be small, but ongoing. Be careful to not include this in an accepted baseline of activity.

Summary

Organizations will need to be able to monitor for anomalous behavior (in addition to traditional security events and logs) to detect the various threats to IAM systems that are present and potentially harmful. An initial assessment should be performed to understand current capabilities with a plan to improve an organization's capability to collect, analyze, detect, and respond to indicators of attack and compromise.

Conclusion

America's critical infrastructure is a prime target for a broad spectrum of threat sources including advanced and ongoing attacks from nation state and terrorist organizations attacks. These threats are real, ongoing, and evolving and the cybersecurity community is especially concerned about certain credible threats to IAM and SSO. IAM weaknesses are frequently exploited in the most insidious threats, APTs, which have led to catastrophic data breaches. The use of SSO without a good MFA foundation and secure design selections, exacerbates the damage of attacks that an organization may be vulnerable to such as password cracking and authenticator hijacking.

The intent of this paper was to provide a clear understanding of how various mitigations counter the threats and to provide actionable recommendations on what organizations should do now. This includes:

- Assess your current IAM capabilities and risk posture.
- For areas that need improvement: select, layer, integrate, and properly configure secure solutions following the best practices provided herein and in referenced guidance.
- Maintain the appropriate level of security to manage risk during continued operations.
- Maintain awareness of correct IAM usage and of risks.

Ultimately every organization has the obligation to ensure their IAM and SSO capabilities are secure to protect not only their own assets but that of their partners and consumers as

Appendix I: Actions to Take Now Checklist

Environmental Hardening

- Take an inventory of all assets within the organization. If there is something missing, or if there are additional assets that are unknown, determine the cause of the discrepancy.
- Identify all the local identities on the assets in order to know who has access to which assets.
- Understand what security controls are in the enterprise environment now and what security gaps persist in an organization's enterprise environment.
- Develop a network traffic baseline that can be used to detect security anomalies in the network. Any compromise to any component in a network has the potential to threaten more critical enterprise systems, including IAM.

Identity Federation/Single Sign-On

- Assess your organization's internal on-premises applications/devices/platforms and your cloud providers ability to connect using single sign-on.
- Determine if your single sign-on integration can collect user context during single sign-on logins including location, device, and behavior.

Multi-Factor Authentication

- Determine the MFA solution best suited in your organization's operating environment. Implement MFA as part of an enterprise SSO solution.
- Maintain a robust inventory of the MFA authenticators deployed in your organization's operating environment.
- Routinely test and patch your organization's MFA infrastructure.

IAM Auditing and Monitoring

- Establish baseline expectations of activity levels and policy and monitor privileged user behavior for both acceptable and suspicious activity. Avoid automatic response actions to suspicious behavior that could be important and legitimate (e.g. system administrator that flags as unusual activity due to logging in from a remote location on a weekend however could be responding to an emergency network problem). Include manual procedures to confirm the legitimacy of these actions before determining how to respond. For example, if the activity includes setting up new accounts or changing privileges a first step would be to determine if there are indications that this may be a malicious insider attack versus preparing for the startup of a new program.

- Monitor general user behaviors in both good and bad terms such as how many successful access attempts versus unsuccessful, what hours typically worked, whether remote access allowed, what systems accessed and amounts of data downloaded.
- Monitor activity between applications and systems and associated network traffic for changes in connectivity, level of activity, and types of data. If an attacker is attempting to move laterally within your network, this may include accesses and traffic that are unusual.
- Monitor external traffic that may include new interactions with previously unknown sites or different types and levels of interactions. Remember that data exfiltration attacks may be 'low and slow' so a change may be small, but ongoing. Be careful to not include this in an accepted baseline of activity.