Secure Cloud Business Applications

Hybrid Identity Solutions Architecture

March 2023

Cybersecurity and Infrastructure Security Agency

CONTENTS

FIGURES

TABLES

1 INTRODUCTION

Identity management for a traditional on-premises enterprise network is usually handled by an on-premises directory service (e.g., Active Directory). When agencies leverage cloud solutions and attempt to integrate them with their on-premises systems (creating a "hybrid" environment), identity management can become significantly more complex.

On-premises identity management solutions need to securely and efficiently integrate with those applied in the cloud to achieve interoperability. This document seeks to help agencies understand potential options for identity management interoperability between on-premises and cloud-based solutions, the challenges involved in each, and how to address those challenges.

1.1 Background

Identity management vulnerabilities have played a key role in several recent high-profile cybersecurity incidents.[1,2] In light of these and other incidents, industry stakeholders, vendors, and other key partners continue to encourage a transition from on-premises to cloud-based identity solutions and phishing-resistant multifactor authentication (MFA).

For a variety of reasons, it is not likely that all agencies will completely abandon on-premises identity services. This will result in a future state in which agencies must securely architect, deploy, maintain, and update on-premises and cloud-based identity services in a manner that integrates across these environments.

These modernization efforts are critical to fulfilling agency mission. They present an opportunity to enable zero trust across the enterprise.[3,4,5] They must be tightly coupled with broader plans to adopt zero trust architectures.

Executive Order (EO) 14028 initiated a government-wide cybersecurity modernization effort to migrate to zero trust architectures, realize the benefits of cloud services, and mitigate associated risks.[6] The Office of Management and Budget (OMB) Memorandum M-22-09 noted that agencies should make use of the rich security features present in cloud infrastructure.[7]

Just as M-22-09 addresses on-premises, cloud, and hybrid systems, this solutions architecture document will take a similar approach for identity services. This document seeks to provide guidance for hybrid identity solutions and cloud-first identity strategies (i.e., enterprise sole or primary reliance on cloud-based identity services).

1.2 Scope

For the purposes of this document, hybrid identity refers to the deployment of integrated on-premises and cloud-based identity services. This solutions architecture document presents potential approaches for addressing identity management in a hybrid environment (1) in which on-premises identity services are deployed as an agency's primary identity solution or (2) in which cloud-based identity services are deployed as an agency's primary identity solution. This document presents various considerations associated with each approach and

---

[1] "Emergency Directive 21-01 – Mitigate SolarWinds Orion Code Compromise," DHS CISA, last modified April 15, 2021, https://www.cisa.gov/emergency-directive-21-01.

[2] "Emergency Directive 21-02 – Mitigate Microsoft Exchange On-Premises Product Vulnerabilities," DHS CISA, last modified April 13, 2021, https://www.cisa.gov/emergency-directive-21-02.

[3] Office of Management and Budget, *OMB. M-22-09: Moving the U.S. Government Towards Zero Trust Cybersecurity Principles,* January 26, 2022, https://www.whitehouse.gov/wp-content/uploads/2022/01/M-22-09.pdf.

[4] DHS CISA. *Zero Trust Maturity Model,* June 2021, https://www.cisa.gov/sites/default/files/publications/CISA%20Zero%20Trust%20Maturity%20Model_Draft.pdf.

[5] NIST. Special Publication 800-207. *Zero Trust Architecture*, August 2020, https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf.

[6] "Executive Order 14028: Improving the Nation's Cybersecurity," White House, May 12, 2021, https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/.

[7] Office of Management and Budget, *OMB. M-22-09: Moving the U.S. Government Towards Zero Trust Cybersecurity Principles, January 26, 2022,* https://www.whitehouse.gov/wp-content/uploads/2022/01/M-22-09.pdf.

proposes potential solutions for several challenges. This document does not seek to provide a comprehensive discussion of each deployment approach or exhaustively cover every possible edge case. Rather, it seeks to provide a basic toolset that agencies can use to better understand and begin approaching identity management challenges. The motivating use case for this document is based on the need for agencies to authenticate and authorize users and entities to access business applications hosted in the cloud. Agencies should keep in mind that these solutions can also facilitate access to traditional on-premises applications.

1.3 Assumptions and Constraints

This document relies of the following assumptions and constraints:

1. **Vendor Agnostic:** This document is intended to be vendor agnostic. Some terms and phrasing may imply agencies should seek a particular vendor or offering; however, that is not intended, nor is it the goal. Technology and offerings are changing rapidly, and this document tries to balance changes with a lack of standard terminology for services and solutions.
2. **Agency Priorities:** Agencies have different missions, priorities, and resources that could restrict their abilities to implement certain solutions. Therefore, this solutions architecture document is not intended to be prescriptive or set requirements for agencies.
3. **Tradeoff Solutions:** Each solution offered involves tradeoffs; agencies must consider their specific mission needs, priorities, and requirements when determining how best to implement a given solution. Agencies may also find that while a solution is not feasible for them enterprise-wide, it may complement other solutions and provide added functionality and security benefits for individual subsections of the agency.

2 IDENTITY AND ACCESS MANAGEMENT

The following section provides an overview of key components for identity and access management and explains how these are structured in both traditional and modern identity architectures. The section describes authentication and authorization in the context of a hybrid identity model (shown in Figure 1), along with considerations for modernizing identity services.

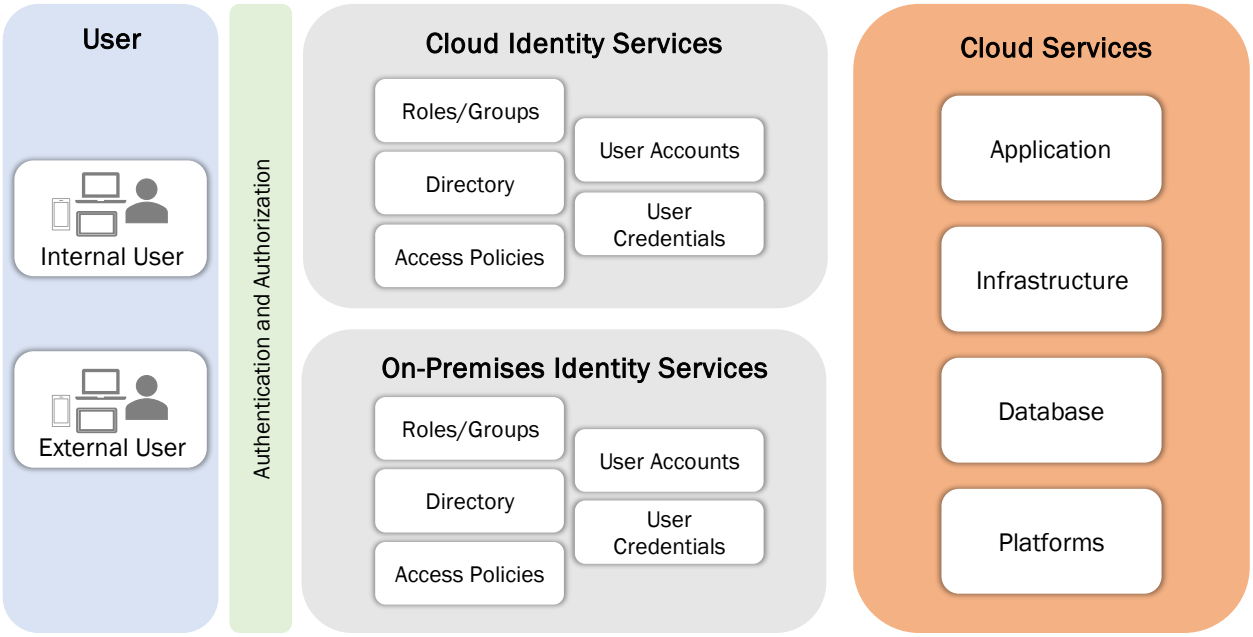

Figure 1: Hybrid Identity Architecture

2.1 Overview

Figure 1 provides a high-level overview of some of the key components of identity and access management involved in the user's access to cloud services. In this document, the term *user* may refer to a human interacting with the system or a non-human entity such as an automated administrative service. User types include:

- **Internal Users:** Users whose identities originate from within a given agency's system through a direct relation such as employment.
- **External Users:** Users whose identities originate from outside a given agency's system; they may be partners, contractors, stakeholders, or members of the public interacting with the agency.

These designations are in relation to the system rather than physical location

Authentication is the process of "verifying the identity of a user, process, or device, often as a prerequisite to allowing access to resources in an information system."[8] **Authorization** is "the process of verifying that a requested action or service is approved for a specific entity."[9] In a hybrid environment, users' identities are federated between on-premises identity services and cloud-based identity services to enable access to cloud services. Identity services include:

- **Roles and Groups:** Labeled collections of users with shared attributes and responsibilities managed by policies set by agencies.

[8] NIST, *Special Publication 800-63-3. Digital Identity Guidelines*, June 2017, https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-63-3.pdf.

[9] NIST, *SP 800-152. A Profile for U.S. Cryptographic Key Management System*, October 2015, https://csrc.nist.gov/publications/detail/sp/800-152/final.

- **Directory:** On-premises or cloud-based hierarchal system for managing objects within a system.
- **Access Policies:** Rules set by administrators that govern user and entity access to system resources.
- **User Accounts:** Collections of information attached to users' identities and that facilitate access to system resources.
- **User Credentials:** An object or data structure that authoritatively binds an identity to at least one authenticator possessed and controlled by a user.[10]

These identity services work together as users are assigned roles and placed into groups. Individual user identities, roles, and groups are managed within a directory, which applies access policies to users and groups. User accounts are tied to their respective identities and, along with user credentials, create the method through which users request and obtain access to systems and resources.

The directory service is a critical component of an enterprise identity architecture, as a compromise of the service can create significant consequences for the organization. This solutions architecture uses the term *Identity Source* to refer to this service. In some situations, however, other components of an agency's identity architecture may also attest to the validity of a given identity and could be exploited to realize similar impacts.

Many different configuration and architectural options facilitate user and entity access to cloud applications. The identity and access management solutions an agency uses should enforce the concept of least privilege, ensuring that the right users are granted just enough access to the right resources, just in time, and for the right purpose. No one identity solution is likely to address an agency's specific needs entirely. Various components must work in tandem to achieve the desired outcomes. Although not addressed in this solutions architecture, these components can include reverse proxies, Secure Access Service Edge (SASE) offerings, Zero Trust Network Access solutions, and host posture assessment capabilities. Agencies also must invest in the requisite training and expertise to implement, maintain, and update solutions over time to keep pace with technological advances, adapt to changing threats, and meet evolving standards/requirements.

2.2 Architectures

Traditional on-premises identity architectures are centered around a directory service, such as Microsoft Active Directory, Red Hat Directory Server, or Oracle Unified Directory. These provide a single view for identities, user accounts, endpoints, applications, and other systems or resources for the enterprise. The architecture may include optional federation tools (e.g., PingFederate or Microsoft's Active Directory Federation Services) as well as other tools that enable single sign-on (SSO) solutions, either via federation or integration with an agent or proxy (e.g., Ping Access for web access), especially for access to software-as-a-service (SaaS) applications in the cloud. In addition, the architecture may have processes for onboarding identities, such as a human resources management system (HRMS), entrance on duty (EOD) system, and others.

On-premises identity architectures traditionally incorporate reusable passwords and SSO service protocols such as Kerberos for access. Some enterprises, particularly in the federal space, rely on a digital certification such as an X.509 certificate on smart cards in lieu of passwords for traditional endpoints. An additional reliance may include a combination of smart cards and one-time password (OTP) systems. An example of an OTP system is a physical security device that can regularly generate access tokens for both on-premises and remote access. The Kerberos protocol is still often used for SSO between applications and services. Current federal guidelines requiring phishing-resistant authentication[11] that will render OTP authentication obsolete for many use cases.

As agencies adopt cloud services, especially cloud-based identity services, they have an opportunity to modernize their identity architecture. Security was not a critical aspect of common traditional directory services, but some recent offerings were developed with guiding security principles in mind. Such offerings typically also allow integration with existing federated identity services and provide modern options for authentication including MFA, SSO, and passwordless options. Agencies must assess their identities and accounts for existing and future needs to decide which will be deployed exclusively in cloud infrastructure, only on-premises, or

---

[10] NIST, *Special Publication 800-63-3: Digital Identity Guidelines,* June 2017, https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-63-3.pdf.

[11] Office of Management and Budget, *OMB. M-22-09. Moving the U.S. Government Towards Zero Trust Cybersecurity Principles,* January 26, 2022, https://www.whitehouse.gov/wp-content/uploads/2022/01/M-22-09.pdf.

synchronized across both cloud and on-premises environments. Section 3 further explores various options to support such implementations.

2.3 Authentication and Authorization

When transitioning to a hybrid environment, an agency's authentication and authorization processes must change to allow access to both on-premises and cloud resources. Figure 2 shows an example instance of a user attempting to access an agency's applications in a hybrid identity model. The user, who may be on-premises (also known as "on-prem") or external, provides their credentials through the agency's cloud identity service. Depending on the agency implementation and the location of the agency's source of identity (i.e., their primary directory service), the user's credentials are used by the cloud identity service and/or on-premises identity services. Typically, authorization policies are enforced in the cloud identity service, while authentication policies may be applied either in the cloud or on-premises. Once the user's identity is validated and their account is authorized, access to agency applications is granted.

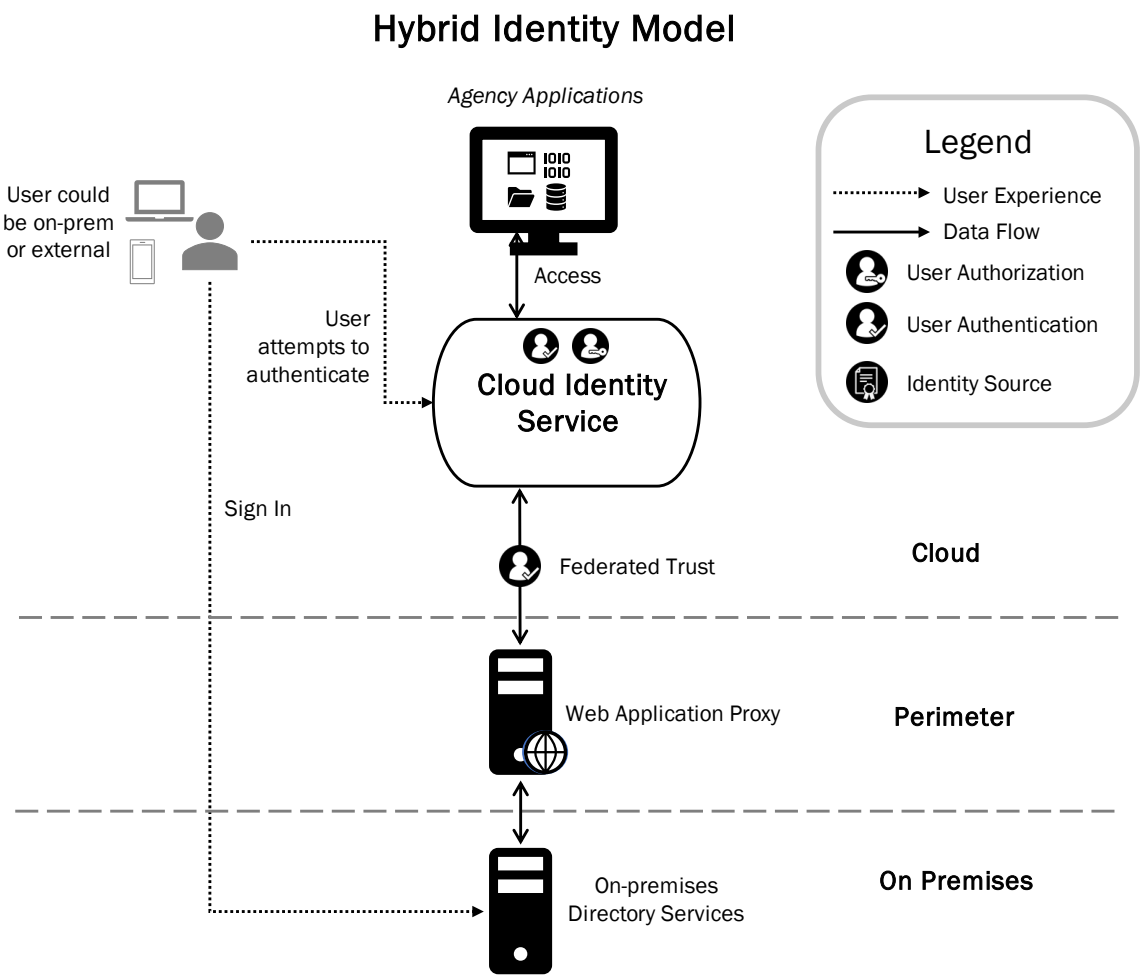

Figure 2: Hybrid Identity Model

Agencies have a range of options for facilitating the authentication and authorization of users in a hybrid identity model. One option is to use one account to facilitate access to both on-premises and cloud resources with the same password through synchronization. In this case, authentication and authorization would take place in the cloud. Another option is for the user to provide credentials to the cloud authentication service, which then passes the information to an on-premises service for validation. In this case, authorization takes place in the cloud and authentication takes place on-premises. Most likely, an agency will adopt a federated authentication model in which authentication generally takes place on-premises, except in certain circumstances, while

authorization takes place in the cloud. Agencies have additional options to shift their identity services to the cloud without needing to federate with on-premises directory services. See Section 3 for further details.

The transition from on-premises to hybrid and cloud-based environments provides agencies the opportunity to leverage more modern authentication and authorization solutions. Agencies should incorporate cloud-based services where possible to leverage modern security features and facilitate a shift away from relying on passwords. Services such as context-based access control allow more granular control over access control decisions, and the more user-friendly SSO both simplifies and strengthens the authentication process. Agencies can conveniently incorporate these solutions during their continuing transitions to hybrid environments with modern, secure protocols; data-driven decisions for access and authorization; and phishing-resistant methods of authentication.

3 HYBRID IDENTITY SOLUTIONS

Agencies are encouraged to review and select the hybrid identity architecture solution that best suits their specific needs and risk tolerances. Technologies in the following subsections can serve as components of an agency's identity architecture and may support overlapping cybersecurity outcomes.

The Cybersecurity and Infrastructure Security Agency (CISA) recommends agencies plan to migrate to cloud-based passwordless authentication via either (1) their existing investments in public key infrastructure (PKI) and Personal Identity Verification (PIV) or Common Access Card (CAC) to authenticate to the identity services, or (2) by leveraging Fast Identity Online Version 2 (FIDO2) and the Web Authentication standard. Agencies may find value from using each option for different use cases or in combination with each other.

3.1 Authentication Options

The following subsections provide introductory information for four different representative models for implementing authentication within a hybrid identity architecture. Each option includes overview information, high-level diagrams, architecture components, and security considerations. While the intention of these sections is not to address all possible edge cases for implementation, the guidance should provide a starting point for agencies to assess their enterprise identity and access management solutions and identify opportunities to modernize via the suggested architecture options.

3.1.1 Federation

Overview

Agencies can use federation services with their on-premises identity services to enable on-premises-based authentication for cloud-based services. Figure 3 introduces this architecture.

When a cloud identity service is federated with an agency's on-premises identity management, the agency's user authentication process can continue to take place on-premises. Since these domains would have pre-established trust, authentication on-premises serves as an acceptable authentication for the cloud service, allowing the user to log on once and access on-premises and cloud-based resources. This authentication setup is beneficial when an agency needs all authentication to be kept on-premises or does not want to transmit their related authentication policies externally. All authentication transactions are handled in one place, ensuring that each account only needs one record and facilitating central policy management and logging for all authentication attempts. Other configurations require hosting logon data redundantly to authenticate users' access from any platform, which also disperses authentication logs.

While there are security benefits to handling all authentication centrally, agencies should consider the potential impacts of increased latency times for their remote users, the ongoing sustainment costs of operating on-premises identity services, and especially the security posture management of their on-premises identity services (directory, federation server, etc.), all of which must be managed as highly sensitive assets. Agencies should also be aware that they can transition from a federation approach to a cloud primary authentication approach over time (see Section 3.1.4).

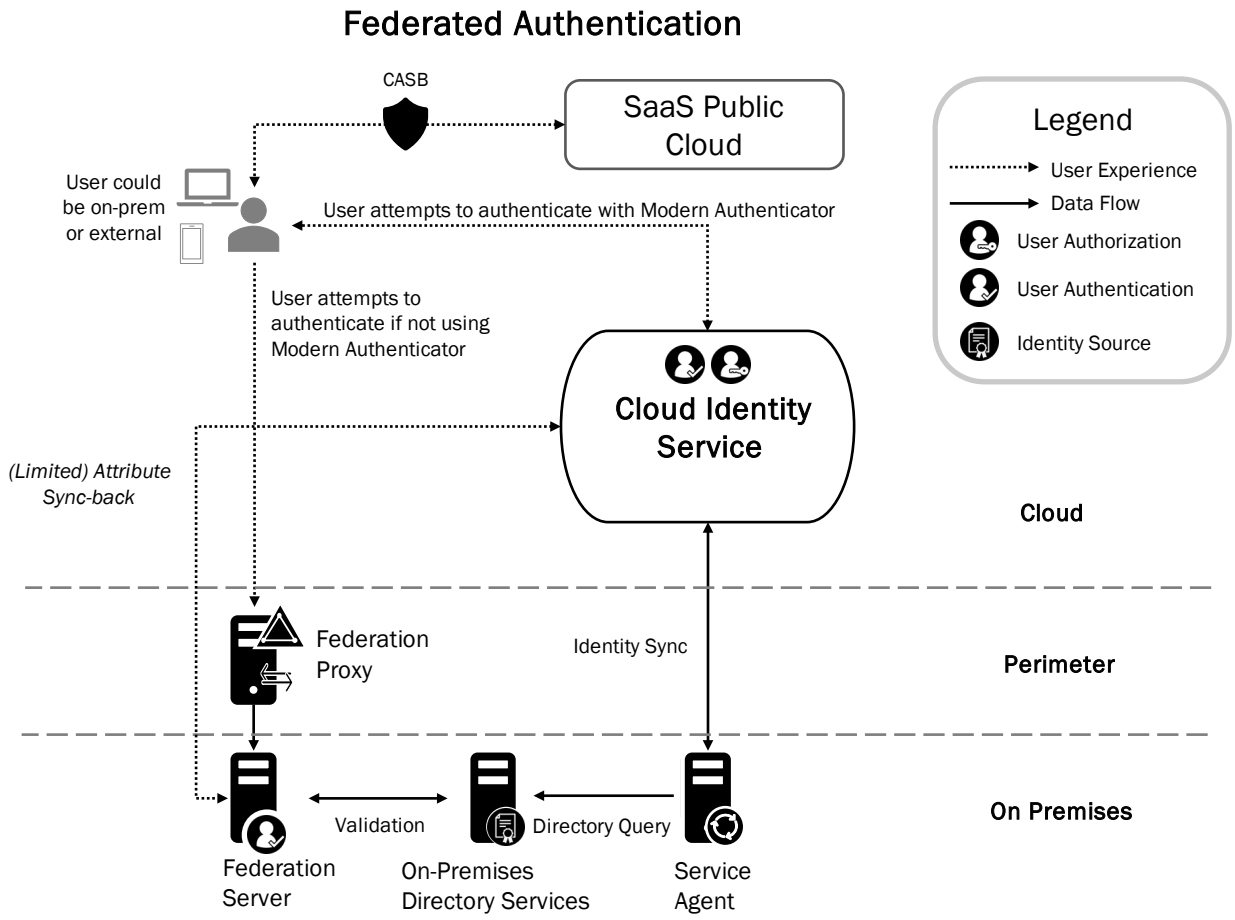

Figure 3: Federated Authentication

Solution Architecture

User Attempts to Authenticate: The user attempts to authenticate via the cloud identity service or, if they are not using an authenticator natively supported by the cloud service (e.g., authentication application, biometric, or other device that provides security guarantees), via a federation proxy. Some implementations, however, may still route traffic through the federation proxy even when using modern authenticators. In either case, this begins the user validation process.

Public Cloud Access: A user's access to the SaaS public cloud is governed by additional protections that provide visibility and enforcement for agency security policies. Agencies have flexibility in determining what solutions will best fit their security needs. Such protections should include capabilities common to a cloud access security broker (CASB) such as a reverse proxy, as well as logging mechanisms, malware protection, etc.

Identity Sync: If authenticated by the cloud identity service, the user's identity is then passed on to an on-premises service agent to check the user's authentication against the on-premises directory service and sync their identity with the cloud.

Federation: A user authenticates via a federation proxy that then validates their credentials with the federation server. The user is authenticated by communicating with the on-premises directory service, logging the user into the resource requested. The federation server can be configured to synchronize select attributes (i.e., as much or as little information pertaining to the authentication request as agencies choose to provide) to the cloud identity service to enable further interoperability of services.

Security Considerations

In this setup, the source of identity is the on-premises directory service. If this service is compromised, an unauthorized entity could pivot and gain access to the cloud identity service. In this way, agencies should

understand that while their on-premises and cloud-based identity services may be architected and secured in different ways, well-known exploits for traditional on-premises-based directory services pose risks that allow attackers to bypass additional security layers protecting an agency's cloud-based identity services in this model.

The federated authentication solution can enable centralized logging of all authentication attempts, improving the accessibility of logon events during security investigations. However, by centralizing all authentication, the increased latency may impact the productivity of remote users that must connect to the on-premises services every time authentication is required.

When deploying this configuration, CISA recommends keeping privileged accounts in the cloud identity service and SaaS public cloud applications as "cloud-only" accounts that do not trust the on-premises federations. This limits lateral movement from a compromised on-premises environment to the cloud to accounts with lower levels of privileges.

3.1.2 Pass-through Authentication

Overview

An alternative model, described in Figure 4, is pass-through authentication to on-premises identity services. In this model, agencies maintain their on-premises identity services and configure cloud services to leverage these resources when users attempt to authenticate. Typically, an additional service (i.e., an agent), is stood up within the agency's on-premises environment to validate these users directly with the cloud service. This allows the agency to keep authentication on premises and enforce established security policies.

Multiple options to configure and implement this model of authentication exist depending on chosen provider and organization-specific needs. For example, an agency may need to conform to security requirements that limit or prohibit installing software directly on the on-premises directory services (e.g., the domain controller). In this case, agencies are more likely to rely on their cloud service provider for an additional service or offering that can handle the authentication request and interface with the on-premises directory services. Regardless, authentication happens on-premises while authorization takes place in the cloud.

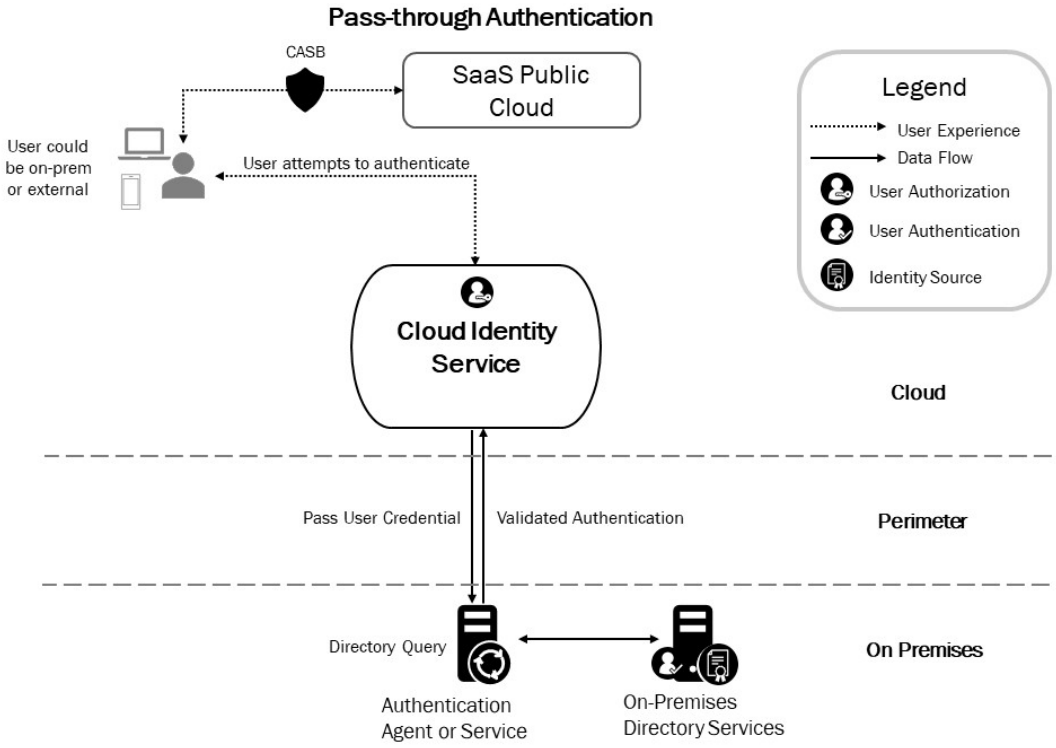

Figure 4: Pass-through Authentication

Solution Architecture

User Attempts to Authenticate: The user attempts to authenticate via the cloud identity service, which begins the process to validate the user.

Public Cloud Access: A user's access to the SaaS public cloud is governed by additional protections that provide visibility and enforcement for agency security policies. As in the federation model, agencies have flexibility in determining what solutions will best fit their security needs; however, they should include capabilities typically part of a CASB, such as a reverse proxy, as well as logging mechanisms, malware protection, and more.

Pass User Credential and Validated Authentication: The cloud identity service passes the user's credentials to the on-premises authentication agent or service where the user is validated against the on-premises directory service. The validated authentication is then returned to the cloud identity service, at which point the user can be authorized for access to resources based on their account information, policies, and other information.

Guidance and Implementation

Typically, this architecture is implemented via an agent, e.g., an agent installed directly on the on-premises directory services or on an alternative component with access to validate against the directory.[12] The agent can establish a connection out to the cloud identity service and operate in a listening or receiving mode. In either case, when an agency user attempts to authenticate via the cloud service provider (CSP), their supplied credentials are passed through to the on-premises identity services. The authentication request is then processed by the agent directly or by a stand-alone service that passes the request to the directory services. The response is then sent back to the cloud identity service.[13] The provider can then complete any authorization

[12] Agencies should note that there is some variation in the intended meaning of *agent* and *agent-less* across providers.

[13] When using an agent, this response can be sent via the established connection, thus reducing the need to open new connections through a firewall.

steps (e.g., checking roles, privileges, groups) in the cloud. Agencies should also add cloud-based MFA after verifying the user via pass-through authentication (see Section 3.2 for further details on security and policy considerations for appropriate MFA technologies) to limit the potential impact of an on-premises compromise to spread.

Security Considerations

Like the federation model, the pass-through authentication model also uses the on-premises directory service as the source of identity. Both models open agencies to the same risk: Exploitation of their on-premises identity services could lead to unauthorized access of their cloud services. Depending on the configuration, agencies may be able to mitigate some of this risk by authenticating users through the on-premises service and then enforcing MFA through the cloud-based identity service.

The requirement to deploy an agent or connector with access to the directory also poses reliability implications, as it is in the critical path of user sign in. High-availability architectures are possible with multiple agents that can fail over during outages. This solution also allows for central logging of all authentication attempts (since they all reach the same core on-prem identity service), improving the accessibility of log-on events during security investigations. However, by centralizing all authentication, the increased latency may affect the productivity of remote users that must connect to the on-premises identity services each time authentication is required.

3.1.3 Password Synchronization

Overview

Password synchronization, outlined in Figure 5, is another approach to hybrid identity. This approach allows an agency to maintain one account for each user that can be used to access either cloud-based or on-premises resources. Having only one password to access agency resources is more user-friendly and can reduce help desk calls for forgotten passwords. All user activity, regardless of their location within the network, is attributed to the same account and authentication happens at the closest point to the user. When syncing passwords, the use of hashing mechanisms and encryption is paramount to preserving the confidentiality and integrity of account credentials. Users with an active cloud session are not interrupted after updating their password but must authenticate with the new password the next time authentication is required.

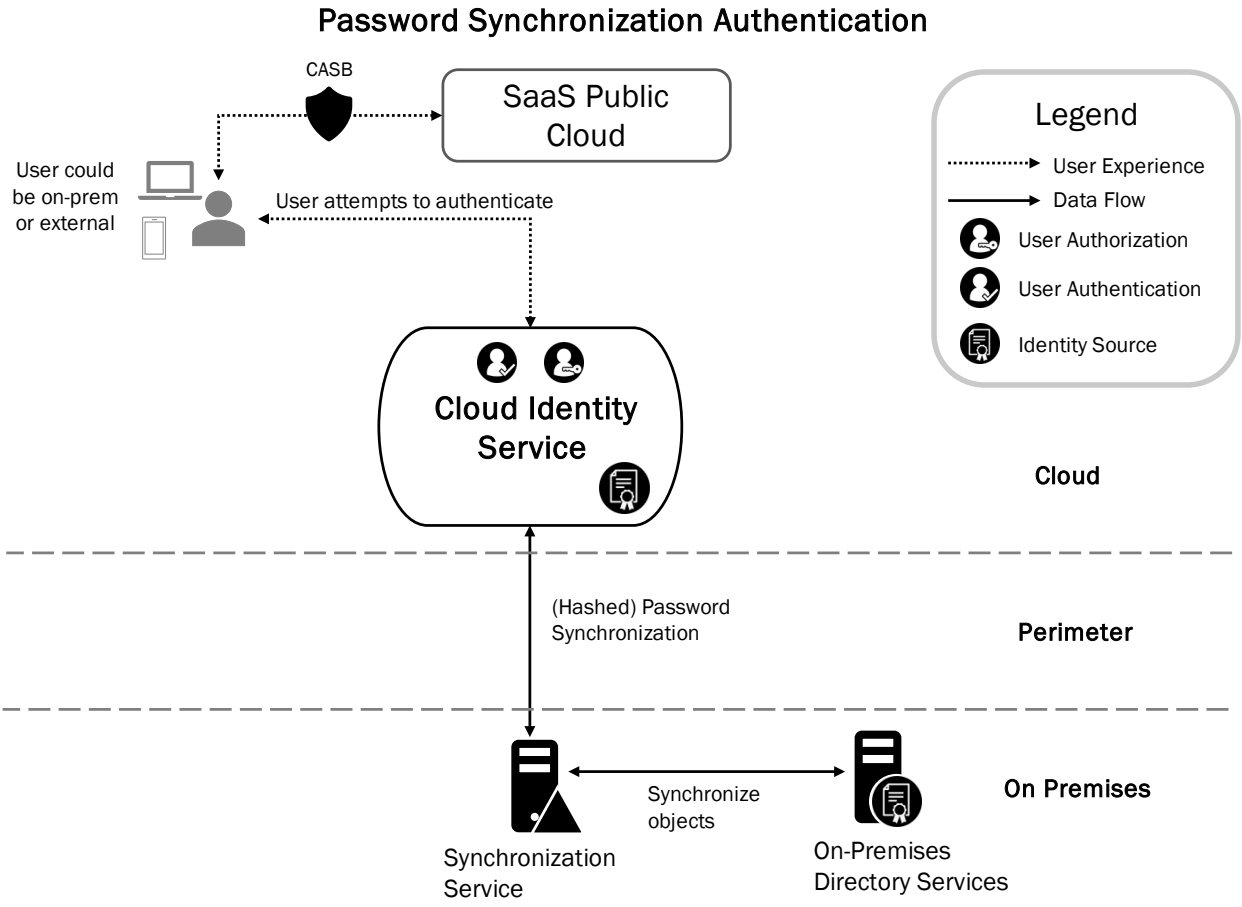

Figure 5: Password Synchronization Authentication

Solution Architecture

User Attempts to Authenticate: The user attempts to authenticate via the cloud identity service, which performs authentication and authorization directly for the user.

Public Cloud Access: A user's access to the SaaS public cloud is governed by additional protections that provide visibility and enforcement for agency security policies. As in each of the previous models, agencies have flexibility in determining what solutions will best fit their security needs; however, they should include capabilities typically included in a CASB such as a reverse proxy and logging mechanisms, malware protection, and more.

Password Syncing: The cloud identity service may synchronize either a hash[14] of the user's password or the password directly with the on-premises authentication agent. This enables the cloud identity service to handle both user authentication and the authorization.

Guidance and Implementation

For this approach, agencies would stand up a service to regularly synchronize either their users' passwords or hashes of their users' passwords stored within their on-premises directory service with a cloud-based identity service. Agencies should note that, depending on their implementation, this synchronization can go in either direction (i.e., they can enable password changes initiated on premises or in the cloud). Agencies should then configure their cloud-based services to leverage this synchronized information when users attempt to authenticate and access cloud resources. This requires agencies to synchronize (hashed) passwords frequently to ensure they are up to date and users have the access they need. Agencies should ensure their cloud-based

[14] Some providers may implement multiple hashes.

identity services are configured to meet their security requirements and that they are consistent with their on-premises policies (e.g., password strength). Agencies should not assume that the policies they have implemented on premises will directly transfer to the policies of their cloud-based solution; instead, they should validate polices that can transfer, and can address gaps where they cannot transfer.

Security Considerations

In this model, the source of identity is the on-premises directory service. Because password synchronization uses the resources and features of the cloud identity service for authentication, attacks leveraged against these mechanisms, such as brute forcing or password spraying, are handled by the cloud service provider's load balancers and resources rather than an agency's on-premises capabilities. Cloud identity providers may be better prepared than agencies to handle such events.

When implementing password synchronization, agencies should be aware that an incompatibility of hash functions could result in a weaker hashing algorithm being selected for operations because the capabilities of on-premises and cloud hashing services differ. Agencies should determine what hashing algorithms are natively available in their configuration and ensure that the most secure algorithms, as applicable, are implemented both on-premises and in their cloud services.

3.1.4 Cloud Primary Authentication

Overview

For the purposes of this document, *cloud primary authentication* refers to a hybrid identity architecture that enables users to authenticate via a cloud-based identity service to access on-premises and cloud-based applications. Figure 6 depicts this architecture. This is distinct from a *cloud native* architecture, in which user authentication is also performed via a cloud-based identity service but only for access to cloud-based applications. In both architectures, a cloud service (or potentially multiple services) handles all aspects of authentication without relying on an on-premises identity service, but the former supports a hybrid enterprise environment while the latter supports cloud environments only.

Agencies have many options to consider when transitioning to a cloud primary authentication model (shown in Figure 6), due in part to the number of service providers, continually expanding offerings, an agency's evolving needs, and more. Therefore, the following section presents a general cloud primary authentication architecture that leverages modern authenticators, supports a passwordless authentication mechanism, and allows for access to on-premises applications.

This approach is intended to be applicable for most agencies. Agencies should leverage existing resources and infrastructure to support such a transition over time; for example, they could gradually configure more of their systems, applications, processes, users, devices, and other resources to leverage their cloud primary authentication architecture. In this vein, agencies that have adopted one of the previously introduced authentication models can use the cloud primary authentication model initially as either a new primary or secondary authentication factor to facilitate their migration. For agencies leveraging a federation model, this cloud primary authentication approach is a logical evolution of their existing hybrid identity architecture. Specifically, it captures the transition from legacy authenticators (which are validated by the on-prem identity services) to modern authenticators (which are validated in the cloud) and the shift of the agencies' identity source[15] to the cloud.

The amount of planning, resources, and effort needed to fully adopt this architecture can be significant. Integrating a cloud identity service with an enterprise for select functions is a different level of effort than transitioning a majority of enterprise services to seamlessly interact with cloud services for identity and access management needs. Migrating applications to the cloud, incorporating modern authentication protocols, and updating operational services (e.g., ticketing systems, antivirus) for this architecture will place more reliance on

[15] For these needs, most agencies in this model are not expected to entirely remove their on-premises identity services, including their identity source. Instead, on-premises identity services play an increasingly secondary role focused on supporting on-premises heavy use cases. Users authenticate directly to Cloud Identity services to access cloud resources and some more modern on-prem resources. Ultimately, agencies' on-premises identity services can act as a relying party to the Cloud Identity services.

cloud identity services. This will take time. However, this is a mature hybrid identity model that agencies should continue to strive to achieve.

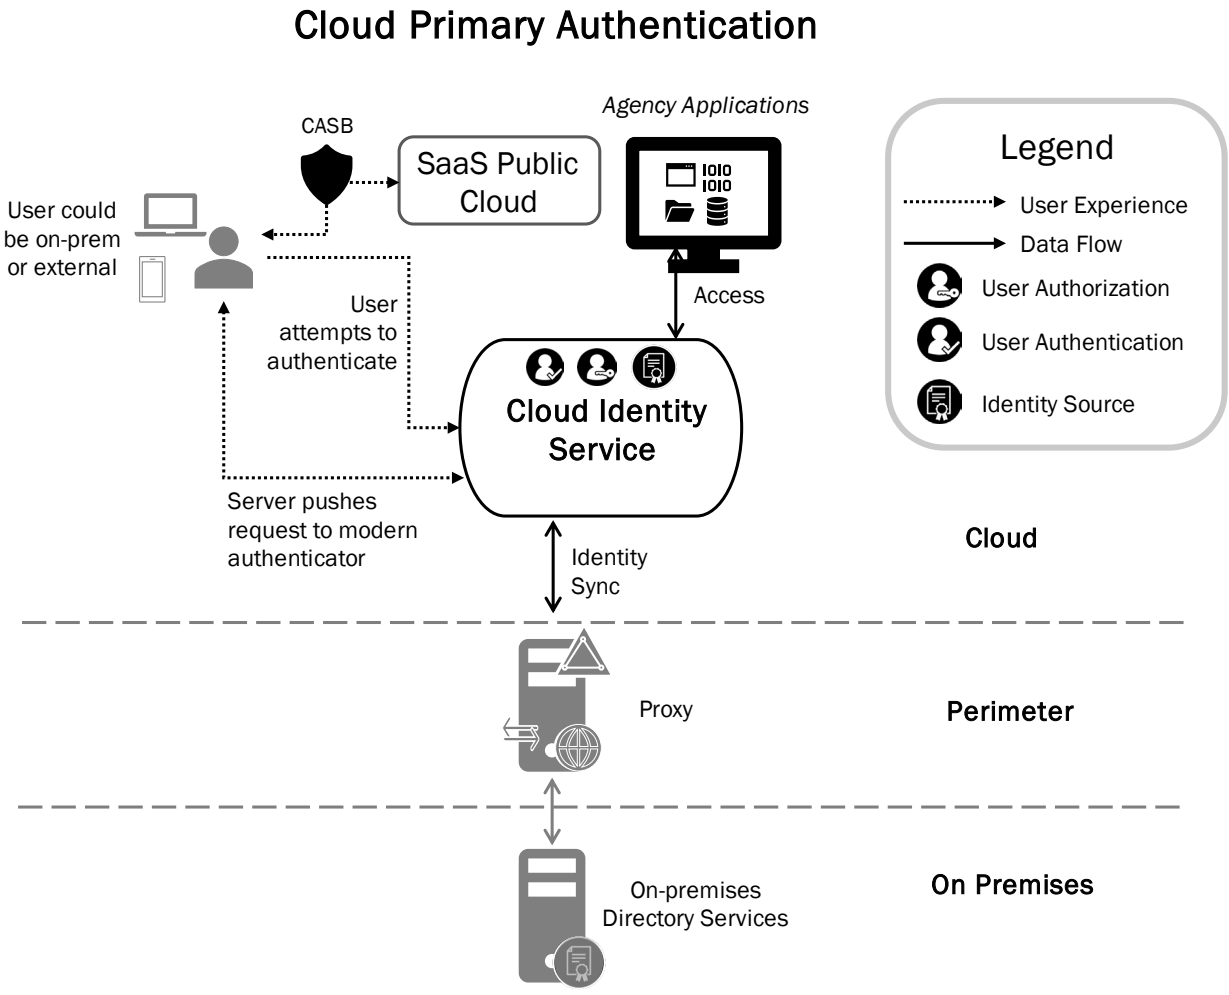

Figure 6: Cloud Primary Authentication

Solution Architecture

User Attempts to Authenticate: The user attempts to authenticate via the cloud identity service, which then begins the process to validate the user via a modern authenticator. In this model, the identity source that governs this login attempt is hosted exclusively in the cloud identity service and leverages a passwordless authentication process (e.g., based on PIV or FIDO2).

Public Cloud Access: A user's access to the SaaS public cloud is governed by additional protections that provide visibility and enforcement for agency security policies. Agencies have flexibility in determining what solutions will best fit their security needs. Such protections should include capabilities common to a CASB, such as a reverse or forward proxy as well as logging mechanisms, malware protection, and more.

Agency Applications: While agencies can configure the Cloud Identity service in previous authentication models to broker access for their on-premises applications, this is purposefully highlighted here because agencies have multiple options for integrating such access, particularly for applications that are incompatible with modern authentication solutions. For example, depending on the agency's current infrastructure, an agency may choose to implement a delegation proxy that facilitates traditional password-constrained access or may configure an on-premises identity service to federate access with the Cloud Identity service via open authentication standards.

Identity Sync: Agencies will likely continue to use existing on-premises services for onboarding new identities to their enterprise. They have different options for synchronizing the associated accounts to the cloud. While most of the agency's identity and access management needs are addressed via the Cloud Identity service, this model,

as mentioned in Agency Applications (above), still allows for the needs of legacy resources and other exceptional use cases to be addressed by on-premises identity services.

Security Considerations

Although this architecture can support traditional alphanumeric password policies, the intended implementation stores credentials in the cloud and never provides the opportunity for users to set passwords for their accounts. Instead, this approach leverages public key encryption, which stores each user's public key in the cloud and private key locally (e.g., on a mobile device, a PIV card, or a security token). Depending on an agency's preferences, this can be combined with additional factors for authentication to support passwordless MFA.

Agencies should carefully assess their options for modern authenticators prior to adoption. Authenticators allow users to securely create credentials and generate assertions when prompted. They are typically distinguished as platform authenticators (i.e., built into a client device), or roaming authenticators (i.e., external to a client device). See sections 3.2 and 3.4 for additional details. By leveraging modern authenticators and open standards for authentication, agencies can more seamlessly benefit from security, monitoring, and other capabilities integrated within their cloud identity services, such as endpoint detection and edge computing.

3.1.5 Authentication Comparison

The four architectures detailed in the preceding sections provide agencies with different options for handling authentication in hybrid environments, some of which will meet their operational and security needs better than others. Agencies should consider where they will maintain their identity source and the accompanying challenges in their selected architecture. For example, agencies should consider how challenges for remote users in the federation and pass-through authentication architectures could impact productivity for users not on-premises. Similarly, if an agency implements password synchronization, they must check the compatibility of hashing algorithms between the on-premises and cloud services to ensure the most secure algorithm is used for all password hashing. The cloud primary authentication architecture will require agencies to address possible challenges associated with the suitability of desired authenticators potentially conflicting with security constraints of select users and/or incompatibility with legacy applications. Table 1 highlights a comparison between these different approaches.

Table 1: Architecture High-Level Comparison

Authentication Architecture	Identity Source	Data Flow	Benefits	Challenges
Federation	On-premises Directory Services	User credentials are passed through the cloud identity service or a federated server to the on-premises identity service for authentication.	Supports legacy authorization and authentication methods and enables a smooth transition between each. Supports passwordless authentication via PIV/CAC.	Latency for remote users Complexity of the architecture can lead to increased operation and maintenance costs. High operational cost to maintain Poor user experience for cloud (including non-modern authenticator and repeated authentication steps).
Pass-through Authentication	On-premises Directory Services	User credentials are passed through the cloud identity service to the on-premises identity service for authentication.	Improved user experience. Enhanced SSO integration.	Latency for remote users. May require the installation of an agent on domain controllers. Reliance on passwords.
Password Synchronization	On-premises Directory Services and/or Cloud-based Identity Services	User credentials are authenticated either at the cloud identity service or the on-premises identity service, depending on the resource's location.	Improved user experience. Enhanced SSO integration.	Authentication logs (or "event" objects) can be stored in multiple locations and systems. Potential for hashing algorithm incompatibility. Susceptibility to brute force and other attacks. May require the installation of an agent on domain controllers. Reliance on passwords.
Cloud Primary Authentication (Passwordless)	Cloud-based Identity Services	User is authenticated by the cloud identity service.	Superior user experience. Full benefit of cloud services. Can be implemented in a federated or non-federated model.	Reliance upon modern equipment/poor support for legacy applications. Hardware asset management may change in complexity. Collaboration with external parties may be complicated due to current lack of standards.

Agencies should begin to shift their primary source of identity from on-premises to cloud-based services along with their authentication and authorization processes (see Figure 7). As mentioned previously, modern cloud-based offerings feature security-focused designs that can help agencies better manage their cybersecurity risk. This transition can facilitate zero trust adoption and an enhanced cybersecurity posture more broadly by readily supporting the cloud-based technology needs of agency users, enabling phishing-resistant MFA solutions (see Section 3.2), providing options for more granular logging, and supporting passwordless authentication solutions across the enterprise (see Section 3.4). Agencies must balance their specific mission needs and risk tolerances (e.g., their visibility into underlying cloud infrastructure and potential concentration of risk within a third-party service provider) in implementing and modernizing their authentication architectures.

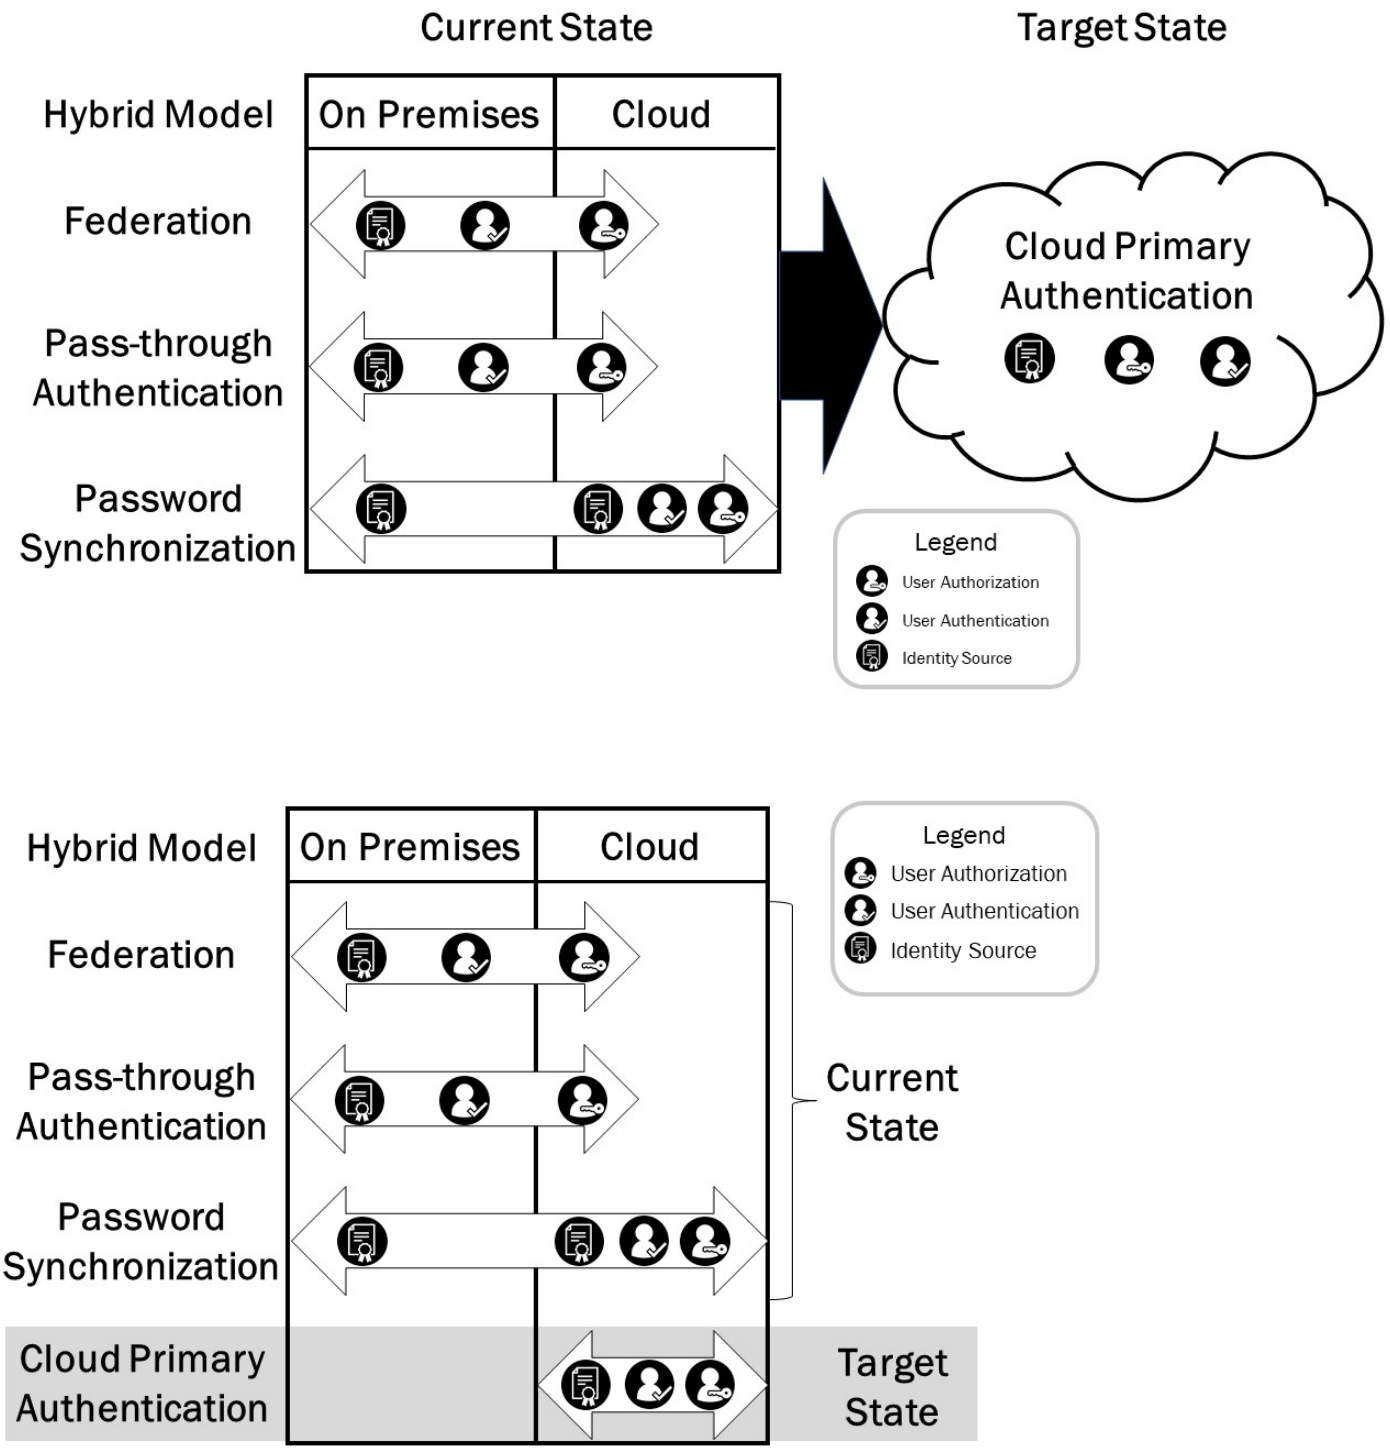

Figure 7: Migration of Identity Architecture to Target State

3.2 Multifactor Authentication Options

Overview

MFA is an important tool for preventing unauthorized access to federal systems, data, and other resources. Per OMB M-22-09 and EO 14028, agencies must implement a minimum of two-factor authentication (MFA with a minimum of two factors) whenever possible, and each authentication method must be phishing-resistant.[16] MFA should be implemented to support a zero trust architecture.

Types of Multifactor Authentication

There are several different types of MFA.[17] The most common are something a user knows (knowledge, e.g., a password or a security question); something a user has (possession, e.g., PIV card, FIDO security key, or smart phone); and something a user is (inherence, e.g., fingerprint, iris, or other biometrics). In some cases, a single authentication (e.g., a PIV card or FIDO security key) may incorporate multiple factors using a local gesture (e.g., touching a button, entering a personal identification number (PIN), or providing a biometric[18]). In other cases, each factor is a separate authenticator presented to the authentication systems.

Agencies should strive to achieve a balance of security and accessibility when implementing MFA solutions. Solutions that are overly burdensome can lead to users avoiding systems or administrators creating exceptions to circumvent MFA enforcement. Agencies should also consider implementing more than one option for MFA based on security and business needs. For example, implementing an MFA solution that requires the use of a mobile device will cause issues for employees who are required to work in secured spaces that do not allow mobile devices. Additionally, agencies may not want to distribute, track, and manage mobile devices to external users of their systems and may choose to use one type of MFA for certain users and another type for other users.

Know (Knowledge)

Knowledge is one of the most common forms of authentication, but it is often the least secure option. The most common example of knowledge-based authentication is a password. Strong passwords are difficult for humans to remember, and computers are particularly good at figuring out weak passwords. Additionally, many password policies are not only complex but also applied differently for different systems. For example, policies can vary greatly in how long a password can be, how complex the password must be (upper case, lower case, special characters, repeatability of characters within passwords, etc.), or require a minimal number of unique passwords before a password can be reused. Due to the difficulty in remembering complex passwords and the constraints of password policies, users are often motivated to reuse passwords directly or with minimal modification where possible. When these passwords are compromised, adversaries can gain access to other services and accounts where that password is used. Passwords are also particularly vulnerable to phishing, and, as such, agencies are highly encouraged to minimize the use of passwords in favor of passwordless authentication mechanisms. Such mechanisms often rely on a PIN (or biometric gesture) as a knowledge factor, usually employed to unlock access to a cryptographic key. In cases where agencies must rely on passwords as a factor for authentication, one option to limit password re-use is to use password managers (see Section 3.5 for more details).

Have (Possession)

Possession types vary the most in implementation and security. Two examples that are less secure are Short Messaging Service (SMS, i.e., texting) OTP authentication and an authenticator application on a mobile device.

SMS OTP is sent through an unencrypted channel; therefore, it is vulnerable to a man-in-the-middle (MITM) attack. SMS OTP also relies on the device's security; if an adversary gains access to the phone or convinces the internet service provider to transfer the number, the adversary will have access to the SMS OTP codes.

---

[16] Office of Management and Budget, *OMB. M-22-09. Moving the U.S. Government Towards Zero Trust Cybersecurity Principles,* January 26, 2022, https://www.whitehouse.gov/wp-content/uploads/2022/01/M-22-09.pdf.

[17] NIST, *Special Publication 800-63-3: Digital Identity Guidelines,* June 2017, https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-63-3.pdf.

[18] "Client to Authenticator Protocol. Proposed Standard June 15, 2021," FIDO Alliance, https://fidoalliance.org/specs/fido-v2.1-ps-20210615/fido-client-to-authenticator-protocol-v2.1-ps-20210615.html.

Authentication applications generate verification codes on a periodic basis, usually within a minute or less. These verification codes use a shared key and time to ensure that an adversary cannot generate the passwords without the key. Authenticator applications may be vulnerable to adversaries who have access to the phone via malware or have physically gained access to a mobile device. Agencies should use an authenticator application that has built-in security measures, such as a PIN or biometric gesture, to access the authenticator application. Both SMS OTP and authenticator applications generate short numbers (usually around six digits) that are easy for users to type in. Due to the short duration of valid codes from an authenticator application, phishing the code is typically more difficult to do than if using an SMS OTP, although it is still possible using real-time MITM proxies. Such attacks are becoming an increasingly common adversary technique.

When choosing a possession-based authentication method, agencies should use the strongest form possible, such as the FIDO2 protocol (see Section 3.4 for more information) or PKI-based authenticators such as PIVs/CACs. Using a PIN to access an authenticator application combines something a user knows with an additional factor of something a user has, which in this case is the authenticator application on the mobile device. Other examples include PIVs/CACs, USB-based dongles, and hardware token cards that require a PIN to unlock the authenticator. Further still, identity service providers may provide custom options for authentication factors not described here.

Are (Inherence)

Biometrics are another form of authentication. Some of the more common types of biometrics include fingerprints, facial scans, and voice recognition. Biometrics rely upon the uniqueness of the body part being scanned and the accuracy of the sensor performing the scanning/detection. Biometrics can either be used locally within an authenticator (e.g., to unlock a cryptographic key) or remotely where a biometric template is stored on the server. Privacy and security properties vary depending on the implementation. Biometrics can be one of the easiest forms of authentication for users to employ.

Biometric authentication can be phishing-resistant when used locally to unlock a cryptographic key. However, depending on the quality of the scanner and algorithm, impersonating someone's biometric data is possible. Another issue with biometrics is the inability to reset or disable a human's biometrics. If there is a vulnerability with the detection scanner or algorithm, then the biometrics portion of the associated MFA solution would need to be turned off to avoid exploitation.

Agencies should be selective in the quality of biometric sensors and algorithms that are implemented. Some biometrics like voice recognition should be avoided as they are not secure due to the potential for impersonation.

Summary

MFA is a powerful tool for increasing security. Phishing attacks can be fully automated to operate inexpensively at scale and obtain passwords, one-time codes, and other information for access. Properly implemented, phishing-resistant MFA can protect federal agencies against various phishing attacks. Additionally, agencies can move into passwordless MFA options, such as FIDO2, to achieve stronger security while easing the authentication experience for their users. Figure 8 provides a snapshot comparison of different MFA options based on the associated security provided.

Federal policy requires agencies move to implement phishing-resistant MFA for employees and contractors. Agencies should have a plan for what form or forms of MFA they should implement, what conditions they should consider for their internal and external users, and what MFA alternatives they might need to provide based on what types of MFA different users require. Agencies also should continue to strive for an MFA implementation that does not place too much of a burden on those using systems and minimizes the need for MFA exemptions the administrators implement. This will require significant commitment and a variety of alternative types of MFA.

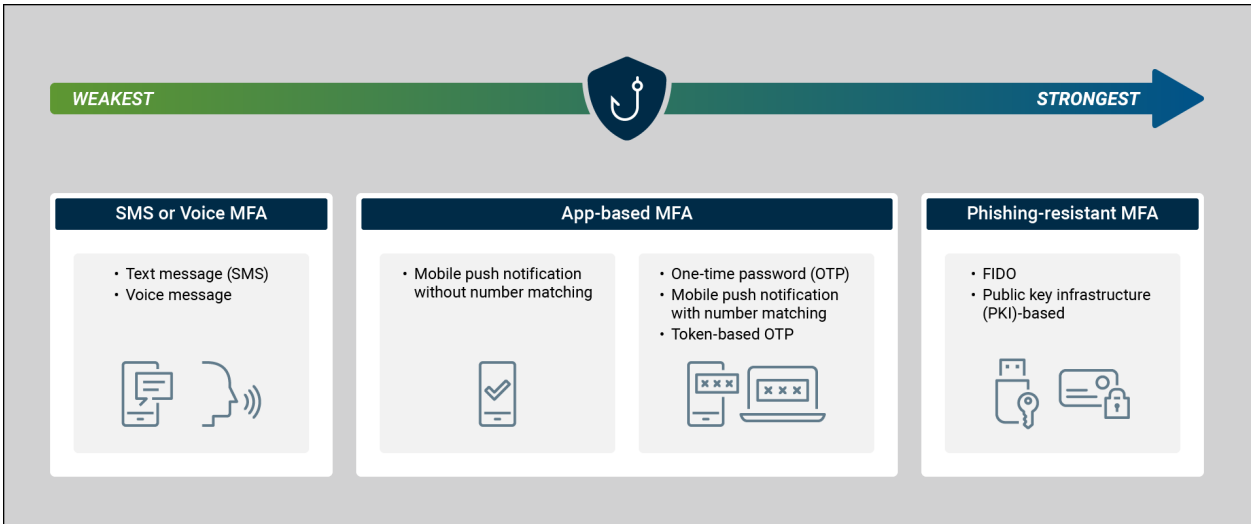

Figure 8: Security Comparison of MFA Options

3.3 Single Sign-on Options

Overview

SSO is a technology that uses federated identity management to authenticate and authorize users across multiple applications on a system by sharing identity attributes. Some identity and access management solutions provide SSO as part of their broader security service. Once the user authenticates to the SSO, it user a password vouching for the user's identity to log them in to authorized services. SSO controls what identity information the system shares with each application.

CISA encourages agencies to implement an SSO with a modern, open standard protocol for increased security and usability as they transition to passwordless authentication. As agencies pursue the opportunity to integrate and leverage cloud-based SSO solutions, they may consult the General Services Administration (GSA) SSO playbook for additional details regarding SSO implementation.[19]

An SSO solution will fit into any of the authentication architectures described above. It is intended to integrate with an agency's identity architecture rather than stand on its own and will, therefore, have dependencies on what the agency enforces for authentication, protocols, MFA, etc. Most SSO implementations use cookies to enable login until the predetermined expiration is reached. These cookies may be specific to session, application, or user and may last longer if the user requests to stay signed in.

Solution Architecture

Figure 9 shows the process of using SSO to manage authentication and authorization between a user, a Cloud Identity service, and multiple SaaS applications. The figure characterizes the process at a high level without detailing the nuances of multiple possible Cloud Identity service configurations to external resources. This document describes some of these nuances with their corresponding components.

[19] "Enterprise Single Sign-On Playbook," General Services Administration (GSA), November 16, 2021, https://playbooks.idmanagement.gov/playbooks/sso/.

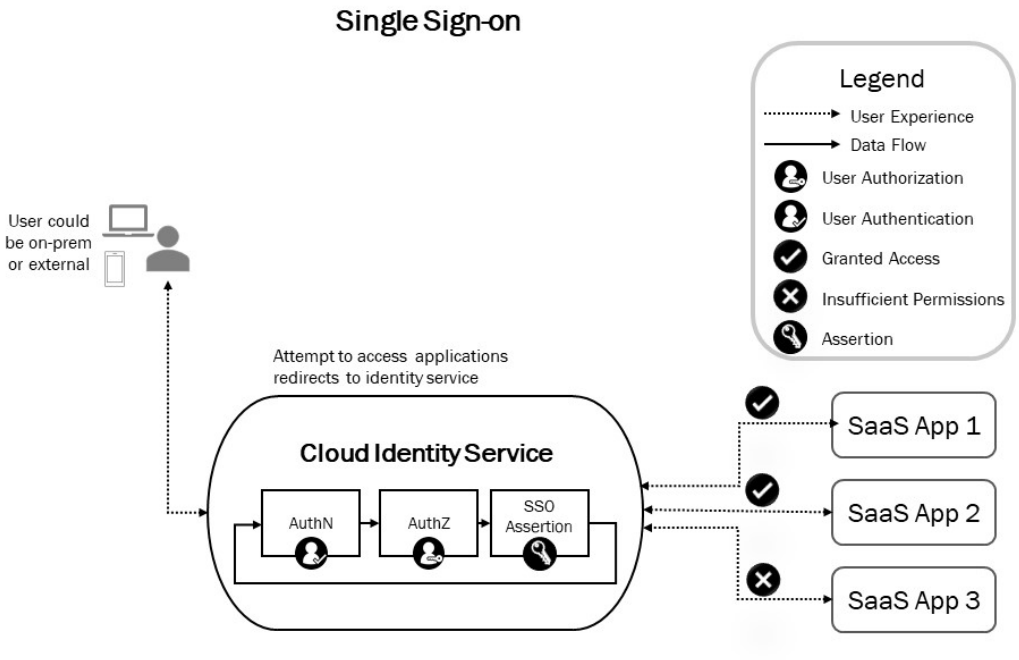

Figure 9: Single Sign-on (SSO) Example

User

A user issues one of two kinds of logon requests: SSO service-initiated or application-initiated.[20] After the user authenticates to the SSO, the user receives proof of their authentication via an assertion that is then passed to an application for access.

- **SSO Service Initiated:** A user visits the SSO or Cloud Identity service's logon portal directly to select an application to access.
- **Application Initiated:** A user visits an SSO-managed application's user portal and is redirected to the SSO or Cloud Identity service to complete authentication and authorization checks.

Cloud Identity Service

The Cloud Identity service may have a built-in SSO that handles all authentication and authorization transactions for application access requests a user makes.[21] These transactions are mostly invisible to the user.

- **Authentication:** The validation of a user's identity using designated authenticators such as passwords, biometrics, OTPs, PINs, etc. The SSO solution may handle the user's authentication request regardless of the authentication approach (e.g., federation, password synchronization, pass through). The SSO or Cloud Identity service provider may also have MFA options to increase security by validating a user's identity through multiple factors. The SSO or Cloud Identity service can apply MFA even when initial authentication is performed on premises.
- **Authorization:** The validation of a user's permission to access an application. The user does not see this part of the transaction since it is handled either by the Cloud Identity service, policy gateway, or the application itself. In some configurations, the Cloud Identity service may request information about user

[20] Terms follow GSA SSO Playbook as defined in Section 4.2 (https://playbooks.idmanagement.gov/playbooks/sso/).

[21] An external SSO solution may also be integrated with the Cloud Identity provider to coordinate authentication and authorization requests before a response is issued to the user.

permissions from an integrated, external authorization provider, such as an on-premises identity store or a context-based access control solution (see Section 3.6).

- **SSO Assertion:** A token, ticket, or assertion,[22] depending on the protocol in use, granted to a user proving their successful authentication. This assertion is provided to any subsequent SSO-managed applications the user attempts to access, removing the need for the user to re-authenticate if they have an active session. The assertion may require the application to check with the SSO or Cloud Identity service to confirm a user's authorization to access the application, but no further action is typically required of the user.

SaaS Applications

If an SSO-managed application receives an access request from an unauthenticated user, it redirects the user to the Cloud Identity service to complete authentication and authorization. When an authenticated user requests access, a user/SSO assertion validating the user's authentication with the Cloud Identity service is sent to the application. If the assertion has expired, or a user's permissions need validation, the application sends a request to the Cloud Identity service for re-authentication or to check a user's authorization.

If a user fails the authorization check, the requested application denies access due to insufficient permissions. This does not affect their ability to access other applications, provided they have appropriate permissions. If a user fails an authentication check, they are unable to access any applications until they complete a subsequent authentication check successfully.

Commonly Used SSO Protocols

SSO can be used with multiple authentication and authorization security standards and protocols. Agencies should select an SSO standard to integrate with the organization's identity management solution and security needs.

- **Security Assertion Markup Language (SAML):** An open authentication standard that uses a specific format to pass authentication information between web applications. It transfers authentication data between identity provider and service provider, called a SAML assertion, allowing authentication across domains. Service providers that set up SAML SSO can avoid storing password data or resetting user passwords by having them authenticate through a trusted identity provider.
- **Open Authorization (OAuth):** An open authorization standard that passes security data between applications without passing user credentials to show user authorization to access an application. The security information passed via OAuth is called an authorization token and is typically sent over HTTPS. Tokens contain information regarding the privileges the user should have on the application and include an expiration.
- **Open ID Connect (OIDC):** An open authentication standard built on top of OAuth 2.0 that can authenticate a user and check their authorization using OAuth 2.0, thereby checking both user identity and permissions before allowing access to a resource. OIDC uses a RESTful HTTP API and JSON format.
- **Kerberos:** An open standard that integrates with Active Directory to authenticate users and give them ticket-granting tickets (TGTs) to access other SSO applications. During the initial authentication process, a user is verified using credentials and then given a TGT for subsequent access. This ticket can be shared with other applications via the operating system, allowing the user to skip the credential verification part of the logon process.
- **Smart Card/PKI:** Smart cards store private keys or X.509 certificates on a physical device used to authenticate users. Smart card authentication often uses Mutual Transport Layer Security (TLS) to provide mutual authentication of server and client certificates. Smart card PINs can be cached to authenticate to a device offline once confirmed by an identity provider and then later used to authenticate multiple applications.

SAML and OIDC are the most-used protocols in enterprise today.

---

[22] Term follows the GSA SSO Playbook as defined in Section 1.1 (https://playbooks.idmanagement.gov/playbooks/sso/). The generalization of the term *assertion* in this document includes tokens, assertions, and tickets.

Considerations

SSO allows users to authenticate to multiple applications with a single set of credentials that are never exposed to compatible applications. SSO can act as a centralized place for controlling account information shared with each application and adding a single MFA checkpoint to access multiple applications. The burden of authenticating a user is removed from the applications and assigned to SSO, a service designed specifically to store and handle user logon data. Password policy enforcement can focus on SSO as a gateway to other applications, increase password complexity requirements because users have fewer passwords to remember, and secure the single point of entry (or move to passwordless authentication). To simplify user access auditing and permissions, SSO can integrate with identity solutions rather than storing user identities itself.

Alternate approaches are used for applications that are not SSO compatible. In some cases, rather than using the open standards discussed above, vendors implement an SSO-like capability by replaying credentials (i.e., passwords) to applications. In this case, the SSO product is essentially acting as a password manager. Considerations around password managers are discussed in Section 3.5. When agency operations are dependent on applications incompatible with SSO, agencies should improve their security posture by modernizing applications to support open authentication standards.

Various implementations of SSO protocols such as SAML or OAuth can and do have vulnerabilities. Robust vulnerability management practices should be adopted where the agency is responsible for patching components of the architecture implementing these protocols. Vulnerabilities can occur either on the SSO provider side or the relying party integrating via SSO. One benefit of cloud-based identity services is that cloud service providers are responsible for patching their own services and usually do so quickly before a vulnerability is publicly disclosed.

SSO can be a single point of failure because one logon is used to access multiple applications. Agencies interested in incorporating SSO or an identity and access management solution with SSO for a hybrid or cloud-based solution should actively mitigate the risks associated with using a single authentication and authorization utility. Requiring evidence of SSO best practices from vendors supplying software that will integrate via SSO can help minimize the risks of a single point of failure and potential violations of the principle of least privileges. Agencies are encouraged to explore compliant vendors' services according to their risk profile with the recommendation to consider options that limit session length to mitigate the risk of compromised user sessions.

Additionally, SSO can be implemented entirely on-premises. For example, in a typical enterprise that uses a federated authentication architecture (see Section 3.1.1), an agency can configure their federation server to provide the SSO directly to their applications. However, for many of the same reasons discussed above, CISA recommends that agencies begin transitioning these solutions to the cloud.

3.4 FIDO2

Overview

Agencies seek both strong authentication solutions as well as easy-to-use and easy-to-implement authentication solutions. FIDO2 is a set of protocols developed in collaboration by the FIDO Alliance and World Wide Web Consortium (W3C). FIDO2 is designed to enable easy, secure, passwordless authentication. FIDO2 is backwards compatible with its predecessor FIDO Universal 2nd Factor (U2F).

Technical Summary of FIDO2

Three parties are involved when FIDO2 is used: an authenticator, a client, and a relying party server. An authenticator is a device or piece of software that provides both encryption key generation management and storage as well as a way for a user to input a gesture. A gesture can be touching a button on a roaming authenticator, biometrics such as a fingerprint reader or facial recognition software, or inputting a PIN submitted from the host, among others. The intent of the gesture is to require a human to be present. Agencies may consider the use of a PIN as a gesture or something resistant to keyboard capture, such as touching a physical device. The authenticator comes in two types: a platform authenticator, which is built into an operating system; and a roaming authenticator, which is a separate physical device (like a PIV card, but usually connected via USB). The client and server within FIDO2 operates in a similar manner to operations in the standard client-server mode (e.g., the client initiates a connection to a dedicated server). In an enterprise context, the server is likely a Cloud Identity service, though the protocols themselves offer compatibility with any type of server and

deployment model (on premises or cloud). W3C's WebAuthn[23] protocol supports this communication between the client and the server. Communication between the authenticator and the client when using a roaming authenticator is facilitated through USB, Near-field Communication (NFC), or Bluetooth Low Energy (BLE) using the FIDO Alliance's Client to Authenticator Protocol (CTAP) protocol.[24] For a generalized overview of how FIDO2 works, see Figure 10. Agencies should be aware that all keys within FIDO2 are unique to each service. The specifications for the relying party server and client/platform shown in the figure are detailed in the W3C's recommendation, while the authenticator should adhere to the proposed standard for CTAP. Agencies must configure their specific implementation so that the client/platform uses the agency's desired protocol to communicate with the authenticator.[25]

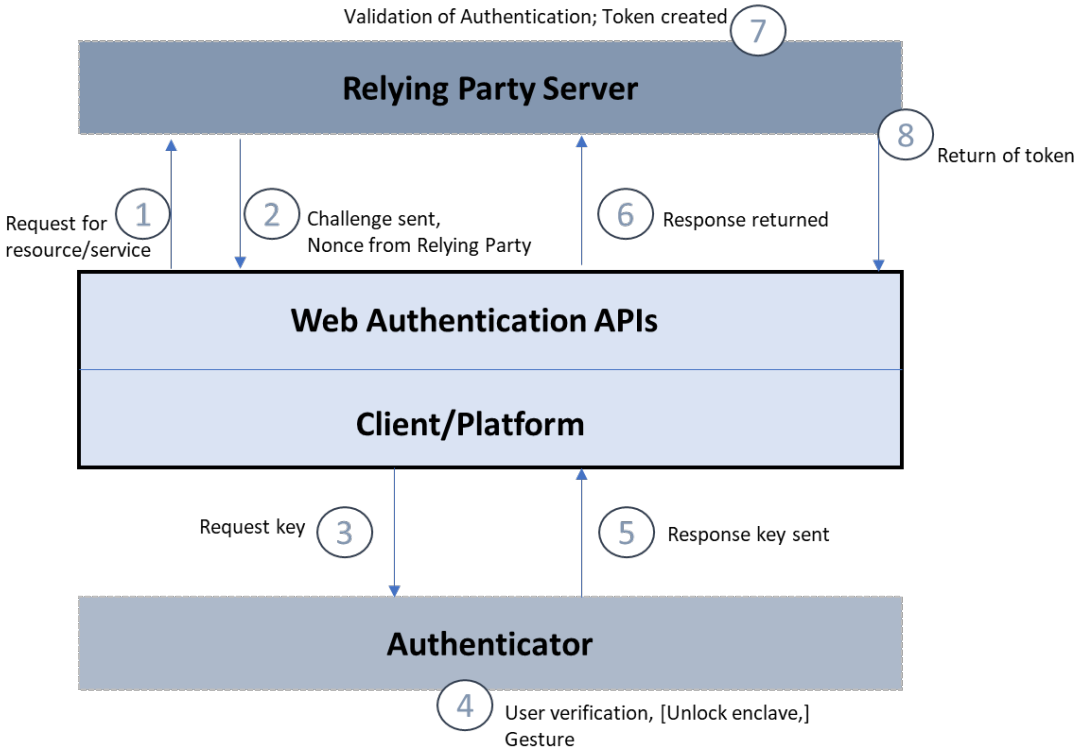

Figure 10: FIDO2 Authentication General Flow

As shown in Figure 10, a sequence of actions takes place during a FIDO2 authentication transaction:

1. Request for resource/service is sent from the client to the relying party server.
2. Challenge sent; nonce from relying party sent back to the client/platform.
3. Response key sent from the client/platform to the authenticator.
4. User verification requires gesture; secure enclave unlocks on the authenticator.
5. Response key sent from authenticator to the client/platform.
6. Response returned from the client/platform to the relying party.

---

[23] "Web Authentication: An API for accessing Public Key Credentials," World Wide Web Consortium, April 8, 2021, https://www.w3.org/TR/2021/REC-webauthn-2-20210408/.

[24] "Client to Authenticator Protocol. Proposed Standard, June 21, 2022," FIDO Alliance, June 21, 2022, https://fidoalliance.org/specs/fido-v2.1-ps-20210615/fido-client-to-authenticator-protocol-v2.1-ps-errata-20220621.html.

[25] Agencies should review offerings to ensure any solution they acquire supports the communication protocols they seek to use (i.e., USB, NFC, and/or BLE).

7. Validation of authentication at the relying party.
8. Return of token from the relying party server to the client/platform.

The Authenticator and Processes

An authenticator is a physical item or a piece of software in the possession of the user that allows them to securely create credentials and generate assertions. A credential is a public key pair generated by an authenticator. The private key is never shared outside the authenticator's secure enclave. During credential registration, the credential ID and public key are shared with the service for which they were created. The two types of authenticators are platform and roaming authenticators:

- **Platform Authenticators:** A type of authenticator built into a client device. This can be implemented with hardware, such as a Trusted Platform Module (TPM) chip, to store the private keys. A gesture can be provided with either a PIN or biometrics, such as fingerprint readers or facial recognition.
- **Roaming Authenticators:** An authenticator external to the client device that can connect to it using protocols such as USB, BLE, and NFC. Usually, keys are encrypted by another key on the roaming authenticator and then stored on the server. This allows the roaming authenticator to theoretically secure an unlimited number of keys.

The processes within FIDO2 include three additional elements:

1. **Registration:** Used when creating a new account or when adding FIDO2 authentication to an existing account. It involves the creation of new keys.
2. **Authentication:** Used when authenticating to an existing account.
3. **Decommissioning:** Used when a relationship between a service (e.g., the Cloud Identity service) and the authentication is to be severed. This can happen either due to a change in the identity (e.g., employee leaves the agency) or due to an authenticator not being used over some period and an automated process triggering its revocation. Regardless of why the decommissioning process starts, it will end with the termination of the given relationship.

Attestation with FIDO2

Attestation is the prosses of establishing to the relying party server certain properties of the authenticator, such as its use of Federal Information Processing Standards (FIPS) 140 validated cryptography, its use of a trusted supply chain, its procurement by a specific agency, or even traceability of issuance to a particular individual at an identity-proofing event. The FIDO alliance asserts: "Compliant FIDO Servers must support all attestation models. Authenticators can choose what attestation model to implement." Agencies should be aware of the various types of attestation, including:

1. **Full Basic Attestation:** In which batches of authenticators of the same model share the same attestation key. This is the most commonly implemented attestation technique in the market today.
2. **Surrogate Basic Attestation:** Uses authenticator to self-sign as surrogate attestation key. Authenticators with minimal protections for an attestation private key can use this model.
3. **Privacy Certificate Authority (CA): "**In this case, the authenticator owns an authenticator-specific (endorsement) key. This key is used to securely communicate with a trusted third party, the Privacy CA. The Authenticator can generate multiple attestation key pairs and asks the Privacy CA to issue an attestation certificate for it."[26]
4. **Direct Anonymous Attestation (DAA):** "DAA credentials are used along with blinding to sign the attestation data."[27]

Due to the nature of work agencies often perform, they may choose to avoid using privacy-preserving mechanisms. The latest version of the FIDO2 standard has added an enterprise attestation capability that will

[26] "FIDO 2.0: Key Attestation Format," FIDO Alliance, September 4, 2015, https://fidoalliance.org/specs/fido-v2.0-ps-20150904/fido-key-attestation-v2.0-ps-20150904.html.

[27] Ibid.

enable increased traceability. When using separate token FIDO2 authenticators, agencies may request unique attestation keys for that agency's batches of devices (or even for each individual device) from the manufacturer. Theoretically, this enables traceability of authenticators issued to personnel. However, such business processes are not currently mature across the public sector and support for these elements of the standards across vendors is currently incomplete.

Authenticator Considerations for Federal Deployment

Agencies can choose any combination of authenticators for their deployment of FIDO2. Agencies should be aware of the particulars of three types of authenticators.

- **Platform Authenticators:** Agencies should be aware that for platform authenticators, authentication is tied to a particular system such as Windows Hello, iOS Face ID, Android fingerprint readers, etc. If that system is upgraded or replaced, the credentials will be lost. Credentials cannot be transferred from one device to another.
- **Roaming Authenticators:** For roaming authenticators, authentication is not tied to a particular system, but to a dedicated device. This makes upgrading and replacing systems independent of the authenticator credentials. This method is more prone to losing or misplacing the authenticator; however, it can be mitigated by requiring registration of two or more authenticators for recovery scenarios or by combining roaming and platform authenticators.[28] Agencies should either have a mix of the two types of authenticators or more than one roaming authenticator to minimize the risk of needing to recover the account.
- **Passkeys:** Passkeys are a recent addition to the FIDO2 ecosystem. Unlike platform authenticators, passkeys are exportable and sharable between devices. [29] While not desirable in the general federal enterprise context, passkeys may be valuable tools for agencies in certain low-assurance situations and other use cases to replace passwords.

FIDO2 Interaction With Existing Infrastructure

Integrating FIDO2 with existing infrastructure happens in two parts: modifying application registration and login pages, and setting up a FIDO server (which may be included in agency Cloud Identity services). Agencies should use open-source developer resources and reference architecture used by other organizations to assist in their implementation of FIDO2[30] and integrate their FIDO2 devices with their primary identity provider either in the cloud or on-premises.

Adhering to Federal Guidelines

CISA assesses that the WebAuthn protocol and FIDO2's roaming authenticators are "phishing resistant" in compliance with OMB's M-22-09. FIDO2 also meets the requirements of National Institute of Standards and Technology (NIST) Special Publication (SP) 800-63 Authenticator Assurance Level (AAL) 2 and may meet the hardware-based key requirements of AAL3, depending on implementation. FIDO2 binds cryptographic signatures to the asserted identity of the relying party server as verified by the client (e.g., the Cloud Identity service's hostname, as verified by the browser), which is subtly different than NIST SP 800-63-3's requirements of *verifier impersonation* resistance. However, OMB M-22-09's adoption of a broader notion of phishing resistance updates federal policy and will be aligned with NIST SP 800-63 in future updates.

Registering FIDO Devices

As discussed above, FIDO2 authenticators are registered to relying parties. In an enterprise context, the relying party to which the FIDO2 authenticator is registered is often the Cloud Identity service, as discussed in Section 3.3, because this is the service that authenticates the user. It could also be an alternative (on-premises or SaaS)

---

[28] FIDO Alliance. *FIDO Alliance White Paper: FIDO Authenticator Lifecycle Management for IT Administrators,* April 2021, https://media.fidoalliance.org/wp-content/uploads/2021/04/FIDO-White-Paper-Lifecycle-Management-for-IT-Administrators.pdf.

[29] For more information, see: https://fidoalliance.org/apple-google-and-microsoft-commit-to-expanded-support-for-fido-standard-to-accelerate-availability-of-passwordless-sign-ins/.

[30] "U.S. General Services Administration's Rollout of FIDO2 on login.gov," FIDO Alliance, March 19, 2020, https://fidoalliance.org/u-s-general-services-administrations-rollout-of-fido2-on-login-gov/.

identity service that federates with the agency's primary Cloud Identity service. The registration process binds the authenticator to an account in the identity services and thus to a valid identity. Note that this registration is conceptually similar in purpose to the registration of an OTP hardware token authenticator, even though the protocols are technically different. This process is unlike the process with PKI-based authenticators such as PIV cards, where the PIV issuer and its certificate authority issue a credential that does not need to be explicitly registered to identity providers, but rather is trusted cryptographically.

FIDO2 authenticators can be integrated into an agency environment in two ways that preserve the key property outlined in FIPS 201[31] of linking the user's identity to their credentials: as a *derived* authenticator and as the *primary* authenticator. Most government usage of FIDO2 today is as a derived authenticator, though in principle both approaches are possible. To use FIDO2 as a derived authenticator, an agency allows end users to register FIDO2 authenticators to the identities maintained in their directory after authenticating via their primary PIV card credential. Agency users can register multiple FIDO2 authenticators of different types, such as roaming authenticators or platform authenticators built into their workstations or mobile phones. All these authenticators can then be used to authenticate to the identity services from different devices. Agencies may define automated processes for removing registrations that have not been used in a defined period (e.g., one year) to unregister unused authenticators. Agencies are encouraged to consider this approach to build experience with FIDO2 technology and address use cases where PIV cards are not suitable as the only authenticator (e.g., access when users leave cards at home, access from devices without smart cards, or access to multiple devices simultaneously).

An alternative is to use FIDO2 authenticators as the primary credential for users and *in lieu of* smart cards. Such a model is not widely deployed today in the government. In such a model, a scheme must be defined to associate the user's initial FIDO2 authenticator with their identity, as established through the PIV process. CISA expects continued developments on these topics during the coming years.

Considerations

Account Recovery

Migrating to an authentication system that incorporates FIDO2 will require minor changes to the account recovery process. Agencies should advise users to register multiple authenticators per account to reduce the need for account recovery.[32] This is not only more convenient for the user but also more secure for the agency. Attackers may use account recovery to circumvent the security benefits of strong MFA. Agencies should treat the risks associated with issuing replacement credentials like the risks associated with issuing initial credentials and should make use of NIST SP 800-63-3.

Additionally, users should have an easy way to report a lost, stolen, or damaged authenticator so it can be deactivated and a new one can be issued.

Security

FIDO2 provides a high degree of resistance to low-skill attacks such as phishing, password stuffing, replay attacks, session hijacking, or MITM attacks. Methods of FIDO2 implementation can vary in security, with roaming authenticators generally being more secure than platform authenticators, particularly if the platform authenticator keys are not handled in a separate cryptographic module. Keeping the authenticator separate from the device with which it pairs helps prevent the private key from being extracted if the physical device falls into the hands of an adversary. Note that any method of FIDO2 is superior to legacy authentication, such as passwords or any other form of authentication that is not resistant to phishing. Platform authenticators are low cost and enable agencies to reduce the issuance of phishable credentials, such as passwords.

Ease of Use

[31] NIST, *FIPS 201-3: Personal Identity Verification (PIV) of Federal Employees and Contractors*, January 2022, https://csrc.nist.gov/publications/detail/fips/201/3/final.

[32] FIDO Alliance, *Recommended Account Recovery Practices for FIDO Relying Parties*, February 2019, https://fidoalliance.org/recommended-account-recovery-practices/.

Proper workforce education is one of the most important steps agencies must take when incorporating FIDO2 into their authentication procedures.[33] This includes user education regarding FIDO's uses and capabilities, how to register a new or pre-existing account, and how to handle account recovery. To better assist users, information technology support employees should also understand the FIDO2 implementation being used by the agency.

Agencies should consider their resources and risk tolerance to decide whether it is best to make FIDO2 authentication mandatory or optional for users. The options an agency provides regarding its implementation will also impact the adoption of FIDO2 by its users. Agencies looking to make FIDO2 authentication mandatory should ensure users have access to roaming authenticators and/or devices that support platform authentication.

Once implemented, FIDO2 is a strong and easy-to-use authentication solution. As is necessary when incorporating other new technologies, adequate planning and education are keys to ensuring a positive user experience and smooth transition.

3.5 Password Managers

Overview

As agencies continue to enhance the security of their enterprise, CISA encourages using SSO and passwordless identity and access management solutions to the extent possible. This will support zero trust migration and reduce reliance on complex password policies.

However, some agencies have applications that are not enterprise managed or are currently unable to perform passwordless authentication. These agencies will find password managers[34] a useful tool to decrease risks associated with weak passwords. Password managers encourage users to use strong passwords by eliminating the need to memorize all passwords. They also provide a secure place to store passwords and can be used to enforce agency password policies. These services can include functionality for storing more types of data than typical passwords (keys, certificates, etc.), but they are referred to as password managers throughout this document.

Storing all passwords for one user in the same location carries inherent risk. A password manager could act as a single point of failure in two ways. First, a breach of the user's primary credentials or vault could expose all saved passwords to unauthorized individuals. Second, if a password manager depends on a database external to a user's device, product availability could be affected by provider outages, keeping users from accessing their accounts. However, features provided by quality password managers, such as using encryption at rest and device-based or cached vaults, can mitigate this single point of failure risk and further increase an agency's security posture while in transition to passwordless solutions.

Implementation

Policy Enforcement and Integration with Identity Services

Password managers designed for enterprises have built-in functionality that can work with internal policies. Some solutions have an internal policy manager that can be updated to reflect corporate policies, while others can be integrated directly with policy-enforcing solutions. Still other password managers can be directly integrated with identity services to support agency security during the transition to passwordless solutions. Agencies using or considering an identity service can check to see if it includes access to a password manager with appropriate features for the agency's risk appetite. Products that do not integrate with existing password policy services may require installation of a tailored product on domain servers or other devices to enforce agency- and product-defined password policies during local password resets.

End-to-End Encryption

---

[33] FIDO Alliance, *FIDO Desktop Authenticator UX Guidelines*, June 2021, https://fidoalliance.org/ux-guidelines/ux-guidelines-desktop-authenticators/.

[34] Although some SSO solutions leverage password managers as an embedded component of their offering, this document treats SSOs and password managers as distinct capabilities that can provide similar outcomes, with the former recommended over the latter.

End-to-end encryption ensures sensitive data, such as user credentials, are never transmitted in plaintext and can only be decrypted by the intended recipient. This prevents attackers from compromising user credentials by intercepting network traffic. Password managers that connect to a vendor's servers or resources external to a user's vault need end-to-end encryption. Using end-to-end encryption paves the way for a zero-knowledge architecture, a model with increased security in which only the user retains the encryption key for their vault.

Zero-knowledge Architecture

Zero-knowledge architecture is a password manager design that ensures the processes external to a user's device (i.e., processes that handle passwords and sensitive data), cannot access a user's primary password or plaintext data. This means the vendor hosting the password manager cannot view a user's unencrypted data because all data stored for the user are encrypted locally and, if the provider offers device syncing or cloud storage, are then securely transmitted to the provider's database. Since the service provider does not keep a record of a user's primary password or encryption key, account recovery can be more difficult if the primary password is forgotten. However, these same processes keep the vault private and increase resiliency against unauthorized access, providing better security for the user's stored credentials. Figure 11 highlights an example implementation of a zero-knowledge architecture designed password manager.

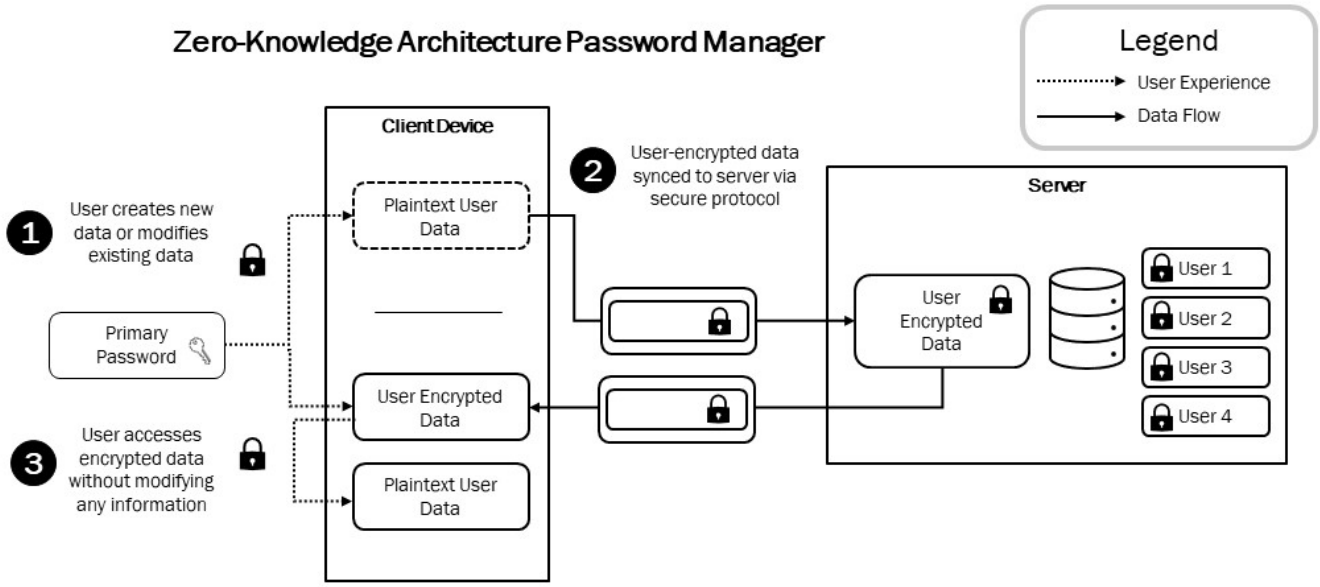

Figure 11: Zero-knowledge Architecture Password Manager

Types of Password Managers

Various types of password managers are available for consideration based on an agency's needs and risk appetite. Each type has different features and usability considerations. For all types except on-premises password managers, agencies are encouraged to check with providers to ensure data are stored on servers physically located in the United States.

Cloud-based: Cloud-based password managers store a user's encrypted vault data in cloud databases. This often allows users to access their accounts from multiple devices. Cloud-based password managers often have an associated mobile or desktop application, a web portal logon, or browser extensions. To ensure availability if internet access or service connectivity is limited, some cloud-based password managers keep cached copies of the vault on the user's device.

On-premises: On-premises password managers are hosted in an agency's environment, allowing users to access vault data even without a stable internet connection. On-premises implementations take longer to deploy, cost more, and have additional overhead that the agency must manage, but they rely less on third parties.

Mobile: Mobile password managers are a functionality built into a device that stores user credentials locally or in the platform's cloud service (e.g., Apple iCloud Keychain, Google Password Manager). Cloud-based storage may allow the user to access credentials on another device by logging in with a platform account. Mobile password managers can fill credentials across the applications and browsers used on the device.

Browser-based: Browser-based password managers rely on data stored on a user's browser account and can usually be accessed on multiple devices using the same browser account (e.g., Firefox, Safari, Chrome, Edge). Compatibility between different browsers is low, and this option does not work well for shared devices.

Features

Enterprise password managers have valuable add-on features with various security and usability considerations. These features are not meant to be a comprehensive list but may be beneficial in understanding the scope of password manager capabilities. Before adopting a potential identity solution, agencies should identify the desired features of a password manager and seek all-in-one or compatible solutions to meet their needs. Password manager features include:

- **Autofill:** Autofill is a feature that allows a password manager to look up and enter stored credentials in logon portals on behalf of the user.
- **Dark Web Monitoring:** Dark web monitoring services scan websites with information from data breaches and generate an alert if they find emails or credentials from a user's vault.
- **Device Sync:** Device sync allows a user to access data stored in their password manager on multiple devices.
- **Encrypted Storage:** Encrypted storage ensures that vault data are not stored as plaintext in the password manager's database or on the user's device. Military-grade encryption (AES-256) for data at rest is best.
- **Multifactor Authentication:** MFA challenges a user to authenticate (also known as identity proofing) through more than one factor (password, biometric, OTP, etc.).
- **Role-based Permissions:** Role-based permissions determine a user's ability to configure an enterprise password manager and access or modify credentials based on their assigned groups.
- **Policy Enforcement:** Password policies can be enforced by password managers to ensure new passwords meet policy complexity and length requirements.
- **Secure Password Generator:** Secure password generators randomly generate a password for an account based on user-specified complexity requirements.
- **Single Sign-on Integration:** SSO complements the capabilities of password managers, increasing security and accessibility by centralizing authentication. SSO can also replace using a primary password to unlock a user's vault.
- **Team Sharing:** Team sharing manages users' ability to share credentials to an account until agencies transition to more secure, passwordless authentication methods.

3.6 Context-based Access Control

Overview

Context-based access control (CBAC) is a method of access control that combines features of role-based access control (RBAC) and attribute-based access control (ABAC) to apply dynamic access policies using device-level signals as cues. CBAC fulfills the fourth tenet of zero trust as defined in NIST SP 800-207, which states, "Access to resources is determined by dynamic policy-including the observable state of client identity, application/service, and the requesting asset and may include other behavioral and environmental attributes." Therefore, each request to access a resource should be evaluated in context (i.e., Who is making the request? What is the user trying to access? From what device is the request being made? Is this a reasonable request from this user, at this time, from this location?) and should be passed through a set of policies that define the access rules of that context. CBAC also aligns with the third tenet of zero trust, which states, "Access to

individual enterprise resources is granted on a per-session basis." Each resource requires its own authentication and authorization request, and access should be granted on the principle of least privilege.

Implementation

In CBAC, attributes such as device location or system version are used to assess the risk of an access request. The requesting user's identity, or role, is at the core of what resources that user will have access to, but the risk derived from the attributes of the request will inform the access decision. For example, managers may have permission to view employee records, but if they attempt to do so from a new device and a new location, they may need to take additional steps to gain access or may not be granted access.

As vendors develop CBAC solutions, the specifics of how this philosophy is implemented into a system's architecture can vary. Different vendors may work better for a given agency's mission needs, risk tolerance, and infrastructure. The key features that compose a strong CBAC solution are:

- Administrators can create and apply policies that allow for high granularity and the application of least privilege.
    - Policies decide how risk is treated in the access decision.
    - CBAC solution determines if the conditions to apply each policy have been met.
- Multiple factors are involved in the analysis of access requests.
    - Typically, high-level user roles are assessed first followed by various attributes, based on the request, to narrow the application of specific attribute-based policies.

Figure 12 shows how a CBAC system operates. First, a user initiates an attempt to access a resource. This may be a cloud resource or an on-premises resource that has been set up to allow remote access. The access request is sent to the CBAC solution along with information, such as the user's identity in the system. Like in an RBAC system, users' roles are what determines if they have privileges to access a specific resource. Additional attributes gathered by the system about the context of the request are included in the access decision. Next, the system processes the access request by applying policies set by administrators. The result of policy application may be to allow access, to require additional authentication or precautions (such as session monitoring), or to deny access.

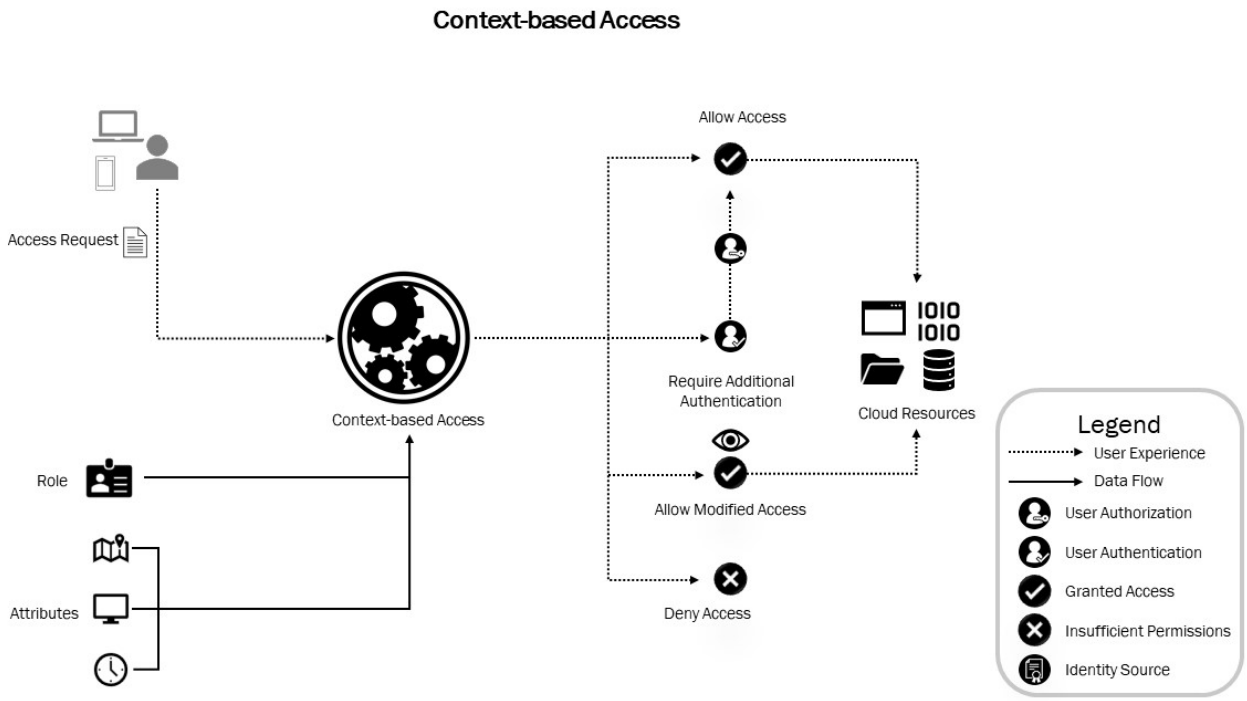

Figure 12: Context-based Access

Additional Considerations

Use of Artificial Intelligence (AI) and Machine Learning (ML)

CBAC solutions may also leverage AI or ML to inform access decisions. Using input from administrators and information collected over time from users, the CBAC solution continuously learns and develops the best application of access policies. As it determines standard user behaviors, what assets are most at risk, and who should be attempting to access specific resources, the response to each login attempt adapts. As this technology advances, it may provide a more accurate and informed risk assessment. However, agencies should prioritize that higher-fidelity data be used in CBAC decisions rather than attempting to analyze weak data using AI/ML.

Fulfilling Federal Requirements

OMB M-22-09 outlines a set of zero trust security goals for agencies that align with the five pillars of CISA's Zero Trust Maturity Model—Identity, Devices, Networks, Applications and Workloads and Data. Within the Identity section of the memorandum, the User Authorization subsection states that a "zero trust architecture should incorporate more granularity and dynamically defined permissions, as ABAC is designed to do." The related task states that "agency authorization systems should work to incorporate at least one device-level signal alongside identity information about the authenticated user." In a context-based solution, an attribute could be device-level signal, such as the originating device of the access request; and identity information would be a role, such as administrator. Agencies should note that the policies their administrators set define what attributes are evaluated and how the associated risk is handled. To fulfill federal requirements, agencies should verify that their policies include at least one device-level signal.

Solution Integration

CBAC solutions may integrate with an agency's existing infrastructure and identity providers. However, some vendors offer CBAC as an add-on to other products rather than as an independent service. Agencies should evaluate what options are compatible with their mission needs, risk tolerance, and infrastructure. Solution options vary in the device-level signals they collect and how the policies set by administrators can use the data. Additionally, the types of signals and methods of analysis may vary depending on the agency's architecture, such as those previously outlined in Section 3.1, and the points at which users can gain access. For example, in an environment where users can access resources from either personal or agency-managed devices, administrators may want to ensure their CBAC solution can distinguish between the two and can enforce access rules accordingly.

4 CONCLUSION

A traditional on-premises enterprise network typically leverages on-premises directory services for identity management. However, attempts to integrate identity management with cloud solutions in a hybrid environment creates new complexity. On-premises identity management solutions should be securely and efficiently integrated with cloud solutions to optimize interoperability. Nonetheless, this integration process can create various challenges.

The guidance presented in this document is intended to assist agencies in understanding the potential options for identity management interoperability between on-premises and cloud-based solutions, the complex challenges involved in each, and how to address these challenges. It explains the difference between traditional and modern identity architectures, discusses different authentication and authorization mechanisms, and explores additional security considerations associated with identity management. It also details various potential components of hybrid identity solutions including authentication, MFA, SSO, FIDO2, password managers, and CBACs.

As agencies make decisions regarding which hybrid identity solution(s) to pursue, they can use this document as reference tool to better understands the risks and benefits associated with each option. Ultimately, agencies must determine which solution(s) best align with their needs and risk posture.