Treasury Board of Canada Secretariat
Secrétariat du Conseil du Trésor du Canada
Canada

Government of Canada (GC) Guideline on Multi-Factor Authentication (MFA): Technical Recommendations for Authenticators to Support MFA Within the GC Enterprise Domain

Published: 2025-05-23

© His Majesty the King in Right of Canada,
represented by the President of the Treasury Board 2025,

Published by Treasury Board of Canada, Secretariat
90 Elgin, Ottawa, Ontario, K1A 0R5, Canada

Catalogue Number: BT48-56/2025E-PDF
ISBN: 978-0-660-76817-5

This document is available on the Government of Canada website, Canada.ca

This document is available in alternative formats upon request.

Aussi offert en français sous le titre : Ligne directrice du gouvernement du Canada (GC) sur l’authentification multifactorielle (AMF) : Recommandations techniques relatives aux authentifiants utilisés pour l’AMF dans le domaine opérationnel du GC

Government of Canada (GC) Guideline on Multi-Factor Authentication (MFA): Technical Recommendations for Authenticators to Support MFA Within the GC Enterprise Domain

On this page

Executive summary

Government of Canada (GC) information systems and networks are lucrative targets to a variety of threat actors and are under constant cyber-attack. This was clearly articulated during a parliamentary committee session in March 2018 when the Chief of the Communications Security Establishment stated, “We’re now blocking over one billion malicious attempts to compromise government systems on average every day. One billion attempts.” [1]

One of the more prevalent techniques used by threat actors to gain unauthorized access to GC systems and information is to compromise user credentials. Attempting to control unauthorized access to user accounts relying solely on single-factor authentication (typically a userid and password) render GC resources particularly vulnerable to compromise, even when proper credential management practices are in place (for example, strong passwords, different password for each account), as these credentials can still be compromised through a number of means such as social engineering and malware. Implementation of multi-factor authentication (MFA) is an essential step towards significantly reducing the risk of account takeover and improving the GC’s overall security posture. Furthermore, **all** users should use MFA when accessing GC resources in order to align with the *Directive on Service and Digital* and zero-trust principles.

The selection of the most suitable authenticator(s) for a given department will depend on a variety of factors, including the ability to leverage existing investments, overall cost, user experience, available technology and more. **Ultimately, the goal is to strike the right balance between security, manageability, interoperability, cost and user experience leading to the deployment of suitable MFA solutions throughout the GC.**

The purpose of this guideline is to establish the technical requirements associated with authenticators, and to recommend specific authenticators that are suitable to support MFA in the GC enterprise domain from a technology and security perspective. The technical requirements are identified in Sections 2.2 and 2.3 of this guideline, corresponding to credential levels of assurance (LoAs) as specified in the Standard on Identity and Credential Assurance.

This guideline recommends that typical users conducting day-to-day business activities must use authenticators that meet all the requirements specified for credential LoA 3. Authenticators that support phishing resistance are strongly recommended (see Section 2.3.1 for additional details regarding phishing resistance requirements). In addition, highly privileged users such as system administrators and high-profile users such as chief financial officers must use authenticators that meet all the requirements for credential LoA 4. At least one of the authenticators must be phishing resistant and must be separate from the computer platform being used to access GC resources.

Figure 1 is a simplified diagram based on Figure 2.2 in the main body that, at a high level, identifies the acceptable and preferred authenticators for typical users as well as privileged and high-profile users. Since this is a simplified representation, the reader is cautioned not to take this diagram out of context; Sections 2.2 and 2.3 should be consulted for specific requirements and recommendations.

Figure 1: high-level summary of acceptable and preferred authenticators for two classes of users

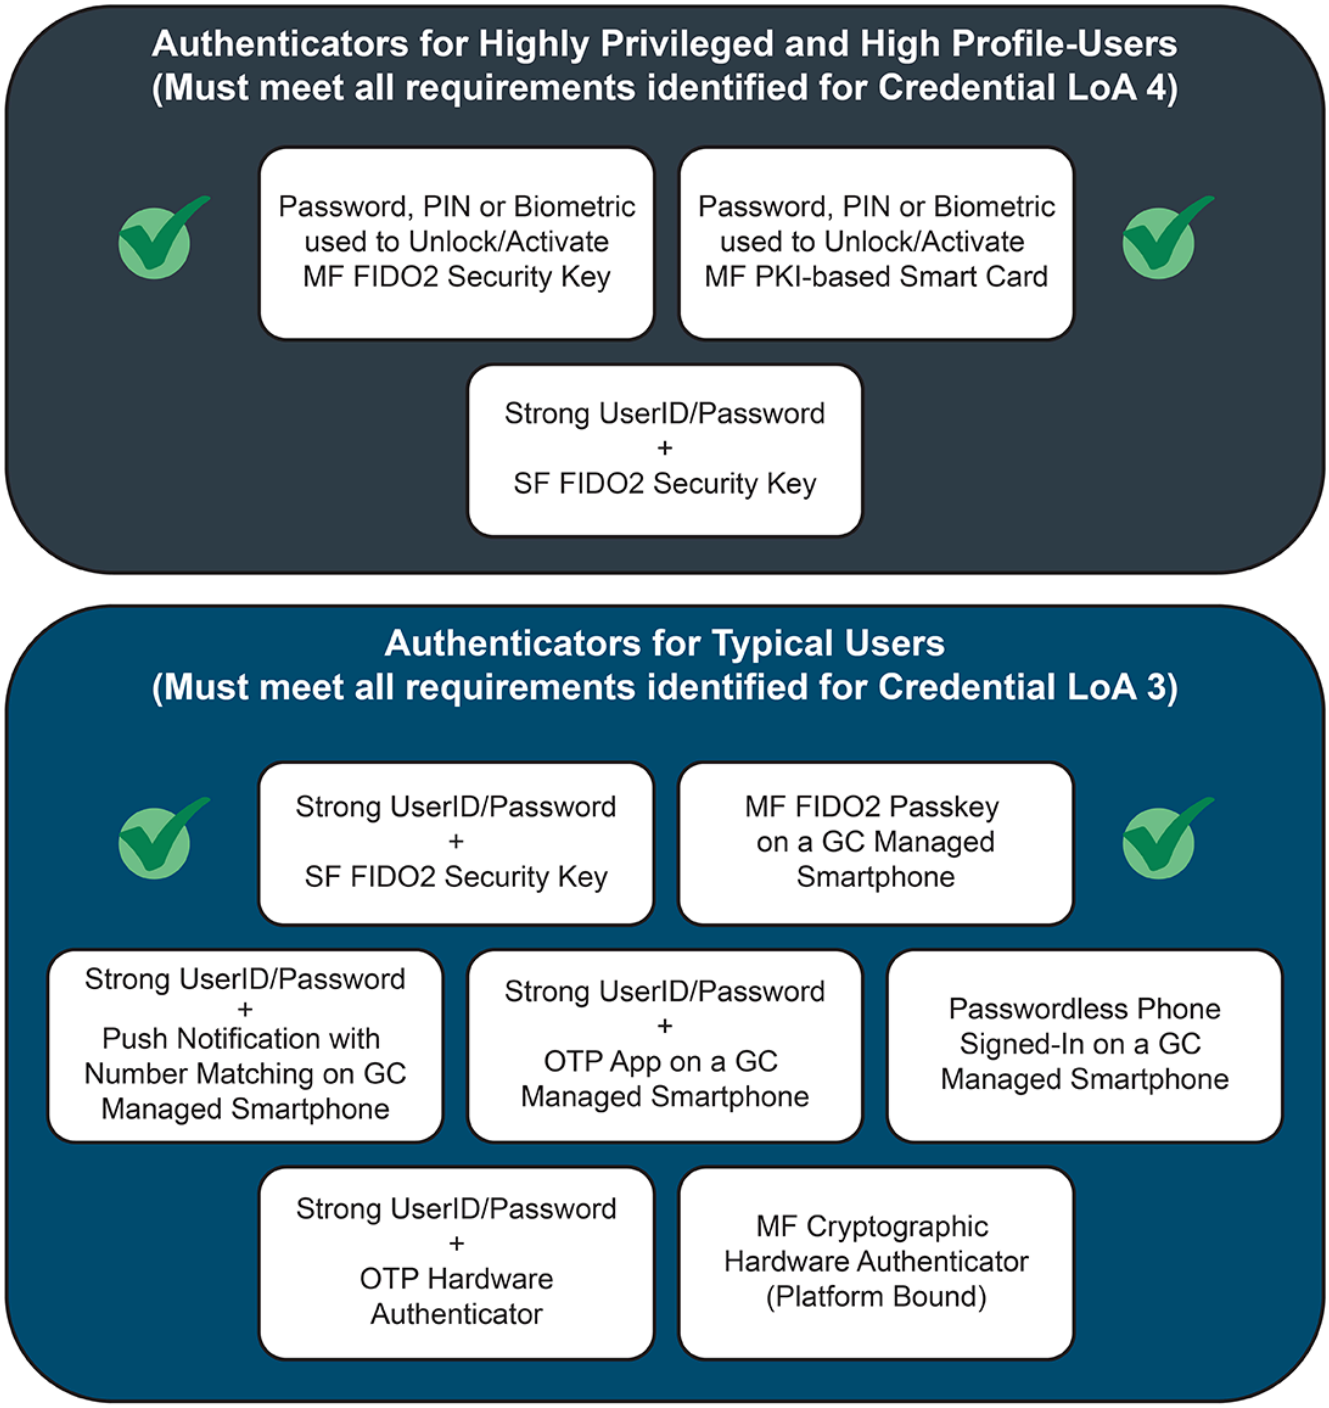

The authenticators tagged with a ✓ indicate the more preferred options for each class of user.

SF = Single-factor
MF = Multi-factor
OTP = One-time Password
PKI = Public Key Infrastructure

▶ Figure 1 - Text version

It should be noted that this guideline addresses a complex and broad topic area, and both the threat landscape and the user authentication technologies and requirements are constantly changing. Therefore, the recommendations made within this guideline are subject to revision over time. Furthermore, references to specific vendors or products, implied or otherwise, are for illustrative purposes only and should not be considered a formal endorsement by the GC.

In terms of obtaining authenticators, departments are expected to use enterprise or shared information technology (IT) solutions, assets, and services to avoid duplication, when available and appropriate, as stipulated in the *Policy on Service and Digital*, Section 4.4.2.3. To that end, departments can leverage enterprise services that support MFA, including leveraging supply arrangements that are established by Shared Services Canada.

1. Introduction

▶ In this section

1.1 Background

GC information systems and networks are lucrative targets to a variety of threat actors and are under constant cyber-attack. This was clearly articulated during a parliamentary committee session in March 2018 when the Chief of the Communications Security Establishment stated, “We’re now blocking over one billion malicious attempts to compromise government systems on average every day. One billion attempts.” [2]

One of the more prevalent techniques used by threat actors to gain unauthorized access to GC systems and information is to compromise user credentials. Attempting to control unauthorized access to user accounts

relying solely on single-factor (SF) authentication (typically a userid and password) renders GC resources particularly vulnerable to compromise, even when proper credential management practices are in place (for example, strong passwords, different password for each account), as these credentials can still be compromised through a number of means such as social engineering and malware.

Implementation of MFA is an essential step towards significantly reducing the risk of account takeover and improving the GC's overall security posture. Furthermore, MFA is a fundamental tenet of the zero-trust security model, which the GC is moving towards.

1.2 Purpose and scope

This document addresses detailed technical requirements associated with authenticators and makes recommendations regarding specific authenticators that can be used to support MFA in the GC Unclassified, Protected A and Protected B enterprise domain. Although some of the considerations identified in this document will also apply to external entities accessing GC public-facing services, this is not within the scope of this document.

Although this document is focused on technology and security considerations, it is recognized that determining the most suitable authenticators for a given department (or target environment) will be based on many additional factors such as user experience, cost, the ability to leverage existing investments, and so on, not just technology. The objective of this document is to assist in this process by identifying suitable authenticators that departments can select from to meet their specific needs.

Recommendations made within this document are based on guidance from several sources, including ITSP.30.031 v3 and its eventual replacement, ITSP.30.031 v4 (currently being drafted), as well as the National Institute of Standards Technology (NIST) Digital Identity Guidelines, particularly NIST Special Publication (SP) 800-63B and NIST SP 800-63B-4 (Revision 4, Initial Public Draft). [3] Similarities and differences between these sources are highlighted where appropriate in order to assist in achieving closer alignment between ITSP.30.031 v4 and the NIST guidelines as part of the ongoing collaboration between TBS and the Canadian Centre for Cyber Security. Once ITSP.30.031 v4 is formally published, the comparisons will no longer be required, and this document can be simplified as a result. Recommendations made within this document will be in complete alignment with ITSP.30.031 v4.

It should be noted that this document addresses a broad and complex topic area, and both the threat landscape and the user authentication technologies and requirements are subject to constant change. It is essential to keep abreast of these ongoing developments and to adjust the relevant guidance accordingly.

This document does not address how the recommendations made within this document are to be enforced; however, it is expected that this will be accomplished through a combination of policy, procedures, user training and awareness, and technology.

Note that this document includes three appendices. These appendices are an integral part of this document and should be read in conjunction with the main body.

This document replaces *Recommendations for Two-factor User Authentication within the GC Enterprise Domain*.

Special note on authenticators and levels of assurance

It is important to recognize that this document addresses technical requirements associated with authenticators only. It does not address the broader considerations associated with the overall user authentication process. In other words, the level of assurance (LoA) of the authenticator (or combination of authenticators) does not necessarily translate to the LoA of the overall user authentication process, which could be lower, depending on the LoAs of the other aspects associated with the authentication (see Appendix A for additional details). Nonetheless, moving to appropriate MFA authenticators will significantly improve the GC's overall security posture and help to significantly reduce the threat of user account takeover.

1.3 Intended audience

This document is intended for IT security practitioners and decision makers responsible for determining suitable MFA solutions within the GC enterprise domain.

1.4 Terminology

Terminology is critical to understanding this topic, particularly when key terms are used inconsistently, depending on the source (for example, marketing material compared to the relevant technical specifications). This document uses a number of key terms that are defined in Section 5, Appendix A (for assurance level definitions and mappings), Appendix B (for FIDO Alliance specific terms) and Appendix C (for trusted platform module

(TPM) types). The definitions for these key terms are extracted or derived from the identified authoritative sources, including relevant standards, specifications and user authentication guidance.

Also note that the phrases “strong password” and “strong, well-managed password” used within this document refer to password composition and management practices that meet the requirements stipulated in the GC’s *Password Guidance*.

1.5 Applicability

The guidance provided within this document is intended for GC departments and agencies with systems that process Unclassified, Protected A or Protected B information.

Ultimately, deputy heads are responsible for “identifying security and identity management requirements for all departmental programs and services, considering potential impacts on internal and external stakeholders” (see the *Policy on Government Security*, Section 4.1.4). The purpose of this document is to provide departments with informed guidance in order to assist in the selection of the most suitable authenticators to support MFA based on each department’s specific business needs and threat context. Exceptions to the recommendations made within this document may be granted subject to GC Enterprise Architecture Review Board (GC EARB) approval.

2. Authenticator considerations

▶ In this section

The purpose of this section is to identify concepts and requirements associated with authenticators used to support MFA within the GC enterprise.

2.1 Authentication factors and types

As discussed in ITSP.30.031 v3 and the NIST Digital Identity Guidelines (see associated definition in NIST SP 800-63-3), an authentication factor is categorized as something you know, something you have, or something you are or do. A few examples of the authenticator types that fall under each authentication factor category are provided in Table 2.1.

Table 2.1: authentication factors and examples

Authentication factor	Authenticator type examples
Something you know	• Memorized-secret authenticators (for example, a password [4] associated with a user account) [5]
Something you have	• Look-up-secret authenticators (for example, static grid or “bingo” cards) • Single-factor (SF) or multi-factor (MF) out-of-band (OOB) authenticators (for example, push notifications with number matching to an OOB mobile device such as a smartphone) • SF or MF one-time-password (OTP) software or hardware authenticators (for example, an OTP app on a mobile device such as a smartphone or an OTP hardware device that displays one-time use authentication codes) [6] • SF or MF cryptographic software or hardware authenticators (for example, software or hardware authenticators that store cryptographic keys and perform cryptographic operations to support the user authentication process)

Authentication factor	Authenticator type examples
Something you are or do	• Biometrics associated with an individual that include both measurement of physical characteristics (for example, fingerprints, iris, facial characteristics) and behavioural characteristics (for example, typing cadence) [7], [8]

Additional details regarding authenticators and authenticator requirements are provided in Sections 2.2 and 2.3 below.

Special note on terminology: authenticator versus token

Prior to the publication of the US NIST Digital Identity Guidelines in 2017, both NIST and the Communications Security Establishment user authentication guidance used the term “token” instead of “authenticator.” The NIST Digital Identity Guidelines changed the term to “authenticator” in order to avoid confusion with the use of “token” in assertion protocols and technologies. This document also uses the term “authenticator” rather than “token,” and it is expected that the next update to ITSP.30.031 will adopt this terminology as well.

Special note on terminology: use of software, hardware and device

NIST has broadened the definition of an OTP device to include hardware and software. Specifically, NIST states: “This category includes hardware devices and software-based OTP generators installed on devices such as mobile phones.” Note that this applies to both an SF OTP device and MF OTP device. This is the only instance where NIST includes hardware and software under the banner of “device.” In all other instances, “device”

means hardware. In order to avoid confusion, this document uses the term “authenticator” rather than “device,” irrespective of whether the authenticator is software or hardware is specified as appropriate. Note that excerpts from the NIST documentation have been changed to reflect this terminology for consistency.

2.2 Authenticators and levels of assurance

Appendix A provides an overview of the four-LoA model used by the GC and how those LoAs correspond to the assurance levels described in the NIST Digital Identity Guidelines. The reader should be familiar with the concepts discussed in Appendix A before proceeding further with this section, particularly with respect to the mappings between the NIST authenticator assurance levels (AALs) and the GC credential LoAs). As defined in the Standard on Identity and Credential Assurance and as discussed in Appendix A, note that credential LoAs are defined as **the level of confidence that the user has maintained control over a credential that has been entrusted to them and that the credential has not been compromised**.

Table 2.2 identifies the credential LoA and AAL of each authenticator type identified in NIST SP 800-63B and draft ITSP.30.031 v4. While there are several differences between NIST SP 800-63B and ITSP.30.031 v3, it is anticipated that the next update ITSP.30.031 will more closely align with NIST SP 800-63B. Entries in this table are subject to change pending formal publication of the ITSP.30.031 update (see associated footnotes and comments).

Table 2.2: authenticators at each credential LoA/AAL

Highest possible LoA and AAL	Authenticators

Highest possible LoA and AAL	Authenticators
Credential LoA 2/AAL1	• Memorized secret • Look-up secret • Single-factor out-of-band authenticator • Single-factor OTP software or hardware authenticator [9] • Single-factor cryptographic software [10] or hardware authenticator

Credential LoA 3/AAL2	• Memorized secret combined with one of the following:
	  ◦ look-up secret
	  ◦ single-factor out-of-band authenticator
	  ◦ single-factor OTP software or hardware authenticator
	  ◦ single-factor cryptographic software authenticator
	• Multi-factor out-of-band authenticator [11]
	• Multi-factor cryptographic software authenticator [12]
	• Multi-factor OTP software or hardware [13] authenticator [14]

Credential LoA 4/AAL3	• Multi-factor cryptographic hardware authenticator • Memorized secret combined with a single-factor cryptographic hardware authenticator [15]

As evident from the entries in Table 2.2, MFA can be supported using a multi-factor authenticator or a combination of two **appropriate** single-factor authenticators. However, it should be noted that authenticators and authenticator combinations are not all created equal, even though they may be at the same credential LoA/AAL. In other words, some authenticators or combination of authenticators are stronger (from a security perspective) than others, even though they considered to be at the same credential LoA/AAL. Also note that acceptable authenticator combinations must not be subject to the same attack vector (that is, method of compromise), and one of the authenticators must be a **something you have** authenticator that cannot be trivially duplicated or

copied. Essentially, this means that the amount of work required to compromise the combination of authenticators must be significantly higher than the amount of work required to compromise just one of the authenticators. In addition, all devices (for example, computer platforms, mobile devices, authenticators) must be managed by (or on behalf of) the GC. This is to ensure that the devices used to access GC resources are in their proper state (for example, the device is running approved and up-to-date software, malware detection is installed and up to date, there is no evidence that the device has been compromised) and approved for the intended use.

Figure 2.1 identifies several different authenticator and authenticator combinations at each credential LoA/AAL and illustrates the relative improvements in security they have in relationship to each other at each level of assurance. Note that the associated credential LoA/AAL represents the **highest possible level** that can be achieved with the indicated authenticator or combination of authenticators and assumes that all applicable requirements at that credential LoA/AAL are met, as stipulated in Section 3.3. In addition, the comparisons are based on the merits of the authenticators themselves, independent from any other security controls that may be implemented to support their use. Additional rationale for their relative placement is provided in Table 2.3. Also note that this is not an exhaustive set of examples but is meant to represent authenticators that might be more applicable within the GC enterprise domain, particularly at credential LoA 3/AAL2 and LoA 4/AAL3.

Figure 2.1: comparison of authenticators

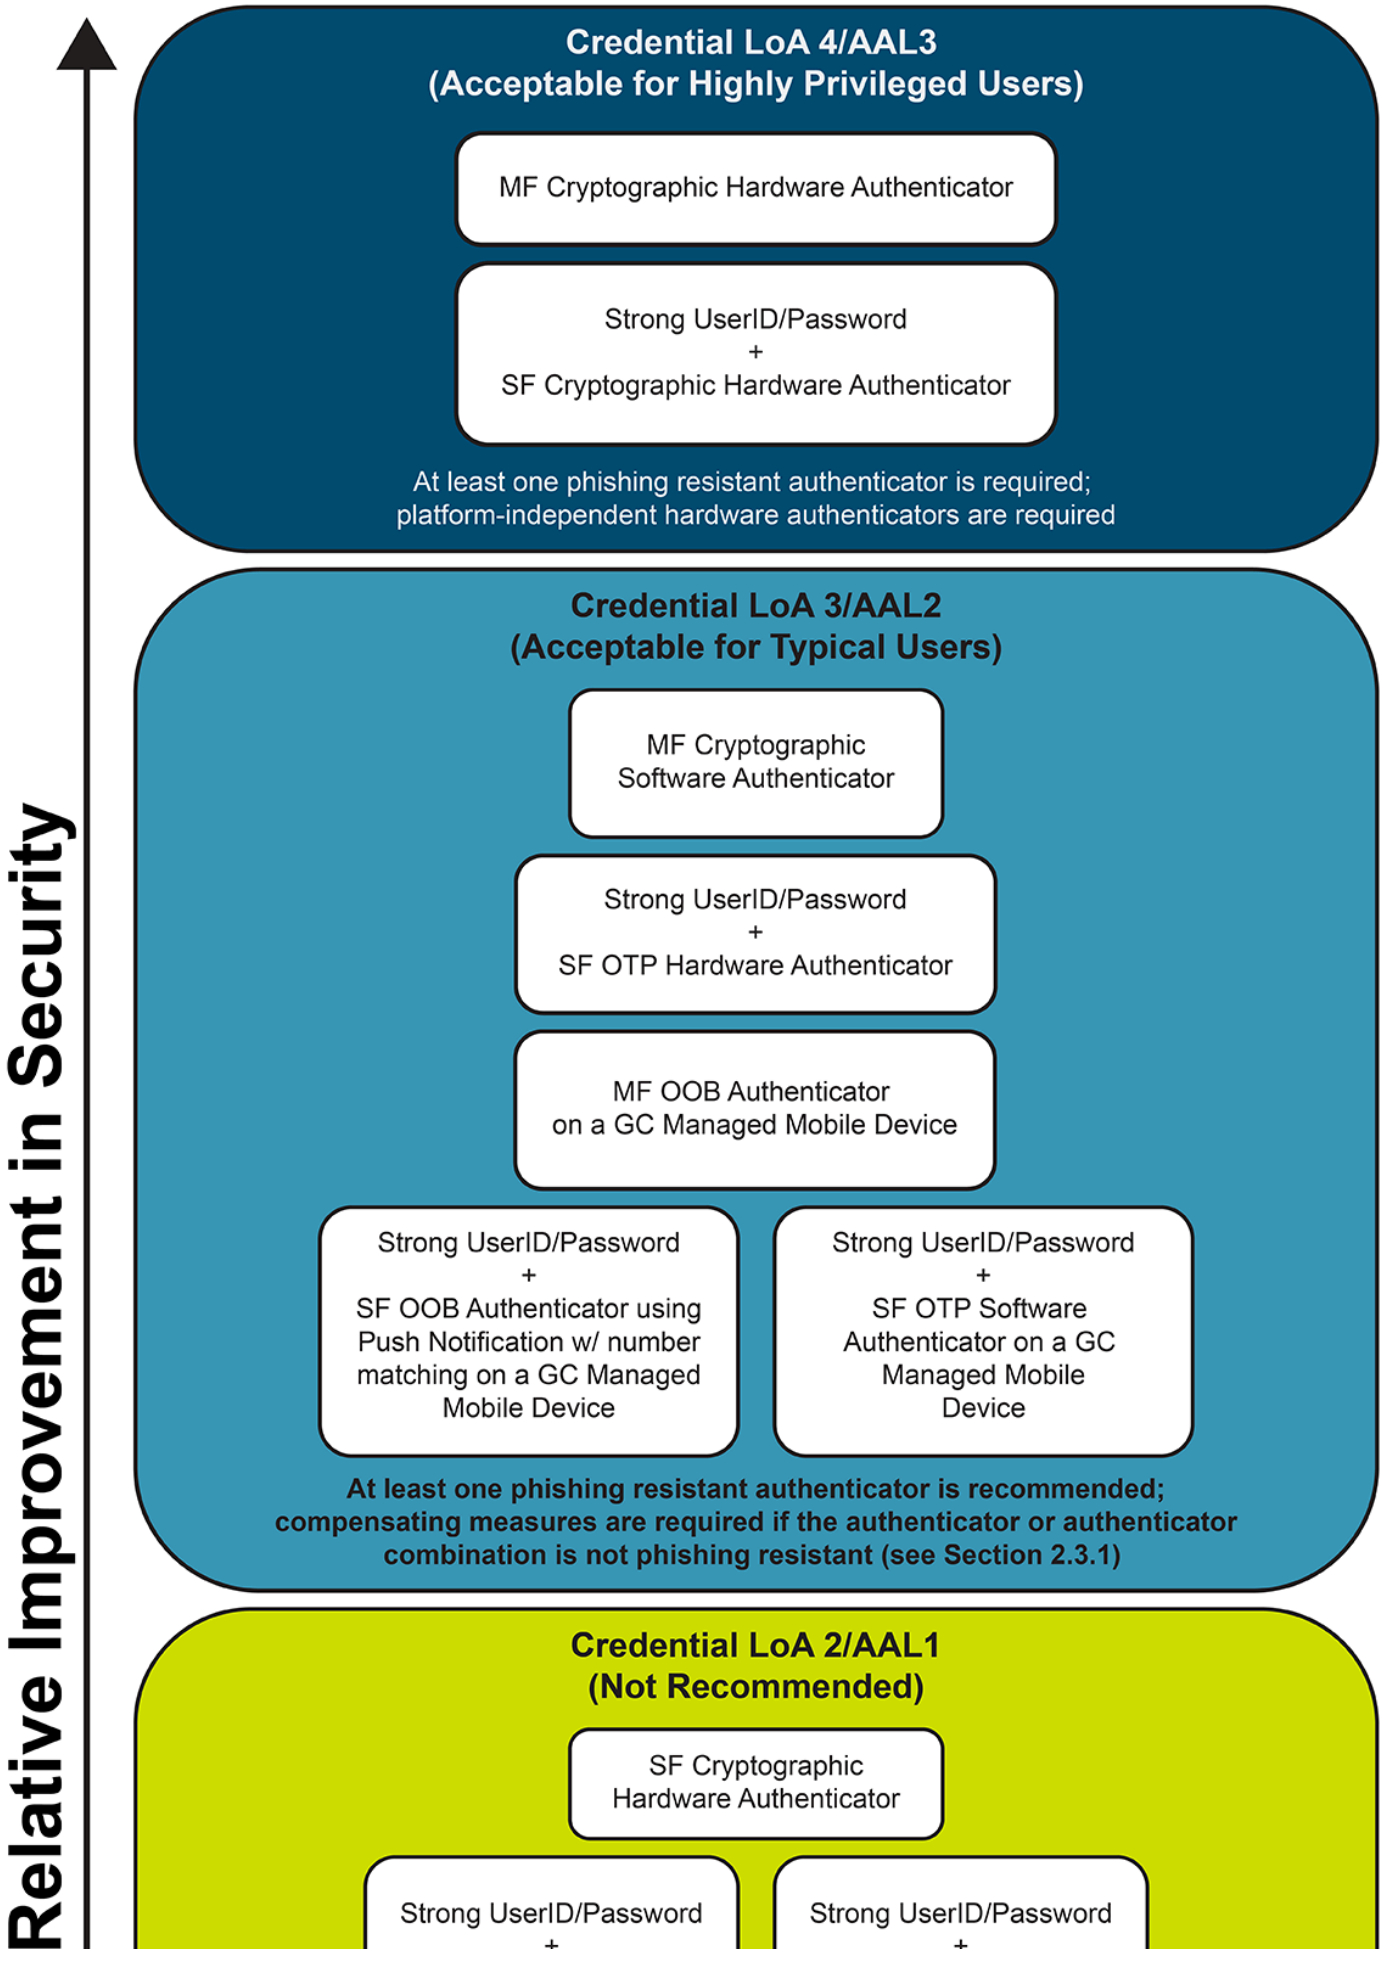
Relative Improvement in Security
Credential LoA 4/AAL3
(Acceptable for Highly Privileged Users)
MF Cryptographic Hardware Authenticator
Strong UserID/Password
+
SF Cryptographic Hardware Authenticator
At least one phishing resistant authenticator is required;
platform-independent hardware authenticators are required
Credential LoA 3/AAL2
(Acceptable for Typical Users)
MF Cryptographic
Software Authenticator
Strong UserID/Password
+
SF OTP Hardware Authenticator
MF OOB Authenticator
on a GC Managed Mobile Device
Strong UserID/Password
+
SF OOB Authenticator using
Push Notification w/ number
matching on a GC Managed
Mobile Device
Strong UserID/Password
+
SF OTP Software
Authenticator on a GC
Managed Mobile
Device
At least one phishing resistant authenticator is recommended;
compensating measures are required if the authenticator or authenticator
combination is not phishing resistant (see Section 2.3.1)
Credential LoA 2/AAL1
(Not Recommended)
SF Cryptographic
Hardware Authenticator
Strong UserID/Password
Strong UserID/Password

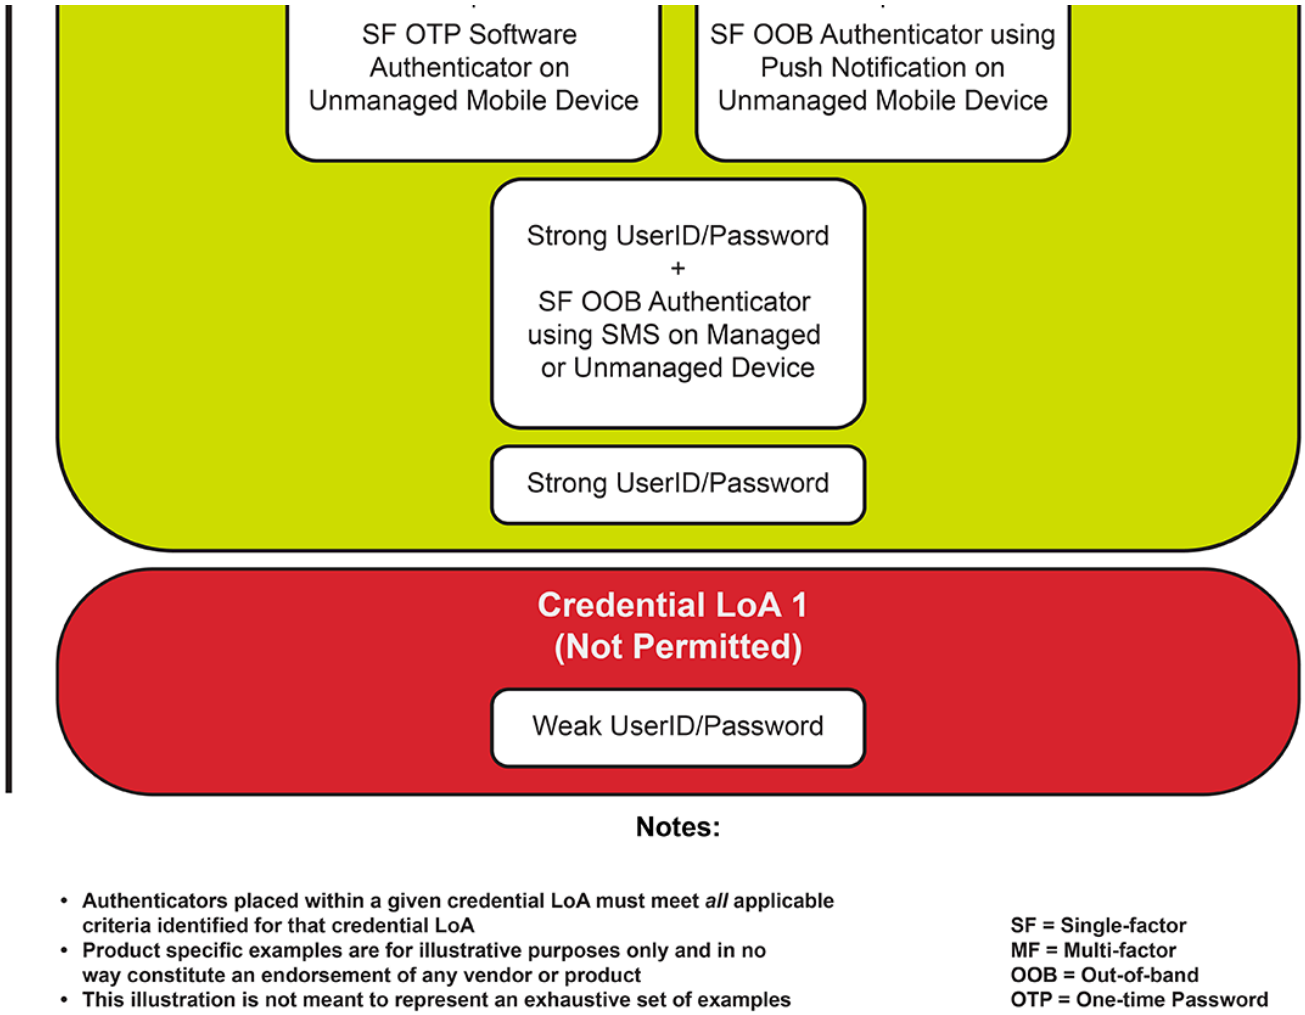

▶ Figure 2.1 - Text version

Figure 2.2 expands on Figure 2.1 to illustrate real-world examples of the different authenticators and authenticator combinations. As noted previously, the product examples are for illustrative purposes only and **do not constitute an endorsement of a particular vendor or product**.

Figure 2.2: comparison of examples of authenticators

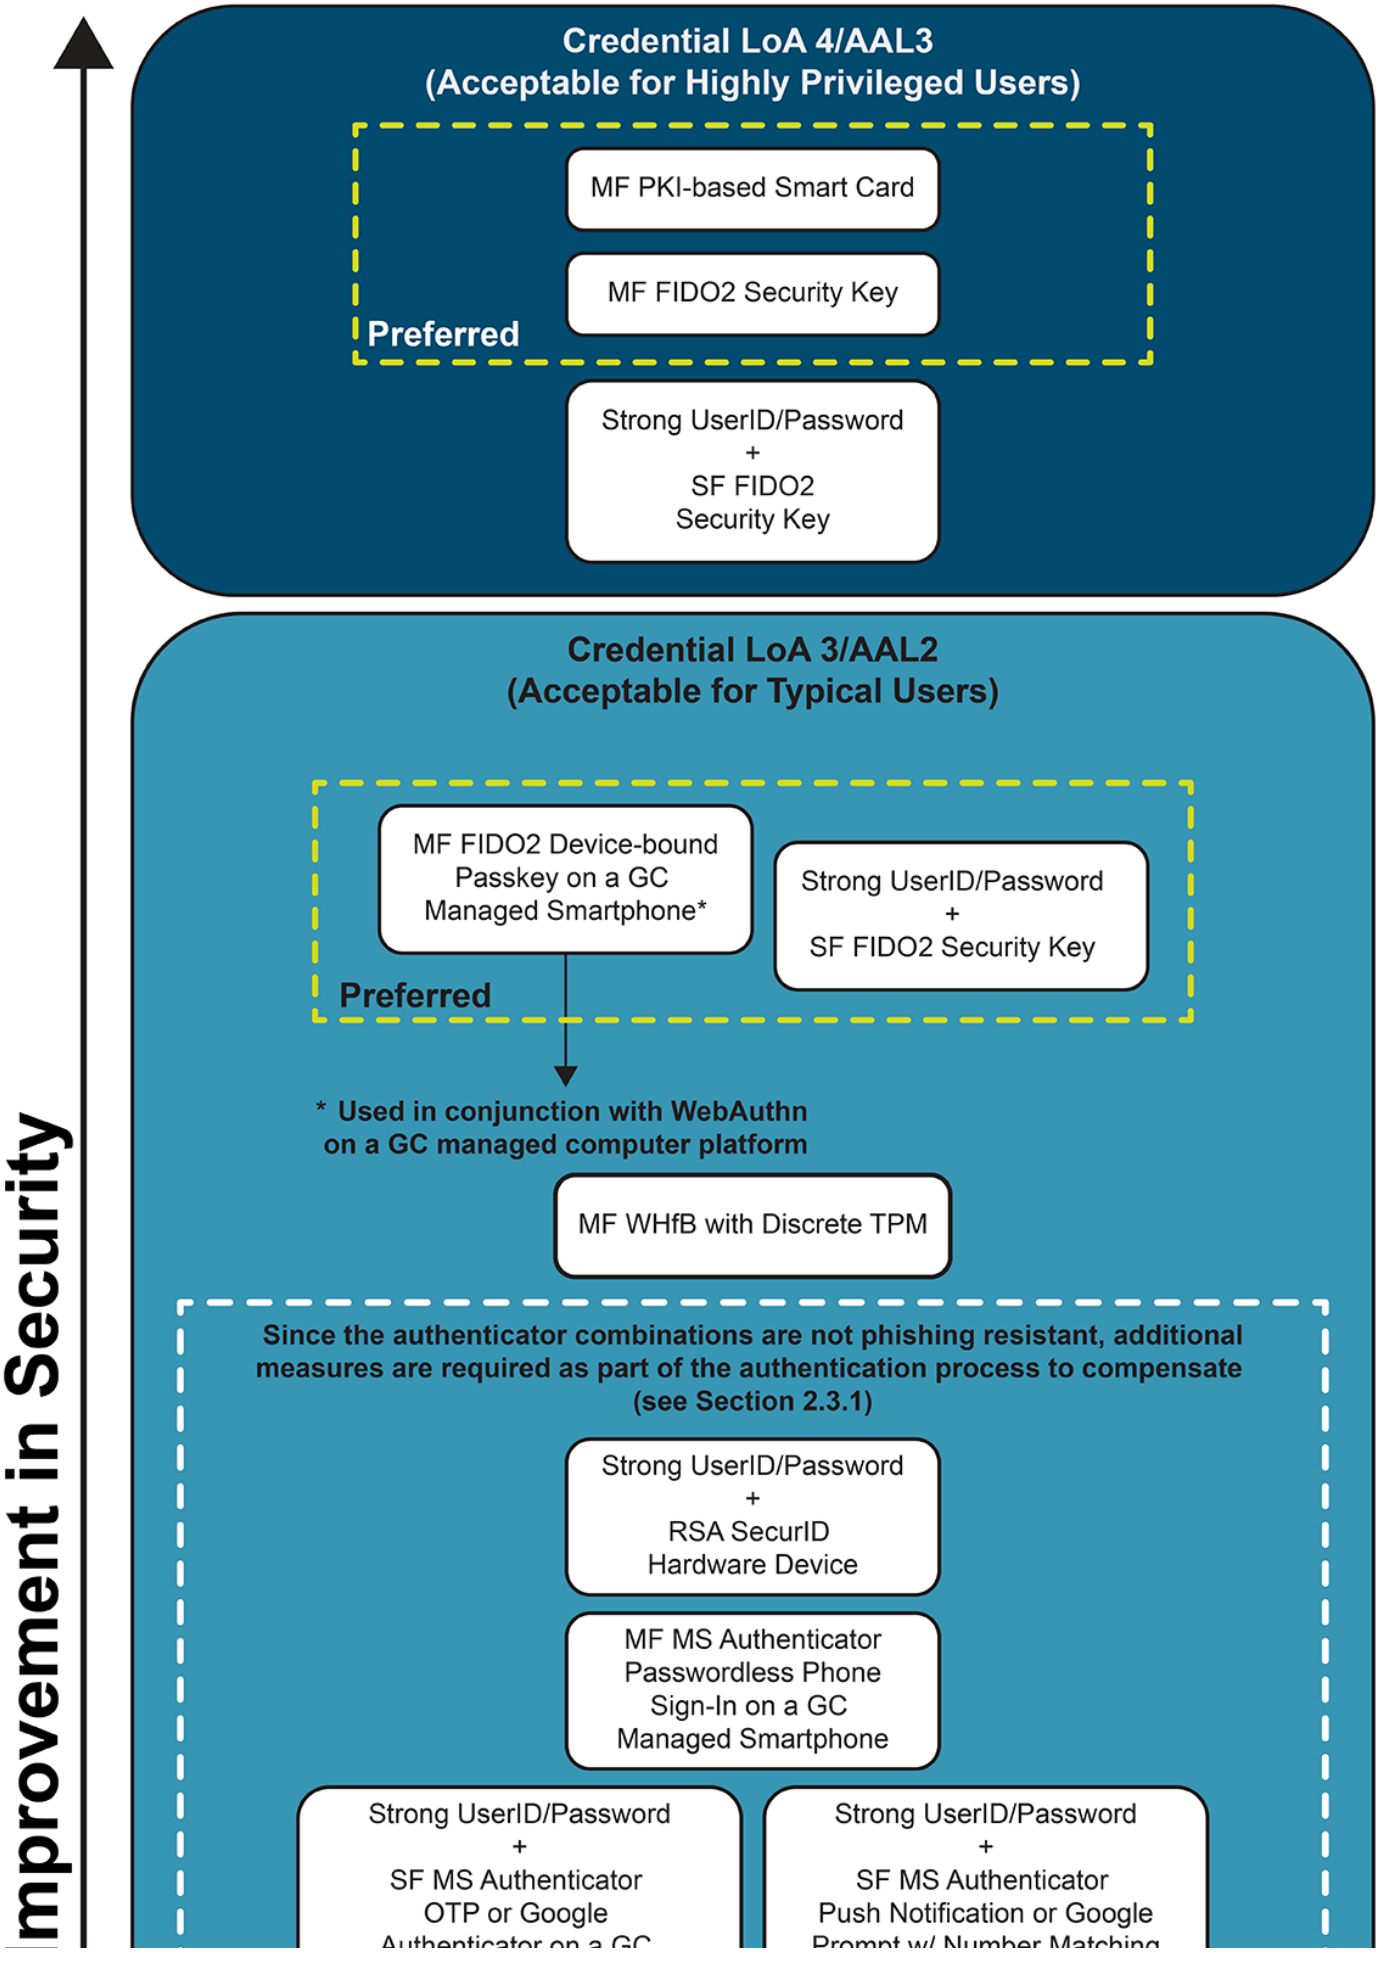
Improvement in Security
Credential LoA 4/AAL3
(Acceptable for Highly Privileged Users)
MF PKI-based Smart Card
MF FIDO2 Security Key
Preferred
Strong UserID/Password
+
SF FIDO2
Security Key
Credential LoA 3/AAL2
(Acceptable for Typical Users)
MF FIDO2 Device-bound
Passkey on a GC
Managed Smartphone*
Strong UserID/Password
+
SF FIDO2 Security Key
Preferred
* Used in conjunction with WebAuthn
on a GC managed computer platform
MF WHfB with Discrete TPM
Since the authenticator combinations are not phishing resistant, additional
measures are required as part of the authentication process to compensate
(see Section 2.3.1)
Strong UserID/Password
+
RSA SecurID
Hardware Device
MF MS Authenticator
Passwordless Phone
Sign-In on a GC
Managed Smartphone
Strong UserID/Password
+
SF MS Authenticator
OTP or Google
Strong UserID/Password
+
SF MS Authenticator
Push Notification or Google

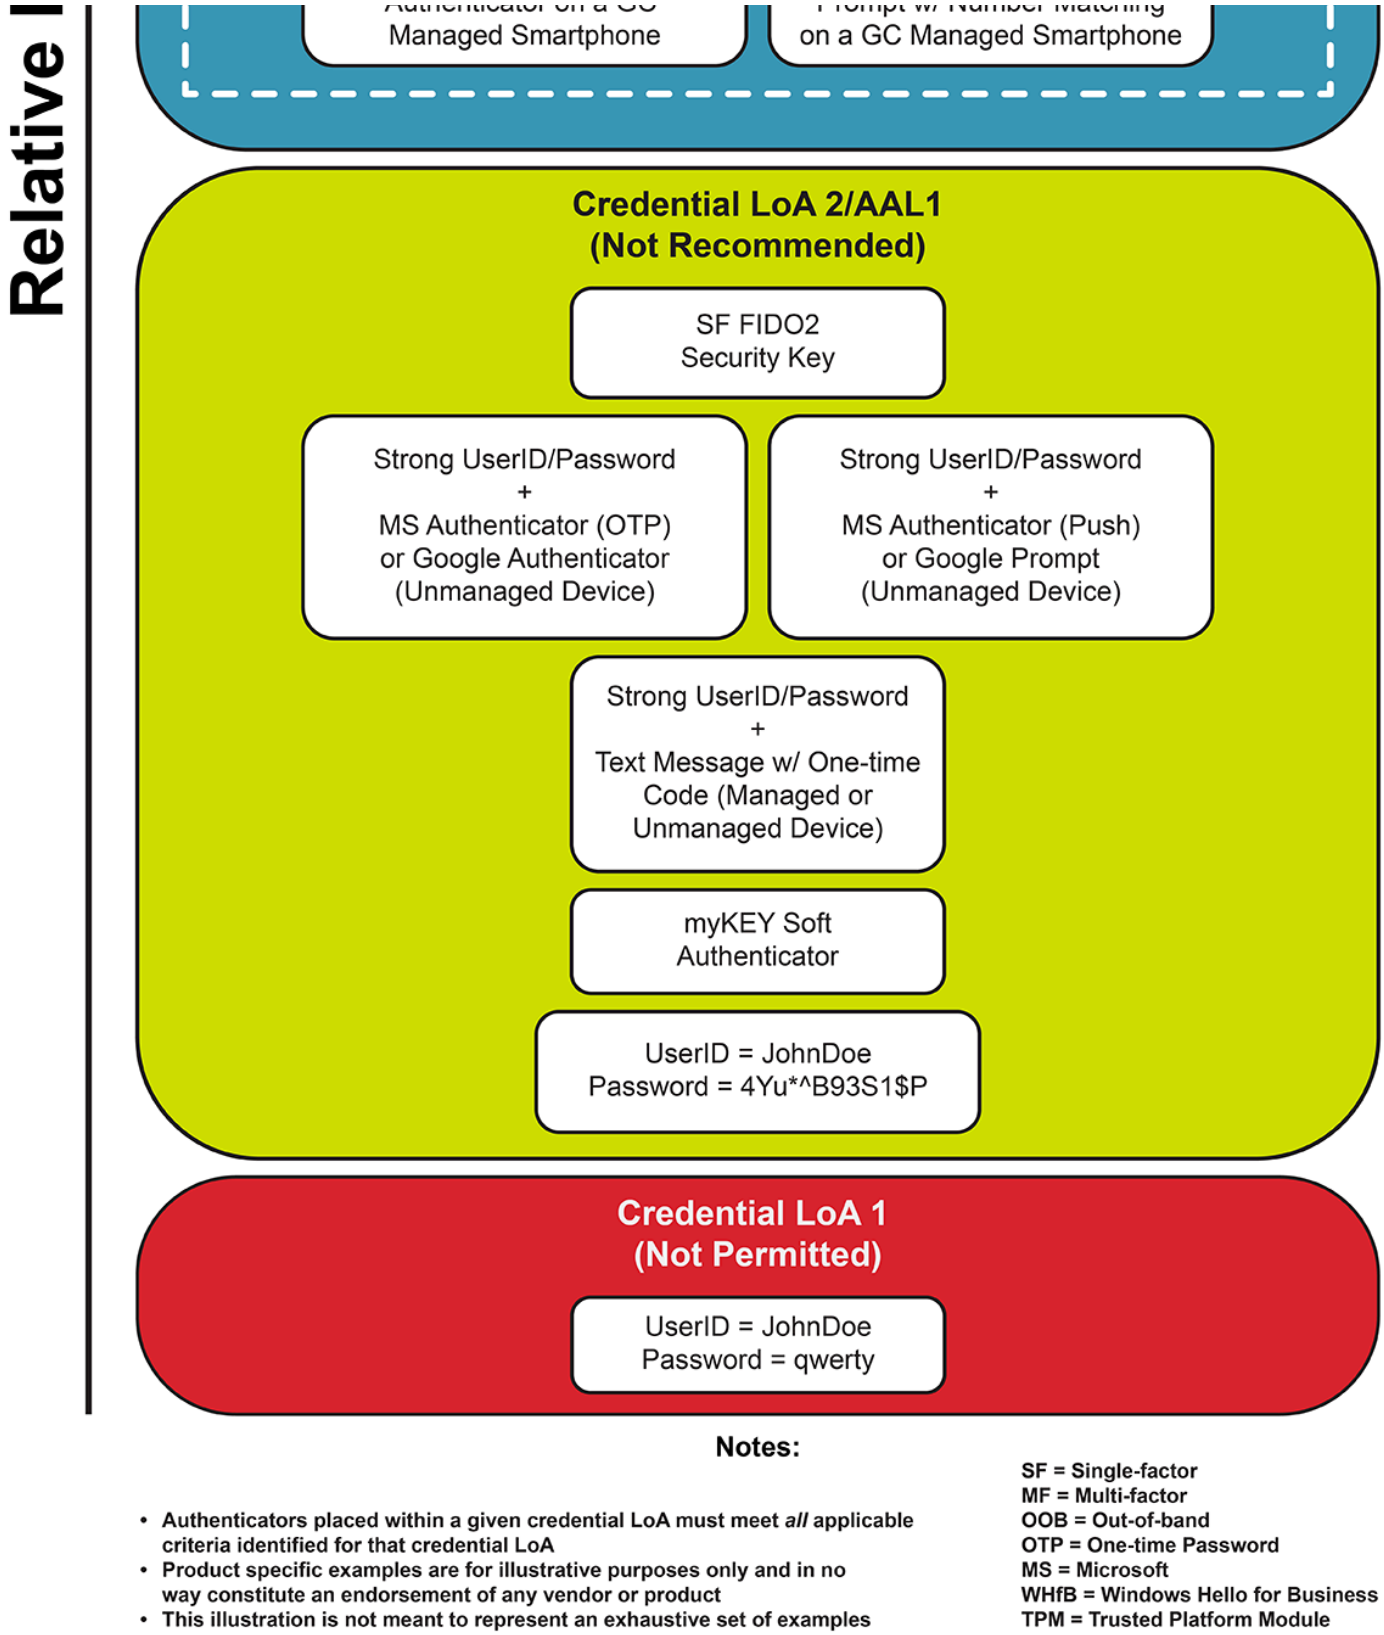

▶ Figure 2.2 - Text version

As illustrated in Figure 2.2, some authenticators do not meet the highest possible level of assurance that the authenticator type to which they belong is capable of achieving as depicted in Figure 2.1. Table 2.3 provides a

rationale for the placement of the authenticator and authenticator combinations provided in Figure 2.1, as well as the examples provided in Figure 2.2.

Note that “credential LoA/AAL” is abbreviated simply as “LoA/AAL” throughout the remainder of Section 2.

Table 2.3: additional rationale for authenticator placements

Authenticator(s)	Additional rationale for placement in Figure 2.1	Examples provided in Figure 2.2
Weak and/or poorly managed password	A weak and/or poorly managed password is the most vulnerable form of authenticator, as it is susceptible to numerous exploits or attacks and provides little to no confidence that the user has maintained sole control over the authenticator or that the authenticator has not been compromised; this type of weak authenticator is LoA 1 at best and should not be used to control access to GC resources.	UserID = JohnDoe Password = qwerty

Strong, well-managed password	Although a strong, well-managed password offers some confidence that the user has maintained control over the authenticator and is therefore considered LoA 2/AAL1, it is the weakest of all LoA 2 authenticators and is therefore positioned at the lowest end of LoA 2 since it is still susceptible to threats such as social engineering and malware. Given the GC's direction with respect to MFA, this type of authenticator should only be used in combination with another appropriate authenticator to achieve MFA and a higher level of assurance (although there may be exceptions such as emergency or "break glass" accounts).	UserID = JohnDoe Password = 4Yu*^B93S1$P

Strong password combined with an SF OOB short message service (SMS (short message service)) to a mobile device (managed or unmanaged)	Although use of SMS to a mobile device is technically an OOB authentication method, it is not considered to be LoA 3 due to number of vulnerabilities such as well-documented weaknesses in the Signalling System 7 (SS7) protocol, as well as other concerns such as subscriber identity module (SIM) jacking and is therefore considered to be only slightly better than a strong password. Note that use of SMS is not recommended as a second factor regardless of whether or not the mobile device is managed by the GC.	A strong password combined with a text message with one-time code (managed or unmanaged mobile device)
Strong password combined with an SF OOB push notification on an unmanaged device	A strong password combined with push notifications on an unmanaged mobile device is considered LoA 2/AAL1 since there is no visibility regarding the state of the mobile device (that is, the device may be compromised); however, it is considered to be slightly better than a strong password combined with SMS and is therefore placed slightly higher in the LoA 2 spectrum.	A strong password combined with Google prompt or Microsoft (MS) authenticator (push notification) on an unmanaged mobile device

Strong password combined with an SF OTP software (app) authenticator on an unmanaged device	A strong password combined with an OTP app on an unmanaged mobile device is considered LoA 2/AAL1 since there is no visibility regarding the state of the mobile device (that is, the device may be compromised); however, it is considered to be slightly better than a strong password combined with SMS and is therefore placed slightly higher in the LoA 2 spectrum.	A strong password combined with Google authenticator or MS authenticator (OTP) on an unmanaged mobile device
SF cryptographic hardware authenticator	An SF cryptographic hardware authenticator is considered to be the best of the LoA 2/AAL1 options (even though it is not MFA) since it is inherently phishing resistant and AitM (adversary-in-the-middle) resistant (see Sections 2.3.1 and 2.3.4) and is therefore positioned at the higher end of the LoA 2/AAL1 spectrum (under the assumption that the binding of the authenticator with the user is done reliably and securely).	SF FIDO2 security key

Strong password combined with an SF OOB authenticator using push notification with number matching on a GC-managed mobile device	A strong password combined with single-factor push notification with number matching on a GC-managed mobile device is considered LoA 3/AAL2; however, it is at the lower end of LoA 3 since this combination of authenticators is not inherently phishing resistant (and therefore compensating measures are required; see Section 2.3.1). Although not represented in Figure 2.1, push notifications without number matching are considered to be weaker than push notifications with number matching; therefore, push notifications without number matching are not recommended. [16] This authenticator combination is considered to be more or less equivalent to a strong password combined with an OTP app on a GC-managed mobile device (see entry directly below).	A strong password combined with MS authenticator (push notification with number matching) or Google prompt (with number matching) on a GC-managed mobile device (that is, smartphone)

Strong password combined with an SF OTP software authenticator on a GC-managed mobile device	A strong password combined with an OTP app on a GC-managed mobile device is considered LoA 3/AAL2; however, it is at the lower end of LoA 3 since this combination of authenticators is not inherently phishing resistant (and therefore compensating measures are required; see Section 2.3.1). This authenticator combination is considered to be more or less equivalent to a strong password combined with a push notification with number matching on a GC-managed mobile device (see entry directly above).	A strong password combined with MS authenticator (OTP) or Google authenticator or on a GC-managed mobile device (that is, smartphone)
MF OOB authenticator on a GC-managed mobile device	An MF OOB authenticator can achieve LoA 3 and is rated slightly better on the LoA 3 spectrum than a strong password combined with either a push notification with number matching or an OTP app since it is resistant to some less sophisticated social engineering attacks; however, the primary channel is not protected from AitM or phishing.	MF MS authenticator passwordless phone sign-in on a GC-managed mobile device (that is, smartphone)

Strong password combined with an SF OTP hardware authenticator	An SF OTP hardware authenticator combined with a strong password is not phishing resistant, but it is placed higher in the LoA 3/AAL2 spectrum than the previous three examples since the OTP hardware authenticator cannot be accessed or circumvented remotely. [17]	A strong password combined with an SF RSA SecurID OTP hardware device

MF cryptographic software authenticator	Some, but not all, MF cryptographic software authenticators can achieve LoA 3/AAL2. In order to achieve LoA 3/AAL2, all applicable requirements at LoA 3/AAL2 as identified in Section 2.3 must be met. MF cryptographic software authenticators that do not meet these requirements are considered to be LoA 2/AAL1. The myKEY MF cryptographic software authenticator is only an LoA 2/AAL1 authenticator since it stores the private keys in a file (.epf) rather than in secure storage and therefore does not meet all the requirements for LoA 3/AAL2; refer to Section 2.3.2. In addition, the Canadian Centre for Cyber Security advises that a myKEY soft authenticator can be easily copied by a threat actor. Once copied, the threat actor only needs to obtain the password or PIN used to unlock the authenticator. This effectively reduces the properties of **something you know** plus **something you have** to simply **something you know** and is therefore only slightly better than a strong password. An MF FIDO2 device-bound passkey on a GC-managed smartphone (used in conjunction with WebAuthn on a GC-	There are two examples provided in Figure 2.2: • myKEY MF cryptographic software authenticator positioned slightly higher than a strong password in the LoA 2/AAL1 spectrum • MF FIDO2 device-bound passkey on a GC-managed smartphone (with a TEE; see explanation in the previous cell to the left) used in conjunction with WebAuthn on a GC-managed computer platform at the higher level of the LoA 3/AAL2 spectrum Although not shown in Figure 2.2, Windows Hello for Business (WHfB) with a firmware TPM on a GC-managed computer platform is also an example of an MF cryptographic software authenticator that would be positioned slightly below the WHfB with a TPM; also note that this is similar to a GC-managed smartphone that is used to access a remote resource with a platform authenticator protected within a TEE on that

	managed computer platform) is at the higher end of the LoA 3/AAL2 spectrum and a preferred authenticator since it is a phishing-resistant authenticator and the authenticator itself is not permanently bound to the computer platform being used to access the resource. Note that the MF FIDO2 device-bound passkey on a GC-managed smartphone represents two possible authenticator types at LoA 3/AAL2, depending on whether the authenticator is supported by a trusted execution environment (TEE) or an embedded secure element (SE). In the case of the TEE, the authenticator is considered to be an MF cryptographic software authenticator, and in the case of the embedded SE, it is considered to be an MF cryptographic hardware authenticator (as reflected below). Although the latter may be considered more secure, both are reasonable choices at LoA 3/AAL2.	smartphone (under the assumption that the computer platforms and mobile devices are protected and managed to the same degree).

Strong password combined with an SF cryptographic hardware authenticator	This combination can achieve LoA 4/AAL3 since it is an MFA solution with at least one phishing-resistant authenticator that is platform independent (assuming all the other applicable requirements at LoA 4/AAL3 are also met); however, the password must be sent to and verified by the relying party (RP), so it is not a preferred solution at LoA 4/AAL3. However, this combination of authenticators may be an attractive option at LoA 3/AAL2 and is therefore represented as one of the preferred options at LoA3/AAL2 in Figure 2.2.	A strong password combined with an SF FIDO2 security key (appears twice)

MF cryptographic hardware authenticator	Some, but not all, MF cryptographic hardware authenticators can achieve LoA 4/AAL3. In order to achieve LoA 4/AAL3, all applicable requirements at LoA 4/AAL3, as identified in Section 3.3, must be met. While WHfB with a discrete TPM is phishing resistant, it is placed at the higher level of the LoA 3/AAL2 spectrum rather than LoA 4/AAL3 since it does not meet the requirement for a separate platform-independent authenticator; [18] this aspect also places it slightly lower than the other platform-independent authenticators at LoA 3/AAL2. An MF FIDO2 device-bound passkey on a GC-managed smartphone used in conjunction with WebAuthn on a GC-managed computer platform is at the higher end of the LoA 3/AAL2 spectrum and is a preferred authenticator since it is a phishing-resistant authenticator and the authenticator itself is not permanently bound to the computer platform being used to access the target resource. However, it is not LoA 4/AAL3 since a smartphone is not dedicated to the authentication function like a FIDO2 security key or PKI-based smart card. Also, as	There are four examples provided in Figure 2.2: 1. WHfB with a discrete TPM (on a GC-managed computer platform) at the higher level of the LoA 3/AAL2 spectrum 2. MF FIDO2 device-bound passkey on a GC-managed smartphone (with an embedded SE; see explanation in the previous cell to the left) used in conjunction with WebAuthn on a GC-managed computer platform placed slightly higher than WHfB in the LoA 3/AAL2 spectrum 3. Password, PIN or biometric used to unlock or activate an MF FIDO2 security key (used in conjunction with WebAuthn on a GC-managed computer platform) positioned in the middle of the LoA 4/AAL3 spectrum 4. Password, PIN or biometric used to unlock or activate an MF PKI-based smart card (used in conjunction with a GC-managed computer platform) at the higher

noted above, an embedded SE on the smartphone is required in order to be considered an MF cryptographic hardware authenticator. Use of a smartphone with a TEE is also permitted at LoA 3, but it would be considered an MF cryptographic software authenticator and therefore placed slightly lower than a smartphone with an embedded SE.

An MF FIDO2 security key is a phishing-resistant platform-independent authenticator that eliminates the need for a userid and password to be conveyed to the RP and is therefore placed higher in the LoA 4/AAL3 spectrum than a strong password combined with an SF security key. An MF FIDO2 security key is one of the preferred authenticators at LoA 4/AAL3 since it is both phishing resistant and platform independent and does not require a password to be transmitted to the RP.

An MF PKI-based smart card is placed at the highest level in the LoA 4/AAL3 spectrum since it is a phishing-resistant platform-independent authenticator and is less susceptible to certain attack vectors than the other LoA 4/AAL 3 options, particularly

end of the LoA 4/AAL3 spectrum

Although not shown in Figure 2.2, a GC-managed smartphone that is used to access a remote resource with a platform authenticator protected within an embedded SE on that smartphone could be considered equivalent to WHfB with a discrete TPM (under the assumption that the computer platforms and mobile devices are protected and managed to the same degree).

	if the activation factor does not traverse the operating system where it could potentially be exposed to a threat actor. An MF PKI-based smart card is one of the preferred authenticators at LoA 4/AAL3 since it is both phishing resistant and platform independent and does not require a password to be transmitted to the RP	

Special note on central versus local password verification

As illustrated within this section, memorized secrets such as passwords can be combined with another single-factor something you have authenticator (for example, a push notification with number matching or OTP app) to achieve a multi-factor solution and a higher level of assurance. In this case, the password (used as the first factor) needs to be sent to and verified by (or on behalf of) the RP. This requires a representation of the password (for example, a salted hash) to be stored centrally on a remote server so that the user's password can be verified. On the other hand, multi-factor authenticators only require that the password (or PIN) used as the activation factor to unlock the authenticator to be verified locally. As discussed in the draft NIST SP 800-63B-4 (Revision 4, Initial Public Draft), Section A.4, the attack surface and vulnerabilities associated with these two scenarios are different in several respects and should be taken into consideration when selecting the most appropriate solutions. In general, multi-factor authenticators should be used in preference to authenticator combinations that require a password to be verified remotely, where possible.

2.3 Implementation guidance

The purpose of this section is to provide implementation guidance with respect to authenticators at each credential LoA/AAL (abbreviated as LoA/AAL). Verifier and credential service provider (CSP) requirements are not included in this document (these details are provided in ITSP.30.031 v3 and NIST SP 800-63B). Definitions for various terms are either copied or derived from the sources indicated.

2.3.1 Phishing (verifier impersonation) resistance

There are many different definitions for the term “phishing.” Some definitions are very narrow and focus on specific types of attacks while others are more general. In general, phishing is an attempt by a threat actor to trick a user into revealing sensitive information such as passwords or bank account numbers, or to do something that they should not, such as clicking on a malicious link that downloads malware.

In the context of authentication, phishing resistance is defined as the ability of the authentication protocol to detect and prevent disclosure of authentication secrets and valid authenticator outputs to an impostor RP without reliance on the vigilance of the user. [19] More specifically, phishing-resistant authenticators provide strong protection against a type of adversary-in-the-middle attack known as verifier impersonation. This is where a threat actor lures a user to a fake website positioned between the user and the legitimate verifier or RP. [20] The threat actor poses as the legitimate verifier website to the user, and as the user to the legitimate verifier, in order to gain unauthorized access to the user’s account or resources the user is attempting to access. Protection against verifier impersonation is provided by channel binding or verifier name binding.

Note that phishing-resistant authenticators use asymmetric cryptography to support verifier impersonation resistance. Authenticators that require the user to manually enter the authenticator output such as passwords, OOB authenticators with number matching or OTP authenticators are not resistant to verifier impersonation since the authenticator output is not bound to the specific session being authenticated.

At LoA 4/AAL3, at least one of the authenticators must be phishing resistant. It is also strongly recommended that at least one of the authenticators at LoA 3/AAL2 should be phishing resistant. However, there are authenticators and authenticator combinations permitted at LoA 3/AAL2 that are not inherently phishing resistant. In these cases, compensating security measures must be implemented as part of the authentication process in order to mitigate the risk of successful phishing attacks. [21] For example, ensuring that the user is authenticating from a GC-managed device, [22] verifying the device is configured properly, and detecting anomalous geolocations can help to compensate for the lack of phishing-resistant authenticators. In addition, user awareness and training is an important mitigation measure that can help to prevent successful phishing attempts. However, note that these examples are not meant to represent an exhaustive list. Furthermore, as the threat landscape evolves, so too must the compensating security measures, and therefore departments should continually monitor and adjust as necessary.

Special note on the use of compensating security measures

It is important to recognize that proper configuration of these additional security measures is critical to ensure their effectiveness, and skilled IT security resources are required to select and configure them correctly. In addition, authentication solutions provided by different vendors may not

support the same compensating security measures or implement them to the same degree. Furthermore, over time, the introduction of new security configuration options or default settings can degrade these compensating measures, requiring constant review and testing to ensure their continued effectiveness. The correct implementation of these measures may result in additional costs (for example, consulting, auditing, recovery from misconfiguration, and so on) that should be taken into account when choosing to implement MFA methods that require compensating measures. To avoid these challenges, it is strongly recommended that at least one of the authenticators should be phishing resistant.

The guidance for phishing-resistant authenticators is summarized as follows:

- LoA 2/AAL1: no stipulation
- LoA 3/AAL2: recommended, otherwise additional security measures as discussed above need to be implemented
- LoA 4/AAL3: required

Additional sources of information regarding phishing resistance include:

- Cybersecurity and Infrastructure Security Agency: Implementing Phishing-Resistant MFA
- Office of Management and Budget M-22-09: Moving the U.S. Government Toward Zero Trust Cybersecurity Principles [23]

2.3.2 Cryptographic key protection

Cryptographic authenticators may be based on symmetric or asymmetric cryptography and may be implemented in software or hardware. In all cases, the secret keys (in the case of symmetric cryptography) or private

keys (in the case of asymmetric cryptography) used to support the user authentication process must be protected from unauthorized use, disclosure or modification.

How cryptographic authenticators protect the cryptographic keys will depend on the type of authenticator (for example, software versus hardware) as well as the LoA/AAL at which it is used, and may be implemented in several ways including the use of a trusted platform module (TPM), [24] secure element (SE) [25] or trusted execution environment (TEE). [26]

Both ITSP.30.031 v3 (see Table 6) and NIST SP 800-63B (see Sections 5.1.6.1 through 5.1.9.1) provide guidance regarding cryptographic key protection. In addition, specific requirements at each credential LoA/AAL include:

- LoA 2/AAL1: cryptographic keys may be stored on disk or other “soft” media as long as they are strongly protected against unauthorized disclosure by the use of access controls that limit access to the key to only those software components on the device requiring access [27]
- LoA 3/AAL2: suitable TPM-, SE- or TEE-required private signing keys must be generated on and must not be exported from the TPM, SE or TEE, [28] and all cryptographic operations to support user authentication must be performed within the boundary of the TPM, SE or TEE; additional guidance includes:
    - software TPMs as defined by the Trusted Computing Group (see Appendix C) must not be used and, if supported on the computer platform or mobile device, must be disabled
    - MF cryptographic software authenticators that store the cryptographic keys on disk or other soft media are not permitted at LoA 3/AAL2 or higher

    - it is recommended that SF or MF OTP software authenticators used at LoA 3/AAL2 or higher operate on a hardware device such as a GC-managed mobile device (that is, smartphone) physically separate from the general-purpose computer platform

- LoA 4/AAL3: separate dedicated hardware such as a discrete TPM (see Appendix C) or embedded SE is required, the cryptographic keys must be generated and securely stored on dedicated hardware and must not be exportable, and all cryptographic operations to support user authentication must be performed within the boundary of the dedicated hardware; additional guidance includes:
    - platform-independent hardware authenticators dedicated to the authentication function are required (for example, a FIDO2 roaming authenticator or PKI-based smart card) [29]
    - the activation factor used to unlock or activate an MF cryptographic hardware authenticator such as a PKI-based smart card should be securely communicated either through a trusted path (that is, does not flow through the computer platform's operating system) or by entering the activation factor directly on the smart card or card reader (for example, a fingerprint scanner on the card or card reader)

In addition, for MF authenticators, each authentication operation requires the input of the associated activation factor to unlock or activate the authenticator. Submission of the activation factor must be a separate operation from unlocking of the host device (for example, computer platform or smartphone), although the same activation factor used to unlock the host device may be used in the authentication operation. [30]

Cryptographic module validation requirements are addressed in Section 2.3.3 below.

2.3.3 Cryptographic module validation

As stipulated in ITSP.30.031 v3 and NIST SP 800-63B, cryptographic modules must be validated to meet the requirements of Federal Information Processing Standard (FIPS) 140 [31] as follows:

- cryptographic authenticators used at LoA 3/AAL2 must be validated to at least FIPS 140 Level 1
- single-factor cryptographic hardware authenticators used in conjunction with another authenticator at LoA4/AAL3 shall be validated at FIPS 140 Level 1 or higher overall, with at least Level 3 physical security
- multi-factor authenticators used at LoA 4/AAL3 shall be hardware cryptographic modules validated at FIPS 140 Level 2 or higher overall, with at least Level 3 physical security

Note: There is no stipulation at LoA 2/AAL1 for authenticators, although see NIST SP 800-63B, Section 4.1.2, for verifier requirements.

Cryptographic authenticators must operate only in FIPS mode. Only cryptographic algorithms recommended in ITSP.40.111 shall be used, and any relevant security protocols shall be configured as recommended in ITSP.40.062.

2.3.4 Adversary-in-the-middle resistance

“Adversary in the middle” (AitM), [32] formerly referred to as “man in the middle” (MitM), is when a threat actor is positioned between two communicating parties in order to intercept and/or alter data travelling between them. In the context of authentication, the threat actor could be positioned between a user and a credential service provider during enrolment or authenticator binding, or between a user and a verifier. [33]

As stipulated in ITSP.30.031 v3 and NIST SP 800-63B, communication between the user and the verifier must be via an authenticated protected channel (for example, TLS) to provide confidentiality of the authenticator output and resistance to AitM attacks at LoA 2/AAL1 and higher. However, the degree of AitM resistance can vary as described in the next paragraph.

A protocol is considered to be weakly resistant to AitM attacks if it provides a mechanism for a user to determine whether they are interacting with the real verifier or RP, but still leaves the opportunity for the non-vigilant user to reveal an authenticator output to an unauthorized party that can then be used to masquerade as the user to the real verifier or RP. For example, sending a password over server-authenticated TLS is weakly resistant to AitM attacks. The web browser allows the user to verify the identity of the verifier or RP; however, if the user is not sufficiently vigilant, the password can still be revealed to an unauthorized party. A protocol is said to be strongly resistant to AitM attacks if it does not rely on the vigilance of the user to prevent disclosure of authenticator outputs to an unauthorized party masquerading as the verifier or RP. An example of such a protocol is client-authenticated TLS, where the browser and the web server mutually authenticate one another using PKI. Authentication protocols that can detect verifier impersonation, as discussed in Section 2.3.1, are strongly resistant to AitM attacks.

Specific AitM resistance guidance is:

- LoA 2/AAL1: at least weak resistance is required
- LoA 3/AAL2: at least weak resistance is required; however, note that at least one phishing-resistant authenticator is recommended, as stipulated in Section 2.3.1
- LoA 4/AAL3: strong resistance is required, at least one phishing-resistant authenticator is required, as stipulated in Section 2.3.1

2.3.5 Replay resistance

An authentication process resists replay attacks if it is impractical to achieve a successful authentication by recording and replaying a previous authentication message. Protocols that use nonces or challenges to prove the “freshness” of the transaction are resistant to replay attacks. Examples of replay-resistant authenticators are OTP authenticators, cryptographic authenticators and look-up secrets. In contrast, memorized secrets are not considered replay resistant because the authenticator output – the secret itself – is provided for each authentication. [34]

The replay resistance guidance provided in ITSP.30.031 v3 and NIST SP 800-63B is adopted (with the exception of the modified LoA 2/AAL1 guidance) as follows:

- LoA 2/AAL1: recommended [35]
- LoA 3/AAL2: required
- LoA 4/AAL3: required

2.3.6 Authentication intent

The latest guidance available from NIST SP 800-63B-4 (Second Public Draft) describes authentication intent as follows:

> An authentication process demonstrates intent if it requires the claimant to respond explicitly to each authentication or reauthentication request. The goal of authentication intent is to make it more difficult for authenticators (e.g., multi-factor cryptographic authenticators) to be used without the claimant’s knowledge, such as by malware on the endpoint. The authenticator itself **shall** establish authentication intent, although multi-factor cryptographic authenticators **may** establish intent by reentry of the activation factor for the authenticator.

> Authentication intent **may** be established in several ways. Authentication processes that require the claimant's intervention can be used to prove intent (e.g., a claimant entering an authenticator output from an OTP authenticator). Cryptographic authenticators that require user action for each authentication or reauthentication operation can also be used to establish intent (e.g., by pushing a button or reinsertion).
>
> The presentation of biometric characteristics does not always establish authentication intent. For example, using a front-facing camera on a mobile phone to capture a face biometric does not constitute intent, as it can be reasonably expected to capture a face image while the device is used for other non-authentication purposes. In these scenarios, an explicit mechanism (e.g., tapping a software or physical button) **shall** be provided to establish authentication intent.

The NIST authentication intent guidance is adopted as follows:

- LoA 2/AAL1: not required
- LoA 3/AAL2: recommended
- LoA 4/AAL3: required

Note that if a single-factor cryptographic hardware authenticator is used in conjunction with a strong password at LoA 4/AAL3, authentication intent must be established by requiring physical input from the user (e.g., the pressing of a button) in order for the authenticator to operate. However, note that the act of pressing the button to confirm physical presence does not constitute an additional authentication factor.

2.3.7 FIDO-specific considerations

An overview of the FIDO Alliance specifications and W3C WebAuthn recommendation, including some of the technical details relevant to this subsection, are provided as Appendix B. The reader should be familiar with the terminology and concepts discussed in Appendix B before proceeding further with this subsection.

2.3.7.1 Credential storage modality

As discussed in Appendix B, private keys can be stored on the authenticator (referred to as client-side discoverable credentials, or simply discoverable credentials), or they can be stored on the server (referred to as server-side credentials) in encrypted form using a strong symmetric encryption algorithm (the symmetric key used to encrypt the private keys is generated on and never exported from the authenticator). The trade-offs between these two credential storage modalities are discussed in Appendix B.

Guidance regarding credential storage modality is:

- LoA 2/AAL1: no stipulation
- LoA 3/AAL2: client-side discoverable credentials or server-side credentials permitted [36]
- LoA 4/AAL3: client-side discoverable credentials required; server-side credentials are not permitted

2.3.7.2 Synced passkeys

As discussed in Appendix B, FIDO credentials can be backed up and copied (or synced) across multiple devices. This is also referred to as multi-device credentials but, more recently, “synced passkeys” appears to be the preferred term.

While the implementation of synced passkeys provides for an improved user experience, the security of the cryptographic keys depends on the implementations provided by third-party service providers and the

method(s) used for authenticator recovery. It also introduces other issues such as the inability to support attestation. However, the next release of the FIDO2 standards (specifically, the WebAuthn Level 3 and Client to Authenticator Protocol (CTAP) 2.2 specification) are expected to include support for attestation as part of the authentication process, not just during registration of the authenticator, which could lead to support for synced passkey attestation in the future. Furthermore, NIST published interim guidance on "synced authenticators" (or synced passkeys) in April 2024, which has since been included as Appendix B to NIST SP 800-63B-4 (Second Public Draft). While subject to further investigation and guidance from CCCS, synced passkeys may be permitted at LoA 3/AAL2 in the future. However, note that synced passkeys are not suitable for higher assurance requirements and therefore will not be permitted at LoA 4/AAL3.

Guidance regarding synced passkeys is:

- LoA 2/AAL1: no stipulation
- LoA 3/AAL2: not permitted at this time (but subject to change as noted above)
- LoA 4/AAL3: not permitted

2.3.7.3 Attestation

As stated in WebAuthn, "attestation is employed to **attest** to the **provenance** of an authenticator and the data it emits." Verifiable attestation statements are conveyed by the authenticator to the RP during registration [37] and provide a means to reliably determine certain properties about the authenticator (for example, make or model). See the FIDO Attestation white paper for additional information regarding attestation.

While several attestation types are defined, only basic attestation, attestation certification authority or enterprise attestation should be used in the GC enterprise domain. Furthermore, relying parties should specify

(or require) direct or enterprise attestation conveyance (see WebAuthn, Section 5.4.7).

Note that enterprise attestation is a special type of attestation that is "intended for controlled deployments within an enterprise where the organization wishes to tie registrations to specific authenticators." [38] Enterprise attestation supports the ability to uniquely identify each authenticator, and this may become the only permitted attestation type in the future, particularly at LoA 4/AAL3. See Appendix B for additional details and references.

At LoA3/AAL2 and LoA4/AAL3, attestation must be used to prevent the registration or use of unapproved authenticators. Guidance regarding attestation requirements is:

- LoA 2/AAL1: no stipulation (attestation is not required)
- LoA 3/AAL2: required
- LoA 4/AAL3: required

2.3.8 Implementation guidance summary

Table 2.4 provides a summary of the recommendations made in the previous subsections. Note that this is an abbreviated summary only; the referenced section should be consulted for additional details. Also note that there are no stipulations with respect to LoA 1 authenticators since they are not permitted.

Table 2.4: implementation guidance summary

	LoA 2/AAL1	**LoA 3/AAL2**	**LoA 4/AAL3**
Phishing-resistant authenticators (see Section 2.3.1)	No stipulation	Recommended	Required

	LoA 2/AAL1	LoA 3/AAL2	LoA 4/AAL3
Cryptographic key protection (see Section 2.3.2)	Disk or other “soft” storage permitted for SF or MF cryptographic software authenticators	Appropriate TPM, SE or TEE required	Dedicated hardware required (for example, discrete TPM or embedded SE)
Cryptographic module validation (see Section 2.3.3)	No stipulation	Overall level 1	MF cryptographic hardware authenticator: overall Level 2 with Level 3 physical security SF cryptographic hardware authenticator (used in combination with another authenticator): overall Level 1 with Level 3 physical security
AitM resistance (see Section 2.3.4)	Required	Required	Required
Replay resistance (see Section 2.3.5)	Recommended	Required	Required
Authentication intent (see Section 2.3.6)	Not required	Recommended	Required
Credential storage modality (see Section 2.3.7.1)	No stipulation	Client-side discoverable credentials or server-side credentials permitted	Client-side discoverable credentials required (device-bound passkeys only)

	LoA 2/AAL1	LoA 3/AAL2	LoA 4/AAL3
Synced passkeys (see Section 2.3.7.2)	No stipulation	Not permitted at this time (subject to change)	Not permitted (device-bound passkeys only)
Attestation (see Section 2.3.7.3)	No stipulation	Required	Required

2.4 Additional considerations

The purpose of this section is to identify additional technical considerations that need to be addressed as part of a comprehensive authentication solution in addition to the technical aspects associated with authenticators as discussed in Section 2.3. It is not meant to represent an exhaustive list, nor is it intended to provide a comprehensive treatment of the identified topics. Additional details will need to be addressed as part of related GC MFA initiatives.

2.4.1 Authenticator binding

Binding authenticators to a specific identity is critical to ensure that the right user is accessing the appropriate resources (for example, accounts, applications, services and information). Binding a user to one or more authenticators occurs during the user enrolment process, and it can also occur later by using previously issued authenticators to bind new authenticators.

Guidance with respect to authenticator binding is provided in Section 6.1 of NIST SP 800-63B. Section 6.1.2.1 of NIST SP 800-63B addresses the use of authenticators at a given level of assurance to bind other authenticators at the same or lower LoA. Section 6.1.2.2 of NIST SP 800-63B addresses the possibility of using a single-factor authenticator to add a different type of second-factor authenticator, potentially raising the level of assurance from

LoA 2/AAL1 to LoA 3/AAL2 (depending on the authenticators). An example is when a user authenticates to a web service using a previously registered userid and password and then registers a **something you have** authenticator (for example, a FIDO authenticator) to be used as a second factor from that point forward. [39]

However, use of an existing userid and password to bind an additional second-factor authenticator does not guarantee that the user account has not already been compromised, so the user attempting to register an additional authenticator could be an imposter. Although NIST recommends that a notification of the event should be sent to the user via an out-of-band method (for example, a registered email account), this may not be sufficient, by itself, to detect account compromise. Therefore, additional measures should be taken to add assurance that the user is whom they claim to be. This can be accomplished if the second factor is already bound to that user (for example, a GC-managed smartphone issued to the user is to be used as the second factor). If the second authenticator has not been previously bound to the user, then other methods should be employed. For example, verification of a one-time passcode sent or provided to the **right** user via a reliable out-of-band means can be used to establish additional confidence that the user involved in the registration is whom they claim to be. The out-of-band mechanism should be time-bound (for example, expires in 10 business days).

2.4.2 Authenticator recovery

Appropriate authenticator recovery mechanisms must be in place to re-establish user access in the event that an existing authenticator is lost, stolen or damaged. Recovery methods will vary depending on the type of authenticator, but in all cases, they must be at least as strong and secure as the method(s) used to establish the original authenticator binding. Section 6.1.2.3 of the NIST SP 800-63B provides additional guidance related

to authenticator recovery. Furthermore, users need to be informed of their responsibilities such as immediately reporting lost or stolen authenticators (see Section 6.2 of NIST SP 800-63B for additional information).

Note that there are other authenticator life-cycle management considerations that need to be taken into consideration as well, including renewal (Section 6.1.4 of NIST SP 800-63B), suspension (Section 6.2 of NIST SP 800-63B), expiration (Section 6.3 of NIST SP 800-63B), and revocation (Section 6.4 of NIST SP 800-63B).

2.4.3 Reauthentication

Reauthentication is used to determine if the previously authenticated user is still present on a given session, either due to user inactivity or after a certain amount of time has passed. The NIST Digital Identity Guidelines (see Sections 4.1.3, 4.2.3, 4.3.3 and 7.2 of NIST SP 800-63B) describe reauthentication requirements at each AAL. Those requirements are adopted with additional clarification as follows: [40]

- LoA 2/AAL1: reauthentication is required at least once every 30 days regardless of user activity; RPs are free to determine their own user inactivity timeout requirements, if any
- LoA 3/AAL2: reauthentication is required at least once every 12 hours regardless of user activity; reauthentication is also required after 30 minutes of user inactivity (although the user may be prompted to see if they are still present before the inactivity threshold is reached [41]); reauthentication of a session that has not yet reached its time limit may require only a memorized secret or a biometric in conjunction with the still-valid session secret
- LoA 4/AAL3: reauthentication is required at least once every 12 hours regardless of user activity; reauthentication is also required after 15 minutes of user inactivity (although the user may be prompted to

see if they are still present before the inactivity threshold is reached [42]); reauthentication requires MFA

If the user does not successfully reauthenticate within one minute of a reauthentication request, the session shall be terminated. If reauthentication is successful, both the session and user inactivity timeout thresholds should be reset.

Note that NIST points out (see NIST SP 800:63 Digital Identity Guidelines – Frequently Asked Questions) that there may be other means of ensuring that the user's device is not used by an unauthorized party, such as locking the user device at the operating system level and requiring local authentication for the user to unlock the device. However, it is not always possible for a RP to know the state of the user's device, and, in these cases, the session timeouts act as a reasonable control. However, if the RP knows that the user is using a managed system that requires screen lock by a system policy, the RP may be able to relax its session timeouts for such devices. Alternatively, an RP could manage most of the user's interactions at a lower AAL and switch to a higher AAL for sensitive operations. In these cases, the lesser-privileged session would be allowed to last much longer than the privilege escalation. Adopting compensating controls such as these are part of the risk assessment process.

[Note: While Section 7.2.1 of NIST SP 800-63B discusses the requirement for an RP to convey the maximum authentication age to an IdP/CSP in a federated scenario, it does not discuss the possibility of forcing reauthentication, which is supported by both the Security Assertion Markup Language (SAML) and OpenID Connect federation protocols. This can be used by RPs to force reauthentication regardless of the status of the authenticated session between the user and the Identity Provider (IdP). Fortunately, the draft Revision 4 of the NIST Digital Identity Guidelines has added the notion of forced authentication in Section 5.6 in SP-800-63C-4.]

2.4.4 Step-up authentication

Step-up authentication, which should not be confused with reauthentication, as described in the previous subsection, is when there is a requirement to elevate the LoA of a given authenticated session (for example, the user is authenticated at LoA 2 and requests access to a resource that requires LoA 3 authentication). RPs should be capable of prompting the user (or redirecting the user to an IdP in a federated environment) to authenticate at the higher level of assurance to support step-up authentication when needed.

2.4.5 Centralized authentication

As noted in Multi-Factor Authentication Considerations and Strategy for GC Enterprise IT Services, a centralized enterprise authentication solution with the ability to support multiple MFA methods based on open, industry-accepted standards is a critical component of the overall GC identity, credential and access management (ICAM) strategy. Authentication should be centralized to the maximum extent possible in order to remove the burden of supporting multiple authentication mechanisms from individual applications and services, provide support for single sign-on and enable federation. Centralized authentication also helps to more easily support continuous, risk-adaptive access control, as discussed in the next subsection.

2.4.6 Risk-adaptive access control

Risk-adaptive access control (or simply adaptive access control) is a dynamic access control model that uses multiple security metrics such as the strength of the user authentication method, the level of assurance of

the session connection between the system and a user, the geolocation of the user, and other considerations in order to render informed access control decisions. [43]

While not meant to be an exhaustive representation, Figure 2.3 helps to illustrate the concept of a risk-adaptive access control model. As illustrated in the diagram, user authentication is only one of many considerations that are used as a basis for access control decisions. For example, device attestation is an important complement to user authentication, as it provides a reliable and secure method to verify the identity and health of the device being used to access GC resources.

Furthermore, access control decisions should be dynamic and adapt to policy changes and real-time threat and risk assessment, as illustrated by the policy and telemetry inputs in Figure 2.3. This model forms the very foundation for supporting continuous, risk-adaptive access control in support of zero-trust principles. [44] It also provides the ability to support compensating measures so that weaknesses in one area can be complemented with other areas to help inform the overall access control decision.

Figure 2.3: conceptual risk-adaptive access control model

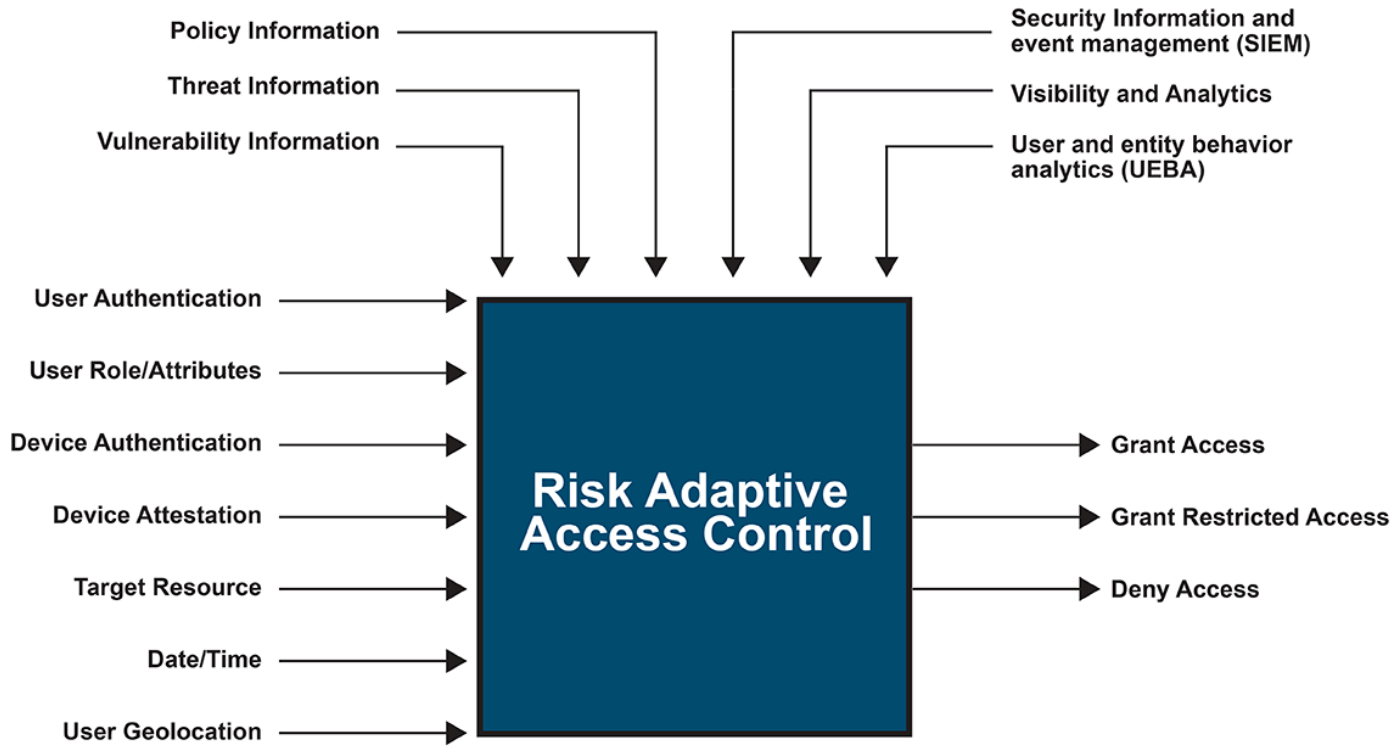

▶ Figure 2.3 - Text version

2.4.7 Use of biometrics

As stated in the NIST Digital Identity Guidelines (see Section 5.2.3 of NIST SP 800-63B), "biometrics SHALL be used only as part of multi-factor authentication with a physical authenticator (**something you have**)." Therefore, biometrics can only be used as a second factor in conjunction with something the user has, such as a multi-factor cryptographic hardware authenticator with an embedded sensor such as a fingerprint scanner. However, when using biometrics as a second factor, certain requirements need to be met, including false match rate and presentation attack detection (PAD) technologies (for example, liveness detection).

It should also be noted that some biometric technologies are designed with user convenience in mind rather than security and, therefore, may not be adequate depending on the level of assurance required. Therefore, caution

must be exercised when using biometrics as an authentication factor. See Section 5.2.3 of NIST SP 800-63B for additional information.

2.4.8 Bring your own authenticator

As noted previously in Section 2.2, all devices used to access GC internal resources should be managed by the GC. This includes mobile devices such as smartphones. However, it is recognized that some departments may have no choice but to implement MFA using personal devices, particularly for contractors. While this does add a measure of additional security over a single-factor solution based solely on a userid and password, the fact that the mobile device is unmanaged makes it more susceptible to compromise than a GC-managed device. However, mobile device management, unified endpoint management (UEM) or mobile application management could potentially be used to manage corporate data and applications on personally owned mobile devices, which would significantly reduce the risk as compared to an unmanaged device. When used in conjunction with a GC-managed computer platform, this could possibly raise the level of assurance from LoA 2 to LoA 3, assuming the appropriate security measures and best practices are in place (for example, separation between the work space and personal space is maintained at all times, rooted and compromised devices are detected, and so on). See End User Device Security for Bring-Your-Own-Device (BYOD) Deployment Models: ITSM.70.003 for additional information.

In any case, the strongest possible methods of authentication should be used. For example, the use of text messages (that is, SMS) is strongly discouraged in favour of more secure mechanisms such as a push notification with number matching or an OTP application. Use of phishing-resistant authenticators such as FIDO2 device-bound passkeys stored on the user’s mobile device or a user-owned FIDO2 security key (either of which could be used to authenticate via WebAuthn on a GC-managed

computer platform) may also be an option. If this approach is adopted, attestation must be used to verify that the authenticator meets the minimum requirements associated with accessing the target resource. Also note that the use of a personal device as an authenticator must not unintentionally introduce the possibility that internal GC resources can be accessed with that device. More specifically, bring-your-own authenticators (BYOAs) are expected to be used in combination with a GC-managed computer platform.

The specific use cases and the need for any additional risk mitigation measures should be carefully considered before adopting any BYOA approach. In addition, privacy considerations must also be assessed.

3. Summary and recommendations

This document provides detailed technical guidance regarding authenticators and recommends specific authenticators that are suitable to provide MFA solutions within the GC enterprise domain as follows:

- For non-privileged users conducting day-to-day business activities, authenticator(s) should meet all requirements identified at credential LoA 3/AAL2. Although some of the acceptable authenticators at credential LoA 3/AAL2 do not support phishing resistance, it is strongly recommended that at least one of the authenticators should be phishing resistant. If the authenticator (or combination of authenticators) does not support phishing resistance, compensating measures, as discussed in Section 2.3.1, must be implemented to help mitigate the risk that the user’s credentials may be compromised through phishing attempts. It is also recommended that one of the authenticators should be physically separate from the computer platform or device being used to access the target resource. In all

cases, the computer platforms, mobile devices and authenticators must be managed by (or on behalf of) the GC. Therefore, the preferred authenticators include:

- phishing-resistant multi-factor FIDO2 device-bound passkeys [45] on a GC-managed smartphone (used to authenticate to a remote RP via WebAuthn on a separate computer platform), or
- a strong, properly managed password [46] (provided to the RP) combined with a phishing-resistant single-factor FIDO2 security key with a push button to verify physical presence (used to authenticate to the RP via WebAuthn on a separate computer platform)

Other acceptable authenticators that may be used to support MFA for typical users include:

- Windows Hello for Business with a TPM (unlocked with an activation factor such as a PIN or biometric entered or presented by the user)
- a strong, properly managed password combined with a single-factor OTP hardware authenticator (for example, an RSA SecurID hardware device) with additional security measures implemented as part of the authentication process to compensate for the lack of phishing-resistant authenticators
- a passwordless multi-factor out-of-band authenticator (for example, the MS Authenticator app using the phone sign-in authentication method) with additional security measures implemented as part of the authentication process to compensate for the lack of phishing-resistant authenticators
- a strong, properly managed password combined with a one-time password (OTP) app (for example, Google authenticator or MS Authenticator with the OTP authentication method), with

additional security measures implemented as part of the authentication process to compensate for the lack of phishing-resistant authenticators
    - a strong, properly managed password combined with a push notification with number matching (for example, Google prompt or MS authenticator with the push notification authentication method), with additional security measures implemented as part of the authentication process to compensate for the lack of phishing-resistant authenticators
- For highly privileged users such as system administrators, and for high-profile users (for example, chief financial officers), authenticator(s) should meet all requirements identified at credential LoA 4/AAL3. At least one of the authenticators must be phishing resistant, and an authenticator that is physically separate from the computer platform or device used to access the target resource must be used. In all cases, the computer platforms, mobile devices and authenticators must be managed by (or on behalf of) the GC. Preferred authenticators include:
    - a phishing-resistant multi-factor PKI-based smart card or
    - a phishing-resistant multi-factor FIDO2 security key

  A strong, properly managed password provided to the RP combined with a phishing-resistant single-factor FIDO2 security key with a push button to verify physical presence (used to authenticate to the RP via WebAuthn on a separate computer platform) is also acceptable, but not preferred, since it involves the use of a password that must be sent to the RP for verification. However, this authenticator combination can also be used for typical users and, in fact, is a preferred solution for that class of users as noted above.

  Note that use of authenticators that meet credential LoA 3/AAL2 requirements may be acceptable for less privileged administrators

such as an application or SaaS administrator [47] (subject to risk assessment); however, at least one of the authenticators must be phishing resistant. Also note that authenticators that meet credential LoA 4/AAL3 requirements may also be used to support credential LoA 3/AAL2 requirements where it makes sense to do so.

While the scope of this document is focused on technical requirements, it should be recognized that the selection of the most suitable authenticator(s) for a given department will depend on a variety of other considerations, including the ability to leverage existing investments, overall cost, user experience and more. Ultimately, the goal is to strike **the right balance between security, manageability, interoperability, cost and user experience leading to the deployment of suitable MFA solutions throughout the GC**.

In terms of acquiring authenticators, departments are expected to use enterprise or shared IT solutions, assets and services to avoid duplication, when available and appropriate, as stipulated in the *Policy on Service and Digital*, Section 4.4.2.3. To that end, departments can leverage enterprise services that support MFA, including leveraging supply arrangements that are established by Shared Services Canada.

Finally, it should be noted that this is a complex and evolving area, and both authentication technologies and threats to those technologies continue to change, sometimes quite rapidly. Therefore, the recommendations made within this document may be subject to change over time. Furthermore, this document identifies several evolving technical considerations that may have an impact on the recommendations in the future, including:

- potential support for synced passkeys at LoA 3/AAL2 (refer to Section 2.3.7.2)
- the potential role and viability of enterprise attestation (refer to Section 2.3.7.3)

- evolving user authentication guidance, particularly ITSP.30.031 v4 (currently being drafted) and Revision 4 of the NIST Digital Identity Guidelines (also being drafted)

Although outside the scope of this document, design and implementation details regarding the related topics identified in Section 2.4 need to be developed, including binding authenticators to individual users, authenticator recovery options, comprehensive risk-adaptive access control solutions and BYOA scenarios.

For enquiries and interpretation of this document, contact the TBS Cyber Security Division at zztbscybers@tbs-sct.gc.ca.

4. References

1. CBC News, "Spy agency chief says new powers would help stop cyberattacks before they happen."
2. Verizon, 2024 Data Breach Investigations Report.
3. CBC News, "Cyberattacks targeting CRA, Canadians' COVID-19 benefits have been brought under control: officials."
4. Treasury Board of Canada Secretariat, *Recommendations for Two-factor User Authentication within the GC Enterprise Domain*; superseded by this document.
5. Treasury Board, *Policy on Service and Digital*.
6. Treasury Board, *Directive on Service and Digital*.
7. Treasury Board, *Directive on Identity Management*: Appendix A – Standard on Identity and Credential Assurance.
8. Treasury Board of Canada Secretariat, *Guideline on Defining Authentication Requirements*.
9. Treasury Board of Canada Secretariat, *Guideline on Identity Assurance*.

10. Canadian Centre for Cyber Security, *User Authentication Guidance for Information Technology Systems* (ITSP.30.031 v3).
11. Canadian Centre for Cyber Security, *User Authentication Guidance for Information Technology Systems* (ITSP.30.031 v4); not yet available – currently being drafted.
12. Canadian Centre for Cyber Security, *Cryptographic Algorithms for Unclassified, Protected A, and Protected B Information: ITSP.40.111*.
13. Canadian Centre for Cyber Security, *Guidance on Securely Configuring Network Security Protocols (ITSP.40.062)*.
14. Canadian Centre for Cyber Security, *Secure Your Accounts and Devices with Multi-Factor Authentication (ITSAP.30.030)*
15. Canadian Centre for Cyber Security, *Annex 2: Information System Security Risk Management Activities (ITSG-33)*
16. US National Institute of Standards and Technology (NIST), *Digital Identity Guidelines: Revision 3* (NIST SP 800-63-3 Series).
17. US National Institute of Standards and Technology (NIST), *Digital Identity Guidelines*Initial Public Draft, Revision 4 (NIST SP 800-63-4 Series).
18. US National Institute of Standards and Technology (NIST), *Zero Trust Architecture* (NIST SP 800-207).
19. W3C, *Web Authentication: An API for Accessing Public Key Credentials – Level 2*.
20. FIDO Alliance, *User Authentication Specifications Overview*.
21. FIDO Alliance, Client-*to-Authenticator Protocol*.
22. Gartner, *Zero Trust is an Initial Step on the Roadmap to CARTA*; available through Gartner subscription.
23. Gartner, *Guidance for Selecting User Authentication Solutions*; available through Gartner subscription.

5. Key terms

Key terms used within this document are provided below. The definitions are copied or derived from the sources indicated in the third column. Text in square brackets is added to some of the definitions to provide additional clarity.

Term	Definition	Source
Access control (logical)	The process of granting or denying specific requests to obtain and use information and related information processing services [An access control decision can be based on numerous inputs, parameters and telemetry; user authentication is just one facet]	NIST Computer Security Resource Center Glossary See Section 2.4.6 in this document
Activation	The process of inputting an **activation factor** into a **multi-factor authenticator** to enable its use for **authentication**	NIST SP 800-63-4
Activation factor	An additional **authentication factor** that is used to enable successful **authentication** with a **multi-factor authenticator**. Since all multi-factor authenticators are physical hardware authenticators, activation factors are either memorized secrets or biometric factors [Something the user enters or presents such as a password, PIN or biometric to unlock a multi-factor authenticator]	NIST SP 800-63-4

Adversary in the middle (AitM)	An attack in which a threat actor is positioned between two legitimate communicating parties in order to intercept and/or alter data travelling between them. In the context of authentication, the threat actor would be positioned between the user and verifier, between the user and CSP during enrolment, or between user and CSP during authenticator binding. [Formerly known as “man in the middle”; also referred to as “attacker in the middle”]	NIST SP 800-63-3 NIST SP 800-63B-4 (Revision 4, Initial Public Draft)
Authentication	Verifying the identity of a user, process or device, often as a prerequisite to allowing access to a system’s resources	NIST SP 800-63-3
Authentication factor	The three types of authentication factors are something you know, something you have, and something you are. Every authenticator has one or more authentication factors.	NIST SP 800-63-3
Authenticator	Something the user possesses and controls (typically a cryptographic module or password) that is used to authenticate the user’s identity Also see the definition for **credential** below.	NIST SP 800-63-3

Credential	An object or data structure that authoritatively binds an identity, via an identifier or identifiers, and (optionally) additional attributes, to at least one authenticator possessed and controlled by a user [Sometimes the term credential and the term authenticator are used interchangeably; however, they are different but related terms. An authenticator, as defined above, is used in conjunction with a credential in order to authenticate a user. A classic example is the use of asymmetric cryptography where the authenticator uses a private signing key unique to the user to digitally sign data (typically a nonce and other data in a challenge-response protocol) and the public key certificate of the user is the credential that is used by the RP to verify the digital signature]	NIST SP 800-63-3
Credential level of assurance (LoA)	The level of confidence the individual has maintained control over a credential that has been entrusted to them and that the credential has not been compromised There are four levels of confidence defined as follows: • LoA 4: very high confidence • LoA 3: high confidence • LoA 2: some confidence • LoA 1: little confidence	Standard on Identity and Credential Assurance

Device-bound passkey	A FIDO2 discoverable credential that is bound to a single authenticator. For example, FIDO2 security keys typically hold device-bound passkeys as the credential cannot leave the device. Device-bound passkeys have been previously referred to as **single-device passkeys**	Passkeys.dev Terms – Device-bound passkey
Embedded hardware security component	A hardware component contained within an endpoint (for example, a computer platform or smartphone) that provides one or more dedicated security services. Embedded hardware security components are separate and apart from the general-purpose computing environment (that is, from the main central processing unit) within the endpoint. They are characterized by dedicated processing and memory resources, have a defined physical boundary, and communicate with other endpoint components through a defined interface and protocol. Depending on their implementation, embedded hardware may be capable of achieving FIPS 140 validation at Level 2 (or higher) overall and Level 3 (or higher) physical security. Examples include a discrete TPM or embedded SE.	ITSP.30.031 v4
FIDO2 security key	A **hardware authenticator** that conforms to the FIDO2 specifications or recommendations	

Hardware authenticator	A purpose-built hardware component that is physically separate and apart from the general-purpose computing environment of a user device. Hardware authenticators are characterized by a physical boundary (device packaging, container), have dedicated processing and storage resources, and communicate with the general-purpose computing environment of a user device over a defined physical interface and protocol. Hardware authenticators can be completely separate from the user device (for example, a smart card), or they can be an embedded component within the user device (for example, a trusted platform module or SE). Hardware authenticators are capable of achieving FIPS 140 Level 3 physical security protections.	ITSP.30.031 v4

Mobile device	A portable computing device that has a small form factor such that it can easily be carried by a single individual; is designed to operate without a physical connection (for example, wirelessly transmit or receive information); possesses local, non-removable data storage; and is powered on for extended periods of time with a self-contained power source. Mobile devices may also include voice communication capabilities, on-board sensors that allow the device to capture (for example, photograph, video, record, or determine location) information, and/or built-in features for synchronizing local data with remote locations. Examples include smartphones, tablets, and e-readers. [For the purposes of this document, references to mobile devices typically denote a smartphone, but there are limited use cases where the device could be a basic cell phone (for example, basic cell phones can support text messaging, but do not have the sophisticated functionality of a smartphone, notably the ability to host and execute mobile apps). Note that the term “smartphone” is also used within this document, particularly when it is used in the context of a specific example or use case.]	NIST SP 800-53 Revision 5 Also referenced in NIST SP 800-124r2

Multi-factor authentication (MFA)	Authentication that requires more than one distinct authentication factor. MFA can be performed using a multi-factor authenticator or a combination of appropriate authenticators that provide different authentication factors [Note that MFA and two-factor authentication (2FA) are sometimes used interchangeably but, in fact, they are not always synonymous. 2FA means the use of exactly two authentication factors whereas MFA means the use of two or more authentication factors.]	NIST SP 800-63-3
Multi-factor (MF) authenticator	An authenticator that provides more than one distinct authentication factor, such as a cryptographic hardware authenticator with an integrated biometric sensor that is required to activate the authenticator [or requires a password or PIN to activate the authenticator]. [The authentication factor used to unlock or activate the MF authenticator is referred to as an **activation factor**]	NIST SP 800-63-3
Nonce	A random or non-repeating value that is included in data exchanged by a [authentication] protocol, usually for the purpose of guaranteeing the transmittal of live data rather than replayed data, thus detecting and protecting against replay attacks.	NIST Computer Security Resource Center: Glossary

Passkey Also see definitions for synced passkey and device-bound passkey	The high-level, end-user-centric term for a FIDO2/WebAuthn discoverable credential. Like “password,” “passkey” is a common noun intended to be used in everyday conversations and experiences (usage: “a passkey” or “passkeys”). Passkeys are designed to be used without additional login challenges. All passkeys can be used with modern sign in experiences like the Autofill UI [user interface] or with a “Sign in with a passkey” button. From the technical side, there are two flavours of passkeys: synced and device-bound. [Note that passkeys are discoverable credentials, and the latest WebAuthn Level 3 W3C Working Draft uses these terms synonymously.]	Passkeys.dev Terms: Passkey WebAuthn Level 3 W3C Working Draft
Phishing resistance	The ability of the authentication protocol to detect and prevent disclosure of authentication secrets and valid authenticator outputs to an impostor RP without reliance on the vigilance of the subscriber	NIST SP 800-63B-4 (Revision 4, Initial Public Draft) (Section 5.2.5)
Platform authenticator	See definition in Appendix B.	FIDO Alliance
Platform-bound authenticator	A generic term for an authenticator that is physically and permanently attached to the computer platform on which user authentication is taking place. Examples include a discrete TPM or embedded SE. Also see the definition for **embedded hardware security component**	Generic term derived from FIDO’s definition of a platform authenticator and as an antithesis to FIDO’s definition of a roaming authenticator

Platform independent authenticator	A generic term for an authenticator that is physically separate (or removable) from the computer platform on which user authentication is taking place. Examples include FIDO2 roaming authenticators and PKI-based smart cards.	Generic term derived from FIDO's definition of a roaming authenticator and as an antithesis to FIDO's definition of a platform authenticator
Relying party (RP)	An entity that relies upon the user's authenticator(s) and credentials or a verifier's assertion of a user's identity, typically to process a transaction or grant access to information or a system	NIST SP 800-63-3
Roaming authenticator	A FIDO authenticator usable with any device the user is trying to sign-in from. Roaming authenticators attach to users' devices in using USB, near-field communication (NFC), and/or Bluetooth. These authenticators are often referred to as security keys. A smartphone can also act as a roaming authenticator using FIDO cross-device authentication	Passkeys.dev Terms: Roaming authenticator
Replay attack	An attempt to achieve successful authentication by recording and replaying a previous authentication message	NIST SP 800-63B (Section 5.2.8)
Secure element (SE)	A tamper-resistant secure hardware component that is used in a device to provide the security, confidentiality, and multiple application environment required to support various business models. May exist in any form factor, such as embedded or integrated SE, SIM, smart card, smart microSD, and so on.	GlobalPlatform Technology: Root of Trust Definitions and Requirements – Version 1.1

Software authenticator	An authenticator that consists of one or more software components implemented on a general-purpose computing environment of a user device (for example, a laptop or smartphone). Software authenticators are characterized by the process separation mechanisms of the operating system. Software authenticators may take advantage of **embedded hardware security components** or a **trusted execution environment (TEE)** to enhance the robustness of the logical authenticator boundary or to store authenticator secrets. Software authenticators are capable of achieving FIPS 140 Level 1 validation.	ITSP.30.031 v4
Synced passkey	A FIDO2 discoverable credential that can reliably be used for bootstrapping sign-in, without requiring other login challenges such as passwords and OTPs. “Reliable” here means that the passkey should be available to, and usable by, the user whenever they need to sign in. This availability can be achieved through different means: for example, passkey providers could sync passkeys in real time across a user’s devices, restore passkeys from a backup whenever a user sets up a new device, offer passkeys across different contexts (a passkey established from an app can be used in the browser when visiting the app’s website), or allow users to exercise passkeys across devices (by, say, using the passkey from a nearby phone when signing in from a laptop).	Passkeys.dev Terms – Synced passkey

Trusted execution environment (TEE)	A secure area of the main processor of a connected device that ensures sensitive data is stored, processed and protected in an isolated and trusted environment. As such, it offers protection against software attacks generated in the Rich Operating System. The TEE's ability to offer safe execution of authorized security software, known as trusted applications (TAs), enables it to provide end-to-end security by protecting the execution of authenticated code, confidentiality, authenticity, privacy, system integrity and data access rights. Comparative to other security environments on the device, the TEE also offers high processing speeds and a large amount of accessible memory. The primary purpose of the isolated execution environment, provided by the TEE, is to protect device and TA assets. An approach in which additional execution modes are introduced to a central processing unit in order to allow for trusted processing. The authenticator boundary is achieved using a combination of temporal and physical separation on the same central processing unit that is used for general-purpose computing. The degree to which resource isolation is achieved depends on the technology implementation.	Introduction to Trusted Execution Environments ITSP.30.031 v4

Trusted platform module (TPM)	Many of the published definitions for a TPM suggest that it is a stand-alone silicon chip attached to the motherboard of a computer platform that is dedicated to security functions such as cryptographic operations and key storage. While this is a legitimate type of TPM (referred to as a “discrete TPM” by the trusted computing group), there are actually different types of TPMs, as discussed in Appendix C.	See Appendix C.
Verifier	An entity that verifies the user’s identity by verifying the user’s possession and control of one or two authenticators using an authentication protocol	NIST SP 800-63-3

6. Abbreviations

AAL	authenticator assurance level
AitM	adversary in the middle
API	application programming interface
Autofill UI	user interface
BYOA	bring your own authenticator
CBC	Canadian Broadcasting Corporation
CRA	Canada Revenue Agency
CSP	credential service provider
CTAP 2.2	FIDO Client to Authenticator Protocol 2.2 Specification
FIDO	fast identity online
FIPS 140	Federal Information Processing Standard 140
GC	Government of Canada
ICAM	identity, credential and access management

AAL	authenticator assurance level
IT	information technology
ITSAP	IT Security Awareness Publication
ITSP	IT Security Publication
LoA	level of assurance
MF	multi-factor
MFA	multi-factor authentication
microSD	Micro SD (secure digital) card
MitM	man in the middle
MS	Microsoft
NFC	near-field communication
NIST	National Institute of Standards Technology
OOB	out of band
OTP	one-time password
PIN	personal identification number
PIV	personal identity verification
PKI	public key infrastructure
RP	relying party
RSA	RSA Security (rsa.com)
SaaS	software as a service
SAML	Security Assertion Markup Language
SE	secure element
SF	single factor
SIM	subscriber identity module
SMS	short messaging service
SP	special publication
SS7	Signalling System 7

AAL	authenticator assurance level
TA	trusted application
TBS	Treasury Board of Canada Secretariat
TEE	trusted execution environment
TLS	transport layer security
TPM	trusted platform module
W3C	World Wide Web Consortium
WebAuthn	Web Authentication: An API for Accessing Public Key Credentials: Level 2
WHfB	Microsoft Windows Hello For Business

Appendix A: Understanding Levels of Assurance

▶ In this section

A-1 Introduction

The purpose of this appendix is to provide an overview of the four levels of assurance model currently used by the Government of Canada (GC) and to compare and contrast that model with the *US National Institute of Standards and Technology (NIST) Digital Identity Guidelines* (Revision 3).

A-2 Government of Canada levels of assurance

The degree of confidence in the GC's identity, credential and authentication is expressed as levels of assurance (LoA) based on a four-tiered model ranging from one to four, where LoA 1 is the lowest level of confidence and LoA 4 is the highest. This model is consistent with ISO/IEC 29115 and, until

June 2017, was also in alignment with user authentication guidance published by the NIST (see the Relationship to NIST User Authentication Guidance section below for additional information).

The GC has developed several policy instruments and guidance documents that address the definitions and requirements associated with each level of assurance, including:

- *Directive on Identity Management – Appendix A: Standard on Identity and Credential Assurance*: defines the four assurance levels associated with identity assurance and credential assurance
- *Guideline on Defining Authentication Requirements*: describes a methodology for determining the minimum assurance level needed for user authentication in a given context
- *Guideline on Identity Assurance*: specifies the minimum requirements to establish the identity of an individual for a given level of assurance
- *User Authentication Guidance for Information Technology Systems (ITSP.30.031 v3)*: provides technical guidance on user authentication requirements at each level of assurance

Note that the GC separates identity assurance from credential assurance. The Standard on Identity and Credential Assurance (published on July 1, 2019) establishes the four levels of assurance associated with identity and credentials as indicated in the following tables.

Table A1: identity assurance

Assurance level	Description
4	**Very high confidence** required that an individual is who they claim to be
3	**High confidence** required that an individual is who they claim to be

Assurance level	Description
2	**Some confidence** required that an individual is who they claim to be
1	**Little confidence** required that an individual is who they claim to be

Table A2: credential assurance

Assurance level	Description
4	**Very high confidence** required that an individual has maintained control over a credential that has been entrusted to them and that the credential has not been compromised
3	**High confidence** required that an individual has maintained control over a credential that has been entrusted to them and that the credential has not been compromised
2	**Some confidence** required that an individual has maintained control over a credential that has been entrusted to them and that the credential has not been compromised
1	**Little confidence** required that an individual has maintained control over a credential that has been entrusted to them and that the credential has not been compromised

The July 1, 2019, version of the Standard on Identity and Credential Assurance replaced the now archived version of the Standard on Identity and Credential Assurance (archived), which was published on February 1, 2013. The original February 1, 2013, version included the notion of harm or injury in the assurance level definitions, which is no longer included in the updated version. [48] However, the older definitions continue to be reflected in the *Guideline on Defining Authentication Requirements*, which was published on November 30, 2012, and has not been updated since.

The *Guideline and Defining Authentication Requirements* describes a technology-independent process that business owners can use to determine the minimum authentication assurance level required to achieve program objectives, deliver a service or properly execute a transaction. Once the minimum assurance level has been determined, requirements for achieving that assurance level can be derived from the appropriate sources. Specifically, the *Guideline on Identity Assurance* identifies identity assurance requirements at each LoA, and the *User Authentication Guidance for Information Technology Systems (ITSP.30.031 v3)* addresses credential assurance requirements as well as the supporting authentication process requirements.

As specified in ITSP.30.031 v3, the overall authentication level of assurance is the "low watermark" of the level of assurance associated with each of the following areas:

- identity proofing, registration and issuance
- authenticator [49] requirements
- authenticator and credential management
- authentication protocol
- assertion/federation requirements
- event logging
- security assurance

For example, if it is determined that the overall authentication assurance must be LoA 3, all of these areas would have to meet the established requirements for either Level 3 or Level 4. [50] If any one of these areas were only Level 2, then the authentication would only be at assurance Level 2 even if the rest were Level 3 or higher. Figure A11 and are provided to help illustrate this concept.

Table A1: Overall level of assurance – Example 1

Identity Proofing, Registration and Issuance	1	2	3	4
Authenticator Requirements	1	2	3	4
Authenticator and Credential Management	1	2	3	4
Authentication Process/Protocol	1	2	3	4
Assertions/Federation	1	2	3	4
Event Logging	1	2	3	4
Security Assurance	1	2	3	4
Overall Authentication Assurance		2		

▶ Table A1 - Text version

Table A2: Overall level of assurance – Example 2

Identity Proofing, Registration and Issuance	1	2	3	4
Authenticator Requirements	1	2	3	4
Authenticator and Credential Management	1	2	3	4
Authentication Process/Protocol	1	2	3	4
Assertions/Federation	1	2	3	4
Event Logging	1	2	3	4
Security Assurance	1	2	3	4
Overall Authentication Assurance			3	

▶ Table A2 - Text version

As alluded to above, the requirements at each LoA associated with identity proofing, registration and issuance are specified in the *Guideline on Identity Assurance*, and the requirements at each LoA for all the other areas is specified in ITSP.30.031 v3.

A-3 Relationship to NIST user authentication guidance

The US National Institute of Standards and Technology (NIST) published its first user authentication guidance in 2004 based on the four levels of assurance established by the US White House Office of Management and Budget Memorandum OMB M-04-04. The user authentication guidance was referred to as the Electronic Authentication Guideline, and it evolved with various releases over time, as illustrated in Figure A3. CSE/CCCS also published user authentication guidance as illustrated in Figure A3 (the dashed arrows indicate which version of the NIST guidelines were used as a reference).

Figure A3: evolution of NIST and CSE/CCCS user authentication guidance

Evolution of US and GC User Authentication Technical Guidance

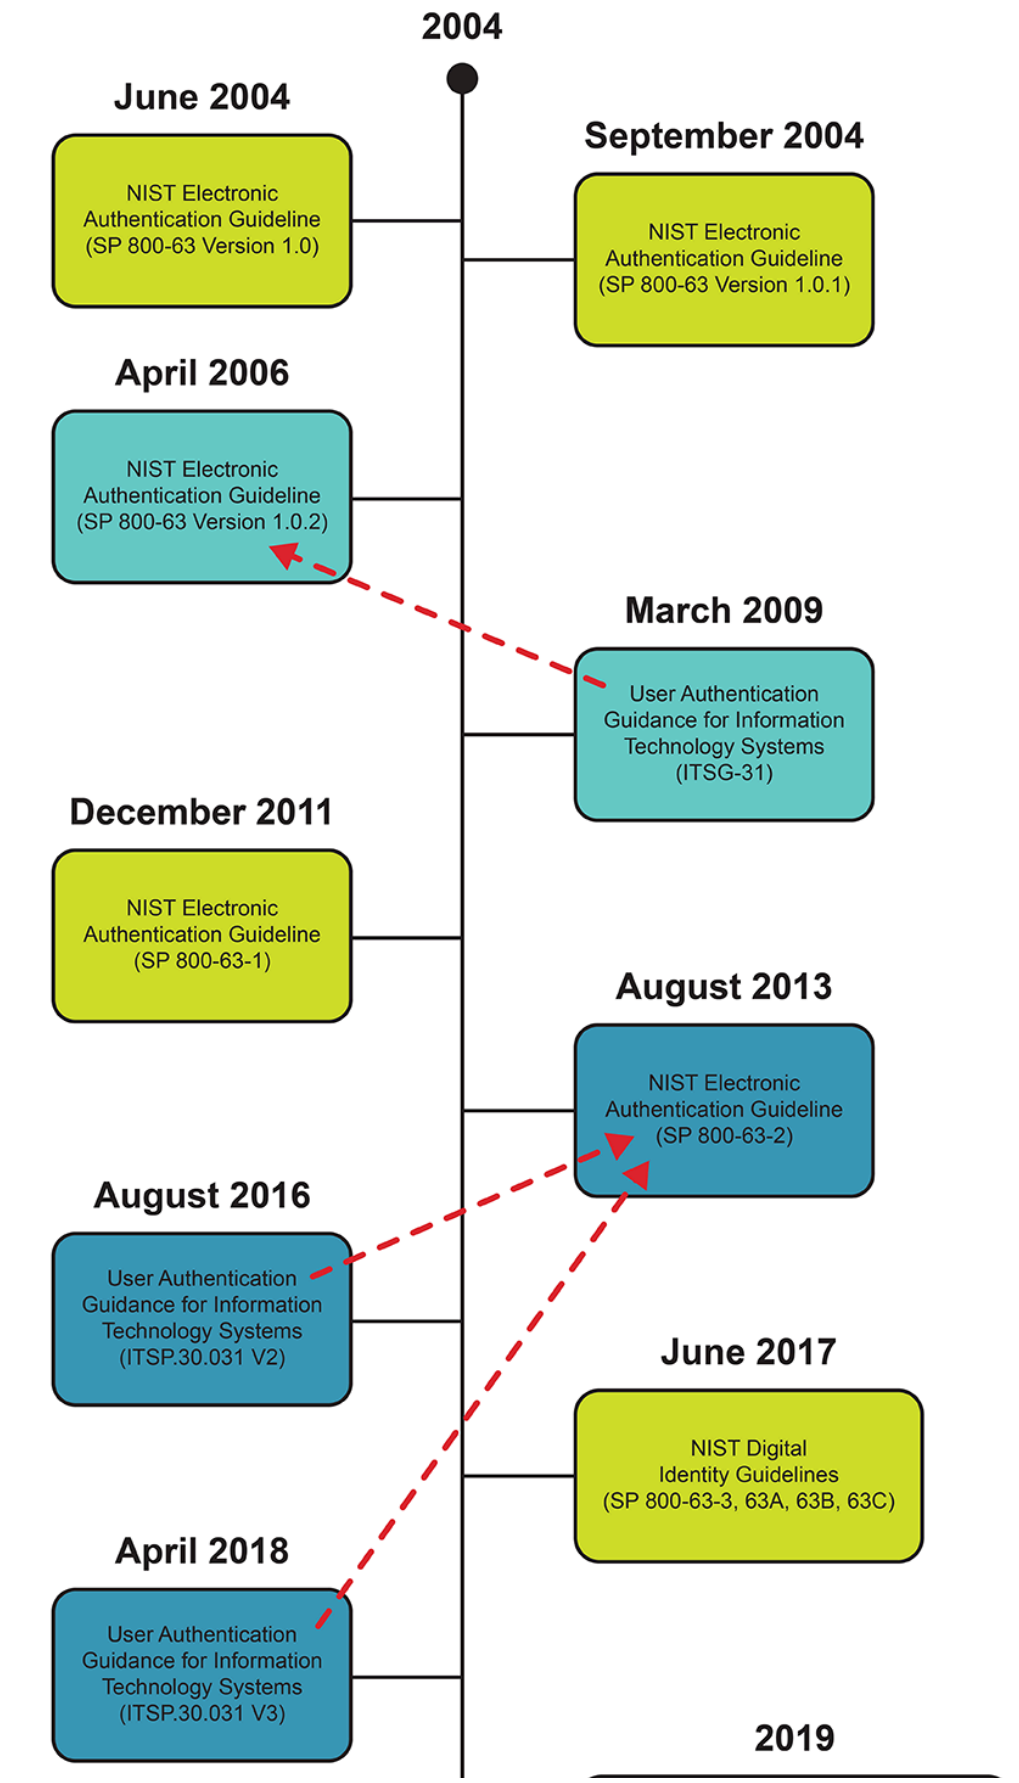

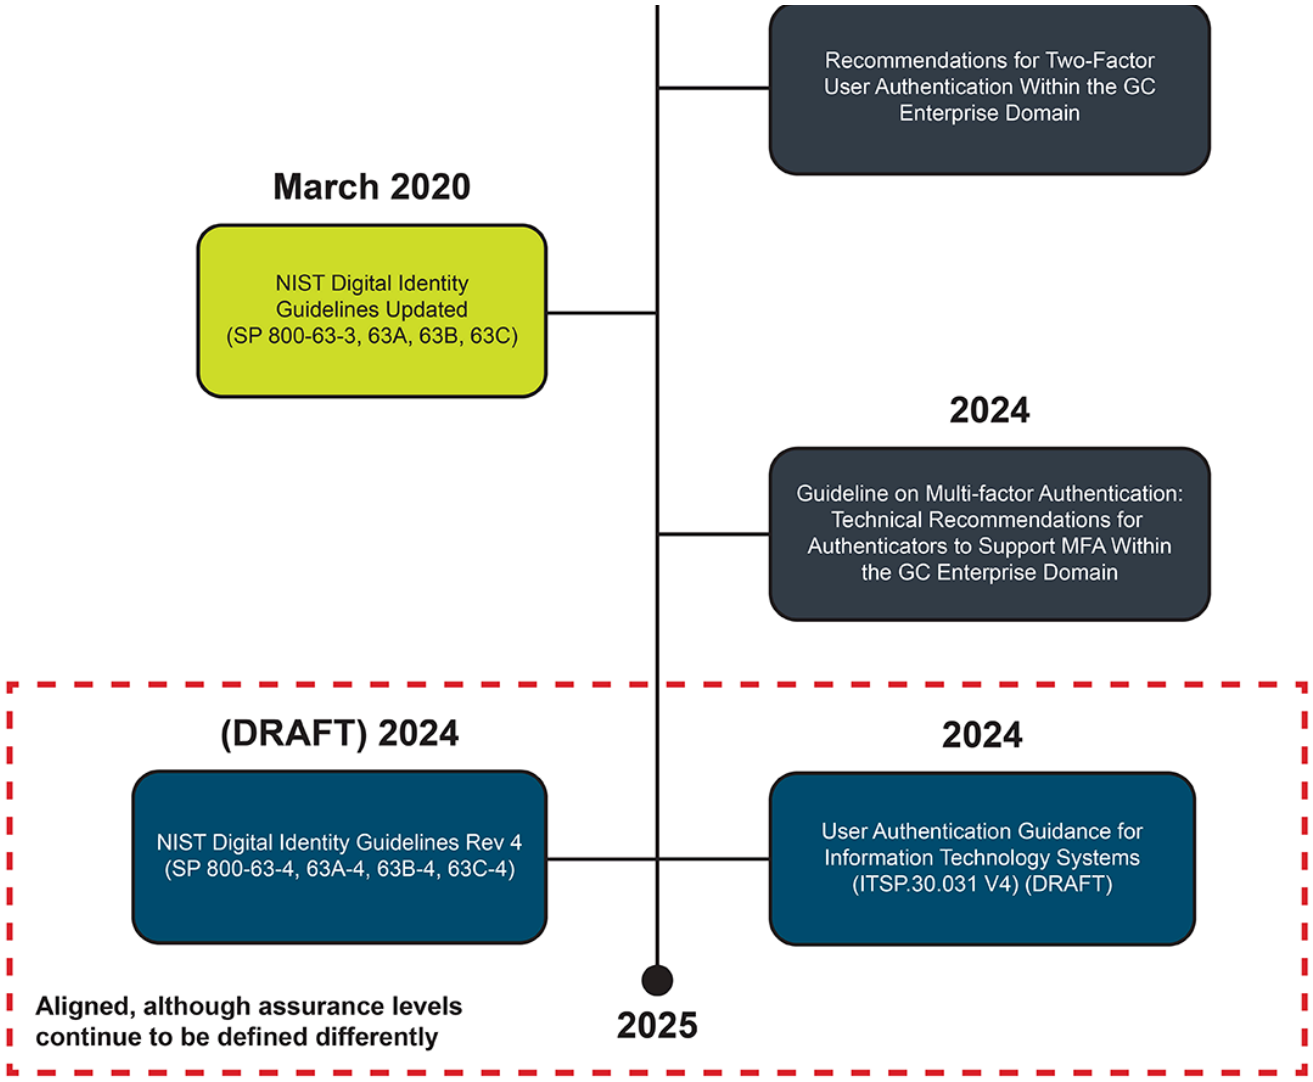

▶ Figure A3 - Text version

Until June 2017, the guidance documents produced by CSE) and CCCS were, with a few exceptions, in alignment with the NIST guidance, and both were based on the four levels of assurance model. However, in June 2017, NIST published a new and substantially different version of its guidance and renamed it to the Digital Identity Guidelines. These guidelines supersede the Electronic Authentication Guideline as follows:

- Sections 1 to 4 of NIST SP 800-63-2 are superseded by SP 800-63-3
- Section 5 of NIST SP 800-63-2 is superseded by SP 800-63A
- Sections 6 to 8 of NIST SP 800-63-2 are superseded by SP 800-63B
- Section 9 of NIST SP 800-63-2 is superseded by SP 800-63C

Perhaps the most significant difference is the new NIST Digital Identity Guidelines are no longer based on the four levels of assurance. [51] Instead, three assurance levels are defined for three different areas: identity, authentication/authenticators and federation. In other words, the NIST Digital Identity Guidelines identify an explicit assurance level associated with each of these three areas rather than a single level of assurance value to represent the overall authentication assurance level. However, the GC policy instruments and related technical implementations continue to be based on the original four levels of assurance. It remains to be seen if, or when, this will change. It is therefore important to understand the relationships between the two different models.

Table A3 describes the three assurance levels for each area defined in Revision 3 of the NIST Digital Identity Guidelines [52] and identifies equivalent (or approximate) mappings with the GC levels of assurance. As there is not always a one-to-one equivalency, some mappings are approximations and may possibly span more than one level. Note that the equivalent GC federation assertion levels are derived from Section 7 of ITSP.30.031 v3.

Table A3: level of assurance mappings

NIST SP 800-63-3 assurance levels		**GC equivalent**
Identity assurance level (IAL)		
IAL1	There is no requirement to link the applicant to a specific real-life identity. Any attributes provided in conjunction with the authentication process are self-asserted or should be treated as such (including attributes a credential service provider, or CSP, asserts to a relying party (RP)).	**Identity LoA 1**

NIST SP 800-63-3 assurance levels		GC equivalent
IAL2	Evidence supports the real-world existence of the claimed identity and verifies that the applicant is appropriately associated with this real-world identity. IAL2 introduces the need for either remote or physically present identity proofing. Attributes can be asserted by CSPs to RPs in support of pseudonymous identity with verified attributes.	**Identity LoA 2/3** [53]
IAL3	Physical presence is required for identity proofing. Identifying attributes must be verified by an authorized and trained representative of the CSP. As with IAL2, attributes can be asserted by CSPs to RPs in support of pseudonymous identity with verified attributes.	**Identity LoA 4**
Authenticator assurance level (AAL)		
AAL1	AAL1 provides some assurance that the claimant controls an authenticator bound to the subscriber's account. AAL1 requires either single-factor or multi-factor authentication using a wide range of available authentication technologies. Successful authentication requires that the claimant prove possession and control of the authenticator through a secure authentication protocol.	**Credential LoA 2**
AAL2	AAL2 provides high confidence that the claimant controls authenticator(s) bound to the subscriber's account. Proof of possession and control of two distinct authentication factors is required through secure authentication protocol(s). Approved cryptographic techniques are required at AAL2 and above.	**Credential LoA 3**
AAL3	AAL3 provides very high confidence that the claimant controls authenticator(s) bound to the subscriber's account. Authentication at AAL3 is based on proof of possession of a key through a cryptographic protocol. AAL3 authentication **shall** use a hardware-based authenticator and an authenticator that provides verifier impersonation resistance; the same device **may** fulfill both these requirements. In order to authenticate at AAL3, claimants **shall** prove possession and control of two distinct authentication factors through secure authentication protocol(s). Approved cryptographic techniques are required.	**Credential LoA 4**

NIST SP 800-63-3 assurance levels		GC equivalent
Federation assurance level (FAL) [54]		
FAL1 [55]	Allows for the subscriber to enable the RP to receive a bearer assertion. The assertion is signed by the IdP using approved cryptography.	**Assertion LoA 1**
FAL2 [56]	Adds the requirement that the assertion be encrypted using approved cryptography such that the RP is the only party that can decrypt it.	**Assertion LoA 2**
FAL3	Requires the subscriber to present proof of possession of a cryptographic key [57] referenced in the assertion in addition to the assertion artifact itself. The assertion is signed by the IdP and encrypted to the RP using approved cryptography.	**Assertion LoA 4**

A-4 Summary

This document provides an overview of the GC's four-tiered level of assurance model and illustrates how the four levels of assurance map to the NIST Digital Identity Guidelines IAL, AAL and FAL assurance levels. The mappings are based on Revision 3 of the NIST Digital Identity Guidelines. However, it is recognized that Revision 4 of the NIST Digital Identity Guidelines is currently under development, and it may have a substantial impact on some of the current definitions and mappings. This document will be updated to reflect any changes once Revision 4 is officially published as the successor to Revision 3.

Appendix B: Fast Identity Online (FIDO) Overview

▶ In this section

B-1 Introduction

“The Fast Identity Online (FIDO) Alliance is a multi-stakeholder, public-private, industry standards development organization comprised of more than 250 companies and government agencies from around the world dedicated to the creation of standards and certification programs for Multi-Factor Authentication (MFA) and passwordless authentication, as well as remote identity verification.” [58]

The FIDO Alliance User Authentication Specifications Overview provides a high-level description of the FIDO Alliance specifications, and the detailed technical specifications are available at Download Authentication Specifications. The purpose and relationship between these specifications can be somewhat confusing, so this overview is provided to help understand the FIDO Alliance (and related W3C) technical specifications landscape and identify some of the more relevant technical considerations in the context of the GC enterprise.

While the focus of this summary is on the technical specifications, it should be recognized that the FIDO Alliance addresses much more. Readers interested in additional details are encouraged to explore the FIDO Alliance website to gain a better appreciation of the available documentation, including the technical specifications, certification program, security requirements and case studies.

It should be noted that the information provided within this overview is based on available information at the time of this writing and is subject to change as the technical specifications and product implementations continue to evolve. It should also be recognized that some of the available online information related to FIDO technical implementation details is sometimes contradictory or misleading. This overview identifies some of these areas and, wherever possible, uses the relevant specifications as the

definitive source. Furthermore, any mention of specific companies or products are for illustrative purposes only and are not to be interpreted as any form of endorsement.

B-2 Technical overview

FIDO authenticators use asymmetric cryptography to support user authentication. A different public/private key pair is generated by the authenticator when the user registers with each relying party (RP) (for example, a website), and that key pair is used to facilitate the user authentication process with that RP thereafter (see Relying Party Identifier (RP ID), *Web Authentication: An API for accessing Public Key Credentials, Level 2*. Each key pair is bound to a specific RP (sometimes referred to as scoped or bound to an origin), so it cannot be used to authenticate to another RP. These properties make FIDO authenticators both phishing resistant and adversary-in-the-middle (AitM) resistant.

FIDO authenticators can be bound to a user device such as a laptop (referred to as **platform authenticators**) or they can be on a separate device such as a security key or mobile phone (referred to as **roaming authenticators**).

The FIDO Alliance has published three sets of technical specifications as follows:

1. **Universal Second Factor (U2F):** a set of FIDO Alliance specifications that define a second factor authentication solution used in conjunction with a first factor (typically a userid or password). As noted below, several U2F specifications have been superseded, but U2F authenticators are still supported.
2. **Universal Authentication Framework (UAF):** a set of FIDO Alliance specifications that define a multi-factor passwordless authentication solution (mostly used as SDKs for native apps and biometric plugins).

3. **Client-to-Authenticator Protocol (CTAP):** a FIDO Alliance specification that describes a protocol for communication between roaming authenticators (for example, FIDO security keys or **passkeys** on mobile devices) and WebAuthn-enabled browsers or platforms over several different transport technologies such as USB, NFC and BLE. [59] It should also be noted that the CTAP 2.1 specification actually describes two protocol versions: CTAP1 (previously referred to as U2F) and CTAP2, which supports passwordless, second-factor and multi-factor authentication. Also note that the CTAP 2.1 specification supersedes the U2F Raw Message Format, U2F HID Protocol, U2F NFC Protocol and the U2F Bluetooth specifications. [60] The current version of CTAP is Version 2.1, and CTAP Version 2.2 is under development (the latest draft at the time of this writing is CTAP 2.2 dated March 21, 2023). Note that CTAP 2.2 defines a new hybrid transport to standardize the protocol for using passkeys on smartphones to authenticate to RPs being accessed on another device (for example, a laptop or tablet). This is referred to as **cross-device authentication** (CDA).

In addition, W3C has published a complementary specification (in collaboration with the FIDO Alliance) referred to as W3C Recommendation Web Authentication: An API for Accessing Public Key Credentials or WebAuthn. WebAuthn "defines a standard web API that is being built into browsers and platforms to enable support for FIDO Authentication." [61] WebAuthn replaces the U2F Javascript API specification but is backward compatible while offering additional functionality. Backward compatibility with the U2F Javascript API ensures continued support for U2F authenticators. Note that all major browsers support the WebAuthn API (see FIDO2: Web Authentication). The current version of WebAuthn recommendation is Level 2, with Level 3 currently under development (the

latest Level 3 editor's draft at the time of this writing is Web Authentication: An API for Accessing Public Key Credentials Level 3 (Editor's Draft dated March 13, 2024)).

Note that the nomenclature in this area tends to vary and can be confusing. The term "FIDO2" (stemming from the "FIDO2 Project") is essentially a term used to collectively refer to the newer FIDO Alliance CTAP specification and W3C WebAuthn recommendation. U2F and UAF are not considered part of "FIDO2," although this distinction is sometimes blurred in various online publications and posts. Furthermore, FIDO2 is an extension of and backward compatible with U2F and may eventually subsume U2F, which adds to the confusion. For the purposes of this document, a "FIDO authenticator" denotes the broadest sense of the term (that is, any authenticator that is FIDO certified). This includes UAF, U2F and FIDO2 authenticators. When necessary to provide additional context or clarity, the specific type of authenticator should be identified. In addition, CTAP and WebAuthn should be identified separately when necessary.

Key takeaways include:

- FIDO authenticators are both phishing resistant and AitM resistant
- the CTAP specification describes how external FIDO authenticators (that is, roaming authenticators physically separate from the user platform) communicate or interface with WebAuthn-enabled browsers and platforms
- WebAuthn is an API that enables FIDO authentication in a web browser or platform
- a U2F or CTAP1 authenticator is used as a second authentication factor in conjunction with a first authentication factor (typically a userid/password)

B-3 Additional technical details

Additional technical details (and areas for further analysis) are discussed below.

B-3.1 Passkeys (synced and device-bound)

The meaning of the term **passkey** has evolved over the last few years. Certain vendors introduced the term as a synonym for multi-device credentials that can be backed up and copied to multiple devices. However, the term **passkey** actually applies to both multi-device and single-device credentials (see FIDO Alliance: Passkeys). In other words, **passkey** is a generic term for FIDO **discoverable credentials** that can be multi-device (copyable) or single-device (not copyable). More recent documentation uses the terms **synced passkeys** and **device-bound passkeys** (for example, see FIDO Deploying Passkeys in the Enterprise: Introduction).

Syncing passkeys is an approach where the private keys used to authenticate a user can be backed up and synchronized (copied) to more than one device, thus eliminating (or reducing) the need to register multiple credentials with the same online web service. Synced passkeys also seamlessly support credential synchronization to new devices. This improves the user experience with respect to authenticator recovery (see Recovery section below) as well as providing seamless authentication for users who use multiple devices to access the same web services. However, there are trade-offs with respect to user convenience and the potential security concerns that synced passkeys may introduce.

The US National Institute of Standards and Technology (NIST) recently published new interim guidance on “synced authenticators” (or synced passkeys). The interim guidance identifies security requirements that must be met so that they can potentially be used at LoA 3/AAL2. This could very well be a game-changer in terms of improving the user experience without

sacrificing security. However, note that there are some products that support only device-bound passkeys in order to meet higher assurance requirements (for example, the Yubikey 5 Series FIDO2 implementation).

Additional information regarding the trade-offs between synced passkeys and device-bound passkeys is provided in the FIDO Alliance white papers FIDO Authentication for Moderate Assurance Use Cases and High Assurance Enterprise FIDO Authentication.

B-3.2 Credential storage modality

There seems to be contradictory (or misleading) information posted on various websites regarding where the private keys used for user authentication are stored. For example, the FIDO Alliance website itself (see FIDO Authentication) actually states, “FIDO2 cryptographic login credentials are unique across every website, never leave the user’s device and are never stored on a server.” However, it is clear from reading the FIDO2 specs (which are, by definition, the authoritative source) that the private signing keys used for authentication can be stored on the device or on the server (in encrypted form). Specifically, see Section 6.2 Authenticator Taxonomy and Section 6.2.2 Credential Storage Modality of the WebAuthn spec and related definitions of client-side discoverable credentials and service-side credentials. [62] Although U2F may not be considered part of “FIDO2,” see also Section 7 of the FIDO Alliance Universal 2nd Factor (U2F) Overview and Section 4.3 of the FIDO Alliance U2F Raw Message Formats spec (for example, “U2F tokens may wrap the generated private key and the application id it was generated for, and output that as the key handle”).

Although it is recognized that specific products may only support discoverable credentials and/or specific solutions such as passwordless authentication may actually require discoverable credentials, blanket statements such as “the private key never leaves the device” are not

entirely accurate when referring to the FIDO Alliance specifications and the W3C WebAuthn recommendation. This is also a misnomer in terms of synced passkeys (see below for more information).

The primary benefit of storing the private keys on the server side is to minimize the amount of memory required on the device, thereby reducing the cost of the authenticator and, in theory, support an unlimited number of key pairs. However, the protection of the private keys when stored on the server may be a concern and will depend on the correct implementation of an approved encryption algorithm and associated key lengths used to protect the private keys. For example, Yubico uses AES-256 in CCM mode (for U2F, see Key generation). Furthermore, the number of passkeys that can be stored on FIDO2 security keys has been increasing. For example, the Yubikey 5 FIDO2 implementation has gone from supporting only 25 device-bound passkeys to 100. In addition, Google has recently launched new Titan Security Keys that can store more than 250 passkeys (see The latest Titan Security Key is in the Google Store).

Also note that the current industry trend appears to be heading towards passkeys which are, by definition, discoverable credentials and therefore cannot be stored on the server (refer to Table B22 for additional context). Nonetheless, server-side credential storage modality should not be dismissed altogether, as there are products available that do implement it.

B-3.3 Recovery

As with any authentication method, recovery options for FIDO-compliant authenticators need to be available in the event that the primary authenticator is lost, stolen or damaged. According to the FIDO Alliance (see Recommended Account Recovery Practices for FIDO Relying Parties), it is recommended that each user should have at least two FIDO compliant authenticators – one of which would be used as a backup in case anything

happens to the other primary authenticator. However, this approach may not be suitable for the GC as it will result in increased costs (due to the need for each user to have multiple authenticators), may result in a poor user experience, and will require users to securely store the backup authenticator(s). Therefore, other recovery methods should be explored (for example, implementation of a centralized authentication approach might support other suitable approaches). Synced passkeys (as discussed above) or high assurance hybrid flows may also provide alternatives.

B-3.4 Enterprise attestation

The original FIDO Alliance specs were designed for a mass consumer market, namely, for users accessing online services on the Internet. While that continues to be a major focus, recent additions to the WebAuthn recommendation and CTAP specification introduce the notion of "enterprise attestation (EA)" in order to render FIDO authenticators more suitable for an enterprise domain. However, how EA actually works in implementation practice requires further investigation.

At the time of this writing, synced passkeys do not implement attestation. [63] This is due to the fact that the current CTAP 2.1 specification and WebAuthn Level 2 recommendation only support attestation during the registration process, so there is no way to assert a new attestation statement when authenticating from another device with a synced passkey. However, support for this new capability is expected to be added to the CTAP 2.2 specification and WebAuthn Level 3 recommendation. What this will mean in implementation practice, and what level of support will be provided by the vendor community, remains to be seen.

Also note that EA was introduced as part of FIDO2 and is not supported by U2F.

See FIDO TechNotes: The Truth About Attestation and FIDO Attestation White Paper for additional information regarding attestation. Additional details regarding EA are also available at Enterprise Attestation, and several white papers related to enterprise deployments are available at Learn how FIDO authentication fits into your enterprise environment.

B-4 Key terms

Key terms used within this document are defined below in Table B11. The definitions are copied or derived from the sources indicated in the third column. Text in square brackets is added to some of the definitions to provide additional clarity.

Table B1: key terms

Term	Definition	Source
Adversary-in-the-middle (AitM)	An attack in which an attacker [threat actor] is positioned between two communicating parties in order to intercept and/or alter data travelling between them. In the context of authentication, the attacker [threat actor] would be positioned between claimant [user] and verifier, between registrant [user] and CSP during enrollment, or between subscriber [user] and CSP during authenticator binding. [Formerly known as man-in-the-middle, also referred to as attacker-in-the-middle]	NIST SP 800-63-3
Authenticator	Something the claimant [user] possesses and controls (typically a cryptographic module or password) that is used to authenticate the claimant's [user's] identity	NIST SP 800-63-3

Term	Definition	Source
Client-to-authenticator protocol (CTAP)	An application layer protocol for communication between a roaming authenticator and another client or platform, as well as bindings of this application protocol to a variety of transport protocols using different physical media	CTAP 2.1 specification
Credential storage modality	Determined by where the credential private key is stored – either the client side (discoverable credential) or server side (non-discoverable credential) In the case of server-side credential storage modality, the credential private key is encrypted (wrapped) using a strong symmetric encryption algorithm (for example, AES-256) such that only this authenticator can decrypt (that is, unwrap) it. The encrypted credential private key is stored by the relying party (as the credential ID parameter in the public key credential source) and returned to the authenticator via the allowCredentials option of get(), which allows the authenticator to decrypt and use the credential private key.	Derived from WebAuthn, Section 6.2.2: Credential Storage Modality
Cross-device authentication (CDA)	FIDO cross-device authentication (CDA) allows a passkey from one device to be used to sign in on another device. For example, your phone can be linked to your laptop, allowing you to use a passkey from your phone to sign into a service on your laptop. CDA is powered by the FIDO client-to-authenticator Protocol (CTAP) using “hybrid” transport. CTAP is implemented by authenticators and client platforms, not relying parties. [The hybrid transport has been added to CTAP 2.2 (see Section 11.5), which is currently a FIDO Alliance draft specification.]	See Cross-Device Authentication (CDA)

Term	Definition	Source
Discoverable credential	A discoverable credential (previously known as a “resident credential” or “resident key”) is a FIDO2/WebAuthn credential that is entirely stored in the authenticator (private key, credential ID, user handle, and other metadata). The relying party (RP) also stores a copy of the **public** key and credential ID. Also known as client-side discoverable credential, a discoverable credential has a client-side credential storage modality and can be “discovered” [or retrieved] without supplying the relying party with user-specific information as a prelude to authentication.	See Discoverable Credential Derived from WebAuthn
Device-bound passkey	A FIDO2 discoverable credential that is bound to a single authenticator. For example, FIDO2 security keys typically hold device-bound passkeys, as the credential cannot leave the device. Device-bound passkeys have been previously referred to as **single-device passkeys**. [Device-bound passkeys are FIDO credentials that never leave the device and therefore cannot be synced or copied.]	See Device-bound passkey
Multi-device credential (also see synced passkey)	A FIDO credential that can be backed up and copied to multiple devices. Also referred to as “synced passkeys” Examples include: • Passwordless login with passkeys • Use passkeys to sign in to apps and websites on iPhone	Derived from FIDO Deploying Passkeys in the Enterprise: Introduction

Term	Definition	Source
Non-discoverable credential	Also known as server-side credential, a non-discoverable credential has a server-side credential storage modality and cannot be “discovered” (or retrieved) without supplying the relying party with user-specific information	Derived from WebAuthn
Passkey (also see synced passkey and device-bound passkey)	The high-level, end-user-centric term for a FIDO2/WebAuthn discoverable credential. Like “password,” “passkey” is a common noun intended to be used in everyday conversations and experiences. Usage: “a passkey” or “passkeys.” Passkeys are designed to be used without additional login challenges. All passkeys can be used with modern sign in experiences like the Autofill UI or with a “Sign in with a passkey” button. From the technical side, there are two flavours of passkeys: synced and device-bound. [Note that passkeys are synonymous with discoverable credentials.]	See Passkey
Phishing resistance	The ability of the authentication protocol to detect and prevent disclosure of authentication secrets and valid authenticator outputs to an impostor relying party without reliance on the vigilance of the subscriber [user]	NIST SP 800-63B-4 (Section 5.2.5)
Platform authenticator	An authenticator that is physically bound to a specific WebAuthn enabled client device such as a TPM on a laptop. Communication between the authenticator and the WebAuthn client is typically via a platform-specific API A FIDO authenticator that is built in to a user’s device.	Derived from WebAuthn and other sources See Platform authenticator

Term	Definition	Source
Roaming authenticator	A FIDO authenticator usable with any device the user is trying to sign in from. Roaming authenticators attach to users' devices using USB, NFC, and/or Bluetooth. These authenticators are often referred to as "security keys." A smartphone can also act as a roaming authenticator using FIDO Cross-device authentication. An authenticator that is not physically bound to a specific WebAuthn-enabled client device such as a FIDO2 security key or FIDO2 passkeys on a mobile phone. A roaming authenticator can be used to authenticate to multiple platforms via several different transport protocols such as USB, BLE and NFC. Roaming authenticators are sometimes referred to as cross-platform or external authenticators.	See Roaming authenticator Derived from WebAuthn and other sources

Term	Definition	Source
Synced passkey	A FIDO2 discoverable credential that can reliably be used for bootstrapping sign-in, without requiring other login challenges such as passwords and OTPs. “Reliable” here means that the passkey should be available to, and usable by, the user whenever they need to sign in. This availability can be achieved through different means: for example, passkey providers could sync passkeys in real time across a user’s devices, restore passkeys from a backup whenever a user sets up a new device, offer passkeys across different contexts (a passkey established from an app can be used in the browser when visiting the app’s website), or allow users to exercise passkeys across devices (by, say, using the passkey from a nearby phone when signing in from a laptop). [Synced passkeys are FIDO credentials that are backed up and made available across multiple devices.]	See Synced passkey

Table B2 summarizes the relationships between various terms/properties associated with FIDO credentials to help place these terms and concepts into context.

Table B2: FIDO credential terminology relationships

	Discoverable credential?	Considered a passkey?	Can be synced?
Multi-device credential	Yes	Yes (“synced passkey”)	Yes
Device-bound credential	Yes	Yes (“device-bound passkey”)	No

Server-side credential	No (non-discoverable)	No [64]	No

Implementation note (subject to change): According to Yubico, Android is "the only mainstream mobile operating system that can generate single device credentials." This is in contrast to Apple iOS devices where passkeys are stored using iCloud Keychain and are synced across a user's registered devices. Furthermore, the passkeys can be shared with others using AirDrop. See the following for additional information:

- About the security of passkeys
- Use passkeys to sign in to apps and websites on iPhone
- Share passkeys and passwords securely with AirDrop on iPad

Appendix C: Trusted Platform Module Overview and Comparison

The latest version of the Trusted Platform Module (TPM) specification, which is published by the Trusted Computing Group, is TPM 2.0. TPM 2.0 includes a number of enhancements as compared to its predecessor, TPM 1.2. Trusted Platform Module (TPM) 2.0: A Brief Introduction provides a high-level overview of TPM 2.0 and summarizes the five different types of TPMs as follows:

- **Discrete TPM** provides the highest level of security, as might be needed for a TPM used to secure the brake controller in a car. The intent of this level is to ensure that the device it's protecting does not get hacked via even sophisticated methods. To accomplish this, a discrete chip is designed, built and evaluated for the highest level of security that can resist tampering with the chip, including probing it and freezing it with all sorts of sophisticated attacks.

- **Integrated TPM** is the next level down in terms of security. This level still has a hardware TPM, but it is integrated into a chip that provides functions other than security. The hardware implementation makes it resistant to software bugs; however, this level is not designed to be tamper resistant.
- **Firmware TPM** is implemented in protected software. The code runs on the main CPU, so a separate chip is not required. While running like any other program, the code is in a protected execution environment called a trusted execution environment (TEE) that is separated from the rest of the programs that are running on the CPU. By doing this, secrets like private keys that might be needed by the TPM but should not be accessed by others can be kept in the TEE, creating a more difficult path for hackers. In addition to the lack of tamper resistance, the downside to the TEE or firmware TPM is that now the TPM is dependent on many additional aspects to keep it secure, including the TEE operating system, bugs in the application code running in the TEE, and more.
- **Software TPM** can be implemented as a software emulator of the TPM. However, a software TPM is open to many vulnerabilities, not only tampering but also the bugs in any operating system running it. It does have key applications: it is very good for testing or building a system prototype with a TPM in it. For testing purposes, a software TPM could provide the right solution and approach.
- **Virtual TPM** Many IoT systems include sensors and cloud processing, which means virtualization. In a cloud environment, one clever way to implement a TPM is through a virtual TPM. The virtual TPM is part of the cloud-based environment, and it provides the same commands that a physical TPM would, but it provides those commands separately to each virtual machine.

Table C1 provides a comparison of the TPM types. Table C1 is based partially on the table provided in Trusted Platform Module (TPM) 2.0: A Brief Introduction. The type of authenticator and level of assurance (LoA) columns have been added. Note that the highest possible LoA assumes that **all** requirements at that LoA are met, including appropriate levels of FIPS 140-3 certification and other applicable considerations. Note that these additions are based on a preliminary high-level assessment and a more thorough analysis is recommended, particularly with respect to the firmware and virtual TPM LoAs.

Table C1: TPM comparison

Type of TPM	Relative security	Execution environment	Type of authenticator	Highest possible LoA
Discrete TPM	Highest	Dedicated, tamper resistant hardware chip	Multi-factor or single-factor cryptographic device	LoA 4
Integrated TPM	Higher	Separate hardware, but not dedicated solely to security functions and not tamper resistant	Multi-factor or single-factor cryptographic device	LoA 3
Firmware TPM	High	TEE on CPU	Multi-factor or single-factor cryptographic software	LoA 3
Software TPM	None	Operating system	Multi-factor or single-factor cryptographic software	LoA 1

Virtual TPM	High	Hypervisor	Multi-factor or single-factor cryptographic software	LoA 3

Note that at LoA 4, a discrete TPM must be used to generate and securely store the private signing key(s) and also perform the cryptographic operations to support the user authentication process (for example, the private key is used to digitally sign a nonce in a challenge-response authentication protocol). In addition, the private keys must not be exportable (that is, the private keys never leave the discrete TPM). Also note that Software TPMs are limited to testing/prototyping purposes and must never be used beyond that.

Footnotes

1 Source: CBC News, "Spy agency chief says new powers would help stop cyberattacks before they happen."

2 Source: CBC News, "Spy agency chief says new powers would help stop cyberattacks before they happen." Although this reference is from 2018, it helps to illustrate that GC systems and information are lucrative targets. It is also reasonable to assume that the number of attempts has only increased.

3 NIST released the Second Public Draft of the Digital Identity Guidelines, including NIST SP 800-63B-4 (Second Public Draft), after this document was submitted to the formal publication process. While there are important changes from the initial public draft, these changes do not have a significant impact on the recommendations made within this document. Some, but not all, of the relevant changes have been identified in this document.

4 This is often represented as “userid and password.” However, the userid is not necessarily, and typically isn’t, known only to the user. In any case, the combination of the userid and password is considered single-factor authentication, as the two components must be used together.

5 Pre-registered knowledge authenticators (for example, pre-established answers to a set of challenge questions) are no longer considered acceptable and are therefore not included as an example of something you know. This is reflected in NIST SP 800-63B and is also expected to be removed in the next update to ITSP.30.031.

6 Some SF OTP products require the user to press a button to display the one-time use code, but the act of pushing the button does not constitute a second authentication factor. Also, there are SF OTP products that do not require the user to press a button (that is, the display is always on, and the codes change automatically after a certain amount of time has passed).

7 Note that biometrics can only be used for local authentication, not for remote authentication. For example, a user’s fingerprint can be used to unlock or activate a multi-factor cryptographic hardware authenticator. See Section 5.2.3 of NIST SP 800-63B for additional restrictions and information.

8 This definition is derived from Section 5.2.3 of NIST SP 800-63B.

9 Note that NIST SP 800-63B, Section 5.1.4, includes both hardware devices and software-based OTP generators installed on mobile devices under this category. ITSP.30.031 v3 does not explicitly address software-based OTP apps.

10 NIST SP 800-63B, Section 5.1.6, introduces a new single-factor cryptographic software authenticator, and this is also expected to be added to the next release of ITSP.30.031.

11 Revision 4 of the NIST Digital Identity Guidelines introduce a new multi-factor out-of-band authenticator. NIST defines this as “a physical device that is uniquely addressable and can communicate securely with the verifier over a distinct communications channel, referred to as the secondary channel.”

12 Section 4.3.1 of ITSP.30.031 v3 considers an MF cryptographic software authenticator to be LoA 2/AAL1; however, this was based on deficiencies related to the myKEY software authenticator. While these deficiencies remain, and therefore a myKEY software authenticator will remain at LoA 2, there are examples of an MF software authenticator that could be considered LoA 3. It is therefore expected that the forthcoming update of ITSP.30.031 will reinstate this type of authenticator (but not a myKEY software authenticator) to LoA 3 as the highest **possible** LoA that can be achieved.

13 The NIST Digital Identity Guidelines no longer consider a multi-factor OTP hardware authenticator, by itself, sufficient to achieve LoA 4/AAL3 (since it is not phishing resistant).

14 NIST SP 800-63B, Section 5.1.5, includes both hardware devices and software-based OTP generators installed on mobile devices under the “OTP device” category. This has been clarified in NIST SP 800-63B-4 (Second Public Draft). ITSP.30.031 v3 does not explicitly address software-based OTP apps.

15 Revision 3 of the NIST Digital Identity Guidelines identified several other authenticator combinations at LoA 4/AAL3 that are omitted here as they are no longer included in NIST SP 800-63B-4 (Second Public Draft) and were never included as examples in Figures 2.1 and 2.2 in this document anyway.

16 Push fatigue, MFA fatigue or push bombing occurs when a threat actor bombards a user with mobile application push notifications until the user either approves the request by accident or out of annoyance. Push notifications with number matching help to ensure that the user is actually engaged in the authentication process. See Cybersecurity and Infrastructure Security Agency: Implementing Number Matching in MFA Applications for additional information.

17 In addition, although both a smartphone and an OTP hardware authenticator can be stolen, an OTP hardware authenticator can be more easily protected from theft (for example, attach the authenticator to the same lanyard as the user's building pass), thus making it more difficult to steal. It might also be argued that a smartphone is generally a more attractive target, as it has value beyond the context of authentication whereas an OTP hardware authenticator does not.

18 The requirement for a separate, platform-independent authenticator is based on guidance provided by CCCS.

19 Source: Section 5.2.5 of NIST SP 800-63B-4 (Revision 4, Initial Public Draft).

20 The method used to lure the user to the fake website is irrelevant in this context.

21 This is not meant to imply that additional security measures should not be taken if one of the authenticators is inherently phishing resistant. See Section 2.4.6.

22 This can be accomplished using device certificates (preferably with associated private keys stored securely in FIPS 140 validated hardware).

23 Note that this memorandum embraces FIDO phishing-resistant authenticators as an alternative to personal identity verification (PIV) cards.

24 Trusted Platform Module (TPM) 2.0: A Brief Introduction provides a high-level overview of TPM 2.0 and summarizes the five different types of TPMs. Also see Appendix C of this document.

25 Introduction to Secure Elements provides a high-level introduction with respect to SEs.

26 Introduction to Trusted Execution Environments provides a high-level overview of TEEs.

27 Source: Derived from NIST SP 800-63B, Sections 5.1.6.1 and 5.1.8.1.

28 This is not meant to preclude the possibility of supporting server-side credential storage modality at LoA 3/AAL2 (see Section 2.3.7.1). In addition, support for synced passkeys at LoA 3/AAL2 may be possible in the future; see Section 2.3.7.2.

29 Platform-independent authenticators must be properly handled by the user in order to realize the additional security benefits that they provide over platform-bound authenticators. This includes things such as maintaining positive control of the authenticator by the user at all times and, if physically connected during use, removal of the authenticator from the computer platform when not in use.

30 Source: NIST SP 800-63B-4 (Revision 4, Initial Public Draft), Sections 5.1.8.1 and 5.1.9.1

31 At the time of this writing, the current version of the FIPS 140 standard (*Security Requirements for Cryptographic Modules*) is FIPS 140-3. However, cryptographic modules validated against its predecessor, FIPS 140-2, remain valid. Therefore, references to FIPS 140 in this document should be interpreted as FIPS 140-2 or FIPS 140-3. See Cryptographic Module Validation Program for additional details.

32 This is also referred to as attacker in the middle.

33 Source: Derived from the MitM definition in NIST SP 800-63-3.

34 Source: Derived from Section 5.2.8 in NIST SP 800-63B.

35 ITSP.30.031 v3 (see Section 6.3) stipulates replay resistance is required even at LoA 1 (which stems from the now obsolete NIST *Electronic Authorization Guideline*). However, NIST SP 800-63B states that replay resistance is not required at LoA 2/AAL1. For now, we have made this recommended at LoA 2, but this will align with whatever ITSP.30.031v4 specifies.

36 If server-side credentials are supported, the private keys must be encrypted with strong symmetric cryptography as recommended in ITSP.40.111, and the cryptographic module must be FIPS 140 validated as stipulated in Section 2.3.3.

37 As noted above in Section 2.3.7.2, this may be changed to include attestation during authentication as well.

38 Source: Client to Authenticator Protocol (CTAP), 2.1, Section 7.1.

39 This could also be used to replace the userid and password with a single-factor phishing-resistant passwordless authenticator.

40 Shorter timeout thresholds may be required in some cases.

41 Prompting the user to force activity before the inactivity threshold is reached may avoid unnecessary reauthentication and therefore result in less friction for the user. This could be a simple “are you still there” prompt. However, note that a response from the user only resets the user inactivity timeout, not the session lifetime timeout, since user reauthentication did not occur.

42 Ibid.

43 Source: Derived from NIST definition in Risk Adaptive (Adaptable) Access Control.

44 Gartner describes a similar model referred to as continuous adaptive risk and trust assessment (CARTA); see Gartner’s *Zero Trust is an Initial Step on the Roadmap to CARTA.*

45 At the time of this writing, these are expected to be device-bound passkeys; however, synced passkeys may be allowed in the future. Refer to Sections 2.3.7.2 and 2.3.7.3 for additional details.

46 The phrases “strong password” and “strong, well-managed password” used within this document refer to password composition and management practices that meet the requirements stipulated in the GC Password Guidance.

47 Refer to the Government of Canada Cloud Guardrails and associated references for more information.

48 Note that the original US OMB M-04-04 (which has been rescinded) talked about “potential impact,” and this is now briefly addressed in NIST SP 800-63-3, Section 5.3. But when it comes down to authenticators, we are really talking about the level of confidence that they are properly bound to and under the control of the user and the level of confidence that the authenticator(s) has or have not been compromised.

49 ITSP.30.031v3 actually uses the term “token” rather than “authenticator,” which is aligned with the now superseded NIST SP 800-63-2. The new NIST Digital Identity Guidelines changed the terminology such that “tokens” are now referred to as “authenticators.” It is expected that the next update to ITSP.030.031 will incorporate this new terminology.

50 However, the *Guideline and Defining Authentication Requirements* recognizes that the target LoA may not always be achievable in practice and therefore allows for compensating controls and risk management when necessary.

51 In addition, OMB M-04-04 was rescinded in May 2019 with the publication of OMB M-19-17, thus officially moving the US federal government away from the four levels of assurance model.

52 It should be noted that Revision 4 of the NIST Digital Identity Guidelines is under development and is expected to have an impact on the at least some of these definitions and mappings.

53 Appears to span both in Revision 3; however, the new draft Revision 4 of the NIST Digital Identity Guidelines moves IAL1 to IAL0 (no identity proofing) and separates IAL2 into IAL1 and IAL2, so the mapping to the GC LoAs may become more straightforward. This should be revisited once Revision 4 stabilizes and is formally published as the successor to Revision 3, but the mappings **may** be to the following: IAL0 maps to Identity LoA 1, IAL1 maps to Identity LoA 2, IAL2 maps to Identity LoA 3, and IAL3 maps to Identity LoA 4.

54 There appear to be substantial differences between Revision 3 and draft Revision 4 with respect to FALs. This should be revisited once Revision 4 is formally published as the successor to Revision 3.

55 NIST SP 800-63C indicates that this is equivalent to the OIDC Basic Client profile (that is, back-channel authorization code flow) and the SAML Web Browser SSO profile using back-channel artifact bindings.

56 NIST SP 800-63C states that FAL2 or higher is required for front channel (that is, SAML redirect binding or OIDC implicit flow).

57 This is referred to as “holder-of-key” assertions in the SAML 2.0 Web Browser SSO Profile.

58 Quote from multiple FIDO Alliance sources.

59 Source: Derived from Client to Authenticator Protocol (CTAP) specification and FIDO Alliance website

60 Source: Section 1.1 of the CTAP specification

61 Source: FIDO Alliance website.

62 Private keys stored on the device were formerly referred to as “resident keys” or “resident credentials” but are now referred to as “client-side discoverable credentials” or “discoverable credentials.” Private keys stored (in encrypted form) on the server were referred to as “non-resident keys” but are now referred to as “server-side credentials” or “non-discoverable credentials.”

63 For example, the FIDO Alliance white paper FIDO Authentication for Moderate Assurance Use Cases, dated June 2023, states, “At the time of publication, synced passkeys do not implement attestation, which means they are not an appropriate solution for scenarios with highly privileged users that require higher levels of assurance or for organizations that want to implement Enterprise Attestation.”

64 For examples, see Discoverable vs non-discoverable credentials and Discoverable Credential.

© His Majesty the King in Right of Canada, represented by the President of the Treasury Board, 2025,
ISBN: 978-0-660-76817-5

Date modified: 2025-05-23

xx
yy