



FIDO Security Reference

FIDO Alliance Implementation Draft 20 February 2018

This version:

<https://fidoalliance.org/specs/fido-uaf-v1.2-id-20180220/fido-security-ref-v1.2-id-20180220.html>

Previous version:

<https://fidoalliance.org/specs/fido-uaf-v1.2-rd-20171128/fido-security-ref-v1.2-rd-20171128.html>

Editor:

[Rolf Lindemann, Nok Nok Labs, Inc.](#)

Contributors:

[Davit Baghdasaryan, Nok Nok Labs, Inc.](#)

[Brad Hill, PayPal, Inc.](#)

[Dr. Joshua E. Hill, InfoGard Laboratories](#)

[Douglas Biggs, InfoGard Laboratories](#)

Copyright © 2013-2018 [FIDO Alliance](#) All Rights Reserved.

Abstract

This document analyzes the security properties of the FIDO UAF and U2F families of protocols.

Status of This Document

This section describes the status of this document at the time of its publication. Other documents may supersede this document. A list of current FIDO Alliance publications and the latest revision of this technical report can be found in the [FIDO Alliance specifications index](https://www.fidoalliance.org/specifications/) at <https://www.fidoalliance.org/specifications/>.

This document was published by the [FIDO Alliance](#) as a Implementation Draft. This document is intended to become a FIDO Alliance Proposed Standard. If you wish to make comments regarding this document, please [Contact Us](#). All comments are welcome.

This Implementation Draft Specification has been prepared by FIDO Alliance, Inc. Permission is hereby granted to use the Specification solely for the purpose of implementing the Specification. No rights are granted to prepare derivative works of this Specification. Entities seeking permission to reproduce portions of this Specification for other uses must contact the FIDO Alliance to determine whether an appropriate license for such use is available.

Implementation of certain elements of this Specification may require licenses under third party intellectual property rights, including without limitation, patent rights. The FIDO Alliance, Inc. and its Members and any other contributors to the Specification are not, and shall not be held, responsible in any manner for identifying or failing to identify any or all such third party intellectual property rights.

THIS FIDO ALLIANCE SPECIFICATION IS PROVIDED "AS IS" AND WITHOUT ANY WARRANTY OF ANY KIND, INCLUDING, WITHOUT LIMITATION, ANY EXPRESS OR IMPLIED WARRANTY OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Table of Contents

- 1. [Notation](#)
 - 1.1 [Key Words](#)
- 2. [Introduction](#)
 - 2.1 [Intended Audience](#)
- 3. [Attack Classification](#)
- 4. [FIDO Security Goals](#)
 - 4.1 [Assets to be Protected](#)
- 5. [FIDO Security Measures](#)
 - 5.1 [Relation between Measures and Goals](#)
- 6. [FIDO Security Assumptions](#)
 - 6.1 [Discussion](#)
- 7. [Threat Analysis](#)
 - 7.1 [Threats to Client Side](#)
 - 7.1.1 [Exploiting User's pattern matching weaknesses](#)
 - 7.1.2 [Threats to the User Device, FIDO Client and Relying Party Client Applications](#)
 - 7.1.3 [Creating a Fake Client](#)
 - 7.1.4 [Threats to FIDO Authenticator](#)
 - 7.1.5 [Threats to Relying Party](#)
 - 7.1.5.1 [Threats to FIDO Server Data](#)

- 7.1.6 Threats to the Secure Channel between Client and Relying Party
 - 7.1.6.1 Exploiting Weaknesses in the Secure Transport of FIDO Messages
- 7.1.7 Threats to the Infrastructure
 - 7.1.7.1 Threats to FIDO Authenticator Manufacturers
 - 7.1.7.2 Threats to FIDO Server Vendors
 - 7.1.7.3 Threats to FIDO Metadata Service Operators
- 7.1.8 Threats Specific to Second Factor Authenticators (UAF / U2F)
 - 7.2 Acknowledgements
- A. References
 - A.1 Informative references

1. Notation

Type names, attribute names and element names are written as `code`.

String literals are enclosed in "", e.g. "UAF-TLV".

In formulas we use "||" to denote byte wise concatenation operations.

UAF specific terminology used in this document is defined in [FIDOGlossary].

1.1 Key Words

The key words "must", "must not", "required", "shall", "shall not", "should", "should not", "recommended", "may", and "optional" in this document are to be interpreted as described in [RFC2119].

2. Introduction

This document analyzes the security properties of the FIDO UAF and U2F families of protocols. Although a brief architectural summary is provided below, readers should familiarize themselves with the FIDO Glossary of Terms [FIDOGlossary] for definitions of terms used throughout. For technical details of various aspects of the architecture, readers should refer to the FIDO Alliance specifications in the Bibliography.

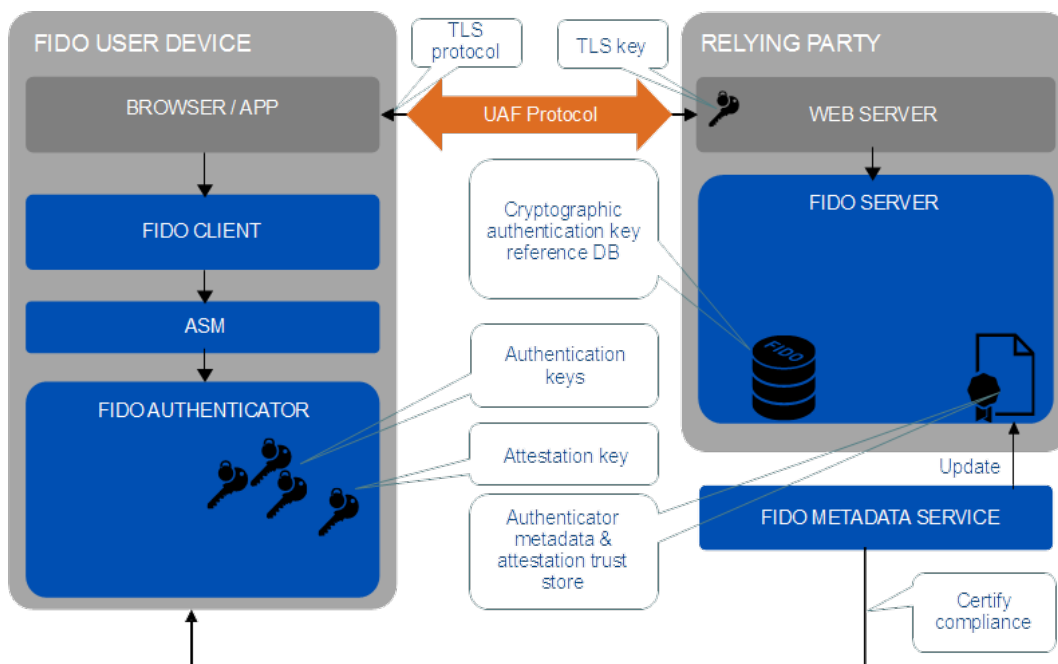


Fig. 1 FIDO Reference Architecture

Conceptually, FIDO involves a conversation between a computing environment controlled by a Relying Party and one controlled by the user to be authenticated. The Relying Party's environment consists conceptually of at least a web server and the server-side portions of a web application, plus a FIDO Server. The FIDO Server has a trust store, containing the (public) trust anchors for the attestation of FIDO Authenticators. The users' environment, referred to as the FIDO user device, consists of one or more FIDO Authenticators, a piece of software called the FIDO Client that is the endpoint for UAF and U2F conversations, and User Agent software. The User Agent software may be a browser hosting a web application delivered by the Relying Party, or it may be a standalone application delivered by the Relying Party. In either case, the FIDO Client, while a conceptually distinct entity, may actually be implemented in whole or part within the boundaries of the User Agent.

2.1 Intended Audience

This document assumes a technical audience that is proficient with security analysis of computing systems and network protocols as well as the specifics of the FIDO architecture and protocol families. It discusses the security goals, security measures, security assumptions and a series of threats to FIDO systems, including the users' computing environment, the Relying Party's computing environment, and the supply chain, including the vendors of FIDO components.

3. Attack Classification

The following attacks all result in user impersonation if successful. However, they have distinguishing characteristics which we use as the basis for attack classification:

1. Automated attacks not focused on the users systems, which affect the user.
2. Automated attacks which are focused on the users' device and which are performed once and lead to the ability to impersonate the user on

an on-going basis without involving him or his device directly.

3. Automated attacks which involve the user or his device for each successful impersonation.
4. Automated attacks to sessions authenticated by the user.
5. Not automatable attacks to the user or his device which are performed once and lead to the ability to impersonate the user on an on-going basis without involving him or his device directly.
6. Not automatable attacks to the user or his device which involve the user or his device for each successful impersonation.

Counter Measures

Examples

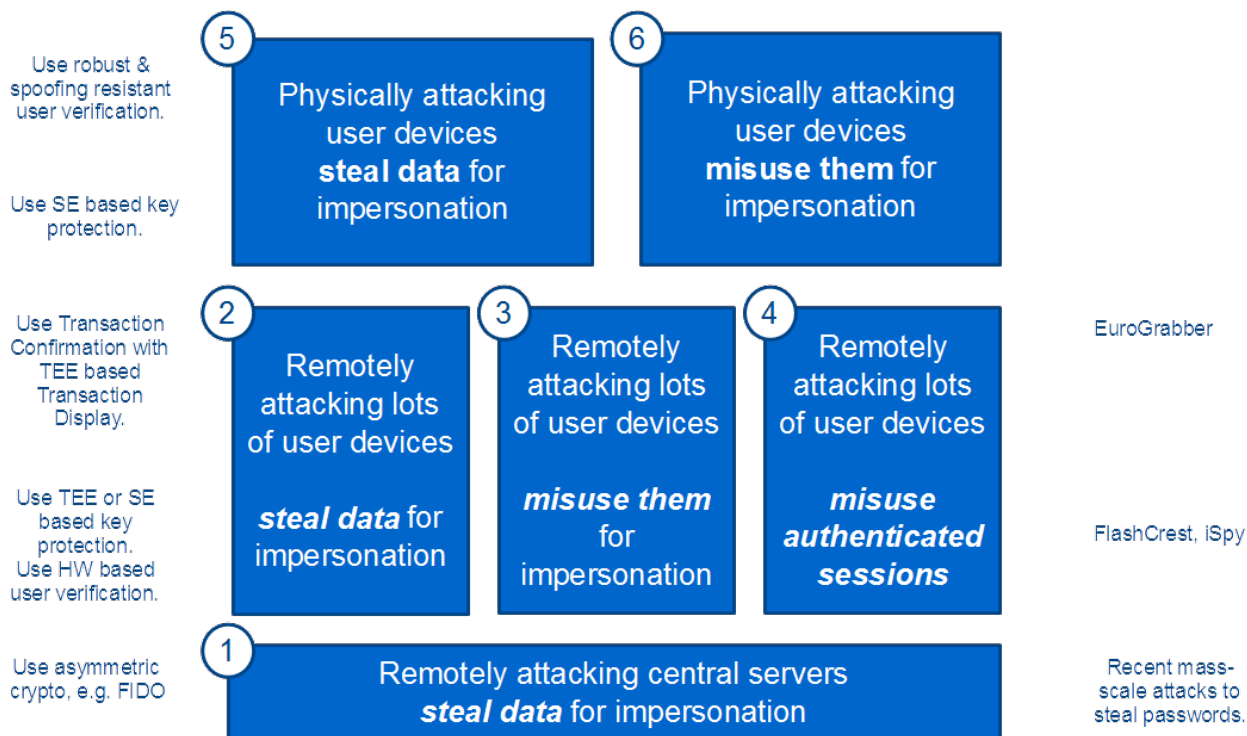


Fig. 2 Attack Classes

The first four attack classes are considered scalable as they are nominally automatable. The attack classes 5 and 6 are not automatable; they involve some kind of manual/physical interaction of the attacker with the user or his device. We will attribute the threats analyzed in this document with the related attack class (AC1 – AC6).

NOTE

1. FIDO uses asymmetric cryptography to protect against AC1. This gives control back to the user, i.e. when using good random numbers, the user's authenticator can make breaking the key as hard as the underlying factoring (in the case of RSA) or discrete logarithm (in the case of DSA or ECDSA) problem.
2. Once counter-measures for this kind of attack are commonly in place, attackers will likely focus on another attack class.
3. The numbers at the attack classes do not imply a feasibility ranking of the related attacks, e.g. it is not necessarily more difficult to perform (AC4) than it is to perform (AC3).
4. The user has almost no influence on the feasibility of attack class (AC1). This makes this attack class really bad.
5. The concept of physical security (i.e. "protect your Authenticator from being stolen"), related to attack classes (AC5) and (AC6) is much better internalized by users than the concept of logical security, related to attack classes (AC2), (AC3) and (AC4).
6. In order to protect against misuse of authenticated sessions (e.g. MITB attacks), the FIDO Authenticator must support the concept of transaction confirmation and the relying party must use it.
7. For an attacker to succeed in impersonating the user, any attack class is sufficient.

Attack Classes

We define the term scalable attack as any attack where the marginal cost of adding an additional target is near zero and which leads to violations of the FIDO security goals.

NOTE

The first four attack classes (AC1, AC2, AC3, and AC4) are considered scalable. The last two attack classes (AC5 and AC6) are not scalable and are performed as one-off user/Relying Party style compromises. We will attribute the threats analyzed in this document with the related attack class (AC1 – AC6).

AC1

Attacks not focused on the users' devices and which lead to violations of FIDO security goals. (e.g., compromise of a Relying Party FIDO database and successful decryption of wrapped keys within the database, phishing, MITM attacks, etc.).

AC2

Scalable attacks involving the Authenticator which, once performed, lead to the ability to violate FIDO security goals on an ongoing basis without later involving the users or their devices directly (e.g., a scalable attack on FIDO Authenticators that recovers the user private keys, allowing the attacker to impersonate the users on an ongoing basis).

AC3

Scalable attacks which involve the user or his device for each instance where the FIDO security goals are violated (e.g., a scalable attack that requires the Authenticator for each successful impersonation).

AC4

Scalable attacks on sessions authenticated by the user which violate FIDO security goals.

AC5

Non-scalable attacks involving the Authenticator which, once performed, lead to the ability to violate FIDO security goals on an ongoing basis without later involving the users or their devices directly (e.g., a non-scalable attack on FIDO Authenticators that recovers the user private keys, allowing the attacker to impersonate the users on an ongoing basis).

AC6

Non-scalable attacks which involve the user or his device for each instance where the FIDO security goals are violated (e.g., a non-scalable attack that requires the Authenticator for each successful impersonation).

NOTE

At this time we are not explicitly addressing classes of physical attacks on the authenticator that may lead to reduced security if the legitimate user uses the authenticator *after* the attacker having physical access to it.

4. FIDO Security Goals

In this section the specific security goals of FIDO are described. The FIDO UAF protocol [[UAFProtocol](#)] and U2F protocol [[U2FOverview](#)] support a variety of different FIDO Authenticators. Even though the security of those authenticators varies, the UAF protocol and the FIDO Server should provide a very high level of security - at least on a conceptual level. In reality it might require a FIDO Authenticator with a high security level in order to fully leverage the FIDO security strength.

NOTE

In certain environments the overall security of the explicit authentication (as provided by FIDO) is less important, as it might be supplemented with a high degree of implicit authentication or the application doesn't even require a high level of authentication strength.

The FIDO U2F protocol [[U2FOverview](#)] supports a more constrained set of Authenticator capabilities. It shares the same security goals as UAF, with the exception of [[SG-14](#)] Transaction Non-Repudiation.

The FIDO protocols have the following security goals:

[SG-1]

Strong User Authentication: Authenticate (i.e. recognize) a user and/or a device to a relying party with high (cryptographic) strength.

[SG-2]

Credential Guessing Resilience: Provide robust protection against eavesdroppers, e.g. *be resilient to physical observation, resilient to targeted impersonation, resilient to throttled and unthrottled guessing.*

[SG-3]

Credential Disclosure Resilience: Be *resilient to phishing attacks* and real-time phishing attack, including resilience to online attacks by adversaries able to actively manipulate network traffic.

[SG-4]

Unlinkability: Protect the protocol conversation such that any two relying parties cannot link the conversation to one user (i.e. be *unlinkable*).

[SG-5]

Verifier Leak Resilience: Be *resilient to leaks from other relying parties* i.e., nothing that a verifier could possibly leak can help an attacker impersonate the user to another relying party.

[SG-6]

Authenticator Leak Resilience: Be resilient to leaks from other FIDO Authenticators. i.e., nothing that a particular FIDO Authenticator could possibly leak can help an attacker to impersonate any other user to any relying party.

[SG-7]

User Consent: Notify the user before a relationship to a new relying party is being established (*requiring explicit consent*).

[SG-8]

Limited PII: Limit the amount of personal identifiable information (PII) exposed to the relying party to the absolute minimum.

[SG-9]

Attestable Properties: Relying Party must be able to verify FIDO Authenticator model/type (in order to calculate the associated risk).

[SG-10]

DoS Resistance: Be resilient to *Denial of Service Attacks* i.e. prevent attackers from inserting invalid registration information for a legitimate user for the next login phase. Afterward, the legitimate user will not be able to login successfully anymore.

[SG-11]

Forgery Resistance: Be resilient to *Forgery Attacks (Impersonation Attacks)*. i.e. prevent attackers from attempting to modify intercepted communications in order to masquerade as the legitimate user and login to the system.

[SG-12]

Parallel Session Resistance: Be resilient to *Parallel Session Attacks*. Without knowing a user's authentication credential, an attacker can masquerade as the legitimate user by creating a valid authentication message out of some eavesdropped communication between the user and the server.

[SG-13]

Forwarding Resistance: Be resilient to *Forwarding and Replay Attacks*. Having intercepted previous communications, an attacker can impersonate the legal user to authenticate to the system. The attacker can replay or forward the intercepted messages.

[SG-14] (not covered by U2F)

Transaction Non-Repudiation: Provide strong cryptographic non-repudiation for secure transactions.

[SG-15]

Respect for Operating Environment Security Boundaries: Ensure that registrations and private key material as a shared system resource is appropriately protected according to the operating environment privilege boundaries in place on the FIDO user device.

[SG-16]

Assessable Level of Security: Ensure that the design and implementation of the Authenticator allows for the testing laboratory / FIDO Alliance to assess the level of security provided by the Authenticator.

NOTE

For a definition of the phrases printed *in italics*, refer to [[QuestToReplacePasswords](#)] and to [[PasswordAuthSchemesKeyIssues](#)]

4.1 Assets to be Protected

Independent of any particular implementation, the FIDO protocols assume some assets to be present and to be protected.

1. Cryptographic Authentication Private Key. Typically, private keys in FIDO are unique for each tuple of (relying party, user account, authenticator).
2. Cryptographic Authentication Key Reference. This is the cryptographic material stored at the relying party and used to uniquely verify the Cryptographic Authentication Key, typically the public key corresponding to the authentication private key.
3. Authenticator Attestation Key (as stored in each authenticator). This should only be usable to attest a Cryptographic Authentication Key and the type/model and manufacturing batch of an Authenticator. Attestation keys are either ECDAA keys [FIDOecdaaAlgorithm] or the attestation keys and certificates are shared by a large number of authenticators in a device class from a given vendor in order to prevent their becoming a linkable identifier across relying parties. Authenticator attestation certificates may be self-signed, or signed by an authority key controlled by the vendor.
4. Authenticator Attestation Authority Key. An authenticator vendor may elect to sign authenticator attestation certificates with a per-vendor certificate authority key.
5. Authenticator Attestation Authority Certificate. Contained in the initial/default trust store as part of the FIDO Server and contained in the active trust store maintained by each relying party.
6. Active Trust Store. Contains all trusted attestation root certificates for a given FIDO server.
7. All data items suitable for uniquely identifying the authenticator across relying parties. An attack on those would break the non-linkability security goal.
8. Private key of Relying Party TLS server certificate.
9. TLS root certificate trust store for the users' browser/app.

5. FIDO Security Measures

NOTE

Particular implementations of FIDO Clients, Authenticators, Servers and participating applications may not implement all of these security measures (e.g. Secure Display, [SM-10] Transaction Confirmation) and they also might (and should) implement additional security measures.

NOTE

The U2F protocol lacks support for [SM-5] Secure Display, [SM-10] Transaction Confirmation, has only server-supplied [SM-8] Protocol Nonces, and [SM-3] Authenticator Class Attestation is implicit as there is only a single class of device.

[SM-1] (U2F + UAF)

Key Protection: Authentication key is protected against misuse. Misuse means any use violating the FIDO specification or the details given in the Metadata Statement. Before a key can be used, it requires the User to unlock it using the user verification method specified in the Authenticator Metadata Statement (Silent Authenticators do not require any user verification method).

[SM-2] (U2F + UAF)

Unique Authentication Keys: Cryptographic authentication key is specific and unique to the tuple of (FIDO Authenticator, User, Relying Party).

[SM-3] (U2F + UAF)

Authenticator Class Attestation: Hardware-based FIDO Authenticators support authenticator attestation using an attestation key using one of the FIDO specified attestation types and algorithms. Each relying party receives regular updates of the trust store (through the FIDO Metadata service).

[SM-4] (UAF)

Authenticator Status Checking: Relying Parties can download latest known status of authenticators included in the FIDO Metadata Service. The FIDO Server should take this information into account. Authenticator manufacturers should notify the FIDO Alliance about compromised authenticators. In the case of FIDO certified authenticators, such notification might even be mandatory.

[SM-5] (UAF)

User Consent: FIDO Client implements a user interface for getting user's consent on any actions (except authentication with silent authenticator) and displaying RP name (derived from server URL).

[SM-6] (U2F + UAF)

Cryptographically Secure Verifier Database: The relying party stores only the public portion of an asymmetric key pair, or an encrypted key handle, as a cryptographic authentication key reference.

[SM-7] (U2F + UAF)

Secure Channel with Server Authentication: The TLS protocol with server authentication or a transport with equivalent properties is used as transport protocol for UAF. The use of https is enforced by a browser or Relying Party application.

[SM-8] (UAF)

Protocol Nonces: Both server and client supplied nonces are used for UAF registration and authentication. U2F requires server supplied nonces.

[SM-9] (U2F + UAF)

Authenticator Certification: The FIDO Metadata Service includes the Authenticator certification status.

[SM-10] (UAF)

Transaction Confirmation (WYSIWYS): Secure Display (WYSIWYS) (optionally) implemented by the FIDO Authenticators is used by FIDO Client for displaying relying party name and transaction data to be confirmed by the user.

[SM-11] (U2F + UAF)

Round Trip Integrity: FIDO server verifies that the transaction data related to the server challenge received in the UAF message from the FIDO client is identical to the transaction data and server challenge delivered as part of the UAF request message.

[SM-12] (U2F + UAF)

Channel Binding: Relying Party servers may verify the continuity of a secure channel with a client application.

[SM-13] (UAF)

Key Handle Access Token: Authenticators not intended to roam between untrusted systems are able to constrain the use of registration keys within the privilege boundaries defined by the operating environment of the user device (per-user, or per application, or per-user + per-application as appropriate).

[SM-14] (U2F + UAF)

AppID Separation: A Relying Party can declare the application identities allowed to access its registered keys, for operating environments on user devices that support this concept.

[SM-15] (U2F + UAF)

Signature Counter: Authenticators send a monotonically increasing signature counter that a Relying Party can check to possibly detect cloned authenticators.

[SM-16] (U2F + UAF)

Use of strong, modern Cryptographic Primitives: The FIDO specifications stipulate the use of strong, modern cryptographic primitives helping to ensure the overall security of conformant FIDO implementations. The FIDO Authenticator certification program defines the "Allowed Cryptography List" for allowed cryptographic primitives to be used in FIDO certified authenticators.

[SM-17] (U2F + UAF)

Resistance to Side Channel Attacks.

[SM-18] (U2F + UAF)

Resistance to Injected Faults in Cryptographic Functions. This security measure purely deals with the cryptographic functions, as compared to the much more general [SM-28].

- [SM-19] (UAF)**
Bounded Probability of a Birthday Collision. For randomly generated nonces, the total number of nonces that can be generated is limited to bound the probability of a birthday collision of generated values.
- [SM-20] (U2F + UAF)**
Individual authenticators are indistinguishable provided authenticators sharing attestation keys are manufactured in sufficiently large (e.g. > 100000) per-model batches.
- [SM-21] (U2F + UAF)**
Authentication and replay-resistance (freshness assurance) of externally-stored protected information.
- [SM-22] (U2F + UAF)**
Certified FIDO Authenticators fully described by the vendor, and tested to verify that it functions as specified.
- [SM-23] (U2F + UAF)**
Key Handles containing a key are cryptographically linked with the Authenticator that produced the Key Handle and with the Relying Party associated with the Key Handle.
- [SM-24] (U2F + UAF)**
Design, implementation and manufacture of certified FIDO Authenticators supports Authenticator security.
- [SM-25] (U2F + UAF)**
Depending on the certification level, certified authenticators are required to implement a Trusted Path for all user / Authenticator direct interactions.
- [SM-26] (U2F + UAF)**
Input Data Validation: Malformed or maliciously crafted input data does not result in unexpected Authenticator behavior.
- [SM-27] (U2F + UAF)**
Protection of user verification reference data and biometric data.
- [SM-28] (U2F + UAF)**
Resistance to Fault Injection Attacks.
- [SM-29] (U2F + UAF)**
Resistance to Remote Timing Attacks: No leakage of secret information to remote entities via variation of operation execution time.

5.1 Relation between Measures and Goals

Security Goal	Supporting Security Measures
[SG-1] Strong User Authentication	<ul style="list-style-type: none"> [SM-1] Key Protection [SM-12] Channel Binding [SM-14] AppID Separation [SM-15] Signature Counter [SM-16] Allowed Crypto Primitives [SM-17] Resistance to Side Channel Attacks [SM-21] Authentication and replay-resistance [SM-23] Key Handles cryptographically linked with the Authenticator [SM-25] Trusted path for all user interactions [SM-29] Resistance to Remote Timing Attacks
[SG-2] Credential Guessing Resilience	<ul style="list-style-type: none"> [SM-1] Key Protection [SM-6] Cryptographically Secure Verifier Database [SM-16] Allowed Crypto Primitives
[SG-3] Credential Disclosure Resilience	<ul style="list-style-type: none"> [SM-1] Key Protection [SM-9] Authenticator Certification [SM-15] Signature Counter [SM-17] Resistance to Side Channel Attacks [SM-29] Resistance to Remote Timing Attacks
[SG-4] Unlinkability	<ul style="list-style-type: none"> [SM-2] Unique Authentication Keys [SM-3] Authenticator Class Attestation [SM-20] No Identifying Information
[SG-5] Verifier Leak Resilience	<ul style="list-style-type: none"> [SM-2] Unique Authentication Keys [SM-6] Cryptographically Secure Verifier Database [SM-16] Allowed Crypto Primitives
[SG-6] Authenticator Leak Resilience	<ul style="list-style-type: none"> [SM-9] Authenticator Certification [SM-15] Signature Counter [SM-16] Allowed Crypto Primitives

Security Goal	Supporting Security Measures
[SG-7] User Consent	[SM-1] Key Protection [SM-5] User Consent [SM-7] Secure Channel with Server Authentication [SM-10] Transaction Confirmation (WYSIWYS) [SM-25] Trusted path for all user interactions
[SG-8] Limited PII	[SM-2] Unique Authentication Keys [SM-20] No Identifying Information
[SG-9] Attestable Properties	[SM-3] Authenticator Class Attestation [SM-4] Authenticator Status Checking [SM-9] Authenticator Certification
[SG-10] DoS Resistance	[SM-8] Protocol Nonces
[SG-11] Forgery Resistance	[SM-7] Secure Channel with Server Authentication [SM-8] Protocol Nonces [SM-11] Round Trip Integrity [SM-12] Channel Binding [SM-17] Resistance to Side Channel Attacks [SM-23] Key Handles cryptographically linked with the Authenticator [SM-29] Resistance to Remote Timing Attacks
[SG-12] Parallel Session Resistance	[SM-7] Secure Channel with Server Authentication [SM-8] Protocol Nonces [SM-11] Round Trip Integrity [SM-12] Channel Binding
[SG-13] Forwarding Resistance	[SM-7] Secure Channel with Server Authentication [SM-8] Protocol Nonces [SM-11] Round Trip Integrity [SM-12] Channel Binding
[SG-14] Transaction Non-Repudiation	[SM-1] Key Protection [SM-2] Unique Authentication Keys [SM-8] Protocol Nonces [SM-9] Authenticator Certification [SM-10] Transaction Confirmation (WYSIWYS) [SM-11] Round Trip Integrity [SM-12] Channel Binding [SM-25] Trusted path for all user interactions
[SG-15] Respect for Operating Environment Security Boundaries	[SM-13] Key Handle Access Token [SM-14] AppID Separation

6. FIDO Security Assumptions

In this section, we enumerate the assumptions we are making regarding the security characteristics of the operating environment components on which a FIDO implementation depends.

- [SA-1] The Authenticator and its cryptographic algorithms and parameters (key size, mode, output length, etc.) in use are not subject to unknown weaknesses that make them unfit for their purpose in encrypting, digitally signing, and authenticating messages.
- [SA-2] Operating system privilege separation mechanisms relied up on by the software modules involved in a FIDO operation on the user device perform as advertised. E.g. boundaries between user and kernel mode, between user accounts, and between applications (where applicable) are securely enforced and security principals can be mutually, securely identifiable.
- [SA-3] Applications on the user device are able to establish secure channels that provide trustworthy server authentication, and confidentiality and integrity for messages (e.g., through TLS).
- [SA-4] The computing environment on the FIDO user device and the and applications involved in a FIDO operation act as trustworthy agents of the user.
- [SA-5] The inherent value of a cryptographic key resides in the confidence it imparts, and this commodity decays with the passage of time, irrespective of any compromise event. As a result the effective assurance level of authenticators will be reduced over time.
- [SA-6] The computing resources at the Relying Party involved in processing a FIDO operation act as trustworthy agents of the Relying Party.

6.1 Discussion

With regard to [SA-4] and malicious computation on the FIDO user device, only very limited guarantees can be made within the scope of these assumptions. Malicious code privileged at the level of the trusted computing base can always violate [SA-2] and [SA-3]. Malicious code privileged at the level of the users' account in traditional multi-user environments will also likely be able to violate [SA-3].

FIDO can also provide only limited protections when a user chooses to deliberately violate [SA-4], e.g. by roaming a USB authenticator to an untrusted system like a kiosk, or by granting permissions to access all authentication keys to a malicious app in a mobile environment. Transaction Confirmation can be used as a method to protect against compromised FIDO user devices.

In to components such as the FIDO Client, Server, Authenticators and the mix of software and hardware modules they are comprised of, the end-to-end security goals also depend on correct implementation and adherence to FIDO security guidance by other participating components, including web browsers and relying party applications. Some configurations and uses may not be able to meet all security goals. For example, authenticators may lack a secure display, they may be composed only of unattestable software components, they may be deliberately designed to roam between untrusted operating environments, and some operating environments may not provide all necessary security primitives (e.g., secure IPC, application isolation, modern TLS implementations, etc.)

7. Threat Analysis

In the following tables describing threats, we mention the relevant attack class(es) in the left column if the threat might lead to user impersonation.

7.1 Threats to Client Side

7.1.1 Exploiting User's pattern matching weaknesses

T-1.1.1	Homograph Mis-Registration	Violates
AC3	<p>The user registers a FIDO authentication key with a fraudulent web site instead of the genuine Relying Party.</p> <p>Consequences: The fraudulent site may convince the user to disclose a set of non-FIDO credentials sufficient to allow the attacker to register a FIDO Authenticator under its own control, at the genuine Relying Party, on the users' behalf, violating [SG-1] Strong User Authentication.</p> <p>Mitigations: Disclosure of non-FIDO credentials is outside of the scope of the FIDO security measures, but Relying Parties should be aware that the initial strength of an authentication key is no better than the identity-proofing applied as part of the registration process.</p>	SG-1

T-1.1.2	Homograph Mis-Authentication	Violates
AC3	<p>The user accidentally browses to a fraudulent web site. The attacker tries to act as man-in-the-middle (MITM) and requests the user to authenticate. In the case of username/password based authentication this is a typical phishing attack.</p> <p>Consequences: The FIDO subsystem will determine that either (a) no FIDO authenticator has been registered with the fraudulent site or (b) it will use the FIDO Uauth key registered to the fraudulent site - which is different from the Uauth key for the relying party's site.</p> <p>Mitigations: FIDO inherently ties keys to the relying party (formally identified by the AppID, and authenticated by TLS and the CA infrastructure).</p>	SG-1, SG-4

7.1.2 Threats to the User Device, FIDO Client and Relying Party Client Applications

T-1.2.1	FIDO Client Corruption	Violates
AC3	<p>Attacker gains ability to execute code in the security context of the FIDO Client.</p> <p>Consequences: Violation of [SA-4].</p> <p>Mitigations: When the operating environment on the FIDO user device allows, the FIDO Client should operate in a privileged and isolated context under [SA-2] to protect itself from malicious modification by anything outside of the Trusted Computing Base.</p>	SA-4

T-1.2.2	Logical/Physical User Device Attack	Violates
	<p>Attacker gains physical access to the FIDO user device but not the FIDO Authenticator.</p> <p>Consequences: Possible violation of [SA-4] by installing malicious software or otherwise tampering with the FIDO user device.</p>	

AC3 1.2.2	Mitigations: [SM-1] Key Protection prevents the disclosure of authentication keys or other assets during a transient compromise of the FIDO user device. Logical/Physical User Device Attack	Violates
AC5	A persistent compromise of the FIDO user device can lead to a violation of [SA-4] unless additional protection measures outside the scope of FIDO are applied to the FIDO user device. (e.g. whole disk encryption and boot-chain integrity).	
T-1.2.3 User Device Account Access Violates		
AC3 / AC4	<p>Attacker gains access to a user's login credentials on the FIDO user device.</p> <p>Consequences: Authenticators might be remotely abused, or weakly-verifying authenticators might be locally abused, violating [SG-1] Strong User Authentication and [SG-13] Transaction Non-Repudiation.</p> <p>Possible violation of [SA-4] by the installation of malicious software.</p> <p>Mitigations: Relying Parties can use [SM-9] Authenticator Certification and [SM-3] Authenticator Class Attestation to determine the nature of authenticators and not rely on weak, or weakly-verifying authenticators for high value operations.</p>	SG-1, SG-13; SA-4
T-1.2.4 App Server Verification Error Violates		
AC3	<p>A client application fails to properly validate the remote sever identity, accepts forged or stolen credentials for a remote server, or allows weak or missing cryptographic protections for the secure channel.</p> <p>Consequences: An active network adversary can modify the Relying Party's authenticator policy and downgrade the client's choice of authenticator to make it easier to attack.</p> <p>An active network adversary can intercept or view FIDO messages intended for the Relying Party. It may be able to use this ability to violate [SG-12] Parallel Session Resistance, [SG-11] Forgery Resistance or [SG-13] Forwarding Resistance.</p> <p>Mitigations: The server can verify [SM-8] Protocol Nonces to detect replayed messages and protect from an adversary that can read but not modify traffic in a secure channel.</p> <p>The server can mandate a channel with strong cryptographic protections to prevent message forgery and can verify a [SM-12] Channel Binding to detect forwarded messages.</p>	SG-11, SG-12, SG-13
T-1.2.5 RP App Corruption Violates		
	<p>An attacker is able to obtain malicious execution in the security context of the Relying Party client application (e.g. via Cross-Site Scripting (XSS)) or abuse the secure channel or session identifier after the user has successfully authenticated. This is a client side attack.</p> <p>Consequences: The attacker is able to control the users' session, violating [SG-14] Transaction Non-Repudiation.</p> <p>Mitigations: The server can employ [SM-10] Transaction Confirmation to gain additional assurance for high value operations.</p>	SG-14
T-1.2.6 Fingerprinting Authenticators Violates		
	<p>A remote adversary is able to uniquely identify a FIDO user device using the fingerprint of discoverable configuration of its FIDO Authenticators.</p> <p>Consequences: The exposed information violates [SG-8] Limited PII, allowing an adversary to violate [SG-7] User Consent by strongly identifying the user without their knowledge and [SG-4] Unlinkability by sharing that fingerprint.</p> <p>Mitigations: [SM-3] Authenticator Class Attestation ensures that the fingerprint of an Authenticator will not be unique.</p> <p>For web browsing situations where this threat is most prominent, user agents may provide additional user controls around the discoverability of FIDO Authenticators.</p>	SG-4, SG7, SG-8
T-1.2.7 App to FIDO Client full MITM attack Violates		
AC3	<p>Malicious software on the FIDO user device is able to read, tamper with, or spoof the endpoint of inter-process communication channels between the FIDO Client and browser or Relying Party application.</p> <p>Consequences: Adversary is able to subvert [SA-2].</p> <p>Mitigations: On platforms where [SA-2] is not strong the security of the system may depend on preventing malicious applications from being loaded onto the FIDO user device. Such protections, e.g. app store policing, are outside the scope of FIDO.</p> <p>When using [SM-10] Transaction Confirmation, the user will be presented with the relevant AppID and transaction text and will be able to evaluate whether or not to consent to the transaction.</p>	SA-2
T-1.2.8 Authenticator to App Read-Only MITM attack Violates		
AC3	<p>An adversary is able to obtain an authenticator's signed protocol response message.</p> <p>Consequences: The attacker attempts to replay the message to authenticate as the user, violating [SG-1] Strong User Authentication, [SG-13] Forwarding Resistance and [SG-12] Parallel Session Resistance.</p> <p>Mitigations: The server can use [SM-8] Protocol Nonces to detect replay of messages and verify [SM-11] Round Trip Integrity to detect modified messages.</p>	SG-1, SG-12, SG-13

T-1.2.9	Malicious App	Violates
AC3	<p>A user installs an application that represents itself as being associated with to one Relying Party application but actually initiates a protocol conversation with a different Relying Party and attempts to abuse previously registered authentication keys at that Relying Party.</p> <p>Consequences: Adversary is able to violate [SG-7] User Consent by misrepresenting the target of authentication.</p> <p>Other consequences equivalent to [T-1.2.5]</p> <p>Mitigations: If a [SM-5] Transaction Confirmation Display is present, the user may be able to verify the true target of an operation.</p> <p>If the malicious application attempts to communicate directly with an Authenticator that uses [SM-13] KeyHandleAccessToken, it should not be able to access keys registered by other FIDO Clients.</p> <p>If the operating environment on the FIDO user device supports it, the FIDO client may be able to determine the application's identity and verify if it is authorized to target that Relying Party using a [SM-14] AppID Separation.</p>	SG-7

T-1.2.10	Phishing Attack	Violates
AC2	<p>A Phisher convinces the user to enter his PIN used for user verification into an application / web site disclosing the PIN to the Phisher. In the traditional username/password world this enables the attacker to successfully impersonate the user (to the relying party).</p> <p>Consequences: None as the phisher additionally would need access to the Authenticator in order to pass user verification [SM-1]. In FIDO, the user verification PIN (if user verification is done via PIN) is not known to the relying party and hence isn't sufficient for user impersonation. If user verification is done using an alternative user verification method, this applies accordingly.</p> <p>Mitigations: In FIDO, the Uauth.priv key is used to sign a relying party supplied challenge. without (use) access to that key, no impersonation is possible.</p>	SG-1

7.1.3 Creating a Fake Client

T-1.3.1	Malicious FIDO Client	Violates
AC3	<p>Attacker convinces users to install and use a malicious FIDO Client.</p> <p>Consequences: Violation of [SA-4]</p> <p>Mitigations: Mitigating malicious software installation is outside the scope of FIDO.</p> <p>If an authenticator implements [SM-1] Key Protection, the user may be able to recover full control of their registered authentication keys by removing the malicious software from their user device.</p> <p>When using [SM-10] Transaction Confirmation, the user sees the real AppIDs and transaction text and can decide to accept or reject the action.</p>	SA-4

7.1.4 Threats to FIDO Authenticator

T-1.4.1	Malicious Authenticator	Violates
AC2, AC3	<p>Attacker convinces users to use a maliciously implemented authenticator.</p> <p>Consequences: The fake authenticator does not implement any appropriate security measures and is able to violate all security goals of FIDO.</p> <p>Mitigations: A user may be unable to distinguish a malicious authenticator, but a Relying Party can use [SM-3] Authenticator Class Attestation to identify and only allow registration of reliable authenticators that have passed [SM-9] Authenticator Certification.</p> <p>A Relying Party can additionally rely on [SM-4] Authenticator Status Checking to check if an attestation presented by a malicious authenticator has been marked as compromised.</p>	SG-1

T-1.4.2	Uauth.priv Key Compromise	Violates
AC2	<p>Attacker succeeds in extracting a user's cryptographic authentication private key for use in a different context.</p> <p>Consequences: The attacker could impersonate the user with a cloned authenticator that does not do trustworthy user verification, violating [SG-1].</p> <p>Mitigations: [SM-1] Key Protection measures are intended to prevent this.</p> <p>Each authentication private key is only used for one relying party.</p> <p>Relying Parties can check [SM-9] Authenticator Certification attributes to determine the type of key protection in use by a given authenticator class.</p> <p>Relying Parties can additionally verify the [SM-15] Signature Counter and detect that an authenticator has been cloned if it ever fails to advance relative to the prior operation.</p>	SG-1

T-1.4.3	User Verification By-Pass	Violates

T-1.4.3	User Verification By-Pass Attacker could use the cryptographic authentication key (inside the authenticator) either with or without being noticed by the legitimate user.	Violates
AC3, AC5	Consequences: Attacker could impersonate user, violating [SG-1]. Mitigations: A user can only register and a Relying Party only allow authenticators that perform [SM-1] Key Protection with an appropriately secure user verification process. Does not apply to Silent Authenticators (see FIDOGlossary).	SG-1
T-1.4.4	Physical Authenticator Attack	Violates
AC2, AC5, AC6	Attacker could get physical access to FIDO Authenticator (e.g. by stealing it). Consequences: Attacker could bring the authenticator in a lab in order to use the authentication key (e.g. by-passing user verification and knowing the RP related to this key). If this physical attack succeeds, the attacker could successfully impersonate the user, violating [SG-1] Strong User Authentication. Attacker can introduce a low entropy situation to recover an ECDSA signature key (or otherwise extract the Uauth.priv key), violating [SG-9] Attestable Properties if the attestation key is targeted or [SG-1] Strong User Authentication if a user key is targeted. Mitigations: [SM-1] Key Protection includes requirements to implement strong protections for key material, including resilience to offline attacks and low entropy situations. Relying Parties should use [SM-3] Authenticator Class Attestation to only accept Authenticators implementing a sufficiently strong user verification method.	SG-1
T-1.4.6	Fake Authenticator	Violates
AC2	Attacker is able to extract the authenticator attestation key from an authenticator, e.g. by neutralizing physical countermeasures in a laboratory setting. Consequences: Attacker can violate [SG-9] Attestable Properties by creating a malicious hardware or software device that represents itself as a legitimate one. Mitigations: Relying Parties can use [SM-4] Authenticator Status Checking to identify known-compromised keys. Identification of such compromise is outside the strict scope of the FIDO protocols.	SG-9
T-1.4.7	Transaction Confirmation Display Overlay Attack	Violates
AC6	Attacker is able to subvert [SM-5] Secure Display functionality (WYSIWYS), perhaps by overlaying the display with false information. Consequences: Violation of [SG-14] Transaction Non-Repudiation. Mitigations: Authenticator implementations must take care to protect in their implementation of a secure display, e.g. by implementing a distinct hardware display or employing appropriate privileges in the operating environment of the user device to protect against spoofing and tampering. [SM-9] Authenticator Certification will provide Relying Parties with metadata about the nature of a transaction confirmation display information that can be used to assess whether it matches the assurance level and risk tolerance of the Relying Party for that particular transaction.	SG-14
T-1.4.8	Signature Algorithm Attack	Violates
AC1, AC2, AC3, AC5	A cryptographic attack is discovered against the public key cryptography system used to sign data by the FIDO authenticator. See also T-1.4.10. Consequences: Attacker is able to use messages generated by the client to violate [SG-2] Credential Guessing Resistance. Mitigations: [SM-8] Protocol Nonces, including client-generated entropy, limit the amount of control any adversary has over the internal structure of an authenticator. [SM-1] Key Protection for non-silent authenticators requires user interaction to authorize any operation performed with the authentication key, severely limiting the rate at which an adversary can perform adaptive cryptographic attacks.	SG-2
T-1.4.9	Abuse Functionality	Violates
AC2, AC3, AC5, AC6	It might be possible for an attacker to abuse the Authenticator functionality by sending commands with invalid parameters or invalid commands to the Authenticator. Consequences: This might lead to e.g., user verification by-pass or potential key extraction. Mitigations: Proper robustness (e.g. due to testing) of the Authenticator firmware.	SG-1
T-1.4.10	Random Number prediction	Violates
AC2, AC3, AC5,	It might be possible for an attacker to get access to information allowing the prediction of RNG data. Consequences: This might lead to key compromise situation [T-1.4.2] when using ECDSA (if the k value is used multiple times or if it is predictable).	SG-1

AC6 1.4.10	Mitigations: Proper robustness of the Authenticator's RNG and verification of the relevant operating environment parameters (e.g. temperature, ...). Random Number prediction	Violates
T-1.4.11 Firmware Rollback		
	Attacker might be able to install a previous and potentially buggy version of the firmware. Consequences: This might lead to successful attacks, e.g. T-1.4.9. Mitigations: Proper robustness firmware update and verification method.	<u>SG-1</u>
T-1.4.12 User Verification Data Injection		
AC3, AC6	Attacker might be able to inject pre-captured user verification data into the Authenticator. For example, if a password is used as user verification method, the attacker could capture the password entered by the user and then send the correct password to the Authenticator (by-passing the expected keyboard/PIN pad). Passwords could be captured ahead of the attack e.g. by convincing the user to enter the password into a malicious app ("phishing") or by spying directly or indirectly the password data. In another example, some malware could play an audio stream which would be recorded by the microphone and used by a Speaker-Recognition based Authenticator. Consequences: This might lead to successful user impersonation (if the attacker has access to valid user verification data). Mitigations: Use a physically secured user verification input method, e.g. Fingerprint Sensor or Trusted-User-Interface for PIN entry which cannot be by-passed by malware.	<u>SG-1</u>
T-1.4.13 Verification Reference Data Modification		
AC3, AC6	An attacker gains logical or physical access to the Authenticator and modifies Verification Reference Data (e.g. hashed PIN value, fingerprint templates) stored in the Authenticator and adds reference data known to or reproducible by the attacker. Consequences: The attacker would be recognized as the legitimate User and could impersonate the user. Mitigations: [SM-27] Proper protection of the the verification reference data and biometric data in the Authenticator.	<u>SG-1</u>
T-1.4.14 Read access to captured user verification data		
AC3, AC6	The Attacker gained read access to the captured user verification data (e.g. PIN, fingerprint image, ...). Consequences: The attacker gets access to PII and could disclose it violating [SG-8]. Mitigations: Limiting access to the user verification data to the Authenticator exclusively.	<u>SG-8</u>
T-1.4.15 Compromised the internal PRNG state and the entropy source		
AC1, AC2, AC5	In this threat, an attacker compromises the entropy source prior to the Authenticator initially seeding the PRNG during initialization or otherwise compromises the internal PRNG state, and the attacker is able to know or specify all future entropy inputs to the PRNG. No PRNG is able to recover to a secure status under this threat, but it serves as a useful point for comparison. Consequences: May undermine [SG-1], [SG-2], [SG-3], [SG-4], [SG-11], [SG-14]. Mitigations: This constitutes a complete compromise of the RNG, with no ability to recover, so mitigation for this threat involves reducing the impact of a compromised RNG. This is partially mitigated by using an allowed random number generator that allows secure integration of additional input [SM-16] and introduction of data derived from the RP challenge additional input to the PRNG, which can help so long as the attacker has not additionally compromised the TLS session or the ASM / Authenticator link. Using the deterministic signature generation methods (e.g., RFC 6979) can reduce the risk of compromise of existing keys during the signature process, as can using the private key and hash of the signed message as additional input to the PRNG during signature generation. Prevention of non-scalable versions of this style of attack is at least partially addressed by [SM-17] and [SM-18].	SG-1, SG-2, SG-3, SG-4, SG-11, SG-14
T-1.4.16 Compromised entropy source after successful seeding during initialization		
AC1, AC2, AC5	In this threat, an attacker gains the ability to influence the Authenticator's entropy source, but only after the initial seeding has been conducted (e.g., if initial seeding occurred prior to the attack and / or as per-Authenticator factory injection of entropy). Consequences: May undermine [SG-1], [SG-2], [SG-3], [SG-4], [SG-11], [SG-14]. Mitigations: This is mitigated by using an allowed PRNG which retains PRNG state between power cycles; i.e., which conserves PRNG state even when being reseeded [SM-16]. Prevention of non-scalable versions of this style of attack is at least partially addressed by [SM-17] and [SM-18].	SG-1, SG-2, SG-3, SG-4, SG-11, SG-14
T-1.4.17 Compromised the internal PRNG state, but not the entropy source		
	In this threat, an attacker compromises the entropy source prior to seeding the PRNG or otherwise compromises the internal PRNG state, but then at some point, the attacker no longer can access / control the entropy source.	

T-1.4.17	Consequences: May undermine [SG-1], [SG-2], [SG-3], [SG-4], [SG-11], [SG-14]. Compromised the internal PRNG state, but not the entropy source	Violates
AC2, AC5	Mitigations: This can be mitigated by Authenticators reseeding periodically from an internal entropy source [SM-16]. As a note, this imposes a total number of random number generator requests prior to a required reseed event; in the event that the Authenticator does not have an entropy source internally, this may act as a hard limit on the number of registrations / authentications that such an Authenticator can perform. Prevention of non-scalable versions of this style of attack is at least partially addressed by [SM-17] and [SM-18].	SG-1, SG-3, SG-4, SG-11, SG-14

T-1.4.18	Bad Key Generation	Violates
AC1, AC2, AC5	In this threat, random chance or active attack causes the key generated to be cryptographically flawed; e.g., an RSA key that can be factored using the Pollard p-1 algorithm more quickly than with the General Number Field Sieve. See also T-1.4.21. Consequences: May undermine [SG-1], [SG-2], [SG-4], [SG-11], [SG-14] Mitigations: This is mitigated by requiring use of an allowed random number generator (in the case of certified authenticators), requiring that keys be generated in the way required in the relevant standard specified in the Allowed Cryptography List [SM-16], and making the key generation process resistant to tampering by the attacker [SM-18].	SG-1, SG-2, SG-4, SG-11, SG-14

T-1.4.19	Local external side channel attacks	Violates
AC2 (associated with shared keys), AC5	In this threat, an attacker with possession of the Authenticator may be able to extract keys using timing, power, RF, or near-field analysis. The impact depends on the key or secret recovered. Consequences: May undermine [SG-1], [SG-2], [SG-4], [SG-11], [SG-14]. Mitigations: This is mitigated by the side channel resistance security measure [SM-17].	SG-1, SG-2, SG-4, SG-11, SG-14

T-1.4.20	Internal side channel attacks	Violates
AC2 (associated with shared keys), AC5	In this threat, an attacker controlling a process running on the same hardware environment as the Authenticator may be able to recover keys by using information leaked by hardware or operating system characteristics (e.g., how often the attacker's process is scheduled, the state of the L1, L2 caches, etc.). Consequences: May undermine [SG-1], [SG-4], [SG-11], [SG-14]. Mitigations: This is mitigated by the side channel resistance security measure [SM-17].	SG-1, SG-4, SG-11, SG-14

T-1.4.21	Error injection during key or signature generation	Violates
AC2 (associated with shared keys), AC5	In this threat, an attacker is able to inject an error in the key or signature generation process that leaks part or all of the private key. Consequences: May undermine [SG-1], [SG-4], [SG-11], [SG-14]. Mitigations: This is mitigated by [SM-18] and [SM-28].	SG-1, SG-4, SG-11, SG-14

T-1.4.22	Birthday Paradox Collision	Violates
AC3, AC6	In this threat, a set of randomly generated parameters collide. The probability of this occurrence can be bounded using analysis similar to that associated with the classical Birthday Paradox. Consequences: May undermine [SG-1], [SG-11], [SG-14]. Mitigations: Establishing a bounded number of allowable outputs based on the size of the randomly generated value [SM-19].	SG-1, SG-11, SG-14

T-1.4.23	Privacy Reduction	Violates
AC1	In this threat, a small number of Authenticators share an attestation key which leaks information about the user across Relying Parties. Consequences: May undermine [SG-4]. Mitigations: This is mitigated by [SM-20].	SG-4

T-1.4.24	Covert Channel	Violates
AC1	In this threat, an Authenticator is malicious (either by design, or after having been independently compromised) and it is configured to leak secret or identifying data within apparently normal exchanges, or to other processes on the same hardware platform as the Authenticator. Consequences: May undermine [SG-1], [SG-4], [SG-5], [SG-6], [SG-8], [SG-11], [SG-14]. Mitigations: Note: This is an interesting thought experiment; use of random nonces and other non-deterministic elements make protection against this threat problematic.	SG-1, SG-4, SG-5, SG-6, SG-8, SG-11, SG-14

T-1.4.25	Substitution of Protected Information	Violates
	In this threat, an attacker substitutes protected information, either by modifying it piecemeal, or by completely substituting it with another value. (Some encryption modes allow an attacker to target bit-level changes to the plaintext. Authenticated	

AC1, AC3, AC5, AC6	data may also have been replaced with data that had previously been authenticated in the same way.) Substitution of Protected Information	SG-1 SG-11
	Consequences: May undermine [SG-1], [SG-4], [SG-11], [SG-14]. Mitigations: This threat is mitigated by [SM-1], [SM-16], [SM-21].	SG-11, SG-14

T-1.4.26	Compromise of Protected Information	Violates
AC1, AC2, AC5, AC6	In this threat, an attacker recovers data that should be protected by the Authenticator. Consequences: May undermine [SG-1], [SG-2], [SG-4], [SG-5], [SG-7], [SG-8], [SG-11], [SG-14]. Mitigations: This threat is mitigated by using allowed cryptographic primitives [SM-1], [SM-16].	SG-1, SG-2, SG-4, SG-5, SG-7, SG-8, SG-11, SG-14

T-1.4.27	Signature or registration counter non-monotonicity	Violates
AC1	In this threat, an attacker may be able to cause these counters to be reset, to roll over, or otherwise to decrease in value. Consequences: May undermine [SG-1], [SG-12], [SG-14]. Mitigations: This threat is mitigated by [SM-15].	SG-1, SG-12, SG-14

T-1.4.28	Hostile ASM / Client	Violates
AC3, AC5, AC6	In this threat, the Authenticator support infrastructure is hostile, and can feed arbitrary data to the Authenticator. Consequences: May undermine [SG-4], [SG-5], [SG-7], [SG-8]. Mitigations: This threat is mitigated by [SM-10], [SM-13].	SG-4, SG-5, SG-7, SG-8

T-1.4.29	Debug Interface	Violates
AC2 (associated with shared keys), AC3 (associated with shared keys), AC5, AC6	In this threat, the Authenticator has a hardware or software debugging interface that is not completely disabled prior to distribution of the Authenticator (e.g., pads for a JTAG port). Consequences: May undermine [SG-1], [SG-4], [SG-5], [SG-6], [SG-8], [SG-11], [SG-14]. Mitigations: This threat is mitigated by [SM-18], [SM-22], and [SM-28].	SG-1, SG-4, SG-5, SG-6, SG-8, SG-11, SG-14

T-1.4.30	Fault induced by malformed input	Violates
AC2, AC3, AC5, AC6	In this threat, the Authenticator behaves in an unexpected fashion due to an error in processing malformed input. The result of this style of attack is poorly controllable, absent strong internal segmentation of the Authenticator. Consequences: May undermine [SG-1], [SG-2], [SG-3], [SG-4], [SG-6], [SG-7], [SG-8], [SG-11], [SG-14], [SG-16]. Mitigations: This threat is mitigated by [SM-1], [SM-2], [SM-4], [SM-5], [SM-10], [SM-5], [SM-23], [SM-13], [SM-26].	SG-1, SG-2, SG-3, SG-4, SG-6, SG-7, SG-8, SG-11, SG-14, SG-16

T-1.4.31	Fault Injection Attack	Violates
AC2 (associated with shared keys), AC5, AC6	In this threat, an attacker subjects the Authenticator to conditions that induce hardware faults (e.g., exposure to photons or charged particles, inducing variations in supply voltage or external clock, altering the temperature, etc.) in an attempt to subvert some logical or physical protection. The result of this style of attack is poorly controllable, absent active detection and response functionality within the Authenticator. This is related to T-1.4.21, but applies more broadly. Consequences: May undermine [SG-1], [SG-2], [SG-3], [SG-4], [SG-6], [SG-7], [SG-8], [SG-11], [SG-14], [SG-16]. Mitigations: Mitigated by [SM-1], [SM-2], [SM-4], [SM-5], [SM-10], [SM-5], [SM-18], [SM-23], [SM-13], [SM-26], [SM-28].	SG-1, SG-2, SG-3, SG-4, SG-6, SG-7, SG-8, SG-11, SG-14, SG-16

T-1.4.32	Remote Timing Attacks	Violates
AC2, AC5	In this threat, an attacker may be able to extract keys using a timing attack from a remote location. The impact depends on the key or secret recovered. Consequences: May undermine [SG-1], [SG-2], [SG-4], [SG-11], [SG-14]. Mitigations: This threat is mitigated by the remote timing attack resistance security measure [SM-29].	SG-1, SG-2, SG-4, SG-11, SG-14

7.1.5 Threats to Relying Party

7.1.5.1 Threats to FIDO Server Data

T-2.1.1 FIDO Server DB Read Attack		Violates
	<p>Attacker could obtains read-access to FIDO Server registration database.</p> <p>Consequences:Attacker can access all cryptographic key handles and authenticator characteristics associated with a username. If an authenticator or combination of authenticators is unique, they might use this to try to violate [SG-2] Unlinkability.</p> <p>Attacker attempts to perform factorization of public keys by virtue of having access to a large corpus of data, violating [SG-5] Verifier Leak Resilience and [SG-2] Credential Guessing Resilience.</p> <p>Mitigations: [SM-2] Unique Authentication Keys help prevent disclosed key material from being useful against any other Relying Party, even if successfully attacked.</p> <p>The use of an [SM-6] Cryptographically Secure Verifier Database helps assure that it is infeasible to attack any leaked verifier keys.</p> <p>[SM-9] Authenticator Certification along with [SM-16] should help prevent authenticators with poor entropy from entering the market, reducing the likelihood that even a large corpus of key material will be useful in mounting attacks.</p>	SG-2, SG-5

T-2.1.2 FIDO Server DB Modification Attack		Violates
	<p>Attacker gains write-access to the FIDO Server registration database.</p> <p>Consequences: Violation of [SA-6]</p>	
AC1	<p>The attacker may inject a key registration under its control, violating [SG-1] Strong User Authentication.</p> <p>Mitigations: Mitigating such attacks is outside the scope of the FIDO specifications. The Relying Party must maintain the integrity of any information it relies up on to identify a user as part of [SA-6].</p>	SA-6

T-2.2.1 Web App Malware		Violates
	<p>Attacker gains ability to execute code in the security context of the Relying Party web application or FIDO Server.</p> <p>Consequences: Attacker is able to violate [SG-1], [SG-10], [SG-9] and any other Relying Party controls.</p> <p>Mitigations: The consequences of such an incident are limited to the relationship between the user and that particular Relying Party by [SM-1], [SM-2], and [SM-5].</p> <p>Even within the Relying Party to user relationship, a user can be protected by [SM-10] Transaction Confirmation if the compromise does not include the users' computing environment.</p>	SG-1, SG-9, SG-10

T-2.2.2 Linking through compromised Relying Party database		Violates
	<p>In this threat, a Relying Party is able to access another Relying Party's database (either because the Relying Parties are collaborating or because of the compromise of another Relying Party's database). The malicious party then sends Key Handles (which may contain a wrapped private key) from the other Relying Party's database in an attempt to link the two separate accounts to the same Authenticator (thus user).</p> <p>Consequences: May undermine [SG-1], [SG-4].</p> <p>Mitigations: This threat is mitigated by [SM-1], [SM-2], [SM-5], [SM-23].</p>	SG-1, SG-4

7.1.6 Threats to the Secure Channel between Client and Relying Party

7.1.6.1 Exploiting Weaknesses in the Secure Transport of FIDO Messages

FIDO takes as a base assumption that [SA-3] applications on the user device are able to establish secure channels that provide trustworthy server authentication, and confidentiality and integrity for messages. e.g. through TLS. [T-1.2.4] Discusses some consequences of violations of this assumption due to implementation errors in a browser or client application, but other threats exist in different layers.

T-3.1.1 TLS Proxy		Violates
	<p>The FIDO user device is administratively configured to connect through a proxy that terminates TLS connections. The client trusts this device, but the connection between the user and FIDO server is no longer end-to-end secure.</p> <p>Consequences: Any such proxies introduce a new party into the protocol. If this party is untrustworthy, consequences may be as for [T-1.2.4].</p>	
AC3	<p>Mitigations: Mitigations for [T-1.2.4] apply, except that the proxy is considered trusted by the client, so certain methods of [SM-12] Channel Binding may indicate a compromised channel even in the absence of an attack. Servers should use multiple methods and adjust their risk scoring appropriately. A trustworthy client that reports a server certificate that is unknown to the server and does not chain to a public root may indicate a client behind such a proxy. A client reporting a server certificate that is unknown to the server but validates for the server's identity according to commonly used public trust roots is more likely to indicate [T-3.1.2].</p>	SG-11, SG-12, SG-13

T-3.1.2 Fraudulent TLS Server Certificate		Violates
	<p>An attacker is able to obtain control of a certificate credential for a Relying Party, perhaps from a compromised Certification Authority or poor protection practices by the Relying Party.</p>	
AC3	<p>Consequences:As for [T-1.2.4].</p>	SG-11, SG-12, SG-13

T-3.1.2	Mitigations: As for [T-1.2.4].	Fraudulent TLS Server Certificate	Violates
----------------	---------------------------------------	--	-----------------

T-3.1.3	Protocol level real-time MITM attack		Violates
AC3	<p>An adversary can intercept and manipulate network packets sent from the relying party to the client. The adversary uses this capability to (a) terminate the underlying TLS session from the client at the adversary and to (b) simultaneously use another TLS session from the adversary to the relying party. In the traditional username/password world, this allows the adversary to intercept the username and the password and then successfully impersonate the user at the relying party.</p> <p>Consequences: None if FIDO channelBinding [SM-12] or transaction confirmation [SM-10] are used.</p> <p>Mitigations: In the case of channelBinding [SM-12], the FIDO server will detect the MITM in the TLS channel by comparing the channel binding information provided by the client and the channel binding information retrieved locally by the server.</p> <p>In the case of transaction confirmation [SM-10], the user verifies and approves a particular transaction. The adversary could modify the transaction before approval. This would lead to rejection by the user. Alternatively, the adversary could modify the transaction after approval. This will break the signature in the transaction confirmation response. The FIDO Server will not accept it as a consequence.</p> <p>HTTP Public Key Pinning (RFC7469) can also be used to mitigate this attack (outside the FIDO stack).</p>		SG-11, SG-12, SG-13

7.1.7 Threats to the Infrastructure

7.1.7.1 Threats to FIDO Authenticator Manufacturers

T-4.1.1	Manufacturer Level Attestation Key Compromise		Violates
AC2	<p>Attacker obtains control of an attestation key or attestation key issuing key.</p> <p>Consequences: Same as [T-1.4.6]: Attacker can violate [SG-9] Attestable Properties by creating a malicious hardware or software device that represents itself as a legitimate one.</p> <p>Mitigations: Same as [T-1.4.6]: Relying Parties can use [SM-4] Authenticator Status Checking to identify known-compromised keys. Identification of such compromise is outside the strict scope of the FIDO protocols.</p>		SG-9

T-4.1.2	Malicious Authenticator HW		Violates
AC1, AC2, AC3, AC5, AC6	<p>FIDO Authenticator manufacturer relies on hardware or software components that generate weak cryptographic authentication key material or contain backdoors.</p> <p>Consequences: Effective violation of [SA-1] in the context of such an Authenticator.</p> <p>Mitigations: The process of [SM-9] Authenticator Certification may reveal a subset of such threats, but it is not possible that all such can be revealed with black box testing and white box examination may be economically infeasible. Users and Relying Parties with special concerns about this class of threat must exercise their own necessary caution about the trustworthiness and verifiability of their vendors and supply chain. [SM-24] builds confidence that an Authenticator is not malicious or poorly implemented.</p>		SA-1

7.1.7.2 Threats to FIDO Server Vendors

T-4.2.1	Vendor Level Trust Anchor Injection Attack		Violates
	<p>Attacker adds malicious trust anchors to the trust list shipped by a FIDO Server vendor.</p> <p>Consequences: Attacker can deploy fake Authenticators which Relying Parties cannot detect as such, which do not implement any appropriate security measures, and is able to violate all security goals of FIDO.</p> <p>Mitigations: This type of supply chain threat is outside the strict scope of the FIDO protocols and violates [SA-6]. Relying Parties can verify their trust list against the data published by the FIDO Alliance Metadata Service [FIDOMetadataService] (see https://fidoalliance.org/mds).</p>		SA-6

7.1.7.3 Threats to FIDO Metadata Service Operators

T-4.3.1	Metadata Service Signing Key Compromise		Violates
	<p>The attacker gets access to the private Metadata TOC signing key.</p> <p>Consequences: The attacker could sign invalid Metadata. The attacker could</p> <ul style="list-style-type: none"> • make trustworthy authenticators look less trustworthy (e.g. by increasing FAR). • make weak authenticators look strong (e.g. by changing the key protection method to a more secure one) • inject malicious attestation trust anchors, e.g. root certificates which cross-signed the original attestation trust anchor and the cross-signed original attestation root certificate. This malicious trust anchors could be used to sign attestation certificates for fraudulent authenticators, e.g. authenticators using the AAID of trustworthy authenticators but not protecting their keys as stated in the metadata. <p>Mitigations: The Metadata Service operator should protect the Metadata signing key appropriately, e.g. using a hardware protected key storage.</p> <p>Relying parties could use out-of-band methods to cross-check Metadata Statements with the respective vendors and cross-check the revocation state of the Metadata signing key with the provider of the Metadata Service.</p>		SG-9

T-4.3.1	Metadata Service Signing Key Compromise	Violates
T-4.3.2	Metadata Statement Data Injection	Violates
	<p>An attacker injects malicious Authenticator data into the Metadata Statement.</p> <p>Consequences: The attacker could make the Metadata Service operator sign invalid Metadata Statements. The attacker could</p> <ul style="list-style-type: none"> • make trustworthy authenticators look less trustworthy (e.g. by increasing FAR). • make weak authenticators look strong (e.g. by changing the key protection method to a more secure one) • inject malicious attestation trust anchors, e.g. root certificates which cross-signed the original attestation trust anchor and the cross-signed original attestation root certificate. This malicious trust anchors could be used to sign attestation certificates for fraudulent authenticators, e.g. authenticators using the AAID of trustworthy authenticators but not protecting their keys as stated in the metadata. <p>Mitigations: The Metadata Service operator could carefully review the delta between the old and the new Metadata Statements. Authenticator vendors could verify the published Metadata Statements related to their Authenticators.</p>	<u>SG-9</u>

7.1.8 Threats Specific to Second Factor Authenticators (UAF / U2F)

T-5.1.1	Error Status Side Channel	Violates
	<p>Relying parties issues an authentication challenge to an authenticator and can infer from error status if it is already registered.</p> <p>Consequences: UAF Silent authenticators / U2F authenticators not requiring user interaction for generating a signed response may be used to track users without their consent by issuing a pre-authentication challenge to them, revealing the identity of an otherwise anonymous user. Users would be identifiable by relying parties without their knowledge, violating [SG-7].</p> <p>Mitigations: The U2F specification recommends that browsers prompt users whether to allow this operation using mechanisms similar to those defined for other privacy sensitive operations like Geolocation.</p>	<u>SG-7</u>

T-5.1.2	Malicious RP	Violates
<u>AC1</u>	<p>Malicious relying party mounts a cryptographic attack on a key handle it is storing.</p> <p>Consequences: If the Relying Party is able to recover the contents of the key handle, it might forge logs of protocol exchanges to associate the user with actions he or she did not perform.</p> <p>If the Relying Party is able to recover the key used to wrap a key handle, that key is likely used for all key handles, and hence might be used to decrypt key handles stored with other Relying Parties and violate [SG-1] Strong User Authentication.</p> <p>Mitigations: None. U2F depends on [SA-1] to hold for key wrapping operations.</p>	<u>SG-1</u>

T-5.1.3	Physical Attack on a User Presence Authenticator	Violates
<u>AC5</u>	<p>Attacker gains physical access to U2F authenticator or a UAF authenticator with only user presence check (e.g., by stealing it).</p> <p>Consequences: Same as for [T-1.4.4].</p> <p>Such authenticators have weak local user verification. If the attacker can guess the username and password/PIN, they can impersonate the user, violating [SG-1] Strong User Authentication.</p> <p>Mitigations: Relying Parties can use strong additional factors.</p> <p>Relying Parties should provide users a means to revoke keys associated with a lost device.</p>	<u>SG-1</u>

T-5.1.4	Physical Attack	Violates
<u>AC2</u> (associated with shared keys), <u>AC5</u>	<p>In this threat, keys or other sensitive information is read out by directly accessing it from the authenticator that the attacker has physically compromised.</p> <p>Consequences: May undermine [SG-1], [SG-4], [SG-11], [SG-14].</p> <p>Authenticator with user presence check have weak local user verification. If the attacker can guess the username and password/PIN, they can impersonate the user, violating [SG-1] Strong User Authentication.</p> <p>Mitigations: Mitigated by resistance to injected faults [SM-18] and [SM-28].</p>	<u>SG-1</u> , <u>SG-4</u> , <u>SG-11</u> , <u>SG-14</u>

7.2 Acknowledgements

We thank [iSECpartners](#) for their review of, and contributions to, this document.

A. References

A.1 Informative references

[FIDOEcdaaAlgorithm]

R. Lindemann; J. Camenisch; M. Drijvers; A. Edgington; A. Lehmann; R. Urian. [FIDO ECDA Algorithm](#). Implementation Draft. URL:

<https://fidoalliance.org/specs/fido-uaf-v1.2-id-20180220/fido-ecdaa-algorithm-v1.2-id-20180220.html>

[FIDOGlossary]

R. Lindemann; D. Baghdasaryan; B. Hill; J. Hodges. *FIDO Technical Glossary*. Implementation Draft. URL: <https://fidoalliance.org/specs/fido-uaf-v1.2-id-20180220/fido-glossary-v1.2-id-20180220.html>

[FIDOMetadataService]

R. Lindemann; B. Hill; D. Baghdasaryan. *FIDO Metadata Service v1.0*. Implementation Draft. URL: <https://fidoalliance.org/specs/fido-uaf-v1.2-id-20180220/fido-metadata-service-v1.2-id-20180220.html>

[PasswordAuthSchemesKeyIssues]

Chwei-Shyong Tsai; Cheng-Chi Lee; Min-Shiang Hwang. *Password Authentication Schemes: Current Status and Key Issues* September 2006. URL: <http://ijns.femto.com.tw/contents/ijns-v3-n2/ijns-2006-v3-n2-p101-115.pdf>

[QuestToReplacePasswords]

Joseph Bonneau; Cormac Herley; Paul C. van Oorschot; Frank Stajano. *The Quest to Replace Passwords: A Framework for Comparative Evaluation of Web Authentication Schemes*. March 2012. URL: <http://research.microsoft.com/pubs/161585/QuestToReplacePasswords.pdf>

[RFC2119]

S. Bradner. *Key words for use in RFCs to Indicate Requirement Levels* March 1997. Best Current Practice. URL: <https://tools.ietf.org/html/rfc2119>

[U2FOverview]

S. Srinivas; D. Balfanz; E. Tiffany. *FIDO U2F Overview v1.0*. Draft. URL: <http://fidoalliance.org/specs/fido-u2f-overview-v1.0-rd-20140209.pdf>

[UAFProtocol]

R. Lindemann; D. Baghdasaryan; E. Tiffany; D. Balfanz; B. Hill; J. Hodges. *FIDO UAF Protocol Specification v1.0*. Proposed Standard. URL: <https://fidoalliance.org/specs/fido-uaf-v1.2-id-20180220/fido-uaf-protocol-v1.2-id-20180220.html>