THREAT MODEL

version 8.4.1

INTERNET IDENTITY CARD ™

Security analysis, trust assumptions & residual risks

Issuer	Https Card — Internet Identity Card Ltd
Company Registration	UK Companies House No. 09168431
Trademark	UK IPO UK00003166480 (2016)
Office	124 City Road, London EC1V 2NX, United Kingdom
Disclosures	TDCommons #10079 · #10167 (CC BY 4.0)

This document describes security properties and limitations. It is not a guarantee of security.

1. Scope & objectives

This document describes the threat model of the Internet Identity Card (IIC) v8.4.1: the assets it protects, the adversaries it is designed to resist, the trust assumptions it depends on, and — importantly — the attacks that are explicitly out of scope. It is written for security auditors, integrators, and reviewers performing due diligence.

An IIC card is a single, self-contained HTML file that encrypts identity data on the issuer's device and is later opened, offline, by a recipient who holds the correct passphrase. There is no server, no account, and no key-exchange protocol. This architecture removes entire classes of risk (there is no backend to breach) but concentrates trust in the endpoints and in the integrity of the file itself. This document is explicit about that trade-off.

A threat model is not a security guarantee. It states what the system is designed to do under stated assumptions. Where those assumptions do not hold (for example, a compromised device), the corresponding protections do not apply, and this is documented in Section 8.

2. System overview

The lifecycle of an IIC card involves three phases and two human roles.

Issuer	The person who fills in identity fields and exports the card. Holds the Memorable Word (long-term issuer secret).
Recipient	The person who later opens the card offline. Receives the Card Passphrase through a separate channel chosen by the issuer.
Generation	Runs entirely in the issuer's browser. Keys are derived locally with Argon2id (cards) and PBKDF2 (backups); data is encrypted with AES-256-GCM.
Distribution	The exported HTML file is transferred by any means (email, USB, print-to-file). It carries its own SHA-256 self-integrity check.
Opening	The recipient opens the file in any modern browser, offline. Decryption happens locally after the correct passphrase is entered.

Because no key material ever leaves the device and no network call is required to open a card, the attack surface is the file at rest, the file in transit, and the two browsers involved.

3. Trust assumptions

The security properties in Section 6 and 7 hold only while these assumptions are true. They are stated plainly so a reviewer can decide whether they fit a given deployment.

Assumption	Why it matters
Trusted generation device	At export time, the issuer's device is free of malware, keyloggers, and screen capture. Identity data and passphrases are in cleartext in memory during this step.
Trusted opening device	At open time, the recipient's device is similarly clean. Decrypted identity is rendered in the browser and is therefore visible to anything running on that device.

Assumption	Why it matters
Authentic generator	The recipient implicitly trusts that the file was produced by an authentic IIC generator. A tampered generator could embed a backdoor or weak key derivation (see §8).
Correct browser crypto	The W3C Web Crypto API and the WASM Argon2id implementation behave per specification. IIC does not re-implement primitives; it relies on the platform.
Out-of-band passphrase	The Card Passphrase is delivered to the recipient over a channel separate from the file itself. Sending both together collapses the two-factor separation.
Passphrase secrecy	The Memorable Word and Card Passphrase are kept secret by their holders and are not reused across unrelated systems.

4. Assets under protection

Asset	Description	Priority
Identity plaintext	The 27 identity fields once decrypted — the primary secret.	**PROTECTED**
Card Passphrase	Recipient-side secret that unlocks the card. ~95 bits of entropy when auto-generated.	**PROTECTED**
Memorable Word	Issuer-side long-term secret used for the backup domain.	**PROTECTED**
File integrity	Assurance that the file has not been altered since export.	**PROTECTED**
Issuer attribution	Assurance about who produced the card (ECDSA, where signed).	**MITIGATED**

Note that **issuer attribution** is marked as partially assured: SHA-256 self-integrity proves the file is internally consistent, but it does not prove *who* created it. Only an ECDSA signature verified against a public key obtained from a trusted source establishes authorship.

5. Adversaries & capabilities

We model four adversaries of increasing capability. The first three are within scope; the fourth is largely out of scope and is the honest boundary of the system.

Adversary	Capability & goal
A1 — Network observer	Can intercept the file in transit (email relay, shared network). Goal: read the identity. Cannot run code on either endpoint.
A2 — File holder	Has obtained a copy of the exported file but not the passphrase. Goal: brute-force or tamper. Has offline compute (GPU/ASIC).
A3 — Malicious recipient	A legitimate recipient who tries to forge a modified card, replay it, or pass it off as issued by someone else.
A4 — Endpoint attacker	Has code execution or physical access on the issuer's or recipient's device (malware, keylogger, coercion). Largely OUT OF SCOPE — see §8.

6. STRIDE analysis

Mapping the six STRIDE categories onto the IIC architecture.

STRIDE category	Applied to IIC	Disposition
Spoofing	ECDSA P-256 signing (where enabled) binds a card to an issuer key; without it, authorship is not proven.	MITIGATED
Tampering	SHA-256 page self-integrity with fail-closed lockdown detects any byte change outside the hash zone.	PROTECTED
Repudiation	No central log exists; blockchain timestamping (BTC/ETH) provides external, tamper-evident proof of existence at a date.	MITIGATED
Information disclosure	AES-256-GCM with Argon2id-derived keys protects identity at rest and in transit against A1 and A2.	PROTECTED
Denial of service	A single offline file has no service to deny; loss of the file or passphrase means loss of access, by design.	PROTECTED
Elevation of privilege	Strict CSP and Permissions-Policy disable camera, microphone, geolocation, payment and USB in the exported card.	PROTECTED

7. Concrete attack scenarios

Scenario	Outcome	Status
Passive eavesdropping (A1)	Observer captures the file but not the Card Passphrase. AES-256-GCM keeps identity unreadable.	PROTECTED
Offline brute force (A2)	Argon2id at 96 MiB / t=4 / p=4 makes large-scale GPU/ASIC guessing economically prohibitive against ~95-bit passphrases.	PROTECTED
Post-export tampering (A2)	Any modification outside the hash zone is detected at load; the card refuses to open (fail-closed).	PROTECTED
Hash recompute & reinject (A2/A3)	A capable attacker can recompute the self-hash after editing. SHA-256 integrity is not a substitute for ECDSA signing; attribution requires the signature.	MITIGATED
Replay / parameter confusion (A2/A3)	AES-256-GCM authenticated additional data (AAD) binds each ciphertext to its declared cipher and key-derivation parameters, so a blob cannot be silently decrypted under different parameters; the GCM tag rejects any mismatch.	PROTECTED
localStorage residue (A4-lite)	Quick-mode data may persist in browser storage. Cleared at first run; this resists incidental dumps but not an attacker with full device access.	MITIGATED
Weak user passphrase	If the issuer overrides the auto-generator with a low-entropy phrase, Argon2id slows but does not save a guessable secret.	MITIGATED
Compromised device (A4)	Malware/keylogger reads identity and passphrases directly from memory or input. No client-side scheme can prevent this.	OUT OF SCOPE

Scenario	Outcome	Status
Tampered generator	A backdoored generator could weaken keys or exfiltrate at export. Recipients must trust the generator's provenance; this is why the generator is access-controlled, not publicly distributed.	**OUT OF SCOPE**
Physical coercion	Forced disclosure of the Memorable Word or Card Passphrase defeats any cryptography.	**OUT OF SCOPE**

8. Residual risks & out of scope

A credible threat model names what it does *not* defend against. The following are explicitly out of scope for IIC v8.4.1, by design.

- **Compromised endpoints.** Malware, keyloggers, or screen capture on either device read plaintext before or after encryption. IIC is a client-side scheme and inherits the trust level of the host.
- **Generator authenticity.** If a recipient is given a tampered generator, all guarantees collapse. The generator is therefore distributed only through a controlled onboarding process and is not published.
- **Browser / platform vulnerabilities.** Bugs in the Web Crypto API, the WASM runtime, or the browser sandbox are outside IIC's control.
- **User-chosen weak secrets.** Overriding the auto-generated high-entropy passphrase with a guessable phrase undermines the key-derivation hardening.
- **Coercion & social engineering.** Forcing or tricking a holder into revealing a secret is a human-layer attack outside cryptographic scope.
- **Metadata & traffic analysis.** The existence, size, and timing of a transferred file may leak information even when its contents do not.
- **Quantum adversaries.** AES-256 retains ~128-bit strength under Grover; ECDSA P-256 would be broken by a large fault-tolerant quantum computer. No post-quantum migration is claimed in v8.4.1.

9. Operational recommendations

For issuers

- Generate cards on a trusted, malware-free device.
- Keep the auto-generated Card Passphrase; do not substitute a weak one.
- Deliver the Card Passphrase out-of-band, never alongside the file.
- Enable ECDSA signing when authorship needs to be provable.

For recipients

- Open cards on a trusted device; treat the decrypted view as sensitive.
- Obtain the generator/public key only from the official, controlled channel.
- Verify the file SHA-256 (and signature, where present) before trusting it.

For auditors

- Validate that the trust assumptions in §3 hold in the target deployment.
- Cross-reference the cryptographic parameters against the Technical Specification and the defensive publications (TDCommons #10079, #10167).

10. Disclaimer

This threat model is provided for informational purposes only. It describes the intended security properties, design assumptions, and known limitations of Internet Identity Card v8.4.1 as understood at the time of writing. It does **not** constitute a warranty, guarantee, certification, accreditation, or legal assurance of security, fitness for a particular purpose, or regulatory compliance.

Security depends on factors outside the software's control — including the integrity of the devices used, the secrecy and strength of user-chosen secrets, the authenticity of the generator, and the correctness of the underlying browser and operating system. Where the trust assumptions in Section 3 do not hold, the corresponding protections do not apply.

No software is immune to all attacks, and no threat model is exhaustive. New attack techniques, implementation defects, or platform vulnerabilities may emerge after publication. This document may be revised; users should refer to the most recent version. Internet Identity Card ™ is a private software-based identity tool developed by Https Card — Internet Identity Card Ltd. References to standards (NIST, FIPS, RFC, W3C) indicate the primitives used and do not imply endorsement or validation by those bodies.

To the maximum extent permitted by applicable law, Https Card — Internet Identity Card Ltd accepts no liability for any loss or damage arising from the use of, or reliance upon, this document or the software it describes. Use of the Internet Identity Card is at the user's own risk.