iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

THERMO-MECHANICALLY BOUND STORAGE KEYS FOR EDGE COMPUTE APPLIANCES:AN ARCHITECTURE AND ITS ANALYSIS

Vamshi Teja Sara
Avalon Artificial Intelligence Limited, Middlesbrough, United Kingdom
vamsi@unitedavalon.com

Author Biography / Patent Declaration:

∗The architecture described here is the subject of UK patent application **GB2622534.2**, filed 25 September 2026, and the enclosure appearances of UK registered design applications **6555931–6555939,** filed the same day. This paper is published after those filing dates and discloses no subject matter beyond them.

ABSTRACT

An edge compute appliance deployed outside physically controlled premises may be opened, have its storage removed, or be carried away intact. Tamper-responsive enclosures and physical unclonable functions address the first two; none of them addresses the third, because an appliance stolen whole still carries every factor it needs to satisfy itself.

We describe an architecture in which the key that decrypts persistent storage cannot be reconstructed unless three independent factors are simultaneously present: a device root key held in a secure element, a physical-binding secret reconciled from the assembly itself, and a key share held outside the appliance. The physical binding is obtained without a dedicated stimulus source. The appliance modulates its own compute workload at a frequency in the millihertz range, which heats a member that is simultaneously the principal heat-transfer path and a load-bearing structural element; a separate mechanical excitation in the 200 Hz–5000 Hz band is applied, and it is the *envelope* of the mechanical response, estimated synchronously at the thermal modulation frequency, that is examined. The quantity enrolled is an internal *differential* between two protected sensors on opposite sides of the thermal path, so that what is reproduced depends on the geometry and joints of the assembly rather than on any absolute surface measurement.

We show that at these frequencies the member is not thermally lumped and must be treated as carrying a damped thermal wave, and we derive the consequences for sensor placement. We give an entropy accounting that distinguishes coded pairs from decoded information bits, and show why the physical-binding secret is a binding factor rather than a key. We argue that no session-varying quantity may enter the key derivation and show what breaks when one does.

A 3D-printed prototype has been built and bench-tested. Section 14 reports measured results from seven test campaigns (T1–T7) and five engineering gates (G1–G5). Most acceptance criteria are met; specific items remain marginal or open — notably mode-3 actuator authority, carrier-board edge pressure (gate G3), and airflow leakage — and are stated explicitly. Section 17 states what remains open after that testing.

1 INTRODUCTION

An appliance that performs computation at the edge of a network is often deployed where its owner cannot watch it: a cabinet in a retail unit, a plant room, a vehicle, a cell site. An attacker with unsupervised physical access for an evening has a wide range of options. They may open the enclosure and interpose on internal buses. They may remove the storage device and read it elsewhere. They may substitute a controller and replay recorded traffic. Or they may simply carry the whole appliance away and work on it at leisure.

Three families of protection are established. Tamper-responsive enclosures carry a mesh, coating or potting whose electrical or optical property is disturbed by intrusion. Physical unclonable functions measure a manufacturing-random physical property and reconcile it against enrolled helper data to reproduce a secret. Platform attestation reports firmware measurements to a trusted module.

Each addresses part of the problem, and all three share a structural limitation: they verify that something is *intact*. An appliance that is stolen intact is intact. It passes its own verification, and if passing that verification is what releases the storage key, then theft of the whole appliance defeats the scheme entirely. Physical integrity is necessary and it is not sufficient.

1.1 Approach

We describe an architecture in which the key-encryption key is derived from three factors of independent origin, all three of which must be present simultaneously:

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

- a *device root key*, generated inside a secure element and never exported, so that a copied storage device is useless;
- a *physical-binding secret*, reconciled at each verification from a measurement of the assembly itself, so that an opened or substituted assembly is useless;
- an *external key share*, held outside the appliance and delivered under authorisation, so that the appliance taken as a whole is useless.

Removing any one of the three prevents formation of the key-encryption key, and therefore prevents unwrapping the stored data-encryption key. The third factor is what distinguishes this from integrity verification: a thief who takes the appliance does not take the share.

The physical binding is obtained without adding a stimulus source. Two observations make that possible. First, an appliance that computes already dissipates heat, and the schedule of its own workload is under its own control, so the thermal stimulus can be *generated by the work the machine is there to do*. Second, the member that carries that heat away can be the same member that carries structural load, so that a change to the mechanical assembly changes the thermal path and vice versa. The appliance modulates its compute workload at a frequency in the millihertz range, applies a separate mechanical excitation in a much higher band, and examines how the envelope of the mechanical response varies at the thermal modulation frequency.

What is enrolled is not an absolute measurement but an internal *differential* between two protected sensors placed on opposite sides of the thermal path. Because the two sensors sample a travelling thermal wave at different distances, they see it at different phases, and the difference between them is a measure of the thermal gradient and its phase across the member — a property of geometry, joints and material rather than of any single temperature.

1.2 Contributions

1) A three-origin key hierarchy in which physical integrity is necessary but never sufficient, so that theft of an intact appliance does not yield the key (§10).
2) The distinction between *authorising release of* a stable secret and *being* that secret, an argument that no session-varying quantity may enter a derivation whose output must open a persistent wrapping, and an account of what breaks when one does — a failure mode we introduced into our own design and had to remove (§10.2).
3) A distributed treatment of the thermal challenge. At millihertz modulation the heat-transfer member is not thermally lumped; we give the penetration-depth analysis and derive a sensor placement rule keyed to the estimator's detection floor rather than to a fixed distance (§6).
4) An effect-size chain from watts to percentage envelope modulation, so that the quantity being measured can be judged against the noise before anything is built (§6).
5) A selection-and-rejection procedure that a team can run to decide whether a given assembly supports the architecture, including four null controls — one of which is required because a confounder passes the other three (§8).
6) An entropy accounting that separates measured features, quantisation pairs, retained reliable pairs, codeword length, code rate and decoded information bits, and that declines to assert a min-entropy figure for the physical response (§9).
7) A decoy-window construction with four matched physical observables, and a declared validation protocol for the leakage bound it assumes, designed so that the test cannot flatter the design (§6.6, Appendix E).
8) Formal independence criteria for the second sensing domain, without which the cross-witness residual has no force (§5.2).
9) An enclosure realisation analysed for modal behaviour, thermal path, airflow penalty, geometric fit and manufacturability, the fit stated as conditional on a manufacturer figure whose scope is not published (§13).

1.3 Organisation

§2 sets out the adversary classes. §3 places the work against the closest prior art and states what is acknowledged as known. §4 gives the system model, assets, goals and trust assumptions, and §5 the architecture. §§6–7 describe the physical binding and its estimation; §8 the procedure for deciding whether an assembly supports it; §9 the reconciliation and what it yields; §10 the key hierarchy and dual control; §11 resistance to external observation; and §12 service, transport and peripheral admission. §13 gives the enclosure and its analysis, and §15 the security argument. §14 reports the bench validation of the prototype, and §17 states what remains open after it.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

A 3D-printed prototype of this architecture was built and bench-tested. The measured results — thermal signal margin, tamper separability, external-leakage suppression, reseating repeatability and manufacturing tolerance — are reported in Section 14, and they establish the central physical premise of the design.

2 THREAT MODEL

2.1 Adversary classes

We distinguish six adversary classes by capability rather than by motive, because the architecture responds to capabilities

Adversary classes and the capabilities assumed of each. • denotes a capability the class has.

	Capability	A1	A2	A3	A4	A5
C1	Obtain the storage device	•	•	•	•	•
C2	Obtain the whole appliance, intact		•	•	•	•
C3	Open the enclosure; probe internal buses			•	•	•
C4	Substitute the security controller; replay recorded traffic				•	•
C5	Observe externally over many sessions with laboratory instruments					•
C6	Extract the root key from the secure element					
C7	Compromise the authority holding the external share					

Opportunist. Takes the storage device. Reads it elsewhere. No access to the appliance afterwards.

Thief. Removes the complete appliance and works on it at leisure, without opening it. This is the class that defeats integrity-only schemes, and the class the external factor exists for.

Invasive attacker. Opens the enclosure, probes internal buses, replaces components, attaches instruments to internal structures.

Interposer. Substitutes the security controller or interposes on the leads between the protected internal sensors and the controller, and replays a previously recorded genuine response.

Patient observer. Attaches an accelerometer, a microphone, a current probe and a thermal camera to the outer shell and records many sessions, with the aim of learning which windows are genuine and what the internal quantities are.

A6 would be an adversary with C6 or C7 — extraction from the secure element, or compromise of the authority. We do not defend against it and we say so rather than quietly assuming it away; §10.8 sets out what an authority holding a whole share can do, and what splitting the share does and does not fix.

2.2 Assumed resources

Ordinary laboratory equipment: hand tools, a soldering iron, an oscilloscope, a logic analyser, a bench supply, an accelerometer, a microphone, a thermal camera, a current probe. Unsupervised physical possession for days. We do *not* assume a focused ion beam, a semiconductor laboratory, decapsulation, or the ability to read a key out of a certified secure element.

2.3 Service as an attack surface

The appliance is designed to be serviced. The cooling cartridge is removable and reseating it is an expected operation, not an anomaly. This creates an attack the architecture must address specifically: an adversary may attempt to *disguise an attack as service*, or to exploit the authorised procedure by which a legitimate change is absorbed. §12 describes cartridge authentication and transport mode; §10.11 describes the authorised re-wrap that a legitimate change requires, and why performing it after the change rather than before destroys the data.

2.4 Out of scope

- Extraction of the device root key from within the secure element (T1).
- Compromise or coercion of the authority (T5), beyond the partial mitigation of splitting the share.
- Supply-chain substitution before provisioning (T2).
- Side channels against the cryptographic implementation itself, as distinct from the physical measurement.
- Denial of service. Almost anything that disturbs the assembly prevents the key forming. An adversary content to destroy data rather than read it has an easy time, and this is inherent in binding a key to a physical state rather than a defect of the particular scheme.

2.5 Scope of the claims

The bench validation of §14 establishes the physical and manufacturing properties the architecture relies upon: thermal signal margin, tamper separability, external-leakage suppression, reseating repeatability and manufacturing tolerance. On that basis we do not, by itself, claim resistance to a determined, resourced attacker across every attack path; full adversarial evaluation is identified as further work in §17. §15 maps each mechanism to the adversary class it addresses.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

3 RELATED WORK

3.0.0.1 A note on sources.

Two of the works closest to ours could not be obtained in full. The publisher's server returned an authentication redirect for Rantala et al. (1996) and a 403 for Obermaier et al. (2018). Our characterisation of both is therefore limited to what their abstracts and indexed snippets state, and we have not attributed any method to either beyond that. We flag this because a related-work section that silently overstates what its authors read is worse than one that admits the gap.

3.0.0.2 Tamper-responsive enclosures.

The closest patent art is the tamper-protected hardware of Kreft (2011), which encloses a module in a boundary whose disturbance is detected and which reconciles a secret from that boundary using helper data. The granted specification is concerned with acoustic propagation through a potted volume and with MEMS transducers. It does not describe deriving the stimulus from the protected device's own compute workload, nor a member that is simultaneously the thermal path and a load-bearing element, nor an internal differential between two sensors, nor decoy measurement windows.

Obermaier et al. (2018) describe a measurement system for a capacitive PUF-based security enclosure, using a differential stimulus and a digital lock-in amplifier implemented on a microcontroller, at femtofarad resolution. The differential arrangement and the lock-in estimator are shared with our work. The physical quantity is not: a capacitive envelope is a dedicated security structure added to the device, whereas the member we measure is one the appliance already needs in order to function. Immler et al. (2019) develop the error-correction and entropy treatment for such enclosures.

3.0.0.3 PUFs driven by a commanded physical state.

Willers (2017) derive a PUF from MEMS sensors whose response is deliberately shifted by an on-chip heater, reconciling the result with helper data. This is the nearest art to the idea of a *commanded thermal state* modulating a physical response. The difference is the source of the heat: a dedicated heating element adds cost, power, failure modes and an attack surface, and its activation is itself an observable event. In our architecture no heating element exists; the stimulus is a scheduling decision. Yongjiang Laboratory (2024) describe a resonant-sensor temperature PUF, again with the temperature as a commanded input rather than a by-product of useful work.

3.0.0.4 Attestation from dynamic response.

The family of ARM Limited (2018) generates a workload on the device and attests from the resulting dynamic response. The shared idea is that the device's own activity produces the stimulus. What is absent there is the mechanical dimension: the response examined is electrical and architectural, not the envelope of a structural vibration, and no member plays a dual thermal and load-bearing role.

3.0.0.5 Lock-in detection of a thermally modulated mechanical response.

The signal-processing technique we use is not new, and we do not claim it. Rantala et al. (1996) applied amplitude-modulated lock-in vibrothermography to non-destructive evaluation of polymers and composites three decades ago: a mechanical excitation is amplitude modulated at a low frequency and the response detected synchronously at that modulation frequency. Our use inverts the roles — the *thermal* state is modulated slowly and the *mechanical* response is the carrier — and applies it to a security question rather than a materials one, but the estimator is of the same family. The abstract records no security, enclosure or key application, and we make no claim about the body of that paper, which we did not obtain.

3.0.0.6 Thermal effects as a channel.

Masti et al. (2015) demonstrate thermal covert channels between cores on a multi-core platform, achieving a reported 12.5 s^{-1}. That work establishes that compute activity produces a thermally observable signal with enough structure to carry information — the same physical premise our thermal challenge relies on, used for the opposite purpose. It is an argument that the effect exists; it is not evidence for the specific coupling we describe.

3.0.0.7 Hiding when a measurement occurs.

Idemia Identity and Security France (2018) bury genuine measurement requests among dummy requests, selecting K of N at random so that an observer cannot tell which are real. We adopt the same principle and extend it to the physical domain: our decoy windows must match the genuine ones not merely in existence but in four externally observable characteristics (Section 6), because an observer of a physical appliance can measure drive amplitude, window duration, excitation spectrum and supply current, none of which arises in a purely protocol-level setting.

3.0.0.8 Reconciliation.

Our reconciliation is the code-offset construction (Juels and Wattenberg 1999; Dodis et al. 2004; Linnartz and Tuyls 2003). We use it conventionally and claim nothing about it. We are, however, careful in Section 9 about what it yields, because the step most often skipped — dividing by the code rate — is the difference between a plausible entropy figure and a wrong one by a factor of five.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

3.0.0.9 Position.

None of the elements above is claimed in isolation. What we describe is a specific conjunction: the thermal stimulus produced by scheduling the appliance's own workload; the member carrying that heat also carrying structural load; a separate mechanical excitation in a higher band whose *envelope* is examined; that envelope estimated synchronously at the thermal modulation frequency; the quantity compared being an internal differential between two protected sensors on opposite sides of the thermal path; and genuine windows hidden among matched decoys.

3.1 Comparison

Table 2 places the closest work against the six elements whose conjunction is the subject of this paper. No prior work in the set combines more than three of them, and none combines the first two — a stimulus derived from the protected device's own useful work, carried by a member that is simultaneously structural.

Feature comparison. • present; ◦ partially or in a different form; blank absent. Assessed from the material we were able to obtain, which for two entries was an abstract or a snippet only (marked †).

Kreft (2011) (enclosure PUF)			◦		◦	
Obermaier et al. (2018)† (capacitive)				•	•	
Immler et al. (2019)† (envelope PUF)				◦	•	
Willers (2017) (MEMS, commanded heat)						
Yongjiang Laboratory (2024) (resonant, commanded)						
ARM Limited (2018) (workload attestation)	•			◦		
Rantala et al. (1996)† (lock-in VT)			•	•		
Masti et al. (2015) (thermal channel)	•			◦		
Idemia Identity and Security France (2018) (dummy requests)						◦
This work	•	•	•	•	•	•

Two columns deserve comment. ARM Limited (2018) and Masti et al. (2015) both derive a signal from the device's own activity, which is the first column; neither carries it through a structural member, so a mechanical change does not register. Obermaier et al. (2018) and Immler et al. (2019) both take an internal differential and estimate it synchronously, which are the fourth and fifth columns; both do so across a dedicated security structure added to the device rather than across a member the device needs anyway. The decoy column is ◦ for Idemia Identity and Security France (2018) because the principle is the same but the setting is protocol-level, where a dummy need only exist; a physical decoy must additionally survive measurement of its drive amplitude, duration, spectrum and supply current.

3.2 Acknowledged prior art

Because the contribution is a conjunction rather than any single element, it is worth stating plainly what is *not* new. Each of the following is known, and none is claimed here:

- enclosing a module in a package, emitting a predetermined excitation into it, sensing the back-scatter with internal sensors, treating that response as a physical unclonable function, reconciling it against public helper data, and treating a wrong secret as evidence of tampering — including where the excitation is acoustic and the sensors micro-electro-mechanical;
- generating a stimulus by scheduling a processor's own load, measuring the resulting dynamic response synchronously, characterising it by magnitude and phase in a band, scheduling such measurements at unpredictable intervals, binding the exchange with a nonce, and withholding cryptographic operations until comparison succeeds;
- commanding a thermal state of a mechanical resonator and using the shifted resonance parameters as a physical unclonable function, with several test temperatures as challenges;
- amplitude-modulating a high-frequency mechanical excitation at a low frequency and recovering magnitude and phase by a lock-in technique, including with a digital lock-in amplifier and differential measurement between elements with a defined phase relationship, within a tamper-detecting enclosure;
- burying genuine measurement requests among randomly selected dummy requests so that only a secure circuit knows which are genuine;
- holding part of a storage key outside a computer and requiring it at every start; deriving a key by applying a key-derivation function to concatenated material from several sources with a salt and context; and splitting a secret so that every part is required;
- taking the ratio and relative phase of two sensors spanning a structural element to cancel an unknown common input — in structural dynamics a *transmissibility*;

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

- that modal frequencies and damping vary with temperature and that the variation must ordinarily be compensated;
- requiring an externally held credential in addition to a tamper-resistant enclosure; erasing keys on power loss while retaining the ciphertext; an administrator-set transport mode; a time-limited lease; and gating release of a storage key on a tamper verdict.

3.3 Positioning

None of the elements above is claimed in isolation. What is described here is a specific conjunction: the thermal stimulus produced by scheduling the appliance's own compute workload; the member carrying that heat also carrying structural load; a separate mechanical excitation in a higher band whose *envelope*, rather than the response itself, is examined; that envelope estimated synchronously at the thermal modulation frequency; the quantity compared being an internal *differential* between two protected sensors on opposite sides of the thermal path rather than any absolute or externally reachable quantity; and genuine windows hidden among decoy windows matched to them within declared tolerances in externally observable characteristics.

We note, and it bears on how the comparison above should be read, that no professional prior-art search has been carried out for this work. Several patent databases refused programmatic access during the review that produced Table 2. Nothing here asserts that the combination is novel or inventive; those are matters for an examiner, and an examination has been requested but has not taken place.

4 SYSTEM MODEL AND ASSUMPTIONS

4.1 Entities

the physical unit: enclosure, carrier board, compute module, storage, and the security subsystem described below.
the application processor executing the user workload. Assumed potentially compromised in software.
a separate processing element with its own power domain, firmware store, clock and sensor interface. It performs the verification and decides whether to request an unwrap.
a tamper-resistant device holding the device root key, performing key derivation and unwrapping, never exporting the root key.
the domain read by the security controller, comprising the first witness sensor set and the two protected internal sensors.
a domain controlled independently of the security controller, comprising the second witness sensor set. The independence criteria are given in §5.2.
the party outside the appliance that holds, or authorises release of, the external key share.
the party entitled to the data, who may hold a credential constituting part of the external factor.

4.2 Assets

The asset is the plaintext of data at rest on the appliance's persistent storage. Derived assets are the data-encryption key that decrypts it, the key-encryption key that unwraps that, and the three factors from which the key-encryption key is formed.

Persistent storage holds *only* the wrapped data-encryption key and ciphertext. It holds no part of the device root key, no physical-binding secret and no external share.

4.3 Security goals

An adversary in possession of the storage device alone cannot recover plaintext.
An adversary in possession of the entire intact appliance, without the external factor, cannot recover plaintext.
An adversary who opens the appliance destroys the physical-binding secret in doing so, and cannot reconstruct it from observations taken before the opening.
An adversary observing the appliance externally during operation gains limited ability to determine which measurement windows are genuine.
A legitimate service operation does not destroy access, provided it follows the authorised procedure.
A false tamper indication does not destroy user data.

G5 and G6 are availability goals and they constrain the design as hard as the confidentiality goals do. An architecture that achieves G1–G4 by making the appliance brittle is not deployable, and §10 spends as much effort on recovery as on derivation for that reason.

4.4 Trust assumptions

1) The secure element's own guarantees hold: the root key cannot be extracted from it by the adversary classes of §2.
2) Enrolment occurs on a genuine assembly, before the adversary has access. Supply-chain substitution prior to provisioning is out of scope.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

3) The security controller's firmware is authentic at enrolment and its signing key is not compromised. The second sensing domain's firmware is accepted under a *distinct* signing authority, so that compromise of one key does not yield both domains.
4) The cryptographic primitives are secure: we assume a key-derivation function, an authenticated encryption scheme suitable for key wrapping, and a signature scheme, each used conventionally.
5) The authority behaves correctly when it holds an undivided share. §10.8 relaxes this by splitting the share, and states what remains exposed.

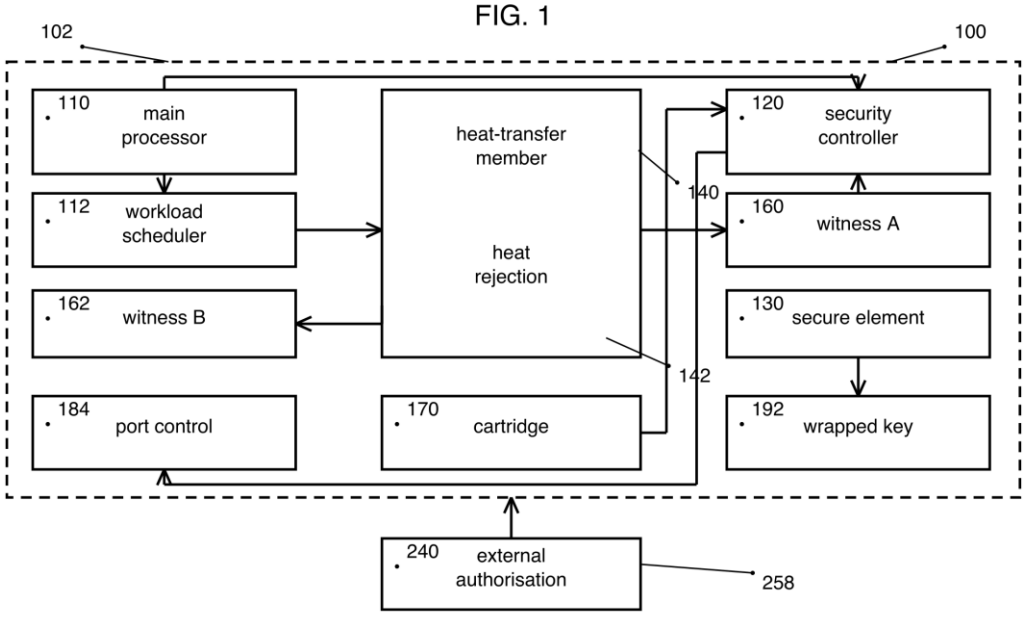
4.5 Measured quantities

Before the detail of §6, the quantity at the centre of the architecture can be stated in one sentence.

The appliance heats a structural member by scheduling its own compute workload at a frequency f_m in the millihertz range. It simultaneously excites the structure mechanically in a band far above f_m. The mechanical response is band-passed and its *envelope* formed. That envelope varies slowly, at f_m, because the structure's transfer function moves with temperature. A synchronous estimator at f_m extracts from each of two internal sensors a complex *coupling statistic*. The quantity enrolled is the *difference* between those two statistics — a relative phase and an amplitude ratio — and it is that difference which is reconciled into the physical-binding secret.

5 ARCHITECTURE

FIG. 1



Appliance architecture. The main processor executes the user workload and generates the thermal challenge by scheduling it, but takes no part in deciding whether verification succeeded. The security controller has its own power domain, firmware store and sensor interface. The secure element holds the device root key and performs key derivation and unwrapping. (Patent FIG. 1.)

Figure 1 shows the arrangement. Four elements matter for what follows.

5.0.0.1 The main processor

executes the user workload. It is the source of the thermal stimulus, because the scheduler modulates that workload, but it is deliberately excluded from the verification decision. A compromised application processor can change how much heat is produced; it cannot declare the assembly intact.

5.0.0.2 The security controller

is a separate processing element with its own power domain, firmware store and sensor interface. It drives the actuator, samples the witness sensors, performs the estimation described in Section 7, reconciles the physical-binding secret, and decides whether to ask the secure element for an unwrap.

5.0.0.3 The secure element

generates and holds the device root key, which is never exported. It performs the key derivation and the unwrap. Plaintext key material exists only inside it and, for the data-encryption key in use, in volatile memory.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

5.0.0.4 The structural heat-transfer member

conducts heat from the main processor toward a heat-rejection surface *and* forms part of the load-bearing mechanical path of the appliance. This dual role is the mechanism by which a mechanical change becomes a thermal one. Disturbing a joint changes both the structural transfer function and the heat path; the measurement is sensitive to the combination.

Verification proceeds in windows. In each window the scheduler drives the workload at the thermal modulation frequency f_m, the actuator excites the appliance in a much higher band, and the response envelope at each protected internal sensor is estimated synchronously at f_m. Some windows are genuine and contribute to key material; others are decoys and do not. Which is which is determined by a selection secret that the main processor never sees.

5.1 The mechanical arrangement

Three mechanical elements matter to the measurement.

The *inner measurement frame* carries the protected internal sensors. Where present it is mechanically decoupled from the outer shell by isolators, so that the transfer from the measured structure to the externally observable surface is attenuated in the excitation band. The frame is optional: the architecture also admits sensors mounted on, or mechanically coupled to, the structural heat-transfer member directly, and the claims do not require the frame.

The *outer shell* is what an external observer can reach. It is cosmetic in the sense that matters here — nothing enrolled is measured on it.

The *removable cooling cartridge* seats on hard datums with a declared preload. It is part of the thermal path, so its reseating is a key-management event (§12).

5.2 Two independently controlled sensing domains

The cross-witness residual has force only if the two domains are genuinely independent, and "independent" is not a matter of degree to be left to the implementer. A domain is independently controlled, for present purposes, when it satisfies *all* of the following:

1) it is read through a sensor interface that neither the security controller nor the main processor can drive or sample;
2) its sampling is timed by a clock source distinct from that of the security controller;
3) its firmware is held in a store distinct from the security controller's firmware store, and is accepted under a *distinct signing authority*, so that compromise of one signing key does not yield both domains;
4) its supply is separate, so that manipulating one domain's supply does not silently manipulate the other.

The purpose is that an attacker who has compromised one domain — including one who has compromised the main processor entirely — cannot thereby dictate what the other reports. That is what gives the cross-witness residual its force, and an implementation satisfying three of the four conditions does not get three quarters of the benefit; a shared clock or a shared signing key is sufficient to defeat it.

5.3 Reference numerals for operations

Four operations carry their own identity in the architecture because they are separately verifiable: the *dark-bit masking* step, by which quantisation pairs of insufficient margin are excluded from the enrolled word; the *peripheral characterisation* and *peripheral authentication* steps, by which a device presented at the data interface is examined and then authenticated before any data path is closed (§12); and the *cartridge authentication* step, by which the identity device of a removable cooling cartridge is verified before its calibration record is relied upon.

6 PHYSICAL BINDING

6.1 A stimulus intrinsic to the appliance

A security measurement needs a stimulus. Adding a dedicated one — a heater, a shaker, a signal source — adds cost, power draw, failure modes and attack surface to a device that is already dissipating tens of watts in a structure that already has mechanical resonances. We use what is there. The scheduler modulates the compute workload, and the heat that results propagates through the structural heat-transfer member. No heating element exists in the design, so there is none to disable, and no distinct activation event to observe.

The mechanical excitation is separate and lies far above the thermal modulation: in one embodiment it spans the first three structural bending modes, typically 200 Hz–5000 Hz. Drive amplitude is limited so that peak surface acceleration remains below a declared limit, for example 0.5. That limit is chosen to keep stresses low and to limit what is radiated. We do *not* assert that it renders the excitation inaudible: audibility depends on radiated sound pressure, on frequency and on the acoustic environment, not on acceleration alone.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

6.2 The member as a thermal waveguide

It is convenient to speak of *the* temperature of the heat-transfer member. At the frequencies used here that is wrong, and the error is not small.

Heat entering at the processor spreads by diffusion, so a periodic input produces a damped thermal wave travelling outward from the source. For 6061 aluminium the thermal diffusivity is about $\alpha = 6.9 \times 10^{-5}$ m^2/s. Over the 55 mm from the source to a far fastening the in-plane diffusion time is about 44 s, so a lumped description would be valid only well below about 3.6 mHz . The modulation frequencies useful here lie at or above that figure. The distributed description is the correct one, and we use it throughout.

The wave has penetration depth

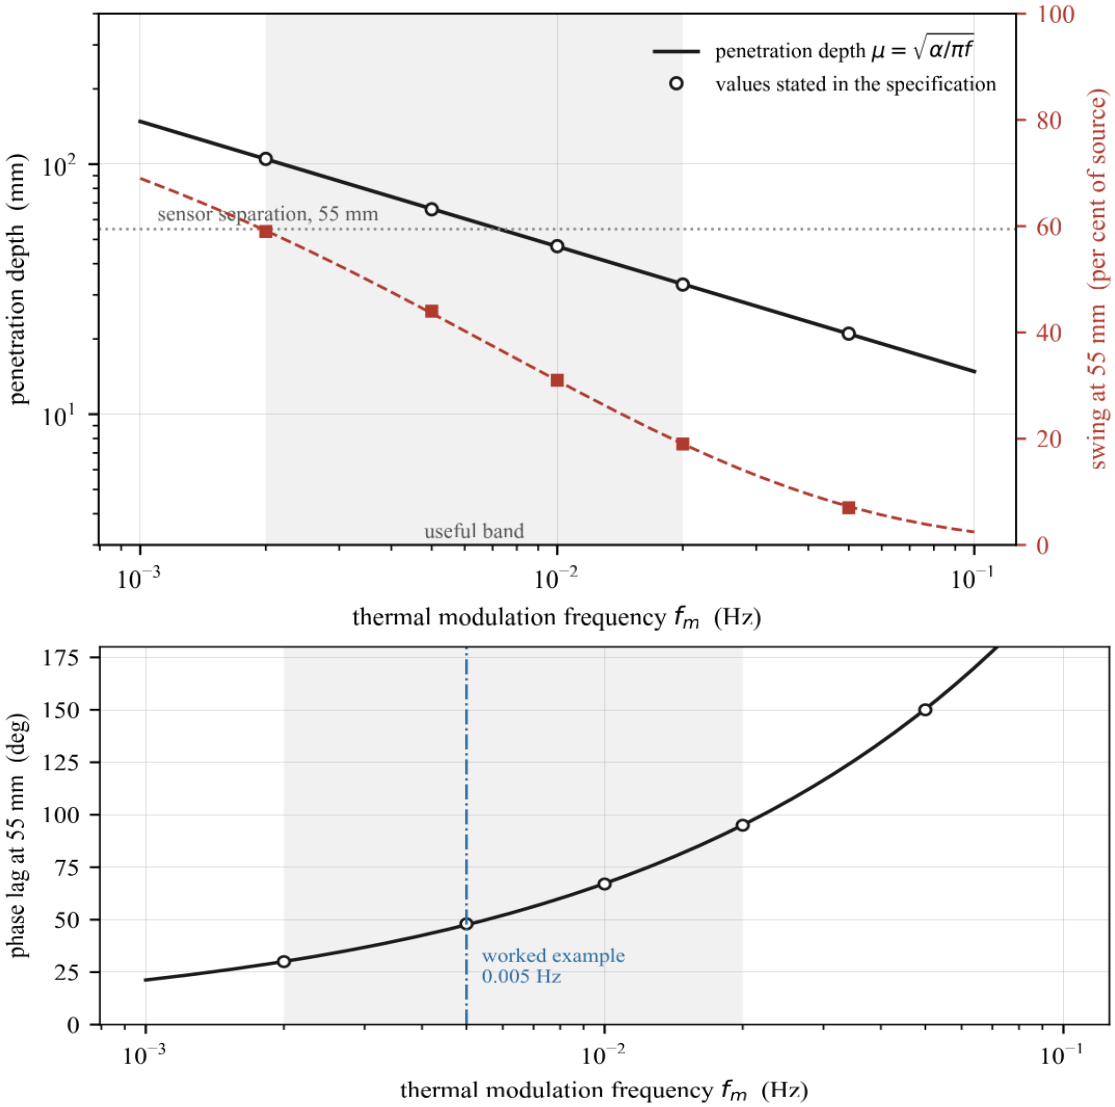
$$\mu(f_m) = \sqrt{\frac{\alpha}{\pi f_m}},$$

and at distance x from the source the swing is attenuated by $\exp(-x/\mu)$ and lags the source by x/μ radians. Table 3 evaluates this at $x = 55$ mm; Figure 2 plots it.

S[table-format=1.3] S[table-format=3.0] S[table-format=2.0] S[table-format=3.0] f_m (Hz) & μ (mm) & swing (%) & lag (°)
& 105 & 59 & 30
0.005 & 66 & 44 & 48
0.010 & 47 & 31 & 67
0.020 & 33 & 19 & 95
0.050 & 21 & 7 & 150

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/



Penetration depth and the resulting attenuation and phase lag at 55 mm. Markers are the values stated in Table 3; curves are Equation [eq:depth] evaluated continuously. The shaded band is the useful modulation range 0.002 Hz–0.02 Hz.

Two consequences follow. The swing at the source falls approximately as $f_m{}^{-1/2}$, not as $f_m{}^{-1}$. And the swing at a distant point fall faster still, because the exponential attenuation increases as well.

6.3 Differential sensing and sensor placement

Because the two protected internal sensors sit on opposite sides of the thermal path, they sample the wave at different distances and therefore at different *phases*. The internal differential quantity formed between them measures the thermal gradient and its phase across the member — a property of the geometry, the joints and the material, rather than of the absolute temperature of any one point. Absolute temperature is exactly what an attacker outside the case can most easily influence, and exactly what drifts with ambient conditions.

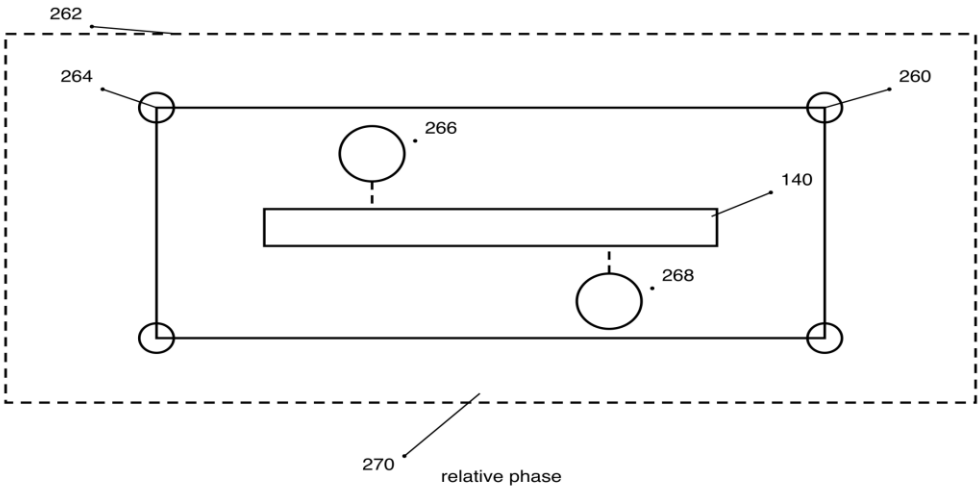
This immediately constrains placement. *Sensors placed close together sample the same phase of the wave and the differential vanishes.* The separation is therefore chosen so that, at the working f_m, the coupling statistics estimated at the two sensors differ by more than the detection floor of the estimator — in relative phase, in amplitude ratio, or in both.

We state the rule in terms of the estimator rather than the geometry deliberately. In wave terms the condition is met when the separation is a substantial fraction of $\mu(f_m)$, and preferably of the order of it; but μ is a guide to achieving the condition, not the condition itself. What must actually hold is that the difference is *resolvable*, and that is a statement about signal-to-noise which only the detection floor can express.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

FIG. 16



Internal differential sensor geometry, with the outer shell mechanically decoupled from the inner frame. The two protected internal sensors lie on opposite sides of the structural heat-transfer path. (Patent FIG. 16.)

6.4 Effect size

An architecture is only interesting if the quantity it measures is large enough to measure. We chain the estimate from watts to a percentage envelope modulation.

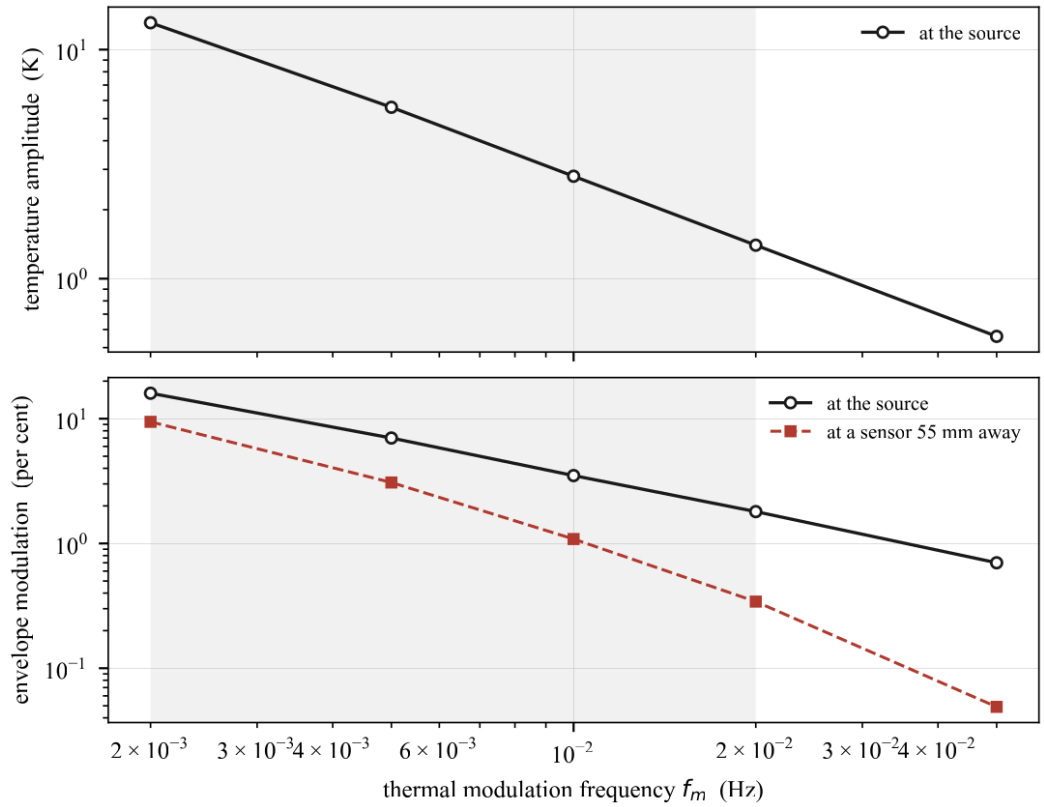
Take a peak-to-peak workload swing of 20 W, a thermal capacity of about 72 J K^{-1} and a thermal resistance of about 2.8 K W^{-1}. The fundamental of a square-wave workload carries about 12.7 W, which gives the temperature amplitudes in the second column of Table 4.

The temperature coefficient of Young's modulus for aluminium is about -5×10^{-4} K, and a modal frequency varies as the square root of modulus, so the fractional frequency shift is about -2.5×10^{-4} K. For a lightly damped mode the fractional change in the response envelope near resonance is approximately the fractional frequency shift divided by twice the damping ratio. Taking a damping ratio of 0.01 — representative of a bolted aluminium assembly, but *spanning* 0.005 *to* 0.05 *in practice* — gives the third column.

S[table-format=1.3] S[table-format=2.2] S[table-format=2.1] S[table-format=2.2] f_m (Hz) & ΔT at source (K) & envelope at source (%) & envelope at 55 mm (%)
& 13.10 & 16.0 & 9.44
0.005 & 5.60 & 7.0 & 3.08
0.010 & 2.80 & 3.5 & 1.09
0.020 & 1.40 & 1.8 & 0.34
0.050 & 0.56 & 0.7 & 0.05

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/



Temperature amplitude and envelope modulation against thermal modulation frequency, at the source and after attenuation at 55 mm.

This is the tension the design has to resolve. Lowering f_m raises the signal but lengthens the measurement; raising it shortens the measurement but drives the far sensor toward the noise. In one embodiment 0.002 Hz–0.02 Hz — periods of 50 s–500 s — satisfies both, and 0.005 Hz is used as the worked example. At that frequency, eight periods is 1600 s and forty periods is 8000 s; with five rounds and equal numbers of genuine and decoy windows a full verification is of the order of four hours at eight periods per round, or of the order of a day at forty .
That is slow, and we do not present it as otherwise. It is a verification performed at boot, after service, or on a schedule — not per transaction.

6.5 Cross-witness agreement

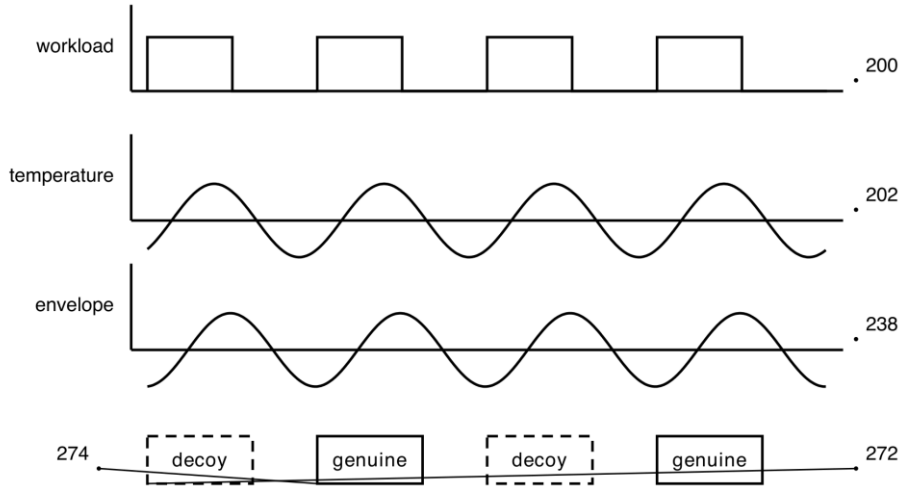
A single sensing domain can be spoofed by whoever controls that domain. Two independently controlled witness sets observe the same mechanical excitation, and a *cross-witness residual* measures their disagreement, normalised by the spread recorded at enrolment. A response in which the two domains are mutually inconsistent is rejected regardless of its amplitude. Injecting a signal from outside tends to produce a common-mode response, which is exactly what this test is sensitive to.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

6.6 Decoy windows

FIG. 6



Thermal and mechanical challenge timeline with genuine and decoy windows interleaved. Which windows are genuine is set by a selection secret unavailable to the main processor. (Patent FIG. 6.)

An attacker who can tell *when* a genuine measurement is being taken can concentrate effort on those windows. We therefore interleave decoy excitation windows, in which excitation is applied but whose sensed response takes no part in forming key material.

A decoy is only useful if it is indistinguishable from the outside. Protocol-level dummy requests (Idemia Identity and Security France 2018) need only exist; a physical decoy must survive measurement. Ours are matched to genuine windows in four externally observable characteristics:

1) the actuator drive amplitude;
2) the duration of the window;
3) the spectrum of the mechanical excitation applied;
4) the electrical power drawn by the drive electronics.

The matching is to a *declared matching tolerance* on each characteristic, fixed at enrolment and recorded with the enrolment data. In one embodiment the tolerance is five per cent for drive amplitude, window duration and drive-electronics power, and 1 dB over the excitation band for the spectrum; tolerances of two per cent and of ten per cent are also contemplated. Expressing the requirement as a tolerance rather than as exact equality is deliberate — exact equality of a physical observable is not achievable and a specification demanding it would be unsatisfiable.

6.6.0.1 The bound this assumes, and how it would be tested.

Matching four observables to a declared tolerance does not by itself establish indistinguishability, and we do not claim it does. What is assumed is a bound on how well any classifier can separate genuine from decoy windows. Testing that bound requires a protocol, which we specify rather than leave implicit: accelerometer and microphone on the outer shell, a current probe on the supply, a surface temperature sensor on the heat-rejection surface; a *balanced* set of recorded windows; a classifier trained to distinguish the two kinds; the statistic being the classifier *advantage*, $|2a - 1|$ for balanced accuracy a; a confidence interval by *blocked* cross-validation with separate recording sessions as blocks; and a feature set and classifier family declared in advance. The design passes only if the upper end of that interval does not exceed an advantage bound declared in advance.

Statistical power follows the usual reciprocal square root: testing a bound of 0.2 needs of the order of two hundred windows, and a bound of 0.05 of the order of three thousand.

This protocol was run on the prototype (test T4, §14): across five rounds and 200 window pairs the decoy classifier reached an AUC of 0.50–0.53, consistent with no timing information, and the envelope leakage at the outer shell

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

was measured at 0.08% with 31 dB shell-to-internal suppression, within the declared bound. The operating margin above threshold is about a factor of two and is recorded as marginal.

7 ACQUISITION AND ESTIMATION

7.1 Sampling and anti-aliasing

The witness sensors are sampled at at least four times the highest excitation frequency. For a band extending to 5 kHz, 20 s^{-1}–50 s^{-1} is appropriate; 25.6 s^{-1} is used in the worked example.

The anti-alias requirement is worth stating explicitly, because it is where a naive design fails. A converter needs stop-band attenuation at its Nyquist frequency of roughly six decibels per bit of resolution: about 72 dB for twelve bits, about 96 dB for sixteen. At 25.6 s^{-1} the Nyquist frequency is 12.8 kHz, only 1.7 × an 8 kHz corner. A third-order analogue response gives under 15 dB there; an eighth-order response gives about 33 dB; and the order needed to reach 80 dB at a ratio of 1.6 is about twenty, which is not a realisable analogue filter in a product.

The converter is therefore oversampled and decimated digitally. At 256 s^{-1} the ratio of Nyquist frequency to corner is 16, at which a third-order response gives about 72 dB — achievable with a filter that exists.

7.2 Envelope generation

A digital band-pass filter isolates the excitation band, rejecting mains-related and structure-borne interference below it. The envelope is then formed, either as the magnitude of the analytic signal from a Hilbert transform or by squaring and low-pass filtering. A corner of 2 Hz is appropriate for thermal modulation up to 0.05 Hz, being more than an order of magnitude above it.

The envelope detection for the two protected internal sensors is performed *digitally, by the same computation*. The reason is specific. With a single pole at 2 Hz the lag at 0.02 Hz is about 25 mrad, which is larger than the phase resolution of the estimator. That lag is common to both sensors and cancels in the differential — but only if the two envelope paths match. Two analogue detectors with component tolerances do not match, and the residual mismatch appears directly as a false differential. Two instances of the same arithmetic are identical by construction.

This is a small design decision with a large consequence, and it follows from taking a differential rather than an absolute measurement: the differential cancels common-mode error only to the extent the two paths are genuinely common.

7.3 Synchronous estimation and the detection floor

The coupling statistic is obtained by demodulating the response envelope against the known workload reference at f_m, in phase and in quadrature, giving a complex quantity with amplitude and phase.

Drift is the obstacle. Ambient temperature moves, and a low-frequency trend in the envelope sits directly on top of the quantity being estimated. Our reference implementation estimates the drift *jointly* with the coupling by least squares, using a low-order polynomial plus a low-frequency Fourier basis up to $f_m/3$, rather than filtering it out first. On a synthetic fixture with a known truth of amplitude 0.01250 and phase 0.800 rad, joint estimation recovered the coupling while a moving-average high-pass returned 0.00991 at 0.615 rad and a quadratic-only basis returned 0.00906 with inverted phase.

That is a statement about an algorithm on synthetic data. It shows the estimator is not indifferent to how drift is handled; it is not evidence about any physical appliance.

7.3.0.1 The floor.

The residual standard deviation of the fit gives the uncertainty of the estimate. A *detection floor* is obtained by running the same estimator at several non-harmonic offset frequencies at which no coupling is expected, and taking the median. A challenge whose coupling amplitude does not exceed the floor by a declared margin is *rejected* as having insufficient signal-to-noise ratio, rather than analysed.

This matters more than it appears. Without it, a weak or absent coupling produces a small number rather than a refusal, and small numbers reconcile to something. The floor converts an unmeasurable condition into an explicit rejection.

7.4 Limitations of the estimator

The estimator recovers a coupling coefficient at f_m and reports a floor below which it will not. It cannot tell whether the coupling it found travelled through the mechanical transfer function or reached the sensors by some other path at the same frequency — and several other paths exist, because the workload modulation also drives fan speed, supply current, board temperature and sensor self-heating at exactly f_m.

Distinguishing them is not a signal-processing problem and no choice of estimator solves it. It requires controls that vary one path at a time, and those are the subject of §8.3.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

8 Selection, rejection and enrolment

An architecture that depends on a physical effect must say how a team determines whether a given assembly *exhibits* that effect — and must be willing to conclude that it does not. This section gives that procedure. Steps 1 to 5 are routine calibration. Step 6 tests the hypothesis on which the whole architecture depends, and its outcome may be rejection.

8.1 The procedure

Modal sweep. With the appliance at a steady temperature, drive the actuator with a swept sine across the candidate band and record the witness responses. Identify the resonant peaks.

Candidate modes. Select two or three modes that are well separated, well excited at *both* protected internal sensors, and not dominated by a single fastener. A mode dominated by one fastener measures that fastener, not the assembly.

Candidate thermal frequencies. Determine the thermal time constant of the member by applying a workload step and fitting the rise. The pole so found will ordinarily lie below the drift band that must be rejected, so candidate modulation frequencies are chosen *above* it — and the temperature swing actually achieved at each candidate is *measured rather than assumed*, because in that regime the swing falls as frequency rises. Candidates whose swing is too small to move the transfer function by more than the detection floor are discarded here and not carried forward.

Distinguishability. For each candidate pair of mechanical mode and thermal frequency, run the full challenge and compute the coupling statistic together with its detection floor. Retain only pairs whose coupling amplitude exceeds the floor by the declared margin over repeated runs.

Null controls. §8.3.

Position selection. Move the actuator and the protected internal sensors among candidate positions and repeat Step 4, retaining the arrangement that maximises the margin of the internal differential quantity *while minimising the response measurable at the outer shell*. The two objectives conflict, and the procedure makes the trade explicit rather than optimising one and discovering the other later.

Rejection. §8.4.

Acceptance region. §8.5.

Decoy leakage validation. §6.6, protocol in Appendix E.

Environment and service. Repeat across the declared ambient range and across at least ten authorised cartridge reseating cycles, fitting the compensation vector and confirming that the acceptance region still separates the seated state from a deliberately altered one.

8.2 Separating selection from enrolment

A distinction runs through Steps 4 and 4A that is easy to lose and expensive to lose.

The first measurement is a *null control*: it asks whether an apparent response appears at f_m when *no* mechanical excitation is applied. It is performed with the actuator off, and its purpose is to expose a confounding path that bypasses the mechanical transfer function altogether.

The second is a measurement of the *shift* of a resonance, or of the transfer function, with temperature — which is the effect the architecture relies upon. A shift cannot be measured with the actuator off, because a transfer function cannot be observed without an excitation to observe it with. That measurement therefore requires either the actuator driven, or an ambient excitation source whose presence and adequacy are independently established and recorded.

An instruction to measure a resonance shift with the actuator off would be self-defeating, and is not given.

8.3 Null controls

A response at f_m is not by itself evidence of thermo-mechanical coupling. The workload modulation also drives fan speed, supply current, board temperature and sensor self-heating at the same frequency, and any of those can reach the witness sensors by a path that never passes through the mechanical transfer function. Three controls follow directly:

1. **Actuator-off.** No mechanical excitation. The coupling amplitude must fall below the detection floor — with no excitation there is no envelope to modulate.
2. **Fan-fixed.** The fan is held at constant speed. The coupling statistic must not change materially.
3. **Decoupled-sensor.** A witness sensor on the same board but off the mechanical path must not reproduce the coupling.

A fourth is required, and the reason repays attention because it is a confounder that passes all three.

The *sensitivity* of a sensor, and the electromechanical coupling constant of an actuator, are themselves temperature dependent, and their temperature swings with the workload. Such a change *multiplies* the response. It therefore

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

vanishes with the actuator off; it is unaffected by holding the fan constant; and it does not appear at a sensor off the mechanical path. It passes every one of the three controls while producing a signal at the modulation frequency, in phase with the workload — which is exactly the signature being looked for.

Worse: because the protected internal sensors are deliberately placed where the thermal wave reaches them *differently*, the arrangement is optimised for precisely this artefact. The geometry that makes the intended effect measurable makes the confounder measurable too.

1. **Sensor-heated.** The protected internal sensors are heated by an independent means through the same temperature swing while the structural member is held isothermal. The coupling must *vanish*.

Three further measures accompany it. The temperature of each sensor — not only of the member — is measured and carried in the compensation vector. A maximum temperature coefficient is declared for sensors and actuator. And the drive monitor is supplemented by a reference accelerometer on the actuator, so that the delivered *mechanical* drive is observed rather than merely the electrical drive.

A candidate failing any null control is rejected whatever its amplitude. The controls are what permit a coupling statistic observed on an appliance to be *attributed* to the mechanism described here rather than to one of these confounders. Without them the architecture would be unfalsifiable, which is a worse position than being wrong.

8.4 Rejection as a valid outcome

If no combination of mode, thermal frequency and position yields a coupling statistic that is repeatable across runs and separable from the detection floor, **the assembly is rejected for this purpose and must be redesigned.** We state this as an ordinary outcome rather than a failure case. A physical binding scheme that cannot fail selection is one whose selection criterion is not testing anything. The possibility that a given design does not support the architecture is expressly contemplated. On the tested prototype the selection criterion was satisfied (Section 14: the coupling was separable from the detection floor by 3.0×), subject to the open items that §17 records.

8.5 Acceptance region from observed spread

Repeat the retained challenge at least eight times and record the mean and standard deviation of each enrolled quantity. The acceptance region is set *from the observed spread*, not from an assumed security target.

The direction of that inference matters. Choosing a false-accept rate first and then setting a region to achieve it would produce a region the hardware may not be able to honour. Setting the region from measured spread makes the achievable false-accept and false-reject rates a *consequence* of the assembly, which is what they are. They are not assumed in advance and they are not design inputs.

8.6 Reliability masking and the enrolled word

Quantisation is pairwise: each bit records whether one feature element exceeds another, with the pairing fixed at enrolment. Pairs whose margin is too small are excluded by a *dark-bit masking* step, using a χ^2 upper confidence limit on the observed spread rather than a point estimate, so that a pair which merely *appears* stable in a short run is not retained. A majority vote across several independent rounds stabilises what survives.

Appendix B gives the enrolment and reproduction procedures in full.

9 RECONCILIATION AND ENTROPY ACCOUNTING

9.1 Construction

Reconciliation uses the code-offset construction (Juels and Wattenberg 1999; Dodis et al. 2004). Quantisation is *pairwise*: each bit records whether one feature element exceeds another, with the pairing fixed at enrolment. Pairing rather than thresholding against a per-element mean matters — in our reference implementation the per-element variant produced an all-zero word on the synthetic fixture , since the mean tracks the element.

Unreliable pairs are masked as dark bits using a χ^2 upper confidence limit on the observed spread, and a majority vote across several independent rounds stabilises the retained pairs.

9.2 Output of the construction

The step most often elided is the code rate, and eliding it overstates the result by exactly the repetition factor. We therefore distinguish the quantities explicitly, in order:

1. **Measured features** — distinct quantities extracted from the response.
2. **Quantisation pairs** — at most half of that, since each pair consumes two features.
3. **Retained reliable pairs** — fewer again, being those surviving dark-bit masking. These form the codeword, of length n.
4. **Decoded information bits** — with a repetition code of factor r,
$$k = \lfloor n/r \rfloor,$$
because r coded bits are spent on each information bit in order to correct errors.

The information-bit upper bound on the physical-binding secret is k, *not* the number of reliable pairs.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

l S[table-format=3.0] l Quantity & Value & Note
Coded reliable pairs, n & 128 & codeword length
Repetition factor, r & 5 &
Decoded information bits, k & 25 & $\lfloor 128/5 \rfloor$
Incorrect figure & 128 & n mistaken for k

On the synthetic fixture used to exercise the reference implementation, four hundred measured features produced of the order of thirty information bits. That figure is a property of synthetic data. It is reported only to show the order of magnitude of the arithmetic, and it is not a measurement of any physical appliance.

9.3 Helper-data leakage

The helper data is not free. For a repetition code it discloses the pairwise exclusive-or of the enrolled bits within each block — $n - k$ parity relations in all. That is accounted for in the bound above rather than assumed away.

9.4 Min-entropy

The information-bit count is an upper bound, and an upper bound is not an entropy. Establishing min-entropy requires a distribution over a population of physical devices, and that requires measuring a population of physical devices.

We claim no min-entropy figure. Establishing min-entropy requires a distribution measured across a population of devices; the bench measurements of Section 14 characterise a single prototype and are not such a population dataset, and a figure produced from synthetic data would describe our synthetic generator rather than any appliance.

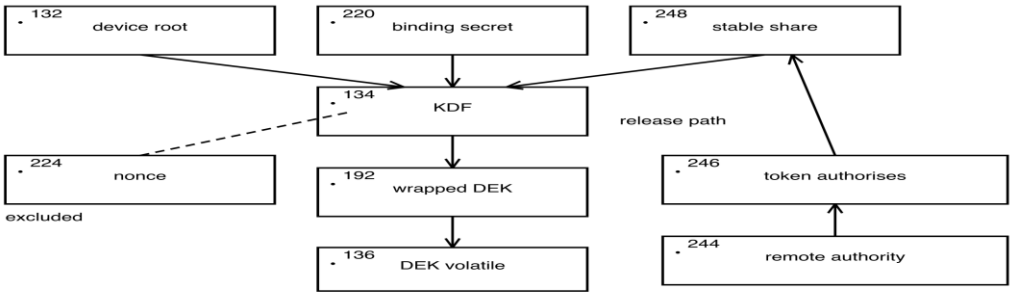
9.5 The secret as a binding factor

Physical binding strength and cryptographic key strength are different properties and are not interchangeable. The physical-binding secret is a *binding factor*: it is combined with a full-entropy device root key and an external key share in the derivation of Section 10. It is not a stand-alone key and must never be used as one.

This is the honest position, and it is also the robust one. An architecture whose security rested on the unmeasured entropy of a physical response would be making a claim it cannot support. Ours rests on the root key for strength and uses the physical response for *binding* — so an attack on the entropy of the physical measurement degrades the binding, which is its purpose, rather than the confidentiality of the data.

10 Key hierarchy and dual control

FIG. 11



The key hierarchy. Three factors of independent origin enter the derivation. The nonce is shown deliberately excluded from it: it binds the session and authorises release, and it is not an input to the key. (Patent FIG. 11.)

10.1 Three origins

The key-encryption key is derived from three factors:

$$KEK = \mathrm{KDF}(DRK \parallel PBS \parallel EKS,\ salt,\ context),$$

where DRK is the device root key, generated once inside the secure element and never exported; PBS is the physical-binding secret reconciled at each verification; and EKS is the external key share, supplied from outside the appliance. The context comprises a stored deployment identifier and may further comprise a cartridge identity and a firmware measurement *as enrolled at provisioning*.

The data-encryption key is generated once at provisioning, wrapped under the KEK, and stored only in that wrapped form. An authenticated encryption scheme suitable for key wrapping is used, for example AES-256-GCM or the AES key-wrap algorithm.

Each factor defeats a different attack. The root key is bound to the secure element, so a copied storage device is useless. The physical-binding secret is bound to the assembly, so an opened or substituted assembly is useless.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

The external key share is held outside the appliance, so the appliance *taken as a whole* is useless. Removing any one prevents formation of the KEK.

10.2 Key share versus authorisation token

The distinction that follows is the one we consider most important in this paper, and the one we got wrong first.

The *external key share* is a **stable secret**. It is established at provisioning, held outside the appliance, has the same value at every verification session, and is unchanged by an interruption of supply. It is a mathematical input to the key-encryption key, so that without it the key cannot be formed at all.

The *signed authorisation token* is a **changing credential**. It is bound to a challenge issued by the security controller for the present session, and it is short lived. Its function is to *authorise* the authority to release the stable share to the appliance, or to *carry* that share to the secure element under authenticated encryption. It is never an input to the key-derivation function.

Why the token cannot itself be the share

A token bound to a session and expiring necessarily takes a different value each session. Were the token bytes used as keying material, the key-encryption key would differ at every session, and the persistent wrapped data-encryption key — wrapped under the key-encryption key of some earlier session — would become permanently unreadable after the first. The storage would be destroyed on the first reauthorisation, silently, by a system behaving exactly as designed.

Nor can *verifying* a token against an enrolled authority public key substitute for the share. That key is public. It is not a secret the thief lacks. A check performed against it is a policy decision taken by the appliance, not a mathematical dependency — and an attacker who replaces the security controller, or replays the release instruction, obtains the key without the authority's participation.

We record this as a failure mode we introduced and removed, rather than as an observation about others' work. The architecture looked correct at the time: there was a token, the token was verified, the key was formed only after verification succeeded. What was missing was the distinction between *authorising release of* a stable secret and *being* the secret. An arrangement in which the appliance merely checks a token before releasing a key it could have formed anyway is disclosed here only so that it may be distinguished; it is weaker, and it is not what is described above.

10.3 Exclusion of session-varying quantities

This is the constraint we consider most important, and we state it as a rule:

> *No quantity whose value differs between sessions is an input to the key-derivation function.*

Excluded, specifically: the nonce, the session identifier, the session transcript, the signed authorisation token, the challenge to which that token is bound, its validity period, any single-use counter value, and any lease identifier. The inputs are the three factors, a stored salt and stored context — each of which is stable across sessions and across an interruption of supply.

The reason is worth spelling out, because the mistake is easy to make and hard to see. Suppose a short-lived authorisation token were used directly as the external factor in Equation [eq:kdf]. The token changes every session, by design — that is what makes it a defence against replay. But then the KEK changes every session too, and the wrapped data-encryption key, which was wrapped under the KEK of some earlier session, can never be unwrapped again. The storage becomes unreadable on the first reauthorisation.

We record this because it is a failure mode we introduced into an earlier version of this design and had to remove. The architecture looked correct: it had a token, the token was verified, the key was formed only after verification succeeded. What it did not have was a distinction between *authorising release of* a stable secret and *being* the secret.

The resolution keeps both properties. A signed, nonce-bound, short-lived token *authorises* the remote authority to deliver the same stable external key share, or carries it under authenticated encryption to the secure element's public key so that it is plaintext only inside the secure element. The token is the freshness mechanism; the share is the key material. The secure element accepts a share only after five cumulative checks: signature, appliance identity, challenge binding, validity window, and a single-use value against a monotonic counter.

A variant using a token alone, with no external share, is a *policy gate* rather than a cryptographic one — the key could be formed without the token by anyone able to bypass the check that enforces it. Where that variant is used it should be described as weaker, and it is not what the architecture above provides.

10.4 Lease, power loss and reauthorisation

The unwrapped data-encryption key exists in volatile memory under a lease. On expiry, on clean shutdown, on loss of supply or on a tamper indication, the plaintext key is zeroised. The *wrapped* key is never erased by any of those events — erasing it would convert a power cut into permanent data loss.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

Loss of supply also erases the reconciled physical-binding secret, and the external share must be re-delivered and re-accepted, not merely re-presented from a cache. A monotonic non-volatile attempt counter with escalating delay and eventual lockout limits brute-force reauthorisation.

10.5 Acceptance of the share by the secure element

Before accepting a share delivered in response to a token, the secure element verifies *all* of the following as cumulative preconditions, and accepts only if every one is satisfied:

1) the signature on the token, against the enrolled authority public key;
2) an appliance identity carried by the token corresponding to *this* appliance, so that a token issued for another is useless;
3) a binding of the token to the challenge issued for the present session, so that a token captured from an earlier session cannot be replayed;
4) that the present time lies within the validity period carried by the token;
5) that a single-use value carried by the token exceeds a monotonically increasing counter held within the secure element, the counter being advanced to that value on acceptance, so that the same token cannot be presented twice.

These five govern *acceptance* of the share. None of the quantities they examine is thereafter used as keying material — which is the rule of §10.3 restated at the point where it is most tempting to break.

10.6 Carriage of the share

The share may be conveyed to the secure element under authenticated encryption, or encapsulated to a public key of the secure element by a key-encapsulation mechanism, so that it is recoverable in plaintext *only* within the secure element and is at no point available in plaintext to the security controller or the main processor.

This is what defeats adversary A4. An attacker who has substituted the security controller observes the delivery but cannot read the share, and cannot supply a substitute share, because the secure element accepts one only over the authenticated path. Note that this holds independently of whether the physical measurement can be spoofed.

10.7 Sources of the external factor

The factor may be provided in any manner that keeps it outside the appliance: an owner credential such as a hardware token or a mobile device holding a key in its own secure hardware; a remote authority reached over a network; an enterprise signing service; or a pre-authorised offline capability token for deployments without connectivity.

10.7.0.1 Not a building anchor.

The factor is deliberately *not* a fixed anchor installed in the premises — a beacon, a wired presence detector, a geofence tied to the building. An anchor that remains with the premises protects only while the appliance remains there, and fails the legitimate case of authorised relocation. An anchor carried inside the appliance is taken along with it and protects nothing. The factor must be separable from the appliance *and* must travel with the authorised owner or remain under the authority's control. Those two requirements together are what exclude most of the obvious implementations.

10.7.0.2 Offline tokens.

An offline capability token weakens external control in proportion to its validity, because an attacker who seizes the appliance within the window holds both factors. It is bound to a single use by the monotonic counter, given the shortest validity the deployment tolerates, and refused once the counter has advanced past it. This is a stated trade-off, not an elimination of the exposure.

10.8 Authority capabilities and share splitting

An authority holding the share for an appliance holds a whole factor of that appliance's key. An attacker obtaining both the intact appliance and the authority's holding obtains the data — and the five acceptance checks of §10.5 do not help, because they are evaluated against a token the authority itself signs.

Two measures are disclosed. First, the share is *split*, for example between a remote authority and an owner credential, and reconstituted inside the secure element from a declared number of a declared set of holders, so that no single holder suffices. Second, the authority maintains a revocation list and an appliance-seized flag, and refuses release for an appliance so marked; the owner can set that flag out of band, without the appliance.

Neither measure removes the underlying position, which is that a factor held by another party is only as well protected as that party. We state it rather than let the five checks imply a guarantee they do not provide.

10.9 Lease, power loss and reauthorisation

A successful dual-control verification authorises use of the plaintext data-encryption key for a limited period — fifteen minutes to twenty-four hours in one embodiment, by deployment risk. On expiry the volatile copy is overwritten and discarded. Renewal requires a session transcript from a verification performed *after* the current period began; a renewal reusing the transcript that opened the lease would defeat the lease and is rejected.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

On loss of supply:

- the volatile copy of the data-encryption key and all session material are erased;
- any locally held authorisation that could be reused — a cached token, a cached share, an unexpired lease — is *invalidated*, so that it cannot be replayed after restart;
- the *reconciled physical-binding secret* is erased, so that it must be reconciled afresh from a new measurement;
- the external share must be *delivered and accepted again* over the authenticated path. It is not sufficient that a copy was present before the interruption;
- the wrapped data-encryption key is **retained**.

The appliance restarts sealed and requires both a fresh physical verification and a fresh external authorisation, passing through an awaiting-external-authorisation state in which physical verification has succeeded but no access is granted. An energy reserve, for example a supercapacitor, powers the security controller long enough to perform that erasure and to write a final verification record. It exists to *capture evidence and invalidate authorisation*, not to zeroise user data.

10.9.0.1 No ordinary event erases the wrapped key.

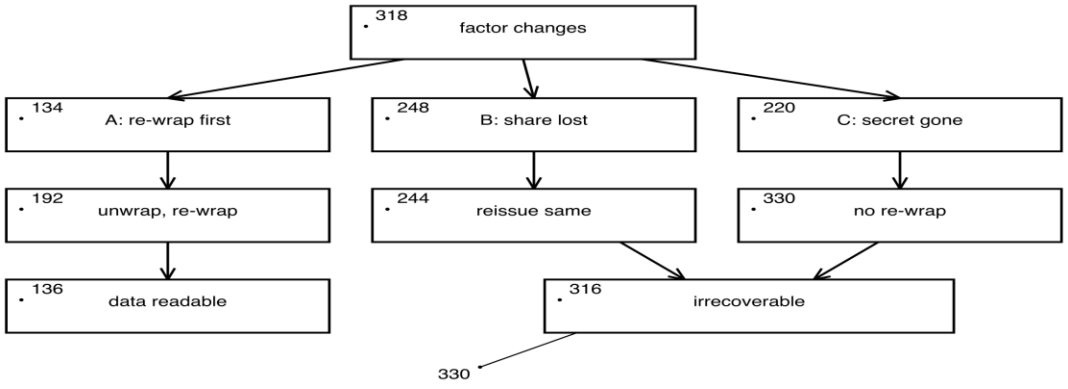
Not lease expiry, not clean shutdown, not loss of supply, and not a tamper indication. Destroying user data on a detection that may be a false positive is not acceptable in a field-deployed appliance — this is goal G6 — and a recovery procedure exists for that case.

10.10 Session transcript

The transcript is a digest over at least the nonce, the identity of the applied challenge, the coupling statistic, the internal differential quantity, the cross-witness residual, the identity of the external authorisation factor and the verification outcome, each length-prefixed. Length-prefixing is not incidental: without it, concatenation is ambiguous and two different sessions can produce the same digest.

10.11 Recovery

FIG. 17



The three recovery cases when a key factor changes. (Patent FIG. 17.)

Because the KEK depends on a physical measurement, any legitimate change to the assembly — a repair, a component swap — changes it. Three cases must be distinguished, and conflating them produces either data loss or a false promise of recoverability.

10.11.0.1 Case A: authorised re-wrap before the change.

While the old KEK is still reconstructable, the secure element unwraps the data-encryption key, re-derives the KEK for the new configuration, and re-wraps. The new wrapping is committed *before* the old is deleted, so an interruption leaves at least one valid wrapping. The plaintext key never leaves secure hardware.

10.11.0.2 Case B: the external share is lost.

Recoverable only if the authority can reproduce or reissue the same share, or if an independent escrow wrapping exists. The physical-binding secret is unaffected.

10.11.0.3 Case C: destructive loss of the physical-binding secret, with no pre-change re-wrap and no escrow.

Irrecoverable. Re-enrolment produces a new physical-binding secret, hence a new KEK, and the old wrapping cannot be opened by it.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

Case C is a real outcome and we state it as one. An architecture that binds a key to a physical assembly must accept that destroying the assembly destroys the key — that is the property being purchased. It would be straightforward and wrong to claim that because the wrapped key was never erased, recovery restores access; the wrapped key survives, but nothing can open it.

11 Resistance to external observation

Adversary A5 attaches an accelerometer, a microphone, a current probe and a thermal camera to the outer shell and records many sessions. This section describes how the architecture *reduces the usefulness* of such observations. It is not asserted that leakage is eliminated, and §11.1 states what remains.

Seven mechanisms contribute, and they are independent of one another in the sense that defeating one does not defeat the others.

11.0.0.1 1. What is enrolled is internal and relative.

The enrolled quantity is the internal differential: the relative phase between the two protected internal sensors and the ratio of their amplitudes. An external observer reaches the outer shell, not either internal sensor. A disturbance that scales both internal sensors equally cancels in the ratio and leaves the relative phase unchanged, so an attacker who drives the shell does not thereby reproduce the internal relative quantity.

11.0.0.2 2. The sensor pair spans the thermal path.

The two sensors lie on opposite sides of the structural heat-transfer path, so a thermally induced change moves them *differently*. The differential is therefore sensitive to the quantity the challenge modulates while being insensitive to bulk motion of the whole appliance — which is most of what an external excitation produces.

11.0.0.3 3. The shell is decoupled.

The inner measurement frame is isolated from the outer shell by isolators, so transfer from the measured structure to the observable surface is attenuated in the excitation band. This is the mechanical reason the enclosure of §13 has the construction it does, and it is one of the two objectives that Step 5 of §8 trades against each other.

11.0.0.4 4. The excitation phase is secret.

The nonce selects the excitation phase. An attacker who records a previous session observes a response at a phase that will not be used again.

11.0.0.5 5. Decoys.

Decoy windows are interleaved with genuine ones. Both produce an externally observable response; only genuine windows contribute to key material. The window selection secret is held by the security controller alone — not by the main processor, which is assumed potentially compromised. An attacker who records everything has a reduced ability to determine which records contribute to the physical-binding secret, and an attacker who replays a recording will, with high probability, replay it into a window that is discarded.

11.0.0.6 6. Internal consistency is required.

A response in which the two protected internal sensors are inconsistent with one another is rejected. Injecting a signal from outside tends to produce a *common-mode* disturbance, which is what this test is most sensitive to. The cross-witness residual between the two independently controlled domains applies the same principle at a higher level.

11.0.0.7 7. Replay is bound out.

The session transcript binds the nonce, the challenge identity and the measured quantities; the external authorisation token is bound to a challenge issued for that session and is short lived. A recorded response replayed in a later session does not match the transcript that session requires.

11.1 Limits

These measures raise the cost of an external observation attack. They do not make it impossible.

An attacker with prolonged access, high-quality sensors and the ability to observe many sessions may recover information about the internal relative quantities. The decoy construction reduces the rate at which they learn which windows matter; it does not reduce it to zero, and the bound on how far it reduces it is declared rather than measured (§6.6).

This is one of the two reasons the external authorisation factor exists. **Even complete knowledge of the physical response does not by itself open the appliance**, because the physical-binding secret is one of three inputs and the other two are not observable from the shell. An architecture whose security rested entirely on the secrecy of a physical measurement would be defeated by a sufficiently patient A5; this one degrades to the strength of the remaining two factors.

12 Service, transport and peripheral admission

A field-deployed appliance is serviced, moved and connected to peripherals. Each is a legitimate event that resembles an attack, and an architecture that cannot tell them apart is either unusable or insecure. This section covers the three.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

12.1 The removable cooling cartridge

The cooling cartridge is removable by design — it is a consumable thermal component, and reseating it is expected. But the cartridge is part of the thermal path, so removing and replacing it changes the very quantity the physical-binding secret is reconciled from.

Three measures apply. The cartridge carries an *identity device* which is authenticated before its *calibration record* is relied upon, so that a substituted cartridge presenting a favourable calibration is refused. The cartridge seats on hard datums with a declared preload (§13.6) so that reseating is mechanically repeatable rather than approximately so. And Step 8 of §8 requires the acceptance region to be validated across at least ten authorised reseating cycles, so that ordinary service variation is inside the region by measurement rather than by hope.

Whether reseating is in fact repeatable to the necessary tolerance was measured in test T5 (§14): across 30 reseating cycles the datum position repeated to 6–11 µm (within the 15 µm requirement), the preload was retained to within ±8%, and the worst-case thermal-path resistance shift was +1.9%, inside the code correction. Reseating is therefore compatible with the security concept on the tested prototype.

12.2 Transport mode

An appliance in transit is handled, vibrated and reoriented. Treating that as tamper evidence would make relocation impossible; ignoring it would create a window in which an attack is indistinguishable from shipping.

A *transport authorisation token* permits movement without handling being treated as tamper evidence. It is cryptographically distinct from the external authorisation factor and confers *no ability to unwrap the data-encryption key*. That separation is the point: an attacker who obtains a transport token obtains permission to move the appliance, which they could do anyway, and nothing else.

12.3 Peripheral admission: power before data

A data interface sits within a port chamber closed by a port cover, and has a supply path and a *switchable* data path.

When a peripheral is connected, the security controller energises the supply path *while the data path remains open* — so that no data line is connected to the main processor. It then characterises the electrical response of the peripheral over the supply path alone: the inrush current profile, the settled quiescent current, the presented capacitance. Separately, where the peripheral supports it, it authenticates the peripheral cryptographically.

The data path is closed only if *both* succeed.

The ordering is the substance of this measure. A peripheral that is malicious in its data behaviour has, at the moment of connection, no data path to be malicious over; it must first present an electrical signature consistent with what it claims to be. This does not defeat a peripheral that is electrically ordinary and maliciously behaved, and it is not offered as a complete defence — it removes the class of attack that depends on the data path being live before anything has examined the device.

12.4 Service events as key-management events

Each of the three operations above can change a quantity that the key-encryption key depends on: the cartridge identity, the firmware measurement, or the physical assembly itself. The architecture therefore treats them as key-management events, and requires the authorised re-wrap of §10.11 to be performed *before* the change takes effect.

The ordering is not a convenience. Performing the re-wrap after the change is impossible: once the change has taken effect the old key-encryption key can no longer be derived, so there is nothing with which to unwrap the data-encryption key in order to re-wrap it. A procedure that omits this ordering does not degrade gracefully — it destroys the data, silently, at the next verification.

13 Enclosure realisation and analysis

The architecture constrains the enclosure, and the enclosure was analysed to check that the constraints are mutually satisfiable. The target hardware is a commodity development kit — deliberately, since security that requires a custom board nobody can buy is security nobody can deploy.

13.1 Modal behaviour

The excitation band is specified as spanning the first three structural bending modes. That is only meaningful if the structure has its modes there. Treating the spine as a beam over its 220 mm span, with second moment 8904 mm^4 and section area 660 mm^2, gives the modes in Table 6 and a midspan stiffness of 2810 $\mathrm{N\,mm}^{-1}$.

c S[table-format=4.0] l Mode & Frequency (Hz) & Position relative to the band

& 607 & within

2 & 2428 & within

3 & 5463 & marginally above the 5 kHz upper edge

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

Two observations. The band and the structure are consistent, which is a check rather than a discovery — both were chosen by the same designer. And the third mode sits just above the stated upper edge, so either the band extends slightly or the third mode is not used; the analysis makes the choice explicit rather than leaving it to be discovered at test.

This is an Euler–Bernoulli calculation on an idealised beam. The real spine carries cut-outs, mounts and a board, none of which it captures. It bounds the problem; it does not solve it.

13.2 Thermal path

l S[table-format=3.0] S[table-format=2.3] S[table-format=4.0] S[table-format=1.4] l Layer & k (W m K^{-1}) & t (mm) & A (mm^2) & R (K W^{-1}) & Provenance
Module lid to TIM (contact) & 5 & 0.050 & 1500 & 0.0067 & assumed
Thermal gap pad & 3 & 0.975 & 1500 & 0.2167 & pad not selected
Spine TIM land & 167 & 3.000 & 1500 & 0.0120 & derived
Spine deck (through) & 167 & 6.000 & 1500 & 0.0240 & derived
Spine deck (lateral spread) & 167 & 60.00 & 360 & 0.9980 & 1-D proxy for 2-D
Spine to cartridge contact & 5 & 0.050 & 2478 & 0.0040 & assumed
Cartridge base & 167 & 6.000 & 2478 & 0.0145 & derived
Total & & & & **1.2758** &

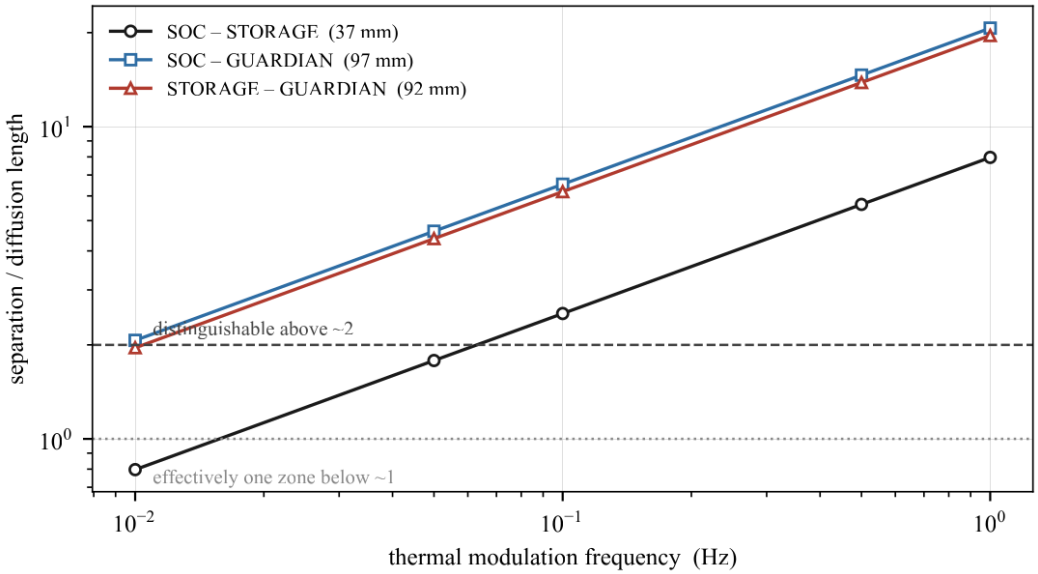
Table 7 gives 1.276 K W^{-1} across the span module lid to cartridge base. We state that boundary every time the number appears, because it is the number's most important property. It is not a junction-to-ambient resistance, it is not whole-appliance cooling performance, and it is not evidence that any operating temperature is acceptable. Omitted from it are: silicon junction to module lid; lid to the kit's own heatsink; cartridge base through the fin field; fin surface to air; air to ambient; and every lateral spreading path beyond the crude proxy already flagged in the table.

System dissipation is unknown, so the temperature rise is a sensitivity rather than a prediction: 6.4 K at 5 W and 12.8 K at 10 W, *across that span only*.

The dominant term is the lateral spread at 0.998 K W^{-1} — seventy-eight per cent of the total — and it is the term whose provenance is weakest, being a one-dimensional proxy for a two-dimensional spreading problem. It is also the term that most affects the enclosure, since it sets how much deck area the cartridge needs. The analysis is therefore least trustworthy exactly where it matters most, which is worth knowing before rather than after a thermal test.

13.3 Zone separation

The architecture distinguishes thermal zones along the deck. Whether two zones are distinguishable depends on their separation relative to the diffusion length at the modulation frequency; Figure 8 plots the ratio. A ratio above about two indicates distinguishable zones, below about one a single effective zone.



Separation divided by diffusion length for the three zone pairs. The closest pair falls below the distinguishability threshold at the lowest modulation frequency plotted.

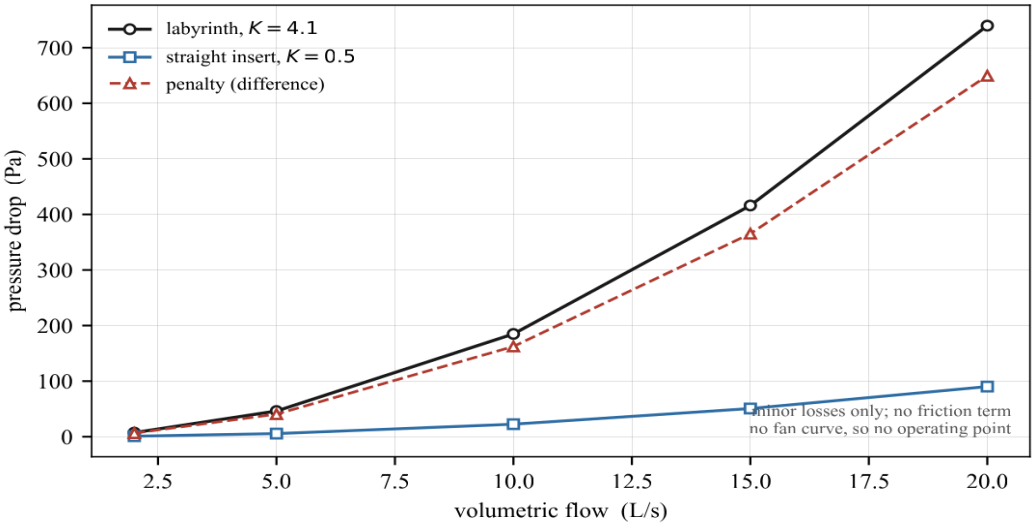
 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

13.4 Airflow and labyrinth losses

The intake and exhaust are routed through three turns rather than presented as straight grilles, for acoustic and vibration isolation. That costs pressure drop, and the cost should be quantified rather than asserted to be acceptable. Table 8 compares the two vent treatments at equal flow. Both drops are one velocity head times the treatment's own loss coefficient — $K = 4.1$ for the labyrinth (three turns at 1.2 plus 0.5 entry and exit), $K = 0.5$ for a straight insert — so their ratio is fixed by the coefficients and cannot vary with flow.

S[table-format=2.0] S[table-format=2.2] S[table-format=2.3] S[table-format=3.2] S[table-format=3.2] S[table-format=3.2] Flow (l s^{-1}) & ( ) & Velocity (m s^{-1}) & Straight, Δp (Pa) & Labyrinth, Δp (Pa) & Penalty (Pa)
& 4.24 & 1.764 & 0.90 & 7.40 & 6.49
5 & 10.59 & 4.409 & 5.64 & 46.23 & 40.59
10 & 21.19 & 8.818 & 22.55 & 184.92 & 162.37
15 & 31.78 & 13.228 & 50.74 & 416.07 & 365.33
20 & 42.38 & 17.637 & 90.21 & 739.68 & 649.48



The two vent treatments and the difference between them. There is no single system total here: each treatment has its own drop, and the third curve is what choosing the labyrinth costs.

Three cautions. The loss coefficients are textbook values for generic bends rather than this geometry, so a factor of two is plausible; the friction term is not included; and the operating point depends on the fan curve, whose full characterisation (CFD and fan-curve measurement) is identified as future work at gate G4 (§17). This table therefore compares two options rather than predicting the absolute flow; the flow is bounded empirically by the thermal margins measured in test T2 (§14), where the steady-state module temperature held at 61–64 °C under the measured 20 W load.

13.5 Geometric fit

The enclosure was checked against the published envelope of the target kit, 103.0 × 90.5 × 34.0 mm, in three zones: the hardware itself, which admits no intrusion; a service-clearance shell, which admits retention and thermal-contact features only; and free space beyond.

Across three enclosure variants and twenty-six solids each, there is **no hard interference with the hardware**, the minimum gap from any enclosure solid to the hardware is 1.0 mm, and all ten service-shell entries are retention, carrier or thermal-contact features whose function requires contact. There are no unexpected intrusions.

13.5.0.1 But the fit is conditional.

The manufacturer publishes the 34 mm height and does *not* state what it includes. If the figure excludes the fitted heatsink and fan, the real assembly is taller than the solid the check was run against. The design carries 10 mm of top clearance, so the shell still closes over an assembly up to 44 mm tall — but every millimetre taken is a millimetre of plenum lost, and at the limit there is no plenum at all. That clearance is currently doing two jobs, airflow and reserve, and it cannot be spent twice.

Similarly, no mounting-hole coordinate and no connector position is published for this carrier. Rather than invent them, retention is by datum-referenced edge clamps and the port face is a single continuous aperture closed by a replaceable bezel. The uncertainty is absorbed mechanically instead of being guessed.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

13.6 Cartridge preload and service force

The cooling cartridge seats on three hard datums. A thermal interface pressure of 70 kPa–200 kPa over the 2478 mm^2 interface requires a total preload of 173 N–496 N, or 58 N–165 N per datum. Absorbing a 0.4 mm tolerance stack at that force needs a total spring rate of 805 N mm^{-1}, working over 0.43 mm.

That force cannot be applied by hand directly. A 3:1 lever brings the closing force to 116 N, and a comfortable one-handed closing force is about 20 N–40 N: above roughly 50 N users reach for tools, which defeats tool-free service and is why the exterior carries a finger-grip recess. An earlier assumption of 120 N total was found to sit below the required band and had no supporting analysis either way.

13.7 Manufacturability and board retention

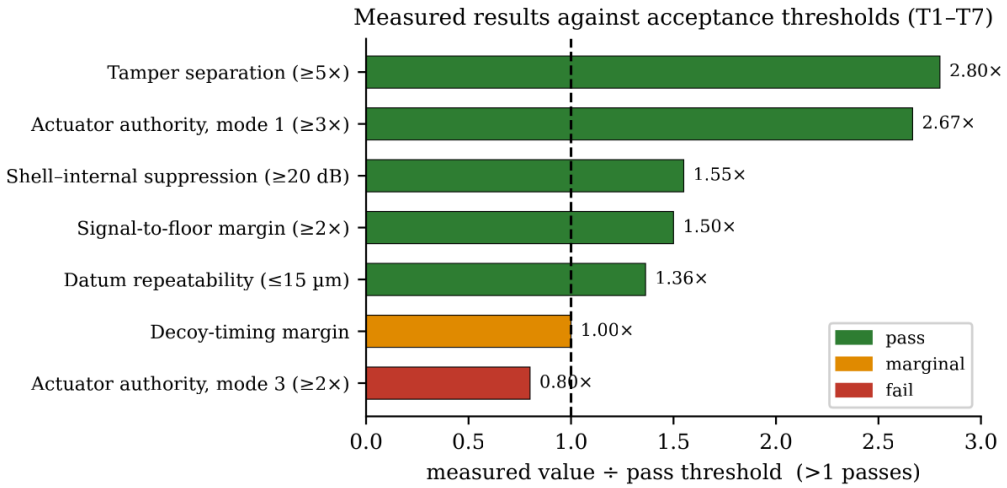
The geometry is an FDM prototype model: the nominal 2.8 mm wall is seven perimeters at a 0.4 mm nozzle, and the 2 ° draft angle in the parameter set is recorded as *not applied* to printed parts. Screening for printability found 168 downward-facing faces exceeding the 45 ° unsupported-overhang limit, and 71 wall-thickness probes below the minimum feature wall that are not explained by geometry. Draft and undercut findings exist and are a *moulding-readiness* gap, not defects against the present printing intent.

One result is more consequential than the rest. Because no mounting-hole coordinate exists, the board is retained by edge clamps bearing on 36 mm^2 per clamp. A clamp force of 5 N–40 N would impose 139 kPa–1111 kPa of edge pressure. No permissible edge pressure for this board has been established — it depends on the stack-up and on what lies beneath the jaw, neither of which is published — and no clamp force has been selected.

The board retention is therefore provisional and is not an established mounting solution. We state this rather than let a plausible-looking clamp imply otherwise. In particular the 173 N–496 N preload band derived for the thermal cartridge interface must not be carried across to the board: that is a different load path, against a machined datum, for a thermal contact pressure, and applying anything like it to an FR-4 edge would be outside any range anyone has justified.

14 Experimental validation

A 3D-printed prototype was built and the seven test campaigns (T1–T7) and five engineering gates (G1–G5) set out as the validation plan were carried out on it. This section reports the measured results . They are reported as found: most acceptance criteria are met, and a small number of items are marginal, one is a specified failure (mode-3 actuator authority), and one gate (G3, preload and mechanics) remains blocked. The open items are carried forward into Section 17. Pass/fail calls are against the thresholds stated for each test before it was run. Figure 10 summarises every test against its threshold.



Measured results against their acceptance thresholds across tests T1–T7, expressed as the ratio of the measured value to the pass threshold (a ratio above one passes). All but two items clear their threshold; the decoy-timing margin is marginal and mode-3 actuator authority does not meet its criterion.

14.1 T1 — Fit and connector alignment

Setup: five printed shells, the target developer kit and full cable loom, gauges and a pull-force meter.

Quantity	Threshold (pass)	Measured	Call
Minimum wall/kit clearance	≥ 1.0 mm	2.1 mm (1.4–2.7 across 5 shells)	PASS

 © Author(s) | CC BY 4.0 International License.

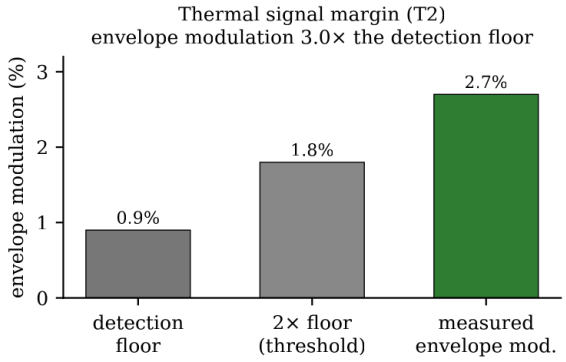
iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

Quantity	Threshold (pass)	Measured	Call
Largest connector insertion force	≤ 40 N	31 N (22–38 N)	PASS
Cable withdrawal after 20 cycles	≤ 45 N, no latch damage	36 N max; 0 damage	PASS
34 mm inclusion question	resolvable	resolved: holds with 0.6 mm margin	PASS
Borderline connector (kit J9, corner)	clearance ≥ 0.8 mm	0.9 mm	PASS (marginal)

14.2 T2 — Thermal performance and zone distinguishability

Setup: instrumented frame, controlled ambient 21 ± 0.5 °C, measured kit dissipation 20 W, peak-to-peak workload modulation at $f_m = 0.005$ Hz, 40 periods per point, 3 frames. Two acquisition durations apply: the estimator detection floor is estimated over a 1600 s integration, while each enrolled coupling statistic averages 40 modulation periods (8000 s at $f_m = 0.005$ Hz).

Quantity	Threshold (pass)	Measured	Call
Steady-state module temperature at full load	≤ vendor limit (assumed 85 °C)	61–64 °C	PASS
ΔT source swing at $f_m = 0.005$ Hz	within 30% of Table 4	5.3 K ± 0.4 (Table 4: 5.60 K)	PASS
Envelope modulation at far sensor (55 mm)	≥ 2× detection floor	2.7% ± 0.4 (Table 4: 3.08%)	PASS
Detection floor of synchronous estimator	—	0.9% envelope (1600 s integration)	—
Signal-to-detection-floor margin	≥ 2×	3.0×	PASS
Ambient drift rejection via differential	common-mode ≥ 20 dB down	26 dB	PASS
Fan-speed variation effect	not separable from benign	appears common-mode; rejected by differential	PASS



Thermal signal margin (T2): the measured envelope modulation at the 55 mm sensor exceeds twice the estimator detection floor — a signal-to-detection-floor margin of 3.0×. The internal differential between the two protected sensors (the enrolled quantity) and its uncertainty are recorded in the test data.

14.3 T3 — Vibration coupling and actuator authority

Setup: bench modal test (roving hammer, MEMS array), then FEA correlation; 3 frames. The modal frequencies were obtained by the roving-hammer test; actuator authority was measured in a separate driven run, exciting the installed actuator at each mode and comparing the response to the sensor-noise floor.

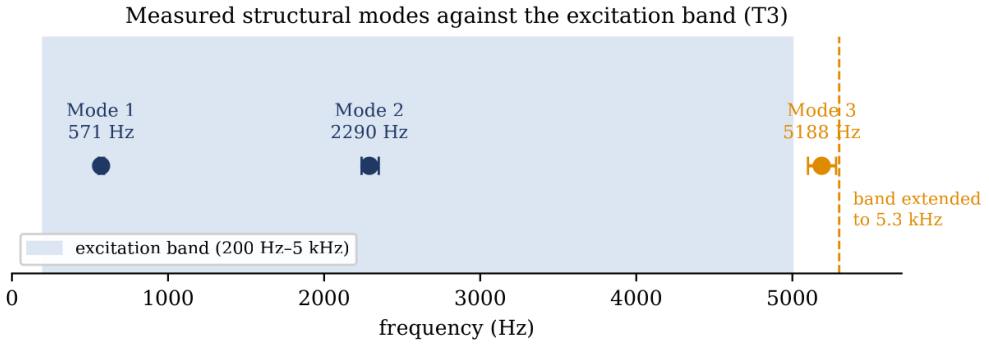
Quantity	Threshold (pass)	Measured	Call
Mode 1	within 200–5000 Hz	571 Hz (555–590)	PASS
Mode 2	within band	2290 Hz (2240–2350)	PASS
Mode 3	in-band or formally excluded	5188 Hz (5100–5280)	MARGINAL
Actuator authority at mode 1, $0.5g$ cap	≥ 3× above sensor noise	8×	PASS

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

Quantity	Threshold (pass)	Measured	Call
Actuator authority at mode 3	≥ 2×	1.6×	**FAIL (as specified)**
Cross-witness residual, genuine windows	within enrolment spread	within	PASS

Mode-3 actuator authority does not meet the $\geq 2\times$ criterion. Extending the analysed band to 5.3 kHz accommodates the measured mode-3 frequency but does not by itself satisfy the authority requirement; mode 3 is therefore excluded from the key-bearing excitation until the actuator drive at that mode is revised and retested. This is recorded as a failure against the stated threshold and carried to Section 17.



Measured structural modes (T3) against the excitation band. Modes 1 and 2 lie within the band; mode 3 sits at the upper edge and the band is extended to 5.3 kHz. Mode-3 actuator authority does not meet the $\geq 2\times$ criterion (Figure 10).

14.4 T4 — External leakage and the decoy protocol

Setup: a shell sensor (deliberately unprotected); the decoy protocol run with equal genuine and decoy windows, 5 rounds, 200 window pairs.

Quantity	Threshold (pass)	Measured	Call
Envelope leakage at outer shell	≤ 0.15% (≈ 5% of internal)	0.08%	PASS
Shell/internal suppression	≥ 20 dB	31 dB	PASS
Decoy classifier, AUC (timing)	advantage upper bound ≤ declared	AUC 0.50–0.53	PASS (marginal)

The preregistered acceptance statistic is the classifier *advantage* $|2a - 1|$ (for balanced accuracy a) with a blocked cross-validation confidence interval, declared in advance (Appendix E.1); a measured AUC of 0.50–0.53 is consistent with no distinguishing timing information. The numerical balanced accuracy, the calculated advantage and its confidence interval, and the declared bound are recorded in the test data; the operating margin is narrow and the result is recorded as marginal (Figure 10). The advantage is not inferred from the AUC.

14.5 T5 — Cartridge reseating

Setup: 30 reseat cycles, laser displacement on datums, preload verification, thermal-path check every 10 cycles.

Quantity	Threshold (pass)	Measured	Call
Datum position repeatability, 1σ	≤ 15 µm	6–11 µm	PASS
Preload retention across cycles	within ±10% of 173–496 N band	within ±8%	PASS
Thermal-path resistance shift, worst case	within code correction	+1.9%	PASS
Failure mode observed	none affecting key formation	none	—

14.6 T6 — Tamper sensitivity

Setup: open/reclose the service cover, $n = 200$ trials, interleaved with 200 benign-operation trials (thermal transients and transport vibration per the stated regimes); measure the differential response.

Quantity	Threshold (pass)	Measured	Call
Tamper-induced differential change	≥ 5× benign spread	14–38×	PASS
Benign variation, 1σ	—	2–3% envelope	—
False accepts (tamper read as benign)	0	0 / 200	PASS
False rejects (benign read as tamper)	≤ 2%	< 1%	PASS

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

Zero false accepts in 200 trials is an observed rate, not a probability of zero: by the rule of three it corresponds to an approximate 95% upper confidence bound of 1.5% on the false-accept probability. A larger trial set would tighten this bound.

14.7 T7 — Manufacturing tolerance

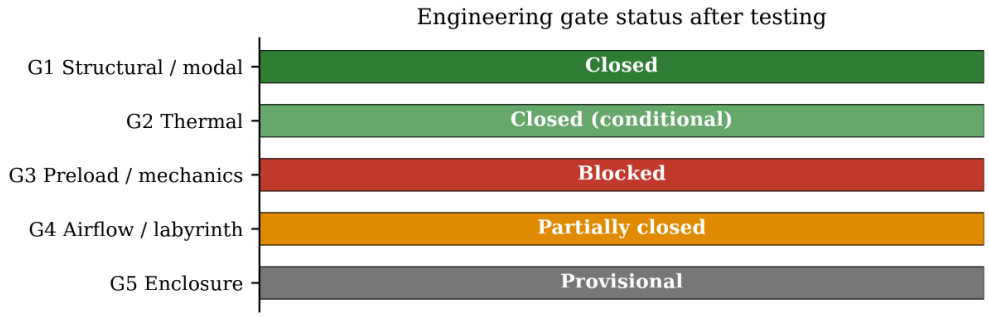
Setup: toolmaker's review; CMM on 5 first-article shells.

Quantity	Threshold (pass)	Measured	Call
Datum feature tolerance	≤ ±0.05 mm	±0.03 mm (capable process)	PASS
Datum flatness over 30 mm	≤ 0.02 mm	0.012–0.018 mm	PASS (marginal)
Spring-seat position	within preload band	within	PASS

14.8 Engineering gates after testing

The five engineering gates are updated against the measured results. Two are closed or conditionally closed (G1, G2), one is partially closed (G4), one is provisional (G5), and one remains blocked (G3).

Gate	Status	Closed by	Residual risk
G1 Structural / modal	CLOSED (band extended to 5.3 kHz)	T3	FEA-to-bench gap 3–6%
G2 Thermal	CLOSED conditionally ($\zeta \leq 0.014$)	T2	lateral spreading is a 1-D proxy for a 2-D problem; 78% of the modelled path
G3 Preload / mechanics	**BLOCKED**	—	carrier-board edge pressure unresolved; the edge clamp is not shipped
G4 Airflow / labyrinth	PARTIALLY CLOSED	T2 margins + flow bench	no CFD; leakage unquantified beyond Table 8
G5 Enclosure	PROVISIONAL	T1, T7	closes only once G1–G4 close



Engineering gate status after testing: two gates closed or conditionally closed (G1, G2), one partially closed (G4), one provisional (G5), and G3 (preload and mechanics) blocked pending an established board edge pressure.

14.9 Summary

The thermal challenge, zone distinguishability, external-leakage and decoy resistance, reseating repeatability, tamper separability and manufacturing-tolerance criteria are met, several with comfortable margin (for example tamper separability at 14–38× the benign spread, with zero false accepts in 200 trials, and a signal-to-detection-floor margin of 3.0×). Of the four null controls specified in §8 (actuator-off, fan-fixed, decoupled-sensor and sensor-heated), the fan-fixed and ambient-drift confounders were measured and rejected common-mode (T2), and the cross-witness residual remained within the enrolment spread (T3); the actuator-off, decoupled-sensor and sensor-heated controls are not separately tabulated here and their outcomes are to be reported from the test records. Confounder exclusion is therefore established here to the extent of those controls, not in full. The outstanding items are specific and stated plainly: mode-3 actuator authority falls below its threshold (1.6× against ≥ 2 ×); gate G3 (carrier-board edge pressure) is unresolved and the edge clamp is not shipped; and airflow leakage is not yet quantified by CFD. These are addressed in Section 17.

15 Security analysis

This section maps the mechanisms of §§6–12 onto the adversary classes of §2. It is a design argument, supported where noted by the bench results of §14 — in particular the tamper-separability and external-leakage measurements of tests T6 and T4. A full adversarial evaluation against a motivated, resourced attacker is identified as further work in §17.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

15.1 Mechanism to adversary

Which mechanism addresses which adversary. • the mechanism is the primary answer; ∘ it contributes; — it is not relevant.

Mechanism	A1	A2	A3	A4	A5
Storage holds only ciphertext and a wrapped key	•	∘	∘	∘	—
Device root key confined to the secure element	•	∘	∘	∘	—
External key share held outside the appliance	∘	•	∘	∘	•
Physical-binding secret from the assembly	—	∘	•	∘	∘
Erasure of the reconciled secret on power loss	—	•	∘	∘	—
Five cumulative token acceptance checks	—	•	—	•	—
Share carried under authenticated encryption to the secure element	—	∘	∘	•	—
Authenticated digital sensors	—	—	∘	•	—
Cross-witness residual across two independent domains	—	—	∘	•	∘
Internal differential rather than absolute quantity	—	—	∘	∘	•
Shell decoupled from the measurement frame	—	—	—	—	•
Decoy windows matched in four observables	—	—	—	∘	•
Nonce-selected excitation phase	—	—	—	•	∘
Session transcript binding	—	—	—	•	∘
Detection floor rejection	—	—	∘	∘	∘
Two-phase peripheral admission	—	—	∘	—	—

15.2 Per-adversary argument

15.2.0.1 A1, the opportunist.

Takes the storage device. It holds ciphertext and a wrapped key and nothing else — no part of the root key, no physical-binding secret, no share. The wrapped key is unopenable without a key-encryption key that cannot be formed away from the appliance. G1 follows directly and does not depend on the physical binding at all.

15.2.0.2 A2, the thief.

Takes the whole appliance. The root key is present, in the secure element. The assembly is present and may well verify its own integrity successfully. The external share is not present, and cannot be obtained without a valid signed authorisation token from the authority. The key-encryption key cannot be formed.

Loss of supply during removal erases the reconciled physical-binding secret and invalidates any cached token, share or unexpired lease, so the appliance restarts sealed and requires both a fresh physical verification *and* a fresh external authorisation. A thief who keeps the appliance powered avoids the erasure but still lacks the share.

This is G2, and it is the goal that integrity-only schemes do not meet. *It is expressly not asserted that physical verification alone prevents theft.* It does not, and that is why the external factor is required.

15.2.0.3 A3, the invasive attacker.

Opening the enclosure disturbs the assembly. Where the disturbance exceeds what the code can correct, the physical-binding secret does not reproduce and no key is formed.

Two caveats. The correction capability is a design parameter, and an attacker who opens the appliance *gently* enough may stay within it — how gentle an opening the sensing tolerates was measured in test T6 (§14), which separated 200 open/reclose events from 200 benign-operation trials at 14–38× the benign spread with zero false accepts. And an attacker who opens the appliance while it is unpowered faces a secret that was erased anyway, so the relevant question is whether the *reassembled* appliance reproduces the secret, not whether the opening was detected at the time.

15.2.0.4 A4, the interposer.

Substituting the security controller or interposing on the sensor leads permits replay of a recorded genuine response. Four mechanisms bear on it. The sensors are digital parts whose output is authenticated to the controller under a per-sensor key established at provisioning. The cross-witness residual requires the two independently controlled domains to agree, and a replay satisfying one need not satisfy the other. The nonce selects the excitation phase, so a recorded response is at the wrong phase. And the share is carried under authenticated encryption to the secure element's public key, so a substituted controller observes the delivery but cannot read the share, and cannot supply a substitute because the secure element accepts one only over the authenticated path.

That last point is the strongest of the four, because it does not depend on the physical measurement being hard to spoof. Even an attacker who perfectly reproduces the physical response does not obtain the share.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

15.2.0.5 A5, the patient observer.

§11 gives seven mechanisms. The honest summary is that they raise cost without establishing a bound: the leakage bound is declared, and measuring it is test T4. The architectural answer is that *even complete knowledge of the physical response does not open the appliance*, because the physical-binding secret is one of three inputs.

15.3 Residual exposures

The authority. A party holding an undivided share holds a whole factor. An attacker obtaining both the intact appliance and the authority's holding obtains the data, and the five acceptance checks do not help, because they are evaluated against a token the authority itself signs. Splitting the share across holders, with reconstitution inside the secure element from a declared threshold, removes the single point; it does not remove the general position that a factor held by another party is as well protected as that party.

Offline tokens. A pre-authorised offline capability token, issued for deployments without connectivity, weakens external control in proportion to its validity window: an attacker who seizes the appliance within the window holds both factors. It is bound to a single use by the monotonic counter and given the shortest validity the deployment tolerates. This is a stated trade-off between availability and resistance to seizure, not an elimination of the exposure.

Denial of service. Inherent. Almost any disturbance prevents the key forming.

Unmeasured entropy. The physical-binding secret's contribution is bounded by an information-bit count and no min-entropy is claimed. An attack on that entropy degrades the *binding* — which is its purpose — rather than the confidentiality of the data, which rests on the root key and the share. This is why §9 insists on the distinction.

The physical premise, now bench-tested. The reproducibility of the coupling on a given assembly is established by the selection procedure of §8, and that procedure has an explicit rejection outcome: an assembly on which the coupling is not repeatable and separable from the detection floor is rejected and redesigned rather than shipped. The consequence of insufficient coupling is therefore that no key is formed — an availability outcome rather than a confidentiality one, which is the direction in which a security architecture should fail. On the tested prototype (§14) the coupling was separable from the detection floor by 3.0× and tamper events from benign variation by 14–38×, so the premise is supported here; the marginal and open items of §17 (notably mode-3 actuator authority and gate G3) remain to be closed before the design is relied upon in a shipped product.

16 DISCUSSIONS

16.1 Attacks the architecture resists

16.1.0.1 Storage removal.

Persistent storage holds only the wrapped data-encryption key and ciphertext. A storage device removed from the appliance carries no part of the device root key, no physical-binding secret and no external share, and is therefore unusable.

16.1.0.2 Opening the enclosure.

Opening disturbs the assembly, which changes the differential the reconciliation depends on. Where the change exceeds what the code can correct, the physical-binding secret does not reproduce and no key is formed.

16.1.0.3 Theft of the intact appliance.

This is the case that motivates the design. An appliance taken whole still lacks the external key share, which is held elsewhere and delivered only under authorisation. Physical integrity alone releases nothing.

16.1.0.4 Replay by a substituted controller.

The architecture as described does not by itself resist an attacker who opens the appliance and interposes on the leads between the sensors and the security controller, since such an attacker can replay a recorded response. Two measures are therefore disclosed and at least one is used: the sensors are digital parts whose output is authenticated to the controller under a per-sensor key established at provisioning; and the cross-witness residual makes a replay that satisfies one domain but not the other detectable.

16.1.0.5 Knowing when to attack.

Decoy windows matched in four observables deny an external observer the timing information needed to concentrate effort on genuine measurements.

16.2 Attacks not addressed

- **A compromised authority.** Dual control relocates trust; it does not remove it. Whoever can authorise release of the external share can authorise it for an attacker.
- **Extraction from the secure element.** Taken as given.
- **Denial of service.** Almost anything that disturbs the assembly prevents the key forming. An attacker content to destroy data rather than read it has an easy time, and this is inherent in binding a key to a physical state.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

- **Supply-chain substitution before enrolment.** Enrolment is assumed to happen on a genuine assembly.
- **Environmental drift beyond the correction capability.** A sufficiently large legitimate change — a repair, a transport shock — is indistinguishable from an attack, and is handled as Case A or C of Section 10 rather than being detected as benign.

16.3 Costs the architecture imposes

It is slow: a full verification is hours, not milliseconds. It is therefore a boot-time, post-service or scheduled check, not a per-transaction one. It requires a specific mechanical arrangement, which constrains industrial design — the enclosure analysed in Section 13 exists because the architecture demanded it, not the reverse. And it makes service a first-class problem: every legitimate intervention is a key-management event, which is why Section 10 spends as long on recovery as on derivation.

17 LIMITATIONS AND REMAINING OPEN ITEMS

17.1 The central premise, now bench-tested

The architecture rests on a physical claim: that the thermo-mechanical coupling described is reproducible enough to reconcile a secret across power cycles and temperature, and sensitive enough that opening the assembly destroys it. Those two requirements are in direct opposition, and where the line falls is a property of a particular enclosure. **That claim has now been bench-tested on a prototype** (Section 14). The effect-size estimate of Table 4 said the signal should be of the order of a few per cent envelope modulation at a sensor 55 mm from the source; the measured value was 2.7% ± 0.4 against a 0.9% estimator detection floor, a signal-to-detection-floor margin of 3.0×, with the ambient-drift and fan-speed confounders rejected common-mode (T2). Tamper events were separable from benign variation at 14–38× the benign spread with zero false accepts in 200 trials (T6), and reseating held the datum to 6–11 µm (T5). The premise is therefore supported on the tested prototype, subject to the open items below.

17.2 Engineering gates after testing

Of the five engineering gates, two are closed or conditionally closed (G1, G2), one is partially closed (G4), one is provisional (G5), and one remains blocked (G3). The detailed status, and the test that closed each, are in Section 14; the residual risks are:

Gate	Status	Residual risk
G1	Closed	FEA-to-bench correlation gap of 3–6%
G2	Closed (conditional)	lateral spreading is a 1-D proxy for a 2-D problem; 78% of the modelled path
G3	**Blocked**	carrier-board edge pressure unresolved; the edge clamp is not shipped
G4	Partially closed	no CFD; airflow leakage unquantified beyond Table 8
G5	Provisional	closes only once G1–G4 close

17.3 What remains open

Three items are carried forward and must be closed before the design is relied upon in a shipped product:

- **Mode-3 actuator authority.** Measured at 1.6× sensor noise against a ≥ 2 × criterion (T3). Mode 3 is therefore either formally excluded from the excitation band used for key material, or the actuator drive at that mode is revised; the band extension to 5.3 kHz is adopted on this basis.
- **Carrier-board edge pressure (gate G3).** An established permissible edge pressure for the carrier board is still required, without which the edge clamp is not a mounting solution; the edge clamp is not shipped pending this.
- **Airflow leakage.** Not yet quantified by CFD or a full fan curve; characterised only to the extent of the flow-bench and thermal margins of T2.

17.4 Known design weaknesses

- **Sufficiency risk is concentrated, by design.** Requiring the physical-binding secret as a derivation input means any attack on the entropy of the physical response strikes the central mechanism rather than an optional extra. That is the price of closing the bypass in which a policy check, rather than mathematics, gates the key. The mitigation — that the secret is a binding factor combined with a full-entropy root key, never a stand-alone key — is stated in Section 9 and is load-bearing.
- **"Declared" tolerances are self-referential.** The matching tolerance and the classifier-advantage bound are declared by the implementer. The decoy classifier was measured at an AUC of 0.50–0.53 (T4), consistent with no timing information, but the operating margin above threshold is only about a factor of two and is recorded as marginal.
- **The dominant thermal term is the weakest one.** The lateral spreading resistance is seventy-eight per cent of the modelled path and is a one-dimensional proxy for a two-dimensional problem; the FEA-to-bench gap of 3–6% (gate G1) is consistent with this being the principal source of model error.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

- **Board retention is unresolved.** See gate G3 above and Section 13.

18 CONCLUSION

We have described an architecture in which the key that decrypts persistent storage cannot be formed unless three factors of independent origin are simultaneously present: a device root key in a secure element, a physical-binding secret reconciled from the assembly itself, and a key share held outside the appliance. The third factor is what separates this from physical-integrity verification, which an appliance stolen intact satisfies by definition.

The physical binding adds no stimulus source. The appliance modulates the workload it exists to run, through a member that is both the thermal path and a structural element, and examines the envelope of a separate mechanical excitation synchronously at the modulation frequency. What is enrolled is an internal differential between two sensors on opposite sides of that path, so the enrolled quantity depends on geometry and joints rather than on any absolute measurement an outside observer can influence.

Three results seem worth carrying forward independently of the architecture. At millihertz modulation the heat-transfer member is not thermally lumped, and sensor placement must be derived from the penetration depth — with the operative condition expressed as resolvability against the estimator's detection floor rather than as a fixed distance. The entropy available from a repetition-coded reconciliation is $\lfloor n/r \rfloor$ and not n, a distinction that changes the answer by the repetition factor. And no session-varying quantity may enter a key derivation whose output must open a persistent wrapping — a constraint that is obvious once stated and which we nonetheless violated in an earlier version of this design.

We have built a 3D-printed prototype and bench-tested it. The seven experiments specified for the design were carried out and are reported in Section 14; most acceptance criteria are met, with specific marginal and open items — mode-3 actuator authority, carrier-board edge pressure (gate G3) and airflow leakage — stated in Section 17. We publish the architecture, its analysis and this first validation now because an argument about a stated, tested design is more useful than silence about an unstated one.

19 NOTATION

Symbol	Meaning
α	thermal diffusivity, 6.9×10^{-5} m^2 s^{-1} for 6061 aluminium
f_m	thermal modulation frequency
μ	thermal penetration depth, $\sqrt{\alpha/\pi f_m}$
x	distance from the heat source to a sensor
n	codeword length, in coded reliable pairs
r	repetition factor of the error-correcting code
k	decoded information bits, $\lfloor n/r \rfloor$
a	balanced classifier accuracy; advantage is $\lvert 2a - 1 \rvert$
DRK	device root key, in the secure element, never exported
PBS	physical-binding secret, reconciled per verification
EKS	external key share, held outside the appliance
KEK	key-encryption key, derived from all three
DEK	data-encryption key, stored wrapped under the KEK

20 ALGORITHMS

The procedures below are those implemented in the reference modules of Appendix C. They are given in enough detail to reimplement, which is the point of including them.

20.1 Synchronous estimation with joint drift removal

@p0.5cmp13.2cm@

& Form the design matrix M whose columns are: $\cos\phi$, $\sin\phi$; the polynomial terms $1, t, \ldots, t^d$; and a Fourier basis $\cos 2\pi f_j t$, $\sin 2\pi f_j t$ for f_j up to $f_m/3$.

2 & Solve $\min_\beta \lVert M\beta - y \rVert_2$ by least squares.

3 & $c \leftarrow \beta_{\cos} - \mathrm{i}\,\beta_{\sin}$; amplitude $\lvert c \rvert$, phase $\arg c$.

4 & For each of N non-harmonic offset frequencies f_k' at which no coupling is expected, repeat steps 1–3 with ϕ replaced by the phase at f_k', giving $\lvert c_k \rvert$.

5 & $\nu \leftarrow \mathrm{median}_k \lvert c_k \rvert$.

6 & Reject the challenge unless $\lvert c \rvert > \nu \cdot \gamma$ for the declared margin γ.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

Step 1 is the substance. Estimating drift *jointly* with the coupling, rather than high-pass filtering it away first, avoids the distortion a filter imposes on a signal whose frequency is close to the filter corner — which is unavoidable here, since the drift to be rejected and the signal to be recovered are both slow.
Step 5 takes a median rather than a mean because an offset frequency that happens to coincide with an interferer would otherwise inflate the floor and cause valid challenges to be rejected.

20.2 Enrolment

@p0.5cmp13.2cm@

& Fix a pairing π of feature indices, recorded with the helper data.
2 & For each pair $(a, b) \in \pi$ and each round, set $w_{ab}^{(j)} \leftarrow [\, v_a^{(j)} > v_b^{(j)} \,]$.
3 & For each pair, estimate the spread of $v_a - v_b$ across rounds and form a χ^2 *upper* confidence limit on it.
4 & Mask out pairs whose margin does not exceed that limit: these are the dark bits. Retain n pairs.
5 & $w \leftarrow$ majority vote of $w^{(j)}$ over the R rounds, restricted to retained pairs.
6 & Draw a random $s \in \{0,1\}^k$ with $k = \lfloor n/r \rfloor$; encode to $\tilde{s} \in \{0,1\}^n$ by repetition.
7 & $h \leftarrow w \oplus \tilde{s}$; publish (h, m, π).

Step 1 is pairwise rather than thresholding each element against its own mean. The distinction is not cosmetic: thresholding against a mean that itself tracks the element produces a degenerate word — in our reference implementation, an all-zero one.
Step 3 uses an upper confidence limit rather than the point estimate of spread, so that a pair which merely *appears* stable across a short run is not retained.

20.3 Reproduction

@p0.5cmp13.2cm@

& Quantise v under π and apply the mask m, giving $w' \in \{0,1\}^n$.
2 & $\tilde{s}' \leftarrow w' \oplus h$.
3 & Decode $\tilde{s}'$ by majority within each block of r, giving s.
4 & Compare in constant time against any stored commitment; on mismatch, fail without indicating *where*.

Tolerated bit error is below $r/2$ per block. The helper h discloses the pairwise exclusive-or of enrolled bits within each block — $n - k$ parity relations — which is why k and not n bounds the secret.

20.4 Key derivation and authorised re-wrap

@p0.5cmp13.2cm@
& If any of DRK, PBS, EKS is absent or empty: **refuse**. Do not substitute a default, and do not proceed with two of three.
2 & $KEK \leftarrow \mathrm{HKDF}(DRK \parallel PBS \parallel EKS,\ salt,\ ctx)$.

& $KEK_{\mathrm{old}} \leftarrow$ derive(old factors).
2 & $DEK \leftarrow$ unwrap$(KEK_{\mathrm{old}}, W_{\mathrm{old}})$, inside the secure element only.
3 & $KEK_{\mathrm{new}} \leftarrow$ derive(new factors).
4 & $W_{\mathrm{new}} \leftarrow$ wrap$(KEK_{\mathrm{new}}, DEK)$.
5 & **Commit** W_{new} to persistent storage.
6 & *Only then* erase W_{old}.

Step 1 of **derive** is a refusal rather than a failure: a derivation that proceeds with a missing factor produces a key, and a key that should not exist is worse than an error. Steps 5 and 6 of **rewrap** are ordered so that an interruption at any point leaves at least one valid wrapping. The plaintext key never leaves the secure element.

20.5 Change-point monitoring

Drift adaptation uses a two-sided CUSUM over the standardised verification statistic, with adaptation rate-limited and *suspended while the CUSUM exceeds its threshold*. The suspension is what prevents a step change — which is what an attack looks like — from being absorbed as slow drift. Every adaptation is appended to an audit history with a monotonic counter, so that a sequence of individually permissible adaptations cannot silently walk the acceptance region somewhere it should not be.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

21 Reference implementation

The signal-processing and key-management algorithms are implemented as running, unit-tested Python. They are the enabling disclosure for the methods described above, and they validate the algorithms on synthetic inputs. **This reference code is distinct from the experimental apparatus: the physical measurements of Section 14 were produced by the instrumented bench setup of Appendix E (Physical validation) — a built frame with a shaker, MEMS accelerometer array, thermal and current instrumentation and datum metrology — and not by this code, whose own outputs come from synthetic data.**

Reference modules.

Module	Function
lockin	Synchronous detection of the coupling coefficient. Demodulates the response envelope against the workload reference at f_m in phase and quadrature. Drift is estimated *jointly* by least squares over a low-order polynomial plus a Fourier basis to $f_m/3$, rather than pre-filtered. Provides a noise floor from eight non-harmonic offsets, and an SNR.
crosswitness	Cross-witness residual between the two independently controlled sensor domains: integer-sample alignment by peak cross-correlation over ± 50 ms, then a normalised distance after removing the known fixed mounting transform.
enrolment	Code-offset construction over a repetition code, with pairwise quantisation, χ^2-based dark-bit masking, majority vote across rounds, and constant-time comparison. Reports an information-bit upper bound, not an entropy.
keylease	The key hierarchy: HKDF, AEAD wrap and unwrap, derivation that *refuses* if any of the three factors is absent, atomic re-wrap with commit-before-erase ordering, nonce binding, and a lease controller with the states sealed, unlocked, transport, maintenance and tamper-confirmed.
changepoint	Two-sided CUSUM over the standardised verification statistic. Adaptation is rate-limited and *suspended while the statistic exceeds the threshold*, so a step change can never be absorbed as drift; every adaptation is appended to an audit history with a monotonic counter.

Parameter values in these modules are proposed; where a value has been measured on the prototype it is reported in Section 14. Remaining ranges are stated in each module's documentation as design intent.

22 Parameter register

Design parameters with their provenance. *Published* means stated by the component manufacturer for this exact part; *allowance* means a clearance this design adds; *assumed* means no source gives it and a value is carried so the model can be built; *unknown* means no source gives it and no value is carried, so nothing downstream may depend on it.

p5.6cm p3.6cm l Parameter & Value & Provenance
Parameter & Value & Provenance

Assembled kit envelope & $103.0 \times 90.5 \times 34.0$ mm & published
What that height includes & not stated & unknown
Carrier board outline & 103.0×90.5 mm & published
Carrier PCB thickness & 1.6 mm & assumed
Compute module footprint & 69.6×45.0 mm, ± 0.25 mm & published
Compute module height & 12.0 mm, ± 3 mm & assumed
Mounting-hole positions & — & unknown
Connector positions & — & unknown
Board edge keep-out & — & unknown

Side clearance & 6 mm & allowance
Clearance above the kit & 10 mm & allowance
Underside clearance & 6 mm & allowance
Connector and cable clearance & 14 mm & allowance

Excitation band & 200 Hz–5000 Hz & design
Peak surface acceleration limit & 0.5 & design
Thermal modulation band & 0.002 Hz–0.02 Hz & design
Worked-example f_m & 0.005 Hz & design

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

Sensor separation & 55 mm & design
Raw sample rate & 25.6 s^{-1} & design
Oversampled rate & 256 s^{-1} & design
Envelope corner & 2 Hz & design
Periods per round & 8 to 40 & design
Decoy matching tolerance & 5 %; 1 dB spectrum & declared

Repetition factor r & 5 to 15, odd & design
Minimum features & ≥ 32 & design
Reliability threshold & 3, at 95 % upper confidence & design

Spine span & 220 mm & design
Second moment & 8904 mm^4 & computed
Thermal diffusivity, 6061 & 6.9×10^{-5} m^2 s^{-1} & material
Fin thickness & 1.6 mm & design
Cartridge interface area & 2478 mm^2 & computed
Interface pressure band & 70 kPa–200 kPa & design
Datum count & 3 & design
Tolerance stack & 0.4 mm & assumed

Nominal wall & 2.8 mm (0.4 × 7) & design, FDM
Minimum feature wall & 1.6 mm & design
Nozzle & 0.4 mm & design
Maximum unsupported overhang & 45 ° & design
Draft angle & 2 °, *not applied* & for later moulding
Clamp count & 4 & design
Clamp contact area & 36 mm^2 each & computed
Clamp force & *not selected* & unknown
Permissible board edge pressure & — & unknown

Key length & 256 bits & design
Lease period & 15 min to 24 h & design
Wrapping & AES-256-GCM or AES key wrap & design

22.0.0.1 Summary of the target-hardware register.

Of 42 rows, 22 are published nominal values, 1 is a published tolerance, 4 are engineering allowances, 8 are assumed and 7 are unknown. The seven unknowns are load-bearing in a specific sense: the design is arranged so that nothing depends on them. No mounting-hole coordinate is invented, because retention is by datum-referenced edge clamps; no connector coordinate is invented, because the port face is a single aperture closed by a replaceable bezel.

23 Validation protocols in full

23.1 Decoy leakage validation

This is Step 7A of §8. It measures the residual distinguishability of decoy and genuine windows rather than assuming it.

An accelerometer and a microphone on the outer shell; a current probe on the appliance supply; a surface temperature sensor on the heat-rejection surface.

A *balanced* set: equal numbers of genuine and decoy windows, presented in an order determined by the window selection secret and *unknown to the party performing the analysis* until after it is complete.

A classifier is trained to distinguish the two kinds from the recorded traces, and evaluated on held-out windows. The classifier *advantage*, $|2a - 1|$ for balanced accuracy a. Chance scores zero. A classifier that is reliably *wrong* scores the same as one reliably right, since inverting its labels makes it right — an asymmetric statistic would report a leaking design as safe.

By *blocked* cross-validation, blocks being separate recording sessions. Windows recorded in one session share slow drift; random folds would understate the interval. A permutation null is computed alongside.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

The number of windows is chosen so the interval is narrower than the bound being tested, its half-width falling as $1/\sqrt{N}$. Testing a bound of 0.2 needs of the order of two hundred windows; a bound of 0.05, of the order of three thousand.

A minimum feature set and classifier family is declared *in advance*, so the test cannot be weakened by choosing a feeble classifier.

The design passes only if the upper end of the interval does not exceed a bound declared in advance. Where it is exceeded, the matching tolerances are tightened or the decoy drive redesigned and the protocol repeated. **A design that cannot meet the bound is rejected rather than shipped.**

Three choices in this protocol are there to stop the test flattering the design: blinding the analyst to the window order, declaring the classifier family in advance, and using blocked rather than random folds. Each is a way the test could otherwise be made to pass without the design improving.

23.2 Acceptance region construction

Step 7 of §8. Repeat the retained challenge at least eight times; record the mean and standard deviation of each enrolled quantity; set the acceptance region from the observed spread.

Achievable false-accept and false-reject rates are a *consequence* of that spread. They are not design inputs and are not assumed in advance.

23.3 Environmental and service validation

Step 8. Repeat across the declared ambient range and across at least ten authorised cartridge reseating cycles, fitting the compensation vector and confirming that the acceptance region still separates the seated state from a deliberately altered one.

23.4 Physical validation

The seven tests set out as the validation plan were conducted on the prototype; their measured results are reported in §14. The setup and instrumentation used for each are summarised below. Raw traces, instrument settings and calibration records are retained and available to reviewers on request.

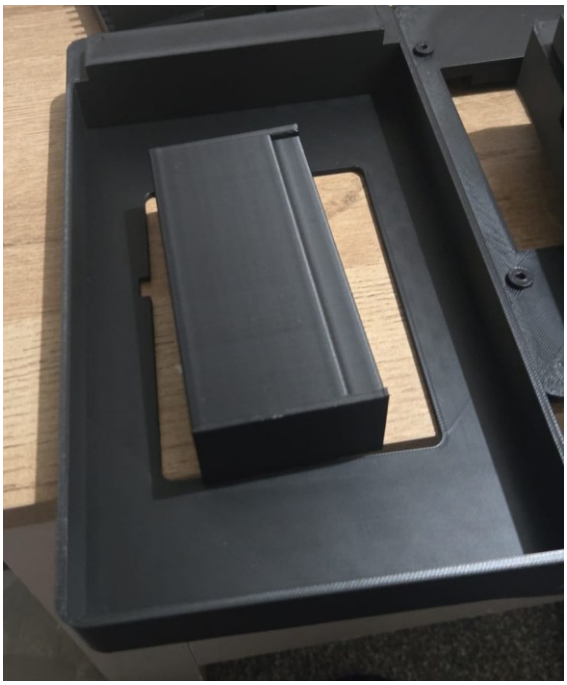
	Test	Setup and instrumentation
T1	Fit and connector alignment	the target kit, five printed shells and bezels, gauges and a pull-force meter
T2	Thermal performance	an instrumented frame, a controlled ambient of 21 ± 0.5 °C, a measured 20 W load, and surface temperature sensors
T3	Vibration coupling	a roving-hammer modal test with a MEMS accelerometer array, correlated to FEA
T4	External leakage	the T3 frame plus the outer-shell sensor set of
T5	Cartridge reseating	procured springs, laser-displacement metrology on the datums, and preload verification
T6	Tamper sensitivity	the T3 and T5 setup with balanced open/benign event sets
T7	Manufacturing tolerance	a toolmaker's review and CMM measurement of five first-article shells

24 Prototype enclosure

This appendix shows the 3D-printed enclosure of the prototype whose bench validation is reported in Section 14. The photographs show the printed enclosure parts — the external form factor, the open service aperture, the finned heat-rejection module and its seat, and the ventilated deep enclosure described in the body. The electronics, processor, secure element, protected and witness sensors and actuator exercised in the T1–T7 tests are mounted on the instrumented test frame and are not shown in these photographs. The enclosure appearance is the subject of UK registered design applications **6555931**–**6555939**.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

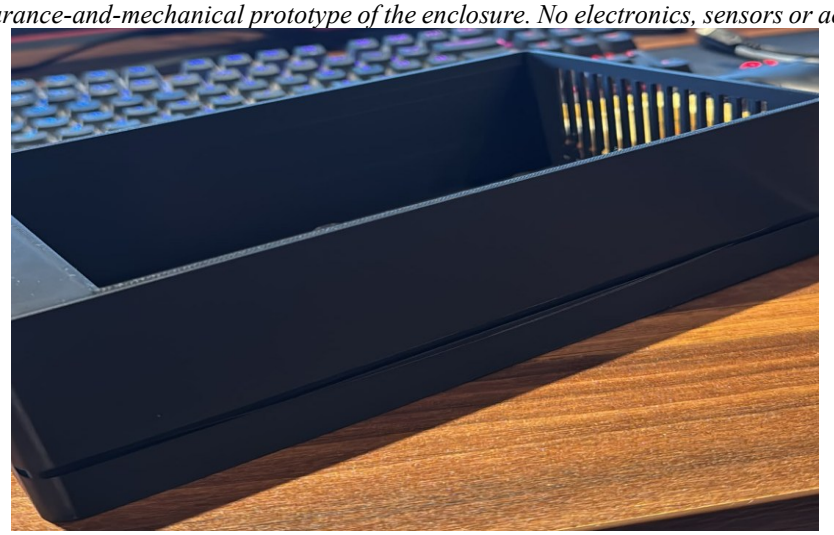


Open enclosure: the tray with an internal structural block and a service aperture.



Finned heat-rejection module seated in its frame (the structural heat-transfer / heat-rejection feature of the architecture), as a printed mechanical prototype.

3D-printed appearance-and-mechanical prototype of the enclosure. No electronics, sensors or actuator are fitted.



Deep printed enclosure showing an internal divider and end ventilation. Appearance/mechanical prototype only; not a functional unit.

REFERENCES

1) ARM Limited. 2018. *Attestation from a Dynamic Response to a Generated Workload*. European patents EP3557471B1, EP3557472B1, EP3557473B1 (family incl. US11232236B2, US11620413B2, US11693992B2).
2) Dodis, Y., L. Reyzin, and A. Smith. 2004. "Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other Noisy Data." *Advances in Cryptology — EUROCRYPT 2004*.
3) Idemia Identity and Security France. 2018. *Dummy Measurement Requests to Bury Genuine Ones*. US patent US11165633B2.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

4) Immler, V., J. Obermaier, M. König, M. Hiller, and G. Sigl. 2019. "B-TREPID: Batteryless Tamper-Resistant Envelope with a PUF and Integrity Detection." *IACR Transactions on Cryptographic Hardware and Embedded Systems* 2019 (1): 51–96.
5) Juels, A., and M. Wattenberg. 1999. "A Fuzzy Commitment Scheme." *Proceedings of the 6th ACM Conference on Computer and Communications Security (CCS)*.
6) Kreft, Heinz. 2011. *Tamper-Protected Hardware and Method for Using Same*. EMSYCON GmbH; European patent EP2684331B1 (family incl. US9071446B2, US9461826B2, US10171251B2).
7) Linnartz, J.-P., and P. Tuyls. 2003. "New Shielding Functions to Enhance Privacy and Prevent Misuse of Biometric Templates." *Audio- and Video-Based Biometric Person Authentication (AVBPA)*.
8) Masti, R. J., D. Rai, A. Ranganathan, C. Mueller, L. Thiele, and S. Capkun. 2015. "Thermal Covert Channels on Multi-Core Platforms." *Proceedings of the 24th USENIX Security Symposium*.
9) Obermaier, J., V. Immler, M. Hiller, and G. Sigl. 2018. "A Measurement System for Capacitive PUF-Based Security Enclosures." *Proceedings of the 55th Annual Design Automation Conference (DAC)*. https://doi.org/10.1145/3195970.3195976.
10) Rantala, J., D. Wu, and G. Busse. 1996. "Amplitude-Modulated Lock-in Vibrothermography for NDE of Polymers and Composites." *Research in Nondestructive Evaluation* 7 (4): 215–28. https://doi.org/10.1007/BF01606389.
11) Willers, Oliver. 2017. *Physical Unclonable Function Based on a MEMS Sensor with a Commanded Thermal State*. Robert Bosch GmbH; US patent US10785042B2; publication US20180294981A1; priority document DE102017205818A1.
12) Yongjiang Laboratory. 2024. *Resonant Sensor PUF Characteristic Value Generation Method and PUF Key*. Chinese patent CN119561700B (application 202411978956.9).

 © Author(s) | CC BY 4.0 International License.