Resolution of the P-INDEP Open Questions and the P-DELEG Principle:

A Mandate Instead of a Key for Autonomous Agents in the Open 6G Ecosystem

Submitted to: A Technical Reference Architecture Framework for an Open 6G Device Ecosystem

Version 0.91 | July 2026 | For IC Activity Review

Abstract. The P-INDEP contribution (June 2026) stated the principle of independence of trust roots by foundation and left three open questions: completeness of the coupling-factor set, independence auditing without privacy loss, and the cost of physically distinct carriers. This document closes all three questions with solutions grounded in verified engineering practice and standards: (1) a transition from a closed checklist to a generative rule over the dependency graph, with precedent-driven evolution of the taxonomy through the ACP registry (PV-3); (2) a two-phase audit - a structural independence audit at the device-class level during certification, plus anonymous attestation of the instance (DAA / property-based attestation) in operation; (3) an inventory of already-existing carriers (eSE/eSIM, a network-side root based on physical-layer authentication) and a minimum-quorum scale k bound to the ACP agent classes. In addition, the document introduces the P-DELEG principle - delegation of authority to an autonomous agent launched from the device wallet without handing the agent any trust roots: authority is conveyed by a one-time evaporating mandate signed by a live quorum of roots. P-DELEG is shown to be a direct consequence of the resolution of question 2: auditing without disclosure and delegation without disclosure are realised by one and the same cryptographic mechanism. Version 0.91 extends the solutions following a diversion analysis and an inversion pass: an experimental audit phase (independence verified by compromise injection), reference quorum profiles (anonymity through uniformity), a witness root on neighbouring mesh devices, an intent model of mandate execution, and protective qualifications for the network-side root against adversarial attacks on classifiers.

Keywords: *6G device ecosystem, trust roots, coupling factors, independence by foundation, anonymous attestation, DAA, property-based attestation, physical-layer authentication, autonomous agents, digital wallet, delegation, IEEE IC25-009-01*

1. Introduction

P-INDEP established that the cost of attacking a quorum of trust roots is determined neither by the most expensive root nor by the sum of the roots, but by the cheapest shared foundation. The principle was proposed as normative, and its Section 7 honestly listed three questions without which the principle remains a declaration: how to make sure the coupling-factor list is complete; how to audit root independence without creating a repository of root data (which would contradict P-AUTH); and what the requirement of physically distinct carriers actually costs.

After the publication of P-INDEP, a fourth question was raised in the activity discussion: how does the independence-of-foundations principle apply to AI agents launched from the device wallet? An agent is a software entity without a body, acting on the user's behalf in the user's absence. Equip the agent with wallet keys and both threats the entire series of contributions was built against are recreated at once: a stored secret (violating P-AUTH-1) and a shared carrier through which compromise of the agent reaches the roots (violating P-INDEP).

This document shows that the fourth question does not add a new problem - it adds a second, independent argument in favour of one specific resolution of question 2. All four answers are assembled from mechanisms that already exist in standards and industrial practice, and from principles already introduced by the contributions of this series (P-AUTH, ACP, P-INDEP). The document introduces no new entities - it introduces the rules for wiring them together.

2. Resolution of Question 1: A Generative Rule Instead of a Closed List

2.1. Why no list can be complete

P-INDEP enumerated seven classes of coupling factors: carrier, power, clock, registration moment, quorum mechanism, vendor, jurisdiction. Any such list is incomplete by construction - the set of possible shared foundations is open, and a rational attacker looks for factor number eight, the one the list does not contain. As long as the audit is defined as a checklist comparison, it is unfalsifiable: passing the audit does not prove independence.

2.2. The generative rule

The resolution is to change the definition. A coupling factor is defined not by enumeration but structurally: it is any node of the device's dependency graph whose single compromise reaches two or more roots counted toward one quorum. Independence auditing thereby becomes graph analysis (fault tree / attack graph) rather than checklist comparison. The seven classes from P-INDEP are retained as a first-pass heuristic - a hint at where to look for graph edges - but not as a criterion of completeness.

This is not an invention but an import of an evolution that has already happened in an adjacent industry. ISO 26262 (dependent-failure analysis, DFA) introduces coupling-factor classes precisely as checklists for different abstraction levels - system, hardware, software, and semiconductor; practice exposed the subjectivity of manual assessment, and the industry moved to automated quantitative DFA: tools analyse the design and identify shared resources as elements appearing in multiple cones of influence. The Independence-as-a-Service approach (OSDI 2014, already cited in P-INDEP as [4]) demonstrates the same for distributed systems: independence is verified by automated analysis of structural dependencies, not by declaration.

2.3. Precedent-driven evolution of the taxonomy

The openness of the factor set stops being a weakness once the taxonomy has a growth mechanism. Such a mechanism already exists in the ecosystem: the ACP registry. Every newly discovered coupling factor - found in an audit, in an incident, or in research - is published as a precedent through PV-3 in anonymised form and becomes a mandatory part of the first-pass heuristic for all subsequent audits. The taxonomy becomes a living one - on the model by which CWE/CVE evolve: the base rule is invariant, the catalogue of cases accumulates.

2.4. Inversion: audit by experiment, not by analysis alone

The dependency graph answers the question “where to look”, but it does not prove independence: an unknown coupling factor is, by definition, absent from the graph as well. Inverting the assumption “auditing is analysis” yields a second, experimental phase: instead of searching for shared foundations on paper - controlled injection of compromise into one root, with measurement of whether it reached the

others. A power glitch, a clock fault, a software compromise - applied to each quorum root in turn; any disturbance that reaches a second root is a discovered coupling factor, including one nobody knew about.

This, again, is not an invention but a transfer of established secure-hardware certification practice: Common Criteria defines a methodology for evaluating resistance to fault attacks, and fault-injection and side-channel testing by independent evaluators is a routine part of the path to EAL 5+ levels for secure elements. P-INDEP extends the subject of the test: what is verified is not the resilience of a single root, but the unreachability of neighbouring roots when a given root is compromised.

2.5. Normative requirements

Requirement R7 (generative rule). For the purposes of the P-INDEP audit, a coupling factor is defined structurally - as a node of the dependency graph whose single compromise reaches more than one quorum root. The independence audit is an analysis of the complete dependency graph of the roots; comparison against the catalogue of known factors is a mandatory but not sufficient step.

Requirement R8 (experimental verification). Device-class certification includes compromise-injection testing: a controlled disturbance (voltage, clock, software compromise, physical access) is applied to each quorum root in turn; any effect on the remaining roots is recorded. Independence is considered confirmed only upon a negative result for all roots. The methodology is aligned with Common Criteria fault-attack evaluation practice.

Requirement R9 (precedents). Discovered coupling factors not present in the current catalogue - whether found by graph analysis, by the R8 experiment, or by an operational incident - are published in the ACP registry (PV-3) in anonymised form and included in the catalogue of the next audit cycle. The absence of new precedents for 24 months is treated under the ACP liveness criterion - as a trigger for revising the audit procedure, not as evidence of catalogue completeness.

3. Resolution of Question 2: Two-Phase Audit - Knowledge About the Class, Anonymity of the Instance

3.1. The contradiction

The auditor must know the structure of the roots (otherwise independence cannot be verified) and must not know it (otherwise a centralised repository of root data is created - a new shared factor and a direct contradiction of the spirit of P-AUTH: what is collected can be seized). The contradiction is resolved by separating the requirements in time and in system level: knowledge pertains to the device class and the moment of certification; anonymity pertains to the instance and the moment of operation.

3.2. Phase 1: structural audit of the class at certification

The full dependency-graph audit (under R7-R8) is performed once - on the design of the device class during certification, on the model of design-level DFA in automotive functional safety. What is audited is the design: which roots, on which carriers, from which vendors, with which shared power and clock domains. Data about specific users and instances does not exist at this phase at all - there are none yet.

3.3. Phase 2: anonymous attestation of the instance in operation

In operation, a device instance must confirm to counterparties one single statement: “I belong to a class that has passed independence certification.” Mature, standardised mechanisms exist for this:

- **Direct Anonymous Attestation (DAA)** - adopted by the Trusted Computing Group in the TPM 2.0 specification, specified in ISO/IEC 20008; the Intel EPID 2.0 implementation is available for licensing together with an open-source SDK; open implementations for TPM 2.0 exist. A DAA signature confirms membership of a group (the certified class) without revealing which instance signed.
- **Property-Based Attestation (PBA)** - the verifier checks a high-level property of the target instead of configuration measurements, and one property may be satisfied by many configurations; PBA variants use zero-knowledge proofs of a configuration's membership in the set of valid ones. In P-INDEP terms: the device proves the property “my roots share no foundation” without revealing which roots it has.
- **A production precedent** - cross-vendor ZK attestation of hardware keys is already deployed on the web: the client sends not a signature but a proof that a signature was generated by a key from an admissible set; the verifier learns only the fact of a valid attestation, not the identity of the device.

The resulting construction: the certification body knows everything about the class and nothing about the instances; a counterparty in operation learns about the instance only the fact of its membership in a certified class; no data about the instance's roots is stored or transmitted. P-AUTH is not violated at either phase.

3.4. Inversion: anonymity through uniformity

An additional privacy layer comes from inverting the original assumption. Instead of “hide the instance's root structure” - “strip the structure of its informativeness”: the standard fixes a closed list of reference quorum profiles (one or two per agent class), and all devices of a class have strictly identical root topology. Knowledge of the topology then says nothing about the instance - the “anonymity loves company” principle on which Tor Browser's fingerprinting resistance is built. Side benefits: the set of valid configurations shrinks, which makes PBA membership proofs cheaper, and any non-standard profile becomes immediately visible - an additional barrier to the “Silent k→1” scenario from the P-INDEP diversion table.

3.5. Normative requirement

Requirement R10 (two-phase audit and reference profiles). The structural independence audit under R7-R8 is performed at the device-class level during certification. The composition and topology of the quorum roots shall conform to one of the reference profiles, the closed list of which is fixed by the standard per agent class. In operation, instance conformance is confirmed by anonymous attestation of membership in the certified class (DAA per ISO/IEC 20008, PBA, or an equivalent with a GlobalPlatform-compatible interface). Requiring an instance to disclose the composition or topology of its roots is prohibited; creating registries linking instances to root composition is prohibited.

4. Resolution of Question 3: Free Carriers and the Quorum Scale

4.1. The principle: do not add carriers - name the existing ones

The question about the cost of distinct carriers proceeded from the implicit assumption that independent carriers would have to be added to the device. A resource inventory shows the opposite: in a typical 6G device and around it there already exist at least three carriers independent by foundation, and using them as roots costs approximately zero in the bill of materials (BOM).

4.2. Four verified carriers

- **eSE/eSIM.** The dedicated secure element is already a separate die in every modern smartphone, with its own vendor and its own certification; modern eSEs protect key material with a hardware PUF and are FIDO-compatible. Using the eSE as a second root adds nothing to the BOM.
- **PUF - with a mandatory qualification.** The technology is mature: PUF IP is deployed in more than a billion devices with certifications at the level of EMVCo, Visa, CC EAL6+, PSA. However, a PUF on the same SoC as another root closes the stored-secret threat but does not provide carrier independence - a precise illustration of P-INDEP's central thesis (difference in nature is not difference in foundation). A PUF counts toward the quorum only on a physically separate die - for example, inside the eSE.
- **Network-side (off-device) root - with a protective qualification.** Physical-layer authentication (PLA) is an active direction in 6G security: location-based authentication relies on RSS, frequency-offset and CSI measurements at 6G's targeted centimetre-level localisation accuracy. The infrastructure for this root - the sub-centimetre positioning of the reference architecture's service level (ToA/AoA from multiple relays) - is already designed into the project. Qualification following the diversion analysis: RF fingerprints are hard to clone physically, but the DL classifiers that recognise them are vulnerable to adversarial attacks (targeted PGD attacks reach 98.9% success; universal perturbations, UAP, cause 80.5% misclassification under grey/black-box conditions; collusion-driven impersonation attacks via surrogate models are documented). The network-side root therefore rests on geometry, not on recognition: consistent measurements from no fewer than three independent nodes; the RF fingerprint and the DL classifier are auxiliary signals, not sole arbiters. Lightweight rogue-device detection based on softmax posterior probabilities and adversarial hardening of classifiers apply in addition.
- **A witness root on neighbouring devices (inversion of ownership).** The most independent carrier is the one the device owner does not own: a quorum of neighbouring mesh devices signs the event as a witness. The witnesses' owner, vendor, power, clock, and firmware are all foreign - a coupling factor with the device is absent by construction. The mechanisms are mature: collective attestation schemes (SANA and relatives) provide publicly verifiable aggregate attestation of large device networks in which full compromise of one device, including its hardware and keys, does not affect the others; in mesh schemes, physical capture of a node is detected collectively - by the node's absence from contact with its neighbours (detection through the absence of a trace - the pheromone-evaporation logic working as an alarm). The cost of such a root is negative: witnessing is one more service of network participation alongside relaying, built into the neighbourhood-hotspot economics of the ICAID.

4.3. The quorum scale by ACP class

The quantitative part of the question ("trade-off per device class") is resolved by binding the minimum k to the ACP agent taxonomy (Principle A-1): the cost and strictness of the quorum grow only together with the stakes - in agreement with the irreversibility-threshold logic of ACP. Table 1 sets the scale and, simultaneously, the P-DELEG mandate width (Section 5).

Table 1. Quorum scale k and mandate width by agent class

Agent class (ACP A-1)	Minimum quorum k	Root composition (free carriers)	P-DELEG mandate width
Executor	k = 1	eSE/eSIM (separate die; PUF inside the eSE is acceptable)	Narrow: fixed list of operations, small limits, short TTL
Adapter	k = 2	eSE/eSIM + network-side (off-device) root: PLA / geometric verification; BOM increase ~ 0	Medium: declared choice space, value and rate limits, TTL of hours to a day
Strategist	k = 3	eSE/eSIM + network-side root + live user presence (P-AUTH)	Wide, but above the irreversibility threshold - only a fresh presence confirmation (human-in-the-loop)

The BOM increase in moving from k=1 to k=2 is close to zero (both carriers already exist); moving to k=3 adds no hardware at all - the third root is the user's live presence under P-AUTH. The requirement of physically distinct carriers thus ceases to be a cost item and becomes a rule for naming resources already paid for.

4.4. Normative requirement

Requirement R11 (quorum scale). The minimum quorum k is bound to the action class under the ACP A-1 taxonomy in accordance with Table 1 and is fixed by the standard, not by the vendor (extending Requirement 3 of P-INDEP). The network-side root is recognised as a full quorum root only under geometric verification by consistent measurements of no fewer than three independent nodes; RF fingerprints and DL classifiers are used as auxiliary signals. A quorum of neighbouring witnesses is recognised as an equivalent root for the Adapter and Strategist classes given no fewer than three witnesses from independent owners. A PUF counts as a root only on a carrier physically separate from the other quorum roots.

5. The P-DELEG Principle: Delegation to an Agent Without Handing Over Roots

5.1. The problem

An autonomous agent launched from the wallet of a 6G device must act on the user's behalf - including when the user is not present. The naive solution - issuing the agent a key or an access token - recreates both threats eliminated by this series of contributions: a stored secret appears whose extraction opens assets (violating P-AUTH-1), and a carrier shared by the agent and the roots turns the agent into the cheapest coupling factor of the system (violating P-INDEP). The question is stated in P-INDEP terms: where must the agent's trust sit so that breaking the agent opens nothing but the agent itself?

5.2. The principle

The answer: trust does not move anywhere. It stays in the roots. The agent does not join the quorum, receives no access to the roots, and carries no authority - it carries a trace of authority: a mandate. A mandate is a one-time commission signed by a live quorum, with explicit bounds: agent class, list of admissible actions, limits, time-to-live (TTL). A mandate is not a secret: its disclosure or theft gives an attacker exactly the remainder of one narrow mandate - and nothing more. When the TTL expires, the mandate evaporates and the agent stops by itself - without a revocation command, without a revocation centre, without a certificate revocation list. Renewal of the mandate is a new quorum event; for the Strategist class - a new user-presence event.

A concrete cryptographic carrier for the mandate exists: macaroons - credentials with caveats built as a chain of HMACs, where adding a caveat is easy and removing one is impossible; the holder can only narrow the authority before passing it on, confining when, where, by whom and for what the credential may be used. Two protections identified by the diversion analysis are built into the technology itself. The first is holder binding (proof-of-possession): the mandate carries the caveat “executable only by a machine holding key K”, where K is the agent's ephemeral key, generated at launch and living only for the session; a stolen mandate without the key is useless. The ephemeral key is not a stored secret in the P-AUTH sense: it opens nothing after the TTL expires and cannot be recovered. The second is a mathematical guarantee of narrowing upon sub-delegation: the HMAC chain makes a sub-agent's mandate strictly no wider than its parent's by construction.

The evaporation logic is not new to this series: it is the same logic by which the Path Quality Tag (PQT) carries TTL-Q, and by which a colony's pheromone trail disappears unless reinforced. Trust in a centre-less ecosystem is everywhere shaped the same way - as a trace with a lifetime, not as a stored artefact.

5.3. The link to the resolution of Question 2

A counterparty accepting an agent's action must verify the agent's right to act - without learning the user, the device, or the composition of the roots. This is exactly the same task as the independence audit of an instance, and it is solved by the same mechanism: the mandate is signed via DAA/PBA of the certified class. The verifier learns: the mandate was issued by a live quorum of a device of a certified independence class, the mandate is in force, the action is within the mandate's bounds - and nothing beyond that. Auditing without disclosure and delegation without disclosure are one cryptographic core. This is precisely why the agent question required no new architecture: it confirmed the one chosen.

5.4. Normative principles

Principle P-DELEG-1 (a mandate instead of a key). The agent does not join the quorum of roots and receives no access to the trust roots. Authority is conveyed to the agent exclusively in the form of a mandate - an attenuable commission (on the macaroon model) signed by a live quorum, with explicit bounds on actions, limits and lifetime, bound to the agent's ephemeral key (proof-of-possession). The mandate is not a secret in the P-AUTH-1 sense: its compromise without the ephemeral key opens nothing, and together with the key - no more than the remainder of the mandate itself. Upon sub-delegation, authority can only narrow, which is guaranteed by the construction of the caveat chain. Rationale: trust placed in the roots must not change its seat upon delegation - otherwise delegation itself becomes the cheapest shared foundation.

Principle P-DELEG-2 (evaporation). Every mandate carries a TTL commensurate with the irreversibility of the permitted actions. TTL expiry stops the agent's authority automatically, without a revocation centre. Renewal is a new quorum event; the standard fixes a budget of automatic renewals (a maximum count or total duration) without a new full quorum event, and every renewal is recorded in the ACP accountability chain - the “convenience eats the quorum” drift becomes visible through the PV layer. For Strategist-class agents and for actions above the irreversibility threshold, renewal is a new verified user-presence event (P-AUTH). Rationale: a stopping mechanism that requires a centre would recreate the centre; a stopping mechanism built into time requires no one.

Principle P-DELEG-3 (anonymous verifiability). The mandate is signed by the class anonymous-attestation mechanism (DAA/PBA per R10). The accepting party verifies the validity of the mandate and the issuing quorum's membership in a certified independence class, receiving no identifiers of the user, the device, or the root composition. The mandate width may not exceed the width of the agent's accountability contract under its ACP A-1 class; the agent's Reachable Anchor (A-3) is inherited from the user whose presence participated in the issuing quorum. Rationale: the chain of accountability for bodiless actions must close on a body - at the point of mandate issuance, not at the point of action.

Principle P-DELEG-4 (execution by the environment). The agent is deprived of the ability to execute - it publishes an intent. Execution is performed by independent nodes of the environment, each of which, before executing, checks the intent against the mandate: action bounds, limits, validity period, quorum signature, binding to the agent's key. A compromised agent can only propose - there is no one to execute the inadmissible. Rationale: the model is operationally proven by smart accounts (intent architectures and session keys: the agent proposes, the policy engine and the wallet approve and execute within hard limits; the agent never receives key material); P-DELEG-4 lifts it from the application level to the level of the ecosystem standard.

5.5. What an attacker gains and does not gain

Compromise of an agent under the P-DELEG model gives the attacker: the ability to propose actions within the bounds of one mandate until its TTL expires, conditional on possession of the agent's ephemeral key, with every intent checked by the executing nodes and continuously recorded in the ACP accountability chain. Compromise of the agent does not give: the roots, the wallet, other mandates, the ability to renew, the ability to widen authority, the ability to execute anything outside the mandate. The answer to the question that opened this section coincides literally with the P-INDEP formula: breaking one place does not open the rest - because in this place nothing but a trace is kept.

6. Coherence of the Resolutions With One Another and With the Series

The four resolutions are not independent - they rest on three shared mechanisms, and this is their test of architectural honesty:

- One cryptographic core. DAA/PBA serves both the instance audit (R10) and the mandate signature (P-DELEG-3). Two different questions are closed by one mechanism - a sign that the mechanism was chosen at the right level.
- One evaporation logic. The TTL of the PQT tag, the evaporation of the pheromone trail, and the TTL of the mandate (P-DELEG-2) are one and the same architectural move: memory and authority are kept in time, not in a centre.

– One scale of stakes. The ACP A-1 classes set the minimum quorum k (R11), the mandate width (P-DELEG-3), and the human-in-the-loop threshold. The cost of trust grows only together with the irreversibility of the action.

No contradictions arise with the earlier contributions: P-AUTH is inherited unchanged (no phase creates a stored secret); ACP gains the missing mechanism for conveying authority from a body to an agent; P-INDEP has its open questions closed by means that themselves satisfy P-INDEP (no resolution introduces a new shared foundation or a new centre).

Compatibility with external standards. P-DELEG does not compete with the emerging IETF stack for agent identity (the AIMS family of drafts composing SPIFFE, WIMSE and OAuth 2.0): that stack requires short-lived credentials with an explicit expiry and treats them as an alternative to explicit revocation mechanisms, declaring static keys an anti-pattern - which coincides with P-DELEG-2. P-DELEG supplies the missing hardware layer underneath that stack: the source of issuance (a quorum of independent roots) and the anonymous verifiability of issuance. Likewise, P-DELEG-4 is compatible with the production practice of session keys and intent architectures in smart accounts, lifting their principles to the level of the ecosystem standard.

Meta-inversion. All the resolutions of this package add up to one shift of goal: not “make compromise impossible” but “make compromise useless and visible”. A root falls - the R8 experiment has already rehearsed it; the structure is known - the reference profiles stripped it of informativeness; the device is captured - the neighbours noticed the absence of a trace; the agent is hacked - it can only ask. This is the same philosophy by which ACP defined its own task: to make the absence of accountability visible and costly. The series is self-consistent at this level as well.

7. Open Questions

By the rule of this series - a principle that cannot survive attack should not be standardised - the document is obliged to name the questions it opens. Two questions of version 0.9 have received confirmed candidates and move to the status “requires specification”: early stopping of an agent is resolved by short TTLs with budgeted auto-renewal (stopping = ceasing to renew) - a scheme confirmed by the practice of short-lived credentials, where explicit revocation mechanisms are recognised as replaceable by short lifetimes; mandate composition is resolved by the macaroon construction, which mathematically guarantees narrowing upon delegation. The following remain open:

– Calibration of TTLs, renewal budgets and irreversibility thresholds. The values are named qualitatively; like the starting TTL-Q values in the PQT proposal, they are subject to refinement on measurements in real networks and are declared a topic for the working group.

– The arms race of adversarial attacks and defences around the DL components of the network-side root. Attacks on RF-fingerprint classifiers and countermeasures to them evolve simultaneously; R11 removes DL components from the role of arbiter, but the boundary of an “auxiliary signal” requires periodic revision based on PV-3 precedents.

– The economics of witnessing. The witness root requires participation incentives and protection against witness collusion (the owner-independence threshold is set in R11, but the incentive model is a separate specification, coupled with the neighbourhood-hotspot economics).

8. Conclusion

The three questions of P-INDEP Section 7 are closed without introducing new entities: completeness - by a generative rule over the dependency graph, by the compromise experiment, and by precedent-driven evolution; audit privacy - by separating knowledge between class and instance, by reference profiles, and by mature anonymous-attestation mechanisms; cost - by an inventory of carriers already paid for, including neighbour witnesses at negative cost, and by a quorum scale bound to the stakes. The fourth question - about agents from the wallet - turned out to be not a new problem but a second proof of the correctness of the resolution of the second one: in an architecture where trust sits in independent roots, an agent can be given everything it needs while being given nothing that can be stolen - and while being deprived of the very ability to execute the inadmissible.

Trust stays in the roots. Everything else is a trace with a lifetime.

Invitation. *This document is submitted as a contribution to IEEE IC25-009-01 for technical review. The authors invite criticism of structural weaknesses - in particular of the proposals for the renewal-budget mechanism and the mandate-composition rules (Section 7) - rather than of implementation details. A principle that cannot survive attack should not be standardised.*

References

A. Project contributions (IEEE IC25-009-01)

[1] IEEE SA Industry Connections Activity IC25-009-01, “A Technical Reference Architecture Framework for an Open 6G Device Ecosystem,” ICAID v1.0, approved 3 December 2025.

[2] “Presence Protocol: A Principle of Authentication Without Stored Secrets for the Open 6G Ecosystem” (P-AUTH-1..P-AUTH-3), contribution to IEEE IC25-009-01, April 2026.

[3] “Accountability Chain Protocol (ACP): A Framework for Verifiable Subjecthood in the Open 6G Device Ecosystem,” contribution to IEEE IC25-009-01, v0.9, April 2026.

[4] “The Independence-of-Foundations Principle (P-INDEP),” contribution to IEEE IC25-009-01, v0.9, June 2026.

[5] “Three Lines in the Standard” (the PQT / TTL-Q / TEE-signature proposal), project article series, March 2026.

B. External sources (verified 2 July 2026)

[6] ISO 26262-9:2018, Road vehicles - Functional safety - Part 9: dependent failure analysis; coupling-factor classes as checklists for the system, hardware, software and semiconductor levels.

[7] E. Zhai, R. Chen, D. I. Wolinsky, B. Ford, “Heading Off Correlated Failures through Independence-as-a-Service,” OSDI 2014 - automated auditing of structural dependencies.

[8] ISO/IEC 20008 (Anonymous digital signatures) and the TCG TPM 2.0 specification - Direct Anonymous Attestation; Intel EPID 2.0 (RAND-Z licensing, open-source SDK). https://en.wikipedia.org/wiki/Direct_Anonymous_Attestation

[9] UBITECH, open-source TPM 2.0 DAA Library. https://github.com/ubitech/daa

[10] G. Arfaoui et al., “Towards a Privacy-preserving Attestation for Virtualized Networks,” IACR ePrint 2023/735 - property-based attestation; zero-knowledge proofs of configuration membership in a valid set. https://eprint.iacr.org/2023/735.pdf

[11] Cloudflare, “Introducing Zero-Knowledge Proofs for Private Web Attestation with Cross/Multi-Vendor Hardware” - production deployment of ZK attestation of hardware keys.

https://blog.cloudflare.com/introducing-zero-knowledge-proofs-for-private-web-attestation-with-cross-multi-vendor-hardware/

[12] Synopsys, Physical Unclonable Function (PUF) Security IP - deployed in 1+ billion devices; EMVCo, Visa, CC EAL6+, PSA certifications. https://www.synopsys.com/designware-ip/security-ip/cryptography-ip/puf.html

[13] Samsung Semiconductor, eSE/eSIM/SIM Security Solution - hardware PUF inside the eSE, FIDO compatibility. https://semiconductor.samsung.com/security-solution/ese-esim-sim/

[14] M. Mitev et al., "What Physical Layer Security Can Do for 6G Security" - location-based authentication at 6G's centimetre-level localisation accuracy. https://arxiv.org/pdf/2212.00427

[15] "Radio Frequency Fingerprint Identification for 5G Mobile Devices Using DCTF and Deep Learning," Entropy 26(1):38, 2024 - RF-fingerprint sources: manufacturing variations of analogue components; uniqueness, stability, cloning resistance.

[16] A. Birgisson, J. G. Politz, U. Erlingsson, A. Taly, M. Vrable, M. Lentczner, "Macaroons: Cookies with Contextual Caveats for Decentralized Authorization in the Cloud," NDSS 2014 (Google Research) - attenuable credentials with an HMAC caveat chain. https://theory.stanford.edu/~ataly/Papers/macaroons.pdf

[17] IETF Internet-Draft draft-klrc-aiagent-auth-00 (AIMS), March 2026 - AI-agent identity composing SPIFFE/WIMSE/OAuth 2.0; short-lived credentials as an alternative to revocation mechanisms; static keys as an anti-pattern.

[18] SPIFFE/SPIRE - short-lived SVIDs with automatic rotation; ceasing issuance as the revocation mechanism. https://spiffe.io/

[19] ERC-4337 / ERC-7715 - smart accounts with session keys: scope- and time-bounded delegations enforced by the wallet; the intent model of execution (the agent proposes - the policy executes).

[20] M. Ambrosin et al., "SANA: Secure and Scalable Aggregate Network Attestation," ACM CCS 2016 - publicly verifiable collective attestation; compromise of one device does not affect the others.

[21] F. Kohnhäuser et al., "Scalable Attestation Resilient to Physical Attacks for Embedded Devices in Mesh Networks," 2017 - mutual attestation of neighbours; collective detection of physical node capture by absence from contact. https://arxiv.org/pdf/1701.08034

[22] B. Yuce, P. Schaumont, M. Witteman, "Fault Attacks on Secure Embedded Software: Threats, Design and Evaluation" - the fault-attack evaluation methodology of Common Criteria certification. https://arxiv.org/pdf/2003.10513

[23] J. Ma et al., "Adversarial Attacks Against Deep Learning-Based Radio Frequency Fingerprint Identification," 2025 - PGD up to 98.9% targeted-attack success, UAP 80.5% misclassification; and "Collusion-Driven Impersonation Attack on Channel-Resistant RF Fingerprinting," 2025. https://arxiv.org/abs/2512.12002

This document was prepared as a contribution to IEEE Industry Connections Activity IC25-009-01 and does not represent an official IEEE position. Comments and critique should be directed to the IC activity mailing list.