IEEE INDUSTRY CONNECTIONS · IC25-009-01 · CONTRIBUTION DOCUMENT

The Independence-of-Foundations Principle (P-INDEP):

Placing Trust So That No Single Hidden Foundation Unlocks the System

Submitted to: A Technical Reference Architecture Framework for an Open 6G Device Ecosystem
Version 0.9 | June 2026 | For IC Activity Review

Abstract. Authentication and routing in the open 6G device ecosystem are being designed around several trust roots of different nature: physical presence (Presence Protocol, P-AUTH), hardware signature (the Path Quality Tag signed by a trusted execution environment), and public visibility (the Accountability Chain Protocol). Each root is chosen so that forging it is expensive. This contribution addresses a complementary and so-far implicit requirement: that compromising any single root must not unlock the system. We show that root heterogeneity by nature is necessary but not sufficient, and that the true determinant of attack cost is the cheapest *shared hidden foundation* - a coupling factor common to several roots (a shared die, power or clock network, the registration moment, the quorum mechanism, or a single vendor). When roots are positively correlated through such a foundation, a k-of-n quorum is provably weaker than one strong root. We propose P-INDEP: a normative principle requiring independence of trust roots by foundation, not merely by nature, together with proactive independence auditing of those roots.

Keywords: *6G device ecosystem, trust roots, common-cause failure, coupling factors, threshold authentication, independence auditing, Presence Protocol, ACP, IEEE IC25-009-01*

1. Motivation

An open 6G device places trust in more than one root. The Presence Protocol grounds it in the irreversibility of a physical moment; the Path Quality Tag grounds it in a hardware signature; the Accountability Chain Protocol grounds it in the visibility of breaches. Each root answers the question “how do we make forgery expensive?” and each answers it well, in its own register.

This contribution starts from a different question. Reliability is not the absence of trust but the deliberate placement of trust where forging it costs the most. The question “whom should we not trust?” is ill-posed. The right question is: where must trust sit so that breaking any one place requires the impossible - and so that breaking one place does not open the rest?

These are two distinct axes. “Expensive to forge” is maximised by the physical nature of a root. “Does not unlock the system when one root falls” is a property not of nature but of how the roots are wired together. A design can be excellent on the first axis and fail on the second. This document is about the second axis.

2. The Problem: Heterogeneity by Nature Is Not Independence by Foundation

Suppose access to a critical action - a wallet transaction, a routing decision, an ACP signature - requires a quorum of k-of-n trust roots of different physical nature. The intended security claim is that the cost of

compromising the system equals the sum of the costs of breaking k heterogeneous roots, rather than the minimum across roots. A rational attacker, however, does not attack head-on; the attacker takes the minimum over all available attack paths. Beyond breaking the roots one by one, the attacker may instead target what the roots share.

Independence of two elements has a precise meaning in the dependent-failure literature: two elements are independent only if there is no interference and no common root cause of failure between them. *The hidden links that violate this are called coupling factors* - shared power or clock networks, a common die, a shared procedure. Heterogeneous roots that ride on a common foundation are not independent, regardless of how different their nature is.

This is not a theoretical nicety. On mobile platforms the secure execution environment is an effective monoculture, and by design the secure world can read all of the normal world's memory - so a single privileged compromise can reach a hardware-signing key, a locally stored physiological baseline, and other on-device roots at once. Three roots of different nature, sharing one die, collapse into one. The same logic applies to a shared registration moment, a centralised quorum collector, or a single trust-anchor vendor.

3. Why a Correlated Quorum Is Weaker Than One Strong Root

The quantitative consequence is sharp. When the assurance levels of authentication factors are positively correlated, the aggregated assurance is lower than in the independent case - and under strong dependence a single authentication is preferable to multi-factor authentication. In other words, a k-of-n quorum whose roots share a foundation does not merely lose part of its benefit; it can fall below the security of one well-built root, while costing more and degrading usability.

This reframes the original question. The cost of attacking the system is determined neither by the most expensive root nor by the sum of the roots, but by the cheapest shared foundation. Reliability is the placement of trust such that the roots have no common floor.

4. The P-INDEP Principle

We propose the following principle for inclusion in the open 6G device reference architecture, at the same level of normative weight as P-AUTH (Presence Protocol) and the ACP principles.

Principle P-INDEP (Independence of Foundations). Where access depends on a quorum of trust roots, no two roots counted toward the quorum shall share a coupling factor - a common physical carrier, power or clock domain, registration moment, quorum-evaluation mechanism, or controlling vendor - whose single compromise would reach more than one root. Independence shall be required by foundation, not merely by the nature of the roots, and shall be verified by proactive auditing rather than assumed. **Rationale:** the cost of compromising a quorum is set by its cheapest shared foundation; absent foundation independence, a heterogeneous quorum can be weaker than a single strong root.

4.1 Normative requirements

1. Roots counted toward one quorum shall reside on physically distinct carriers. Where on-device isolation is unavoidable, at least one root shall be anchored off-device (e.g. network-side geometric verification).
2. The quorum-evaluation mechanism shall be executed locally and on heterogeneous implementations, with no single centralised collector of votes that could itself become a common root.
3. The minimum k for irreversible actions shall be fixed by the standard, not left to the deploying vendor, and shall align with the irreversibility-threshold logic of the ACP Strategist class.
4. Roots shall be independent by controlling party and, where feasible, by jurisdiction, consistent with the ACP Reachable Anchor requirement.
5. The registration moment shall itself require a minimal quorum and shall be published in the ACP registry (PV-1), so that substitution at onboarding becomes visible.
6. Conformance shall include a structural independence audit of the trust roots; the set of coupling factors is treated as open and revisable, not as a fixed checklist.

5. Diversion Analysis: How the Principle Is Defeated If Left Implicit

The following failure scenarios are presented as normative warnings. Each corresponds to a coupling factor that re-introduces a single hidden foundation.

Scenario	Mechanism	Counter-requirement
Carrier collapse	All roots computed/stored on one SoC whose secure world reads all memory; one exploit reaches every root while the quorum appears honest.	Req. 1 (distinct carriers; one root off-device)
Registration attack	Roots substituted at onboarding (T2), before the quorum is established; later checks pass cleanly on an already-compromised identity.	Req. 5 (quorum at registration; PV-1 visibility)
Silent k→1	Vendor lowers default k or allows “remember this device” for usability; the quorum is formal, the reality is a mono-root.	Req. 3 (standard-fixed minimum k)
Vendor / geometry monoculture	Triangulation nodes owned by one operator, or the whole TEE market is one vendor; “different” roots share one exploit.	Req. 4 (independence by party / jurisdiction)
Quorum-mechanism attack	A centralised or identical vote collector; one bug nullifies the whole construction (the cheapest path of all).	Req. 2 (local, heterogeneous evaluation)

Likelihood note. Against the reference class of “mandatory-diversity” security requirements on an immature market - multi-factor schemes degrading to the convenient factor, diversity requirements consolidating to the dominant supplier - the probability that first-generation open 6G devices share at least one coupling factor across roots, absent a normative requirement, is estimated at roughly 72%

(range 60–82%). This is a calibrated class base rate, not a measurement of any specific design; it is the argument for making P-INDEP normative rather than advisory.

6. Relationship to Existing Principles and Standards

P-INDEP does not compete with the existing contributions; it states the wiring condition under which their guarantees survive composition.

Presence Protocol (P-AUTH). P-INDEP inherits P-AUTH unchanged for embodied subjects and adds a constraint the Presence Protocol does not state: the locally stored personal baseline must not share a carrier with another root counted toward the same quorum.

ACP. P-INDEP operationalises, at the trust-root layer, the same anti-centralisation logic ACP applies at the accountability layer; the ACP registry (PV-1) provides the visibility channel that makes foundation-sharing auditable.

Dependent-failure and independence-auditing practice. P-INDEP imports the established notions of common-cause failure and coupling factors from functional-safety practice (e.g. ISO 26262 dependent-failure analysis) and proactive structural independence auditing, applying them to authentication trust roots rather than to safety elements.

Threshold authentication. k-of-n group/threshold schemes provide a mature mechanism for “sum, not minimum”; their known weak point - the initial key-distribution / registration phase - is exactly the moment P-INDEP Requirement 5 protects.

7. Open Questions

- Completeness of the coupling-factor set. The list (carrier, power, clock, registration moment, quorum mechanism, vendor, jurisdiction) is open. P-INDEP therefore mandates auditing rather than a fixed checklist; a canonical taxonomy for the 6G device context remains to be specified.
- Auditability without privacy loss. Structural independence auditing requires dependency data about roots; reconciling this with the no-stored-secret stance of P-AUTH needs specification.
- Cost of distinct carriers. Requiring physically distinct carriers raises bill-of-materials cost; the trade-off against the assurance gain should be quantified per device class.

8. Conclusion

Heterogeneity of trust roots is necessary but not sufficient. The property that protects the system - that breaking one root does not open the rest - is not provided by the differing nature of the roots but by the absence of a shared foundation beneath them. The cost of attack is set by the cheapest coupling factor, and a correlated quorum can be weaker than one strong root. P-INDEP names this condition and requires it to be verified rather than assumed.

Invitation. *This document is submitted as a contribution to IEEE IC25-009-01 for technical review and critique. The authors specifically invite criticism of structural weaknesses - in particular, proposed*

additions to the coupling-factor taxonomy of Section 7 - rather than 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" (principles P-AUTH-1 to 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.

B. Supporting external literature (verified)

[4] E. Zhai, R. Chen, D. I. Wolinsky, and B. Ford, "Heading Off Correlated Failures through Independence-as-a-Service," in Proc. 11th USENIX Symp. on Operating Systems Design and Implementation (OSDI '14), Broomfield, CO, USA, Oct. 2014, pp. 317–334. https://www.usenix.org/sites/default/files/osdi14_full_proceedings.pdf

[5] OWASP Foundation, "Multifactor Authentication Cheat Sheet," OWASP Cheat Sheet Series. "The factors used should be independent of each other and should not be able to be compromised by the same attack." https://cheatsheetseries.owasp.org/cheatsheets/Multifactor_Authentication_Cheat_Sheet.html

[6] National Security Agency (NSA), "Transition to Multi-factor Authentication," Cybersecurity guidance, Sep. 2019. "The strength of MFA decreases when all the factors can be compromised … attackers may try to find and gain persistence on a host where the factors come together." https://media.defense.gov/2019/Sep/09/2002180346/-1/-1/0/Transition%20to%20Multi-factor%20Authentication.pdf

[7] "Risk-link authentication for optimizing decisions of multi-factor authentications," U.S. Patent 10,210,518. "Where the levels of assurance are positively correlated … the aggregated assurance is expected to be less than … the independent case … a single authentication was preferred over multi-factor authentication." https://patents.google.com/patent/US10210518B2

[8] A. Shamir, "How to Share a Secret," Communications of the ACM, vol. 22, no. 11, pp. 612–613, Nov. 1979.

[9] L. Harn, "Group Authentication," IEEE Transactions on Computers, vol. 62, no. 9, pp. 1893–1898, 2013 - (t, m, n) group authentication based on Shamir's (t, n) secret sharing. https://doi.org/10.1109/TC.2012.251

[10] ISO 26262-9:2018, Road vehicles - Functional safety - Part 9: ASIL-oriented and safety-oriented analyses (dependent-failure analysis; coupling factors). Applied here by analogy to authentication trust roots.

Note on sources: references [1]–[3] are the project contributions on which this document builds. References [4]–[10] are external sources introduced by this contribution to substantiate the common-cause-failure argument; they do not appear in [1]–[3]. All external URLs were verified accessible at the time of writing (20 June 2026).

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