

CRA-Capable Digital Product Passports: A Verifiable Evidence Architecture for Cybersecurity, Compliance, Functional Safety, and Product Lifecycle Assurance

Legal-technical design science article

Final manuscript — Evidence cut-off: 15 August 2026

Dr. Carsten Stöcker
Spherity GmbH

Prepared: August 2026



Except where otherwise noted, this work is licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0). © 2026 Dr. Carsten Stöcker and Spherity GmbH.

Abstract

Regulation (EU) 2024/2847, the Cyber Resilience Act (CRA), establishes horizontal cybersecurity requirements for products with digital elements and places manufacturers under lifecycle obligations that extend beyond market entry. The regime requires risk-based secure design, documented software composition, vulnerability handling, secure updates, user information, post-market corrective action, and time-critical reporting of actively exploited vulnerabilities and severe incidents. In parallel, the Ecodesign for Sustainable Products Regulation (ESPR) establishes the Digital Product Passport (DPP) as an interoperable and increasingly operational European product information infrastructure. These regimes are legally distinct: the CRA does not mandate a DPP, and the ESPR DPP is not, by itself, evidence of CRA conformity. Their convergence nevertheless creates a design opportunity. Product identifiers, versioned data, verified economic operators, semantic models, controlled access, and lifecycle events can provide a common infrastructure through which cybersecurity evidence is resolved and governed. This article develops the concept of a CRA-capable DPP as a legal-technical design artefact. It defines such a DPP not as a static certificate or public compliance page, but as a distributed, versioned, and access-controlled evidence architecture that cryptographically binds product identity, hardware and software composition, vulnerability and exploitability status, security updates, conformity evidence, and lifecycle events to the relevant product model, release, physical or logical instance, deployment configuration, and remote data processing service. The proposed architecture separates a control plane for legal-person identity, organisational authority, mandates, and policy enforcement from an evidence plane for product claims, assessment results, SBOMs, VEX statements, update receipts, configuration states, and modification histories. The European Digital Identity Framework and the proposed European Business Wallets framework are examined as a prospective European implementation of the control plane, without treating the proposal as applicable law. Methodologically, the study combines normative analysis of the CRA and ESPR, standards analysis of the 17 ETSI Final Draft European Standards available under Public Enquiry at the evidence cut-off, design science construction of a verifiable evidence graph, and cross-sector evaluation. The article distinguishes binding legal

obligations, candidate harmonised standards, indirect governance requirements, and novel architectural choices. It argues that the principal value of a CRA-capable DPP lies in continuous conformity: the capacity to maintain a verifiable relation between the assessed product, the product actually deployed, and the changing cybersecurity state throughout the support period. The resulting model supports cybersecurity operations, conformity assessment, supply-chain accountability, lifecycle services, and the controlled coordination of cybersecurity and functional-safety change processes.

Keywords: *Cyber Resilience Act; Digital Product Passport; product evidence infrastructure; SBOM; VEX; continuous conformity; evidence graph; European Business Wallet; eIDAS 2.0; zero trust; functional safety; product lifecycle assurance*

Contents

1	Introduction	5
1.1	Regulatory convergence around products with digital elements	5
1.2	The product cybersecurity evidence problem	6
1.3	Research gap	7
1.4	Research objective and central proposition	9
1.5	Research questions	9
1.6	Normative boundaries and terminology	10
1.7	Methodology: legal-technical design science	11
1.8	Expected contributions	13
1.9	Organisation of the article	14
2	Part I - Regulatory and Standards Foundation	14
2.1	Normative hierarchy and permissible inferences	15
2.2	The CRA as a lifecycle evidence regime	16
2.3	The European DPP infrastructure as a reusable evidence substrate	20
2.4	Identity, authority, and the control-plane foundation	22
2.5	CRA standardisation architecture under M/606	23
2.6	Analysis of the 17 ETSI CRA vertical standards	24
2.7	Cross-cutting standards findings for product evidence infrastructure	29
2.8	Direct, indirect, and proposed requirements for a CRA-capable DPP	30
2.9	Derived design requirements for Part II	31
3	Part II - Concept and Reference Architecture	33
3.1	Definition, scope, and architectural boundary	33
3.2	Design principles and assurance objectives	34
3.3	Five-subject product model	35
3.4	Dual-plane reference architecture	38

3.5	Control-plane identity, mandate, and zero-trust authority	40
3.6	Evidence object model and verifiable evidence graph	41
3.7	Evidence assurance vector and policy evaluation	44
3.8	Composition, build provenance, and supplier evidence	45
3.9	Vulnerability, exploitability, and VEX lifecycle	47
3.10	Security update, installed-state, rollback, and recovery evidence	49
3.11	Continuous conformity and lifecycle lineage	50
3.12	Functional-safety and essential-availability interface	53
3.13	Access control, confidentiality, and selective disclosure	54
3.14	Cryptographic binding, status, and long-term verifiability	55
3.15	Interoperability stack and deployment patterns	57
3.16	Conformance profiles and validation rules	59
3.17	Threat model for the DPP evidence infrastructure	60
3.18	Worked example: managed enterprise router with remote controller	62
3.19	Architecture traceability and specification for Part III	64
4	Part III - Operational Value and Sector Implementation	64
4.1	Evaluation design and analytical scenarios	65
4.2	Cybersecurity operations value	66
4.3	Compliance operations value	71
4.4	Business and lifecycle value	73
4.5	Functional safety and essential availability	75
4.6	Reusable implementation patterns	77
4.7	Sector implementation and applicability boundaries	80
4.8	Cross-sector scenario evaluation	85
4.9	Findings and implementation priorities for Part IV	87
5	Part IV - Evaluation, Governance, and Research Implications	88
5.1	Evaluation method and validity conditions	89
5.2	Legal accuracy and regulatory coherence	91
5.3	Interoperability and semantic governance	94
5.4	Scalability, performance, and proportionality	97
5.5	Security and resilience of the evidence infrastructure	99
5.6	Transparency, confidentiality, data protection, and trade secrets	102
5.7	Provider dependency, portability, and long-term continuity	104
5.8	Governance and accountability model	107
5.9	Limits of automated conformity and proof-graph evidence scoring	109

5.10	Integrated evaluation of propositions and design requirements	112
5.11	Limitations	114
5.12	Research, standards, and implementation implications	115
6	Conclusion	118
6.1	Synthesis of findings	118
6.2	CRA requirements heat map and final interpretation	119
A	Requirements-to-evidence matrix and versioned source register	120
A.1	Matrix and source-register schema	120
A.2	Representative versioned records	121
	References	123

1 Introduction

1.1 Regulatory convergence around products with digital elements

Products with digital elements no longer possess a stable cybersecurity state after manufacture or initial release. Their security properties depend on changing software components, firmware revisions, cloud and application-programming-interface dependencies, cryptographic material, configurations, operational environments, and third-party services. A product that was adequately assessed at market entry may later contain an exploitable component, run an unsupported release, depend on a modified remote service, or operate outside the deployment assumptions used in its original risk assessment. This temporal instability changes the character of product compliance. The relevant question is not only whether a product type satisfied defined requirements at a point in time, but whether the manufacturer and other responsible actors can maintain, demonstrate, and communicate an adequate cybersecurity state throughout the period in which the product is supported and used.

The CRA responds to this problem by creating a horizontal European market-access regime for products with digital elements [1]. Its essential cybersecurity requirements cover the design, development, and production of products as well as the processes through which manufacturers identify and document vulnerabilities, prepare software bills of materials, remediate weaknesses, distribute security updates, provide coordinated vulnerability disclosure channels, and communicate relevant information to users. The Regulation also brings certain integrated remote data processing solutions within scope where their absence would prevent the product from performing one of its functions. The regulated object can therefore extend beyond a device or locally installed software package to a product-service configuration whose components evolve at different rates.

The CRA further establishes a post-market evidence obligation. Technical documentation must contain the data needed to demonstrate conformity and must be continuously updated where appropriate during the support period. Manufacturers must determine and disclose the support period, handle vulnerabilities effectively during that period, maintain security updates for the legally specified duration, take corrective action where a product or process is no longer conformant, and report actively exploited vulnerabilities and severe incidents within short statutory timeframes. Article 14 reporting obligations apply from 11 September 2026, while the CRA becomes generally applicable on 11 December 2027 [1]. These dates convert lifecycle cybersecurity from a future design consideration into an immediate operational requirement.

European standardisation is intended to translate the CRA's technology-neutral essential requirements into testable product and process requirements. Commission Implementing Decision C(2025) 618, commonly referenced as standardisation request M/606, identifies 41 horizontal and product-specific standardisation items [2]. As of 15 August 2026, ETSI has made 17 Final Draft European Standards available under Public Enquiry for product categories including browsers, password managers, antimalware products, virtual private networks, network-management systems, security information and event management systems, boot managers, public-key-infrastructure products, network interfaces, operating systems, routers, smart-home products, toys, wearables, virtualisation stacks, and firewalls [3]. These documents are candidate harmonised standards. They do not create a presumption of conformity until final standards are adopted and their references are published in the Official Journal of the European Union. This status distinction is central to any academically and legally accurate architecture.

In parallel, the ESPR establishes the DPP as a European product information mechanism intended to support sustainability, circularity, market surveillance, and value-chain information exchange [4]. The legal and technical implementation of this infrastructure advanced materially

in July 2026 with Commission Implementing Regulation (EU) 2026/1778 on the DPP registry [5]. The registry framework provides for verified economic operators, persistent registration identifiers, model-, batch-, and item-level granularity, versioned data, semantic validation, logs, a semantic repository, links to evidence, and proofs of registration protected by qualified electronic seals and timestamps. It also explicitly retains a decentralised DPP model: the registry stores defined identifiers and registration data, while the substantive product information remains distributed across DPP information systems and service providers. Importantly, automated registry checks concern structure, semantics, and registration completeness; they do not prove the substantive correctness of product data or compliance with applicable Union law [5].

The two regulatory developments are therefore complementary but not interchangeable. The CRA creates cybersecurity obligations and conformity duties for products with digital elements. The ESPR and related product legislation create an infrastructure for persistent product identifiers, machine-readable data, regulated access, and lifecycle information. The CRA does not explicitly require manufacturers to use that infrastructure, and a DPP registered under the ESPR does not constitute a CRA certificate. Nevertheless, the emergence of a common European DPP infrastructure makes it reasonable to ask whether CRA evidence can be represented, linked, governed, and resolved through compatible product-passport mechanisms. This article addresses that design question without converting an architectural opportunity into a fictional legal mandate.

1.2 The product cybersecurity evidence problem

The central implementation problem is a mismatch between the static form of conventional compliance documentation and the dynamic state of contemporary products. A technical file, declaration of conformity, test report, SBOM, or certificate is normally issued for a defined subject and at a defined time. Operational risk, however, is generated by the relationship among several subjects: the product model placed on the market; the hardware, firmware, and software release assessed; the physical or logical instance deployed; the configuration and operating environment of that instance; and the remote service on which one or more product functions depend. Treating these subjects as one undifferentiated product record produces ambiguous and potentially misleading claims.

A model-level conformity assessment, for example, cannot establish that a particular installed device runs the assessed firmware. A release-level SBOM cannot establish that a deployed instance has successfully installed the corresponding security update. A manufacturer advisory can identify an affected release, but it cannot by itself establish exploitability in a specific configuration. A certificate for a component cannot establish that an integrator has preserved its security assumptions within a larger system. A cloud-service change can alter the security properties of a connected product even where the physical device and local software remain unchanged. These distinctions become critical in routers, industrial controllers, battery-energy-storage systems, virtualisation platforms, smart-home products, connected machinery, and software products distributed across heterogeneous environments.

The evidence needed to manage these relations is currently fragmented across engineering, product lifecycle management, vulnerability management, conformity assessment, customer support, incident response, and supplier-management systems. SBOMs describe software composition, but they may be generated at different stages, with different levels of transitivity and inconsistent product identifiers. VEX statements or equivalent exploitability assessments add product-specific context, but their value depends on the exact release, component path, vulnerability source, assessment time, and mitigation state [6]. Security advisories identify affected and fixed versions, while update systems record packages and installations. Technical files docu-

ment applicable requirements and test evidence. Asset inventories record deployed instances and configurations. Supplier portals carry component declarations. None of these artefacts, considered alone, provides a reliable and complete answer to the question: what cybersecurity claim applies to which product subject, on whose authority, on the basis of which evidence, and for what period of validity?

The problem is not solved by placing these artefacts behind a QR code or combining them in a central database. A DPP that merely aggregates unverified data can reproduce the same ambiguity at a new access point. Evidence must be bound to the correct subject and version; claims must be attributable to identified organisations and authorised issuers; status changes must be time-stamped, supersedable, and revocable; and access must reflect the difference between public user information, operator evidence, confidential supplier data, conformity-assessment material, and authority access. A trustworthy architecture must also distinguish declared state from observed state. A manufacturer may declare that release R contains SBOM hash H, while a device attestation establishes that instance I currently runs release R, an operator record establishes configuration C, and a laboratory report establishes that R was assessed against requirement Q. Combining these claims into a single binary status would remove information that is necessary for both security decisions and legal accountability.

Supply-chain composition adds a further governance problem. Final manufacturers depend on software libraries, firmware, hardware components, operating systems, cloud services, and security products supplied by other organisations. Suppliers may issue component SBOMs, vulnerability statements, updates, certifications, or support commitments. The final manufacturer must evaluate those claims in the context of the final product and retains the legal responsibility assigned by the CRA. A product evidence infrastructure must therefore preserve claim provenance without confusing evidence contribution with a transfer of regulatory responsibility. It must also support changes in responsibility where an importer, distributor, integrator, repairer, remanufacturer, or other actor places a substantially modified product on the market under conditions that make that actor the manufacturer of the modified product.

Cybersecurity change can also interact with functional safety. The CRA is not a functional-safety regulation and a CRA-capable DPP cannot replace a safety case, sector-specific certification, or established safety lifecycle. Yet software, firmware, configuration, cryptographic, and network changes can invalidate assumptions used in a safety assessment. An urgent security update may alter timing, diagnostic coverage, fail-safe behaviour, communications, or availability. Conversely, delaying an update for safety reasons may leave an exploitable vulnerability in operation. The evidence problem is therefore not to collapse cybersecurity and safety into one status, but to make their dependencies explicit: which security-relevant change affects which approved configuration, which safety analysis was performed, who authorised deployment, and which rollback or recovery state remains permitted.

The research problem can thus be stated precisely. Europe is establishing legally binding cybersecurity duties and an operational DPP infrastructure, but it lacks a generally defined architecture that translates CRA and candidate ETSI requirements into verifiable, versioned, access-controlled product evidence across model, release, instance, deployment, and service layers. Without such an architecture, manufacturers may digitise documents without achieving operational traceability, and DPP implementations may expose data without establishing authority, freshness, scope, or evidentiary value.

1.3 Research gap

The DPP research literature has advanced the analysis of circular-economy information requirements, lifecycle data sharing, traceability, product identity, and technical implementation. Ex-

isting work identifies requirements for DPP ecosystems, examines data-sharing and governance challenges, and evaluates architectures based on distributed ledgers, decentralised identifiers, verifiable credentials, and industrial data models [7–11]. This literature establishes that a DPP is not merely a consumer label. It is an ecosystem artefact whose value depends on product granularity, semantic interoperability, actor incentives, data quality, access rights, and lifecycle persistence.

However, most DPP scholarship remains anchored in sustainability, circularity, materials, repair, and end-of-life processes. Cybersecurity is commonly treated as a requirement of the DPP platform itself, as a generic data-security concern, or as an additional data category. Less attention has been given to the use of the DPP as the evidentiary infrastructure for a product cybersecurity regime. In particular, the literature has not yet established a rigorous mapping between CRA legal obligations, ETSI product requirements, cybersecurity evidence objects, issuing authorities, temporal state, and the distinct subject layers to which claims apply.

The software supply-chain literature provides a complementary but incomplete perspective. Empirical studies of SBOM adoption identify problems of component naming, tooling, completeness, supplier coordination, business incentives, and operational integration [12, 13]. VEX and advisory formats address the need to contextualise whether a known component vulnerability affects a product. These mechanisms are important building blocks, but they do not by themselves establish product identity, economic-operator responsibility, conformity scope, update deployment, support entitlement, substantial-modification lineage, or controlled authority access. An SBOM is an inventory, not a conformity argument. A VEX statement is an exploitability claim, not proof that an affected instance has been remediated.

Research on decentralised identity and verifiable credentials provides mechanisms for issuer authentication, cryptographic integrity, selective disclosure, status checking, and portable verification [14–17]. Zero-trust architecture establishes the principle that access and action should depend on continuously evaluated identity, context, policy, and evidence rather than network location or implicit trust [18]. These concepts are relevant to cross-organisational product evidence, but they do not prescribe how CRA requirements should be represented, how legal-person mandates should be resolved, or how product evidence should be separated across model, release, instance, deployment, and remote-service subjects.

The CRA, Commission guidance, M/606, and the ETSI drafts provide the normative and technical requirements from the opposite direction. They identify manufacturer duties, essential cybersecurity outcomes, product-specific controls, assessment methods, and evidence expectations. They do not define a shared DPP architecture, a universal evidence graph, or a European organisational-authority layer for issuing and verifying lifecycle claims. This omission is appropriate: legislation and standards should not prematurely mandate one implementation technology. It nevertheless leaves a design space that must be addressed by manufacturers, infrastructure providers, conformity-assessment bodies, data-space operators, and public authorities.

The resulting gap is both conceptual and operational. There is no sufficiently precise definition of a CRA-capable DPP; no common method for separating direct legal information duties from indirect evidence-governance controls; no complete subject model for cyber-physical and software products; and no reference architecture that combines organisational authority with verifiable product evidence while respecting confidentiality and legal boundaries. This article addresses that gap through legal-technical design science.

1.4 Research objective and central proposition

The objective of this study is to design and evaluate a reference architecture through which DPP mechanisms can support CRA implementation without being misrepresented as a statutory requirement or an autonomous proof of compliance. The intended artefact is not a new certificate. It is an evidence architecture that enables a relying party to resolve a scoped claim, its subject, issuer, authority, supporting evidence, assessment method, time, status, and access policy.

Central research proposition. *A CRA-capable Digital Product Passport is not a static compliance document. It is a distributed, versioned, and access-controlled evidence architecture that cryptographically binds product identity, software and hardware composition, vulnerability status, security updates, conformity evidence, and lifecycle events to the relevant product model, release, instance, deployment, and remote service.*

The term distributed is used in an architectural rather than ledger-specific sense. Evidence may remain with authoritative manufacturers, suppliers, laboratories, operators, registries, or DPP service providers, provided that it can be resolved and verified through stable identifiers and governed interfaces. The design does not require a blockchain. Cryptographic binding means that the integrity and provenance of a claim can be verified and that the claim is unambiguously associated with the intended product subject and evidence object. It does not mean that cryptography establishes the truth of the underlying factual assertion. Substantive reliability continues to depend on competent issuers, valid authority, adequate assessment methods, trustworthy source data, and effective oversight.

The architecture is framed as product evidence infrastructure. It links a regulatory requirement to a technical control, an assessment method, an evidence object, an issuer, and a product subject. It further links model-level design evidence to release-level composition and vulnerability evidence, instance-level observed state, deployment-level configuration, and service-level dependencies. This enables continuous conformity, understood here as the controlled maintenance of an evidentiary relationship between the assessed product and the product state that exists after market entry. Continuous conformity is not continuous certification, and it does not imply that software can determine legal compliance autonomously. It denotes the capacity to identify when evidence remains applicable, when a change invalidates an assumption, and when reassessment, corrective action, escalation, or human decision is required.

A dual-plane model structures the proposal. The control plane establishes the identity and authority of legal persons, organisational units, employees, services, and automated agents that issue, update, request, or rely on evidence. It also carries mandates, representation rights, policy rules, and revocation status. The evidence plane carries cryptographically protected product claims, SBOM and HBOM objects, VEX statements, advisories, assessment results, update and rollback receipts, configuration states, support-period information, and lifecycle events. The European Digital Identity framework and the Commission proposal for European Business Wallets provide a plausible European implementation path for the control plane [19,20]. The paper treats the proposed European Business Wallet as prospective infrastructure, not as applicable law or a mandatory dependency.

1.5 Research questions

The study is organised around four research questions:

RQ1. Which CRA obligations and candidate ETSI requirements create direct information and evidence needs that can be represented or resolved through a DPP, and which additional governance controls arise only indirectly from the need to make that evidence reliable?

RQ2. What reference architecture and evidence model can bind cybersecurity claims to the correct product model, hardware or software release, physical or logical instance, deployment configuration, and remote data processing service while preserving provenance, temporal validity, confidentiality, and issuer accountability?

RQ3. How can a dual-plane architecture based on organisational authority and verifiable product evidence support continuous conformity, cybersecurity operations, conformity assessment, lifecycle services, and functional-safety change control?

RQ4. Under which technical, legal, organisational, and sector-specific conditions does the architecture remain interoperable, scalable, secure, and useful, and where are the limits of automated claims and DPP-based assurance?

1.6 Normative boundaries and terminology

The paper adopts a strict hierarchy of normative sources. This is necessary because the expression CRA-capable DPP combines an enacted cybersecurity regulation, a legally established but product-dependent DPP framework, candidate harmonised standards, identity and trust infrastructures at different stages of legal maturity, and a new architectural proposal. The sources cannot be treated as if they carried the same legal weight.

Table 1: Normative layers and their treatment in this study

Layer	Status as of 15 August 2026	Role and boundary in this article
Cyber Resilience Act, Regulation (EU) 2024/2847	Binding Union law. Article 14 applies from 11 September 2026; general application from 11 December 2027.	Defines legal duties, essential cybersecurity requirements, conformity procedures, technical documentation, support, vulnerability handling, and reporting. It does not mandate a DPP.
ESPR and DPP implementing legislation	Binding framework. DPP duties arise for product groups through delegated or sectoral legislation. Regulation (EU) 2026/1778 governs the DPP registry.	Provides product identifiers, granularity, versioning, semantic and registry mechanisms. Registration or structural validation is not substantive proof of CRA conformity.
Harmonised European standards cited in the OJEU	Voluntary technical specifications that can create a presumption of conformity for covered requirements once their references are published.	Future conformity route. Applicability, edition, scope, and partially applied clauses must be recorded precisely.
17 ETSI Final Draft European Standards under Public Enquiry	Candidate harmonised standards. Not yet OJEU-cited and subject to change.	Primary standards-analysis corpus. Requirements are treated as draft technical evidence expectations, not binding law or an existing presumption of conformity.
eIDAS 2.0 / European Digital Identity framework	Binding legal framework for electronic identification, wallets, attestations, signatures, seals, and trust services.	Provides trust mechanisms that may support actor authentication, organisational evidence, seals, timestamps, and verification. It does not by itself define a CRA DPP.
European Business Wallet proposal, COM(2025) 838	Commission legislative proposal; not treated as adopted law in this paper.	Prospective European control-plane implementation for legal-person identity, mandates, representation, official communications, and attestations.

Layer	Status as of 15 August 2026	Role and boundary in this article
Architecture proposed in this article	Non-normative design science artefact.	Translates legal and standards requirements into an evidence model, trust architecture, implementation patterns, and evaluation criteria. It must not be represented as a legal requirement.

Source: Authors' classification based on [1–5, 19, 20].

Three requirement categories are used throughout the study. Direct requirements are information, process, documentation, or evidence obligations that can be traced explicitly to the CRA or a relevant standard. Examples include product identification, risk assessment, SBOM generation, vulnerability handling, update information, support-period disclosure, and conformity documentation. Indirect requirements are controls that are not expressed as DPP duties but are necessary if a DPP is to carry reliable CRA evidence. These include issuer authentication, mandate verification, cryptographic binding, status and revocation, evidence freshness, provenance, access control, confidentiality, retention, and continuity. Proposed architectural choices are design decisions introduced by this paper, such as the five-layer subject model, dual-plane separation, evidence-graph relations, disclosure tiers, and policy-engine integration.

The term DPP is used broadly enough to include software-only and service-dependent products where a physical data carrier is absent. The relevant access mechanism may be a product identifier, package identifier, device certificate, registry reference, API, or other resolvable link. A CRA-capable DPP may share infrastructure and identifiers with an ESPR or sectoral DPP, but it can contain additional cybersecurity evidence under different access policies. The term passport should not be read as implying that all evidence is stored in one object, made public, or controlled by one platform.

Functional safety is treated as an adjacent assurance domain. The paper examines how cybersecurity evidence can preserve, challenge, or trigger reassessment of safety assumptions. It does not claim that CRA conformity establishes functional safety, nor that a DPP replaces safety standards, notified assessment, or a safety case. Similarly, the use of European Business Wallets, decentralised identifiers, verifiable credentials, Asset Administration Shells, SPDX, CycloneDX, CSAF, VEX, JSON-LD, RDF, or other technologies is evaluated as an implementation option unless a legal act or standard makes a more specific requirement.

1.7 Methodology: legal-technical design science

The study follows a legal-technical design science methodology. Design science is appropriate where the research objective is to construct and evaluate an artefact that addresses a defined organisational and technical problem [21–23]. The artefact in this study consists of a set of design requirements, a conceptual data and trust model, a dual-plane reference architecture, an evidence-graph model, and cross-sector implementation patterns. The methodology does not treat legal rules as system requirements without interpretation. It combines doctrinal and normative legal analysis with standards engineering and system design, while preserving the differences among legal validity, technical conformity, operational security, and architectural preference.

The first research activity is normative legal analysis. The primary legal corpus comprises the CRA, the ESPR, relevant DPP implementing legislation, the European Digital Identity framework, the Commission proposal for European Business Wallets, and Commission CRA implementation guidance. The analysis identifies regulated actors, product scope, lifecycle duties, documentation requirements, conformity routes, reporting obligations, access and retention

conditions, and legal exclusions. Each proposition is classified by legal source and status. Commission guidance is used as interpretive material rather than as legislation. The EBW proposal is analysed as a prospective institutional and technical framework and is not presented as binding law.

The second activity is standards analysis. The corpus includes M/606 and all 17 ETSI Final Draft European Standards under Public Enquiry available on the evidence cut-off date. An additional ETSI work item, EN 304 642 on network functions of telecommunications systems, is tracked but excluded from the complete clause-level corpus because it had not reached the same public final-draft status. For every normative or assessment-relevant requirement extracted from the ETSI corpus, the analysis records the standard and version, clause and requirement identifier, covered CRA requirement, product subject, expected evidence, assessment method, temporal character, responsible issuer, and potential DPP relevance. This method prevents a selected example from being misrepresented as complete standards coverage and allows later versions of a draft to supersede earlier mappings explicitly.

The third activity is requirements synthesis and artefact construction. Legal and standards findings are transformed into design requirements only after their source and status have been recorded. The synthesis asks: what must be identified; what must be asserted; what must be observed; what must be assessed; who is legally or operationally responsible; which actor is competent to issue the claim; to which subject layer the claim applies; how long it remains valid; what event supersedes it; and who may access the underlying evidence. The resulting architecture separates identity and authorisation functions from evidence functions, and separates product model, release, instance, deployment, and remote-service state. Evidence is represented as a graph because the same component, vulnerability, advisory, test, or update can relate to multiple products and lifecycle events without being copied into one monolithic record.

The fourth activity is analytical evaluation across representative implementation patterns. The evaluation compares software-release passports, connected-device and fleet passports, remote-service passports, industrial system-of-systems passports, and substantial-modification or second-life passports. Sector scenarios include enterprise and cybersecurity software, routers and telecom equipment, smart-home products, industrial machinery and robotics, energy systems and batteries, charging infrastructure, wearables, semiconductor security components, and AI or edge-computing systems. Sectoral exclusions and overlapping regimes are treated as boundary conditions rather than ignored. The evaluation examines whether the architecture can preserve product identity, evidence traceability, supplier provenance, confidentiality, update state, support status, and safety-relevant configuration across these patterns.

Table 2: Research design, analytical operations, and intended outputs

Research activity	Primary corpus or input	Analytical operation	Output
Normative legal analysis	CRA, ESPR, DPP registry legislation, eIDAS 2.0, EBW proposal, Commission guidance	Identify duties, actors, scope, timelines, legal status, exclusions, and documentation obligations	Legally bounded requirement set and terminology
Standards analysis	M/606 and 17 ETSI Public Enquiry drafts	Extract clauses, requirement identifiers, assessment methods, evidence expectations, and CRA mappings	Versioned standards-to-evidence matrix
Requirements synthesis	Legal and standards findings	Classify direct, indirect, and proposed requirements; assign subject, issuer, time, access, and status	Design requirements for a CRA-capable DPP

Research activity	Primary corpus or input	Analytical operation	Output
Artefact construction	Identity, zero-trust, DPP, SBOM/VEX, and semantic-web building blocks	Design control plane, evidence plane, subject model, evidence graph, and disclosure tiers	Reference architecture and conceptual data model
Cross-sector evaluation	Software, devices, remote services, industrial systems, second-life scenarios	Apply architecture to representative lifecycle and incident cases	Implementation patterns, limitations, and sector findings

Source: Authors' research design, adapted from design science principles in [21, 22].

The evaluation criteria are legal accuracy, traceability, evidentiary completeness, interoperability, security, privacy and confidentiality, temporal correctness, scalability, operational usefulness, and safety relevance. Legal accuracy requires that the architecture neither omits a material duty nor implies that a technical mechanism creates compliance automatically. Traceability requires every material claim to resolve to its subject, issuer, authority, source, time, and status. Interoperability concerns identifiers, schemas, semantics, APIs, and portable verification. Security includes resistance to unauthorised modification, fraudulent issuance, replay, stale evidence, provider compromise, and denial of access. Operational usefulness concerns the ability to support vulnerability triage, targeted patching, procurement, supplier management, conformity assessment, market surveillance, maintenance, repair, and second-life decisions.

The evidence cut-off for this paper is 15 August 2026. This date is material because the ETSI documents remain under Public Enquiry, CRA guidance continues to develop, and the European Business Wallet instrument remains a legislative proposal. All standards mappings therefore retain the source identifier, edition, date, clause or requirement identifier, and legal status. Later drafts, adopted standards, OJEU references, implementing acts, or legislative amendments may change the analysis and must be incorporated through controlled versioning rather than silent replacement. Appendix A specifies the fields used for this requirements-to-evidence mapping and source-status register.

The present research is principally analytical and design-oriented. It does not claim statistical evidence of industry-wide effectiveness. Scenario evaluation can establish internal coherence, legal fit, and technical feasibility, but it cannot by itself demonstrate adoption, cost reduction, or risk reduction in production. Later iterations of the research may add prototype implementation, performance testing, conformity-assessment review, expert interviews, and field validation. Such evidence should be reported separately from the normative and architectural argument.

1.8 Expected contributions

The article is designed to make six contributions:

- 1. Conceptual definition.** It defines the CRA-capable DPP as a product evidence infrastructure and distinguishes it from a static document, a public compliance label, a technical file repository, and a legal presumption of conformity.
- 2. Normative traceability.** It creates a disciplined mapping among CRA obligations, candidate ETSI requirements, evidence objects, assessment methods, and proposed architectural controls, while preserving the status of each source.
- 3. Subject separation.** It establishes model, release, instance, deployment, and remote-service layers as distinct evidence subjects. This resolves a central ambiguity in software-defined and connected products.

4. Dual-plane trust architecture. It separates organisational authority, mandates, identity, and policy enforcement from product evidence, and examines eIDAS 2.0 and the proposed European Business Wallet as a European control-plane implementation path.

5. Continuous conformity model. It connects pre-market conformity evidence with post-market vulnerability, update, configuration, support, incident, repair, and modification evidence throughout the support period without equating monitoring with automatic legal compliance.

6. Cross-domain assurance. It evaluates how the same evidence infrastructure can support cybersecurity operations, regulatory compliance, business processes, and controlled coordination with functional-safety lifecycle decisions across different product and sector patterns.

1.9 Organisation of the article

Part I establishes the regulatory and standards foundation. It distinguishes CRA legal obligations, ESPR and DPP infrastructure requirements, candidate ETSI requirements, indirect evidence-governance needs, and the architectural choices proposed in this article. Part II defines the CRA-capable DPP and its reference architecture, including the five-subject model, dual-plane trust architecture, evidence graph, SBOM and VEX integration, update and roll-back evidence, issuer authority, and access control. Part III evaluates operational value and implementation patterns across cybersecurity operations, conformity processes, lifecycle services, functional-safety change control, software, connected products, remote services, industrial systems, and second-life applications. Part IV evaluates legal fidelity, interoperability, scalability, security, confidentiality, provider dependency, governance, and the limits of automated conformity claims. Appendix A specifies the schema used for the requirements-to-evidence mapping and source-status register, including version and supersession fields that permit later legal and standards changes to be incorporated explicitly rather than silently. Section 6 concludes the article and presents a heat map indicating where CRA-capable DPPs can most directly support CRA implementation.

2 Part I - Regulatory and Standards Foundation

This Part answers RQ1 and establishes the normative foundation from which the reference architecture is derived. It analyses the Cyber Resilience Act (CRA), the Ecodesign for Sustainable Products Regulation (ESPR) and its Digital Product Passport (DPP) implementing framework, the European digital identity framework, Commission standardisation request M/606, and the 17 ETSI CRA vertical standards available as Public Enquiry drafts at the evidence cut-off. The purpose is not to assemble a broad catalogue of cybersecurity controls. It is to identify which legal and standards requirements create evidence that can be represented, resolved, or governed through a CRA-capable DPP, and to preserve the legal status of every inference.

The analysis applies three tests to each candidate requirement. First, the source-status test asks whether the requirement arises from binding Union law, an OJEU-cited harmonised standard, a draft candidate harmonised standard, a legislative proposal, or the authors' design. Second, the evidence test asks what information, document, test result, operational observation, or actor statement is needed to demonstrate or operate the requirement. Third, the subject test asks whether that evidence concerns a product model, hardware or software release, physical or logical instance, deployment configuration, remote data processing solution, economic operator, or lifecycle event. Only after these tests are applied is an architectural requirement derived.

2.1 Normative hierarchy and permissible inferences

The expression "CRA-capable DPP" is an analytical construct. The CRA does not require a Digital Product Passport. It requires manufacturers and other responsible economic operators to meet product cybersecurity, vulnerability handling, documentation, conformity assessment, user-information, and post-market duties. A DPP can support these duties only where its data, trust, access, temporal, and evidence properties are suitable for the specific obligation [1].

The ESPR creates a legally defined DPP infrastructure, but a DPP duty arises only where an applicable delegated act or another Union instrument requires one for a product group. The ESPR therefore supplies reusable infrastructure properties - persistent identifiers, product granularity, machine-readable data, interoperability, controlled access, persistence, and lifecycle linking - without converting the DPP into a universal CRA obligation [4, Arts. 9-11].

A second distinction concerns harmonised standards. Compliance with an OJEU-cited harmonised standard may create a presumption of conformity only for the legal requirements and parts covered by that standard. The six DPP standards cited in July 2026 create such a presumption for relevant ESPR Articles 10 and 11 requirements. They do not create a presumption of CRA conformity. Conversely, the 17 ETSI CRA documents analysed below are draft candidate harmonised standards under Public Enquiry. At the evidence cut-off, they provide technically important requirements and assessment models but no Article 27 CRA presumption of conformity [24]; [1], Art. 27; [3].

Table 3: Source-status categories and permitted claims

Source category	Status at evidence cut-off	Permitted use in this article	Claim that is not permitted
CRA, Regulation (EU) 2024/2847	Binding Union law. Article 14 applies from 11 September 2026; general application from 11 December 2027.	Identify legal duties, actors, timelines, essential requirements, documentation and conformity obligations.	The CRA mandates a DPP or accepts a DPP as automatic proof of conformity.
ESPR and DPP implementing law	Binding framework. Product-specific DPP duties depend on delegated or sectoral law.	Use legally defined identifiers, granularity, access, persistence, registry and interoperability mechanisms.	Every CRA product must possess an ESPR DPP.
EN 18216, EN 18219-18223:2026	OJEU-cited harmonised DPP standards under Decision (EU) 2026/1736.	Claim presumption only for covered ESPR Articles 10 and 11 requirements.	The standards establish CRA cybersecurity conformity.
Commission request M/606	Binding instruction to European standardisation organisations; not a product requirement addressed directly to manufacturers.	Define the planned horizontal and vertical CRA standardisation architecture and requested outcomes.	M/606 itself supplies presumption of conformity.
ETSI EN 304 617-627 and 631-636 drafts	Seventeen candidate harmonised standards in Public Enquiry; not yet OJEU-cited.	Extract product-specific requirements, assessment criteria, evidence expectations and emerging implementation patterns.	A product conforming to the draft already benefits from CRA presumption of conformity.
eIDAS 2.0 and DPP registry identity rules	Binding trust-services and identity law, supplemented by binding registry rules.	Ground legal-person verification, signatures, seals, attestations, status and delegated access.	A valid signature alone proves that the signer had the organisational authority to make the claim.
European Business Wallet proposal	Commission legislative proposal; ordinary legislative procedure ongoing at the evidence cut-off.	Evaluate a prospective control-plane mechanism for legal persons, roles, mandates and secure transactions.	EBW functions are currently binding obligations or universally available infrastructure.

Source category	Status at evidence cut-off	Permitted use in this article	Claim that is not permitted
CRA-capable DPP reference architecture	Authors' design artefact.	Derive requirements, trust relationships, evidence objects and implementation patterns from the analysed corpus.	Architectural choices are existing legal mandates unless an explicit source is identified.

Source: Authors' classification based on [1–5, 19, 20, 24].

Normative boundary. A CRA-capable DPP may organise and make evidence verifiable. It cannot change the responsible economic operator, replace the applicable conformity assessment procedure, or transform a structurally valid passport into proof that the underlying product satisfies the CRA.

2.2 The CRA as a lifecycle evidence regime

2.2.1 Scope, regulated object, and remote dependencies

The CRA applies to products with digital elements made available on the Union market where the intended purpose or reasonably foreseeable use includes a direct or indirect logical or physical connection to a device or network [1, Art. 2(1)]. The regulated category includes hardware and software products and can include a remote data processing solution where the remote processing is designed and developed by, or under the responsibility of, the manufacturer and is necessary for one of the product's functions. This makes product scope dependent on functional and responsibility boundaries rather than on physical enclosure alone [1, Art. 3].

This scope has a direct evidence implication. A connected product may change materially without a change to its physical identifier. A cloud API, policy service, database, model-serving endpoint, or remote management component may be updated independently from the device or local software. A technical file organised only around the commercial model therefore risks obscuring which release, service version, and deployment state was assessed. The CRA does not create the model-release-instance-deployment-service taxonomy used in this paper. That taxonomy is an architectural response to the CRA's requirement that evidence remain attributable to the product and processes actually assessed.

Sectoral exclusions must remain explicit. Certain medical devices, in vitro diagnostic medical devices, and products covered by the motor-vehicle type-approval cybersecurity regime are excluded from the CRA, while general-purpose ICT used in the same sectors can remain within scope [1, Art. 2]. A CRA-capable DPP therefore requires a machine-readable regulatory-scope statement. It must not infer CRA applicability merely because an object has a digital identifier or because a DPP exists under another Union regime.

2.2.2 Lifecycle responsibility, support, and modification

Article 6 links market access to both product properties and manufacturer processes. Products may be made available only where the applicable Annex I Part I requirements are met, subject to proper installation, maintenance, intended or reasonably foreseeable use, and installation of necessary security updates, and where the manufacturer's vulnerability handling processes comply with Annex I Part II [1, Art. 6]. This dual structure already distinguishes a product state from the organisational process that maintains that state.

Article 13 makes the cybersecurity risk assessment a lifecycle artefact. It must be documented, updated as appropriate during the support period, and based on intended purpose, reasonably

foreseeable use, operating conditions, assets to be protected, and expected use time. Manufacturers must document the applicability of Annex I requirements, exercise due diligence for third-party components, systematically record relevant cybersecurity information, and update the risk assessment where appropriate [1, Art. 13(2)-(7)]. These duties require evidence lineage: a later assessment must preserve its relationship to the earlier assessment, the product version to which each applies, and the new facts that caused the change.

The support period must normally be at least five years unless the expected product use is shorter. Its determination must consider expected use, user expectations, intended purpose, relevant product-lifetime law, the operating environment, and support periods of core third-party components. The determination must be documented. Security updates made available during the support period must remain available for at least ten years after issuance or for the remainder of the support period, whichever is longer [1, Art. 13(8)-(9)]. A support-end date is thus not merely commercial metadata. It is an evidence-bearing regulatory state that affects vulnerability handling, user communication, procurement, maintenance, and residual risk.

Responsibility can also change. Importers and distributors become manufacturers where they place a product on the market under their own name or trademark or substantially modify it. Another natural or legal person that carries out a substantial modification and makes the product available can likewise become the manufacturer for the affected part or, where the modification affects the product as a whole, for the entire product [1, Arts. 21-22]. Economic-operator traceability must be retained for ten years [1, Art. 23]. A trustworthy DPP implementation should therefore create a new, attributable lineage node for a substantial modification rather than silently overwrite the original product record.

2.2.3 Direct CRA information and evidence requirements

For this study, a requirement is classified as direct where the CRA explicitly requires information, documentation, a process record, a test result, an identifier, a declaration, a notification, or other evidence. Direct does not mean that the CRA requires the information to be stored in a DPP. The DPP role may be to carry the information, resolve it from an authoritative source, bind it to the correct subject, expose only a status or hash, or provide controlled access to a restricted technical-file object.

Annex I Part I defines the required product properties: risk-appropriate cybersecurity; absence of known exploitable vulnerabilities at market placement; secure defaults; security updates; access control; confidentiality; integrity; data minimisation; availability and resilience; non-interference with other devices and networks; attack-surface reduction; exploit mitigation; security logging and monitoring; and secure removal or transfer of data and settings [1, Annex I Part I]. Annex I Part II then requires documented components and vulnerabilities, a machine-readable SBOM covering at least top-level dependencies, remediation without delay, regular security tests, public information on fixed vulnerabilities, coordinated vulnerability disclosure, a reporting contact, secure update distribution, and free and timely security updates [1, Annex I Part II].

Annex II defines user-facing information. It includes product and manufacturer identification, the vulnerability reporting contact, intended purpose and security environment, reasonably foreseeable risks, the declaration of conformity or its location, support type and end date, and instructions for secure commissioning, operation, updates, decommissioning, data removal, and integration. The full SBOM need not be public; Annex II requires access information only where the manufacturer decides to make the SBOM available to users [1, Annex II].

Annex VII defines the technical-documentation evidence set. It includes compliance-relevant software versions, system architecture and component relationships, vulnerability handling pro-

cesses, the SBOM, disclosure policy, reporting contact, secure update solution, production and monitoring processes, the cybersecurity risk assessment and applicability rationale, support-period determination, standards and specifications applied, test reports, and the EU declaration of conformity [1, Annex VII]. Article 31 requires this documentation to be created before market placement and continuously updated where appropriate, at least during the support period. This provides the strongest legal basis for the continuous-conformity framing used in this article, while "continuous conformity" itself remains an analytical term [1, Art. 31].

Table 4: CRA obligations translated into evidence subjects and DPP treatment

CRA source	Required information or process	Primary evidence subject	Representative evidence object	CRA-capable DPP treatment
Arts. 13(15)-(17); Annex II; Annex V	Unique product and manufacturer identification; contact and responsible actor.	Model/type; batch or instance; legal person.	Product identifier; manufacturer identity; vulnerability contact; declaration metadata.	Public identity view plus verified issuer and economic-operator links.
Art. 13(2)-(4); Annex II(4)-(5); Annex VII(1),(3)	Intended purpose, foreseeable use and misuse, operating environment, protected assets, applicability rationale.	Model and deployment profile.	Security context profile; risk assessment; applicability matrix; residual-risk statement.	Public summary; restricted assessment and machine-readable scope constraints.
Annex I Part II(1); Annex VII(2),(8)	Identify components and vulnerabilities; machine-readable SBOM covering at least top-level dependencies.	Exact software or firmware release/build.	SBOM, component identifiers, supplier/origin, hash and generation metadata.	Restricted composition object cryptographically bound to the assessed release; public access pointer only where required.
Annex I Part I(2)(a); Part II(2)-(4)	No known exploitable vulnerabilities at market placement; remediation and disclosure of fixed vulnerabilities.	Release/build and time-bounded assessment context.	Candidate-vulnerability correlation; exploitability decision; mitigation or VEX; advisory; remediation record.	Versioned vulnerability claim with source cutoff, retrieval time, status, evidence hash and supersession.
Annex I Part I(2)(b); Annex II(8)(a)	Secure-by-default configuration and reset capability.	Model default and deployed configuration.	Secure-default profile; reset method; configuration assessment; approved deviations.	Separate declared baseline from operator-observed configuration; never expose secrets.
Annex I Part I(2)(c); Part II(7)-(8); Art. 13(9)	Secure, timely and where applicable automatic security updates; update availability over time.	Release, update package, installed instance.	Package hash and signature; target version; advisory; installation result; rollback and recovery record.	Link released, delivered, installed and verified states; preserve update and rollback lineage.
Art. 13(8),(18)-(19); Annex II(7)	Support period, support type, end date, end-of-support notice and long-term information availability.	Model and release family; affected instances.	Support determination; support entitlement; end-of-support notice; migration guidance.	Public support status with dated evidence; restricted rationale and component-support dependencies.
Annex I Part I(2)(l); Annex VII(2)(c)	Security-relevant recording and monitoring with user opt-out where applicable.	Product capability and deployment profile.	Logging capability statement; time-source assurance; retention policy; test report; audit-record hash.	Expose capabilities and controlled pointers, not raw logs or personal telemetry.
Annex I Part II(5)-(6); Art. 13(17)	Coordinated vulnerability disclosure policy and contact address.	Manufacturer and product family.	CVD policy; reporting endpoint; acknowledgement and case identifiers.	Public contact and policy; authenticated restricted case linkage.
Art. 14	Staged reporting of actively exploited vulnerabilities and severe incidents.	Affected products, releases, instances and jurisdictions.	24-hour warning; 72-hour notification; final report; affected-version set; mitigation and fixed release.	Prepare structured reporting inputs and affected-population evidence; do not replace the official platform.

CRA source	Required information or process	Primary evidence subject	Representative evidence object	CRA-capable DPP treatment
Arts. 27-32; Annexes V-VIII	Conformity route, applied standards, tests, notified body or certificate, declaration and CE marking.	Model/type and compliance-relevant release.	Requirement-to-control matrix; test reports; certificate; EU declaration; OJEU reference status.	Restricted technical-file index plus public declaration metadata; no autonomous compliance flag.
Arts. 21-23; Annex II(8)(f)	Responsibility changes, substantial modification, integrator information and economic-operator traceability.	Modified product lineage; responsible legal person.	Modification event; affected scope; new risk assessment; new declaration; actor chain.	Create an immutable lineage branch and new authority context rather than overwrite history.
Annex I Part I(2)(m); Annex II(8)(d)	Secure removal or transfer of data and settings; secure decommissioning.	Instance and operator handover event.	Reset attestation; data-removal result; credential revocation; transfer record.	Record verifiable result or hash; exclude deleted data and secrets.
Art. 3; Annex VII(2)(a)	Architecture and relevant remote data processing solutions within manufacturer responsibility.	Remote service release and product-service dependency.	Service identifier; API version; compatibility; incident/advisory link; architecture evidence.	Represent the service as an independent evidence subject linked to product models, releases and deployments.

Source: Authors' mapping based on [1]. VEX and some instance-level evidence objects are standards-derived or architectural representations; they are not terms mandated universally by the CRA.

2.2.4 Reporting, conformity assessment, and presumption of conformity

Article 14 creates time-critical reporting duties. For an actively exploited vulnerability, the manufacturer must submit an early warning within 24 hours, a fuller notification within 72 hours unless already provided, and a final report no later than 14 days after a corrective or mitigating measure becomes available. Severe incidents follow a parallel 24-hour and 72-hour sequence and a final report within one month [1, Art. 14]. These deadlines make identification of affected product versions, market geography, mitigations, and fixed releases an operational evidence problem. A CRA-capable DPP can supply structured inputs and dissemination targets, but the official reporting platform remains the legally relevant channel.

The conformity route depends on product category and the standards or certification instruments applied. Article 32 allows internal control, EU type examination followed by conformity to type, full quality assurance, or an applicable European cybersecurity certification scheme. Important products in Class I may use internal control only under the conditions set out in Article 32; partial or absent use of relevant harmonised standards can trigger third-party assessment. Class II and certain critical products are subject to stronger routes [1, Arts. 7-8, 32; Annexes III-IV]. The passport must therefore record the actual route and scope, not merely the existence of a test report.

Article 27 is decisive for legal claims. Presumption of conformity arises only where a product and the manufacturer's processes conform to a harmonised standard or part whose reference has been published in the OJEU, and only for the Annex I requirements covered. Common specifications and specified cybersecurity certification schemes can provide equivalent legal effects within their defined scope [1, Art. 27]. A DPP may record the standard edition, clauses applied, OJEU status, conformity route, assessor, and evidence. It cannot itself generate the legal presumption.

2.3 The European DPP infrastructure as a reusable evidence substrate

2.3.1 Conditional DPP duties and product granularity

The ESPR does not impose one immediate DPP obligation on every product. Article 9 requires a DPP where the applicable delegated act establishes the information requirement. That act determines the data, carrier, presentation, access rights, update rights, duration, and whether the passport is established at model, batch, or item level [4, Art. 9]. Other Union legislation can use the same infrastructure, as is already reflected in the 2026 registry rules for batteries, construction products, toys, detergents, and future product regimes [5, Art. 1].

The paper therefore uses two meanings with care. A statutory DPP is a passport required under applicable Union product law and conforming to that regime. A CRA-capable DPP is the proposed evidence architecture. It may be implemented on the statutory DPP infrastructure, coexist with it, or use equivalent identifiers and resolvers for software products that do not have a physical data carrier. Only the first meaning carries the legal label and obligations created by the relevant DPP legislation.

2.3.2 ESPR architectural requirements

ESPR Articles 10 and 11 already establish several infrastructure properties required by a CRA evidence system. A DPP must be connected through a data carrier to a persistent unique product identifier. Data must be based on open standards and, as appropriate, be machine-readable, structured, searchable, transferable, and exchanged without vendor lock-in. Data must refer to the legally specified model, batch, or item level. Access is governed by actor rights, and a back-up copy must be made available through a DPP service provider [4, Art. 10].

The technical design must provide technical, semantic, and organisational interoperability; free and easy role-based access; persistence after insolvency or cessation; restricted update rights; authentication, reliability and integrity; and links from a replacement passport to the original passport [4, Art. 11]. These requirements are directly reusable for CRA evidence, especially persistent identifiers, subject granularity, lineage, access policies, continuity, and controlled updates. They do not specify the cybersecurity evidence content analysed in Section 2.2.

2.3.3 Harmonised DPP infrastructure standards

Commission Implementing Decision (EU) 2026/1736 cited six DPP standards in the OJEU on 15 July 2026. Compliance with these standards can create a presumption of conformity with the covered requirements of ESPR Articles 10 and 11. The standards provide a mature interoperability substrate for CRA-capable evidence services, but their legal presumption is confined to DPP infrastructure requirements [24].

Table 5: OJEU-cited DPP standards and reuse in a CRA-capable evidence architecture

Standard	DPP infrastructure function	CRA-capable reuse	Legal boundary
EN 18216:2026	Data exchange protocols.	Protocol-neutral retrieval and exchange of signed evidence objects, status data and controlled references.	Presumption concerns covered ESPR DPP exchange requirements, not CRA security controls.
EN 18219:2026	Unique identifiers.	Persistent identifiers for product models, batches, items, actors and resolvable evidence subjects.	CRA release, deployment and service identifiers may require additional namespaces and mappings.

Standard	DPP infrastructure function	CRA-capable reuse	Legal boundary
EN 18220:2026	Data carriers.	Resolution from a physical product or packaging to the appropriate pass-port and evidence view.	Software-only products and remote services may require non-physical resolution mechanisms.
EN 18221:2026	Data storage, archiving and persistence.	Evidence continuity, historical versions, auditability and resilience after provider or manufacturer cessation.	Retention and confidentiality must still follow CRA, product law, contract and data-protection requirements.
EN 18222:2026	APIs for lifecycle management and searchability.	Programmatic issue, update, retrieval, status, query and lifecycle event interfaces.	API conformance does not validate the factual correctness of evidence.
EN 18223:2026	System interoperability.	Federated manufacturer, supplier, operator, conformity-assessment and authority evidence exchange.	Semantic and trust profiles for CRA evidence remain to be defined by the proposed architecture and relevant standards.

Source: Authors' analysis based on Commission Implementing Decision (EU) 2026/1736 [24].

2.3.4 The DPP registry: structural assurance, not product conformity

Commission Implementing Regulation (EU) 2026/1778 operationalises the DPP registry. It provides secure user and API interfaces, verified economic-operator and value-chain-actor profiles, identification and authorisation schemes, a semantic repository, logging, registration identifiers, and DPP service-provider references. Legal persons can be verified through qualified electronic seals or qualified electronic attestations of attributes under eIDAS, and verified actors may delegate access to third-party users where Union law permits [5, Arts. 3-6].

At registration, the registry automatically checks semantic conformity, mandatory-data coherence where relevant, the legally required model, batch, or item granularity, and specified registry links. It then creates a persistent registration identifier. A downloadable proof of registration identifies the product and verified operator, records the registration time, contains a hash of the DPP version, and is protected by a qualified electronic seal and Commission timestamp. Registration data are versioned and logged [5, Arts. 8-10].

The Regulation draws an important boundary: automated registry verification is not proof of compliance with the Union rules applicable to the product. The registry can prove that a particular DPP version was registered by a verified actor at a certain time and that it passed defined structural and semantic checks. It cannot prove that a vulnerability assessment is technically sound, that an SBOM is complete, that a patch was installed on a specific device, or that the product meets the CRA [5, recital 16; Art. 9]. This distinction is central to the reference architecture.

Table 6: Assurance provided by the DPP registry and residual CRA evidence questions

Registry capability	What it can establish	What it cannot establish	Required additional CRA evidence
Verified operator profile	Identity and establishment of the registered economic operator under the permitted eIDAS mechanisms.	That the individual or workload issuing a specific technical claim had the relevant internal mandate.	Role, mandate, delegation scope, issuer credential and status.
Semantic conformity check	Required fields, structures, types and controlled vocabularies conform to the applicable data model.	That the represented values are factually true or technically sufficient.	Signed source claim, assessment method, evidence hash, provenance and independent verification where required.

Registry capability	What it can establish	What it cannot establish	Required additional CRA evidence
Granularity check	Passport is registered at the legally specified model, batch or item level.	That the granularity is sufficient for release, deployment or remote-service cybersecurity state.	Linked release/build, instance, deployment and service subjects.
Registration identifier and proof	A DPP version was registered by the identified operator at a recorded time; the proof is sealed and time-stamped.	That the product or passport was CRA-conformant at that time.	Requirement-to-evidence mapping, conformity route, test results and current status.
Versioning and registry log	Changes to registered data are recorded and time-stamped.	That evidence outside the registry is immutable or that a later operational event occurred.	Cryptographic evidence object identifiers, external status services and lifecycle event attestations.
Service-provider and backup references	The registered passport can be associated with a provider and backup where required.	Long-term independent verifiability after provider failure unless evidence is portable.	Exportable signed evidence, open verification, continuity and migration controls.

Source: Authors' analysis based on [5].

2.4 Identity, authority, and the control-plane foundation

Cybersecurity evidence is not trustworthy merely because a data object is signed. The verifier must establish the legal person represented by the issuer, the issuer's organisational role, the mandate or delegation under which the claim was made, the scope of products and requirements covered, and whether the authority and credential remained valid at the relevant time. Manufacturer, component supplier, test laboratory, notified body, system integrator, operator, repairer, and market-surveillance claims must remain distinguishable.

The amended eIDAS framework provides relevant legal and technical building blocks. European Digital Identity Wallets can support qualified electronic signatures and seals and are intended to enable powers of representation and e-mandates. Qualified trust services and electronic attestations of attributes can provide identity, integrity, authenticity, status, and revocation functions [19]. The 2026 DPP registry rules already use qualified seals and qualified attestations to verify legal persons, which demonstrates a direct regulatory bridge between product-passport infrastructure and European trust services [5, Arts. 4-6].

The Commission proposal for European Business Wallets extends this direction to economic operators. Its proposed functions include identification and authentication, signing and sealing, document exchange, notifications, roles, attributes, and authorisations. The proposal requires role-attribute mappings to be verifiable, auditable, revocable and traceable to legitimate issuers, and anticipates controls against role conflict, over-delegation and expired authorisations [20]. At the evidence cut-off, these provisions are prospective rather than binding. The architecture therefore treats EBW as a likely European control-plane implementation path, not as a current prerequisite for a CRA-capable DPP.

Table 7: Control-plane assertions required for accountable CRA evidence

Assertion	Possible European trust mechanism	Verification question	Failure prevented
Legal-person identity	Qualified electronic seal, qualified attestation of attributes, verified DPP registry profile, future EBW owner data.	Which legal entity is accountable for the claim?	Anonymous or falsely attributed evidence.
Organisational role	Role or professional attribute credential; authoritative directory or registry.	Is the issuer acting as manufacturer, supplier, laboratory, notified body, integrator or operator?	Role confusion and reliance on claims outside the issuer's competence.
Mandate or delegation	Power of representation, e-mandate, delegated registry access, policy decision.	Was the person, workload or agent authorised to issue this claim for this product and requirement?	Valid signature by an unauthorised employee or service.
Scope and limits	Credential constraints, policy attributes, product and requirement identifiers.	Which products, versions, sites, claims and time period are covered?	Over-broad claims and over-delegation.
Credential and mandate status	Revocation and status services; wallet or registry validity information.	Was the authority valid at issuance and is the claim still acceptable?	Reliance on expired, revoked or compromised authority.
Action accountability	Qualified or advanced seal/signature, signed policy decision, timestamp and audit receipt.	Who approved, issued, changed or withdrew the evidence?	Untraceable changes and repudiation.
Relying-party policy	Zero-trust policy engine using identity, role, evidence and risk context.	Is this issuer and evidence type sufficient for the requested decision?	Authentication being mistaken for authorisation or assurance.

Source: Authors' synthesis based on [5, 19, 20].

2.5 CRA standardisation architecture under M/606

Commission Implementing Decision C(2025) 618, standardisation request M/606, organises CRA standardisation into 41 requested deliverables. Entries 1-14 concern horizontal product cybersecurity properties aligned with Annex I Part I. Entry 15 addresses horizontal vulnerability handling aligned with Annex I Part II. Entries 16-41 concern vertical standards for specific important and critical product categories [2]. The request requires objective and verifiable criteria, attention to intended purpose and reasonably foreseeable use, and consideration of health and safety effects. It also requires coherence between horizontal and vertical standards.

This structure matters for evidence architecture. Horizontal standards define reusable requirement and assessment patterns. Vertical standards adapt them to product functions, use cases, attack surfaces, risk or harm classes, and product-specific evidence. A CRA-capable DPP should therefore not hard-code one vertical schema as the universal CRA model. It should identify the applicable horizontal and vertical profiles, their exact editions and clauses, and the evidence required by each.

The 17 ETSI drafts correspond to a substantial but incomplete subset of the M/606 vertical programme: browsers through routers and switches, plus selected smart-home, toy, wearable, virtualisation, and firewall categories. M/606 also requests standards for identity and privileged-access management, security-relevant processors and controllers, hardware security devices, smart meter gateways, smartcards and secure elements. Those categories are outside the 17-document ETSI Public Enquiry corpus analysed here and may be addressed by other European standardisation work. Completeness with respect to ETSI must therefore not be confused with completeness with respect to all CRA standardisation.

Table 8: M/606 standardisation layers and their evidence implications

M/606 layer	Requested scope	Evidence implication for a CRA-capable DPP	Status consequence
Horizontal entry 1	Risk-based appropriate level of cybersecurity.	Risk context, intended purpose, operating environment, applicability and residual-risk evidence.	A reusable evidence profile is needed across product categories.
Horizontal entries 2-14	Product properties: exploitable vulnerabilities, secure default, updates, access, confidentiality, integrity, minimisation, availability, non-interference, attack surface, exploit mitigation, monitoring, removal/transfer.	Requirement-to-control-to-test graph; product- and deployment-specific assessment evidence.	Product-specific standards may refine applicability and test methods.
Horizontal entry 15	Vulnerability handling processes.	SBOM, vulnerability intake, assessment, remediation, disclosure, update distribution and support evidence.	Dynamic evidence must be updated throughout the support period.
Vertical entries 16-41	Important and critical product categories.	Use-case profiles, component scope, product-specific requirements, test activities, verdicts and evidence lists.	The passport must record exactly which vertical standard and clauses apply.
Cross-standard coherence	Vertical standards implement and further develop horizontal standards.	Common semantic identifiers for requirements and evidence while preserving product-specific extensions.	Draft cross-references and editions must be versioned until the corpus stabilises.
Presumption of conformity	Legal effect arises only after OJEU citation and only for covered requirements.	Store OJEU status, edition, scope, partial application and conformity route as first-class metadata.	Draft conformity must never be presented as legal presumption.

Source: Authors' analysis based on M/606 [2] and CRA Article 27 [1].

2.6 Analysis of the 17 ETSI CRA vertical standards

2.6.1 Corpus, status, and extraction method

ETSI announced 17 Final Draft European Standards for Public Enquiry on 13 August 2026. The public corpus comprises EN 304 617 through EN 304 627 and EN 304 631 through EN 304 636. At the evidence cut-off, the documents were in the approval or Public Enquiry process and had not been cited in the OJEU. The paper therefore refers to them as ETSI drafts or candidate harmonised standards [3].

An additional ETSI work item, EN 304 642 on network functions of telecommunications systems, was at early-draft stage and planned to build primarily on GSMA NESAS and 3GPP SCAS. It is tracked but excluded from clause-level mapping because it was not part of the 17 Public Enquiry drafts [25]. This prevents an unstable work item from being presented with the same evidentiary maturity as the analysed corpus.

The extraction focused on requirements and assessment criteria that materially affect a DPP evidence architecture: product and version identity; component inventories and SBOMs; vulnerability sources and exploitability decisions; secure defaults; security updates, installed-version reporting, rollback and recovery; logging and trustworthy time; remote data processing dependencies; support and end-of-support; test evidence; and safety- or harm-sensitive applicability. The tables report selected high-value clauses. They are not a substitute for a complete confor-

mity assessment against the final published standards.

2.6.2 Browsers, credential and endpoint security, VPN, management, and SIEM

EN 304 617 provides the most explicit vulnerability evidence model in the corpus. Its assessment of known exploitable vulnerabilities uses an SBOM hash, raw component-to-vulnerability correlation output, exploitability sources and retrieval timestamps, a VEX or mitigation document tied to the assessed product version, verification records, remediation tracking, and snapshots of vulnerability and exploitation databases. This is a direct standards basis for time-bounded, release-specific vulnerability claims rather than an undated statement that a product is secure [26, cls. 5.3.1 and 6.3.1].

EN 304 618 combines vulnerability and high-sensitivity audit requirements for password managers. REQ-KEV-NORM-1 addresses exploitable vulnerabilities affecting confidentiality, integrity, availability or authenticity. REQ-LOG-CRIT-1 requires comprehensive tamper-protected audit logs for critical use cases. The DPP implication is twofold: the release requires current exploitability evidence, while the passport should expose only logging capabilities and controlled audit references, never credentials, vault contents or raw sensitive logs [27, cls. 5.3.1.1 and 5.14.3.1].

EN 304 619 addresses antimalware products. Its secure-default controls and assessment of end-of-support notification illustrate that compliance evidence includes configuration behaviour and user communication, not only component vulnerability data. REQ-FRD-01 assesses whether the product clearly identifies and tracks the end of support, informs administrative and end users where feasible, and provides continued-use, upgrade, replacement or migration information [28, cl. 6.15.1].

EN 304 620 makes update authenticity operational. REQ-SU-07 requires cryptographically signed updates and signature verification before installation. Related requirements reject updates with unexpected hashes and hashes not signed by an authorised authority. A CRA-capable DPP should bind the update package identifier, hash, signer authority, target release, verification result and installation outcome rather than record only that an update exists [29, cls. 5.5.8-5.5.10].

EN 304 621 extends the evidence chain for network management systems. Its security-update requirements cover update capability, independent scheduling, signature verification, downgrade prevention, failed-update recovery, postponement, explicit authorisation and auditable metadata for intentional rollback, user notification, component version tracking, and update download logs. This standard demonstrates why release publication, device installation, approved rollback and recovery must be separate lifecycle events [30, cl. 5.5].

EN 304 622 provides product-specific alternatives for the absence of known exploitable vulnerabilities in SIEM products. REQ-KEV-01 recognises testing results, recent disclosure windows, documented risk and mitigation, and vulnerability sources such as EUVD or direct notification. It also treats time correction as a monitored event in relevant use cases. The resulting evidence claim must record the chosen sub-requirement, source context, date, justification and mitigation rather than collapse the alternatives into one Boolean status [31, cls. 5.2.3 and 5.3.2].

Table 9A: Selected ETSI requirements and DPP implications: EN 304 617-622

Draft standard	Selected requirement or assessment	Evidence expected	DPP implication
EN 304 617 - browsers	REQ-KEV-1; assessment cl. 6.3.1.	SBOM hash; raw correlation; source and retrieval time; VEX/mitigation version; verification and remediation records; database snapshots.	Issue a time-bounded release claim linked to exact SBOM and evidence-source snapshot.
EN 304 618 - password managers	REQ-KEV-NORM-1; REQ-LOG-CRIT-1.	Exploitability evidence and tamper-protected critical audit logs.	Separate vulnerability status from logging capability; restrict raw logs and secrets.
EN 304 619 - antimalware	Secure-default requirements; REQ-FRD-01 end-of-support assessment.	Configuration tests; clear, logged and auditable support-end notification; migration guidance.	Model secure baseline and public support state with dated notification evidence.
EN 304 620 - VPN products	REQ-SU-07 to REQ-SU-09.	Signed update; pre-installation signature check; expected hash; authorised signing authority.	Bind update artefact, signer authority, target release and verification result.
EN 304 621 - network management systems	SU_UPDATES-1 to SU_UPDATES-15, especially 8-12.	Version tracking; authorised rollback metadata; recovery; user notice; update logs.	Represent release, install, rollback and recovery as distinct auditable events.
EN 304 622 - SIEM	REQ-KEV-01; REQ-ALC-02.	Testing, recent-disclosure or mitigation rationale; vulnerability source; time-drift correction events.	Record the applicable alternative, time basis, mitigation and time-assurance evidence.

Source: Authors' extraction from [26–31]. Requirement wording and numbering remain subject to change during the standards process.

2.6.3 Boot, PKI, interfaces, operating systems, routers, modems, and switches

EN 304 623 requires boot-manager assessment to be anchored in architecture. Evidence includes a pre-market vulnerability assessment, SBOM and component inventory, a vulnerability-database cutoff date, shipped configuration, and architectural references supporting claims that a component is absent or non-reachable. Annex C further maps threat modelling, boot flow and trust boundaries, cryptographic architecture, memory architecture, interfaces, recovery mechanisms and isolation boundaries. A non-exploitability claim therefore requires a link to the relevant attack-surface and architecture evidence, not merely a VEX status value [32, cl. 6.3 and Annex C].

EN 304 624 applies the same principle to PKI and certificate-issuance software. ACC-PKI-KEV-01 requires product documentation, component inventories or bills of materials, version and patch levels, recognised vulnerability sources, scanning where feasible, and a product-specific determination that each identified vulnerability is not exploitable or is controlled through specific user guidance. The passport must keep the PKI software release distinct from the operated trust service and must never expose private keys [33, cl. 6.2].

EN 304 625 demonstrates that evidence extends beyond inventory. Its assessment for untrusted inputs requires identified input fields, acceptable behaviours, input-generation or coverage-guided fuzzing, static-analysis results, test logs and sufficiency analysis. A CRA-capable DPP should therefore link a requirement to an assessment method and test-result object; an SBOM alone cannot demonstrate robust input handling [34, cl. 6.2.2].

EN 304 626 requires operating-system testing for known exploitable vulnerabilities using public and manufacturer records, product-specific prioritisation, documented test selection, results, residual risk and compensating controls. It also defines remote data processing solutions in the product context. The assessed OS build, component set, deployment state and manufacturer-controlled remote service must therefore remain distinguishable [35, cls. 4.8.2 and 6.2.2].

EN 304 627 contains the clearest instance-binding requirement. REQ-KEV-1-01 requires a machine-readable SBOM for each product version, and the assessment confirms that the version declared in the SBOM matches the installed version reported by the product. REQ-UPDATE-1-03 requires reporting the installed version after an update, while REQ-UPDATE-1-04 addresses automatic update checking or a documented checking process. This directly supports the model-release-instance separation at the centre of Part II [36, cls. 5-6].

Table 9B: Selected ETSI requirements and DPP implications: EN 304 623-627

Draft standard	Selected requirement or assessment	Evidence expected	DPP implication
EN 304 623 - boot managers	Known-exploitable-vulnerability assessment; Annex C architecture artefacts.	SBOM/component inventory; database cutoff; shipped configuration; boot flow, trust boundaries, cryptography, interfaces, recovery and isolation evidence.	Bind exploitability or non-reachability claims to exact architecture and build.
EN 304 624 - PKI issuance software	ACC-PKI-KEV-01.	Inventory or BoM, versions and patch levels, scanning, recognised sources, non-exploitability or user-guidance justification.	Separate software-product evidence from operated PKI service and key material.
EN 304 625 - network interfaces	Cl. 6.2.2 untrusted-input assessment; KEV assessment.	Input analysis, fuzzing/static-analysis results, test logs, acceptable behaviour, mitigation and sufficiency evidence.	Link requirements to assessment methods and result objects, not composition alone.
EN 304 626 - operating systems	AC-KEVT; remote data processing definition.	Product-specific vulnerability list, prioritisation, tests, residual risk, compensating controls and remote-dependency context.	Bind exact OS build and deployed instance to associated remote services and current assessment.
EN 304 627 - routers, modems and switches	REQ-KEV-1-01; REQ-UPDATE-1-03 and -1-04.	Per-version machine-readable SBOM; declared-installed version match; installed-version report; update-check evidence.	Use separate model, release and device-instance subjects with observed installed state.

Source: Authors' extraction from [32–36].

2.6.4 Smart-home assistants and security products, toys, wearables, virtualisation, and firewalls

EN 304 631 makes the interaction between cybersecurity and safety explicit. SU-AUTO requires automated security updates where the product can connect to a public network and automated installation does not create unacceptable risk to essential-function availability or safety. The draft identifies emergency access and alarm functions as examples. Where automatic installation is not appropriate, the decision requires product-specific justification and user information. A CRA-capable DPP should therefore carry a security-update safety decision tied to the relevant deployment profile, not a universal automatic-update flag [37, cl. 5.5.3].

EN 304 632 develops availability evidence for smart-home security products. AVAI-TIME-

RECO-POW requires restoration after loss of power, AVAI-TIME-NETW requires local operation of time-sensitive functions where network access is not necessary, and AVAI-TIME-RECO-NETW requires clean reconnection after network loss. The standard also uses harm classes and reliable time for high-impact logging. The DPP can expose tested fallback and recovery capabilities and the approved configuration without disclosing security-sensitive topology [38, cls. 5.10 and 5.14].

EN 304 633 applies harm-sensitive logging to internet-connected toys. LOG-TIME-HIGH requires real-time timestamps where a function can cause high harm. LOG-STORE-MEDIUM and LOG-STORE-HIGH define differentiated persistence requirements, and the standard addresses factory reset and data portability. A consumer-facing passport can expose support, update and data-removal information, while detailed audit and harm-assessment evidence remains restricted [39, cls. 5.14-5.15].

EN 304 634 permits assessment of known exploitable vulnerabilities in personal wearables using a transitive SBOM or a vulnerability scanner applied to the default configuration. NKEV-MKAV requires a remediation plan, records of identified potential vulnerabilities, and validation of the documentation. This illustrates that top-level SBOM coverage, while the CRA minimum, may be insufficient for a product-specific standard and that default configuration is part of the assessment context [40, cl. 6.3.1].

EN 304 635 contains the broadest composition model. REQ-ASSUR-SBOM-001 requires a machine-readable SBOM with names, versions, supplier or origin information and persistent identifiers where available. The VES and CES extensions cover hypervisor artefacts, management interfaces, lifecycle controllers, scheduling and device virtualisation modules, runtime binaries and services, container engines, image management, networking and storage plugins, orchestrator components, admission controllers, schedulers, secrets-management components, policy enforcement, cluster lifecycle management and hardware-related dependencies. The standard therefore requires a runtime dependency graph rather than a simple package list [41, cl. 5.4.2].

EN 304 636 mirrors the per-version SBOM and installed-version matching pattern of EN 304 627 for firewalls, intrusion detection systems and intrusion prevention systems. It also requires audit events for communication failures with a remote data processing solution. Its scope excludes industrial operational-technology products and points to prEN 50770-1. The passport must record this scope boundary: a general firewall standard cannot be used to imply coverage of an industrial OT appliance outside its scope [42, scope; cls. 5-6].

Table 9C: Selected ETSI requirements and DPP implications: EN 304 631-636

Draft standard	Selected requirement or assessment	Evidence expected	DPP implication
EN 304 631 - smart-home assistants	SU-AUTO automated security updates.	Connectivity and update capability; safety and essential-availability risk decision; user information.	Add a deployment-specific cyber-safety update gate and documented exception.
EN 304 632 - smart-home security	AVAI-TIME-RECO-POW, AVAI-TIME-NETW, AVAI-TIME-RECO-NETW; harm-based logging.	Power and network recovery tests; local fallback; time and log evidence.	Represent safe fallback, recovery capability and tested deployment profile.
EN 304 633 - internet-connected toys	LOG-TIME-HIGH; LOG-STORE-MEDIUM/HIGH; factory reset.	Harm class, trustworthy time, retention behaviour, user controls and reset evidence.	Separate parent/consumer view from restricted safety and audit evidence.

Draft standard	Selected requirement or assessment	Evidence expected	DPP implication
EN 304 634 - personal wearables	NKEV-MKAV.	Transitive SBOM or scanner results for default configuration; remediation plan and validation records.	Support deeper dependency evidence and bind vulnerability status to assessed default configuration.
EN 304 635 - virtualisation/container stacks	REQ-ASSUR-SBOM-001 to -003.	Machine-readable SBOM covering runtime, engines, orchestrators, controllers, policy and hardware-related dependencies.	Represent a runtime dependency graph and actual deployment separately from vendor release.
EN 304 636 - firewalls, IDS and IPS	Per-version SBOM and installed-version match; REQ-RDPS-1-05; OT exclusion.	SBOM, observed version, remote-service failure events and scope evidence.	Bind product instance and remote service; record that industrial OT products require another profile.

Source: Authors' extraction from [37–42].

2.7 Cross-cutting standards findings for product evidence infrastructure

The 17 drafts differ substantially in product scope and drafting style, but they converge on a common evidentiary pattern. Conformity is assessed through a relation among product context, a requirement, an assessment objective, preparation artefacts, activities, a verdict, and retained evidence. This structure is more demanding than a data inventory. It requires the DPP to preserve why a claim was accepted, for which product subject, using which evidence and source state, by which authorised actor, and for how long the result remains meaningful.

Five findings are especially relevant. First, product version is frequently a normative assessment boundary. Second, the assessed composition must be linked to that version and, in EN 304 627 and EN 304 636, to the version actually reported by the installed product. Third, known-exploitable-vulnerability claims are temporal and source-dependent. Fourth, mitigation or non-reachability claims require product-specific verification and often architecture evidence. Fifth, updates are lifecycle events with authenticity, installation, rollback, recovery and safety implications. These findings jointly reject a static model-level passport as the sole CRA evidence object.

The standards also distinguish evidence retention from evidence disclosure. Detailed SBOMs, architecture diagrams, fuzzing records, vulnerability correlation outputs, mitigation tests and operational logs can be necessary for assessment but unsafe or inappropriate for unrestricted publication. The CRA itself differentiates public user information, public fixed-vulnerability notices, technical documentation available to authorities, and sensitive Article 14 reporting. The architecture must therefore provide role- and purpose-based disclosure rather than one public dataset.

Table 10: Cross-cutting evidence patterns derived from the ETSI corpus

Pattern	Standards examples	Required evidence property	Architectural consequence
Version-specific composition	EN 304 617, 627, 634, 635, 636.	SBOM or component set tied to exact assessed release, with hash and provenance.	Release/build is a first-class subject; composition is not attached only to the model.
Observed installed state	EN 304 627 and 636; version tracking in 621.	Device-reported or otherwise observed installed version and update outcome.	Instance state must be distinct from manufacturer release declaration.

Pattern	Standards examples	Required evidence property	Architectural consequence
Time-bounded vulnerability status	EN 304 617, 622-624, 626, 634.	Vulnerability and exploitation source, retrieval time, cutoff date, assessment and remediation state.	Claims require freshness, expiry/review and supersession.
Verified exploitability or mitigation	EN 304 617, 623, 624, 626.	VEX, mitigation test, architecture-based non-reachability or compensating control.	A CVE match and an exploitability verdict must remain separate evidence objects.
Update authenticity and authority	EN 304 620, 621, 627.	Package hash, signature, authorised signer, target version and verification outcome.	Update provenance belongs in both control and evidence planes.
Rollback and recovery	EN 304 621; recovery in 623 and 632.	Authorisation, reason, metadata, restored state and post-event verification.	Rollback is not ordinary installation; it is a controlled exception event.
Configuration context	EN 304 619, 623, 625, 634.	Secure-default or shipped configuration, deployment profile and deviations.	Declared, assessed and observed configuration states must be separated.
Trusted logging and time	EN 304 618, 622, 632, 633.	Tamper protection, time source, retention, opt-out and harm/risk context.	Passport carries capability and evidence references, not raw operational logs.
Remote-service dependency	EN 304 626 and 636; CRA product definition.	Service identity/version, interface, compatibility, failure and advisory linkage.	Remote data processing solution is an independent linked subject.
Safety and essential availability	EN 304 631-633; CRA Art. 13(2).	Safety/availability impact assessment, fallback, recovery and update decision.	Security changes can trigger a controlled functional-safety hand-off.
Scope and profile constraints	Use cases and risk/harm classes across the drafts; OT exclusion in 636.	Applicability profile, exclusions and assessment scope.	A conformity claim is valid only within its documented product and deployment scope.

Source: Authors' synthesis from [26–42].

2.8 Direct, indirect, and proposed requirements for a CRA-capable DPP

The legal and standards analysis supports a three-level classification. Direct requirements are explicit legal or standards requirements for information, processes, technical properties, assessment activities or evidence. Indirect requirements are not stated as DPP requirements but are necessary for the represented evidence to be secure, attributable, current, confidential and usable. Proposed requirements are design choices that integrate the first two levels into a coherent evidence infrastructure. The distinction prevents technical necessity from being mislabelled as statutory text.

Table 11: Requirement classes for a CRA-capable DPP

Class	Definition	Representative examples	Normative treatment
Direct legal	Express CRA or applicable DPP-law information, documentation, process or identity duty.	Product identification; risk assessment; SBOM; support end; CVD contact; user instructions; tests; declaration; Article 14 reports.	Cite legal article or annex. DPP is one possible implementation unless the applicable DPP law requires it.

Class	Definition	Representative examples	Normative treatment
Direct standards	Normative requirement or assessment evidence in an applicable standard.	SBOM hash; VEX/mitigation evidence; installed-version match; signed update; rollback metadata; fuzzing records; trustworthy time.	Cite exact standard edition and clause. Draft status must be visible until OJEU citation.
Indirect governance	Control required to make evidence reliable or safe, derived from the nature of the direct obligation.	Issuer authority; mandate; evidence status; confidentiality; selective disclosure; continuity; independent verification.	State the derivation and avoid presenting it as express CRA wording.
Proposed architecture	Authors' mechanism for combining requirements into product evidence infrastructure.	Dual-plane model; evidence graph; five subject layers; signed lifecycle events; policy-engine decisions; proof-graph scoring.	Evaluate as a design artefact in Parts II-IV.
Excluded claim	Representation that exceeds the legal or evidentiary basis.	Single "CRA compliant" flag; registry proof treated as product conformity; signature treated as authority; static model passport treated as fleet state.	Explicitly prohibit or qualify in the architecture and user interface.

Source: Authors' classification.

2.9 Derived design requirements for Part II

The following design requirements are derived from the preceding analysis. They form the acceptance criteria for the reference architecture in Part II. Each requirement is architectural unless it explicitly restates a legal or standards duty. The labels do not imply that the CRA uses the same terminology.

Table 12A: Derived design requirements: legal status, subject model, and trust

ID	Design requirement	Principal basis
DR-01	Every requirement, claim and standard reference shall carry its source, edition, clause, legal status, scope, and where relevant OJEU citation status. Draft, proposal, harmonised standard and authors' design must remain machine-distinguishable.	CRA Art. 27 and Annex VII; source-status analysis.
DR-02	The architecture shall maintain distinct identifiers and records for product model or type, hardware or software release/build, physical or logical instance, deployment and configuration state, and remote data processing solution.	CRA scope and Annex VII; EN 304 627 and 636 instance matching.
DR-03	Every evidence subject and evidence object shall use a persistent identifier and resolvable relationship. Where an ESPR DPP applies, the architecture shall reuse its legally defined product identifiers and granularity while permitting linked cybersecurity-specific subjects.	ESPR Arts. 9-11; EN 18219 and EN 18220.
DR-04	An SBOM, firmware manifest or relevant hardware/component inventory shall be cryptographically bound to the exact release or build assessed, with hash, generation time, tool and provenance metadata.	CRA Annex I Part II(1); EN 304 617, 627, 635, 636.
DR-05	Manufacturer-declared release and configuration, device-observed installed state, operator-configured state, and independently assessed state shall remain separate claim classes and may be compared by policy.	EN 304 627/636; operational evidence distinction.
DR-06	A manufacturer-responsible remote data processing solution shall be represented as an independent versioned subject with explicit dependency, interface, compatibility, support, vulnerability and incident relationships.	CRA remote data processing scope; EN 304 626/636.
DR-07	Every authoritative claim shall identify the accountable legal person and the role under which it acts, such as manufacturer, supplier, laboratory, notified body, integrator, operator or repairer.	CRA economic-operator duties; eIDAS and DPP registry.

ID	Design requirement	Principal basis
DR-08	The architecture shall verify that the human, service or automated agent issuing or changing evidence had a valid, scoped and non-expired mandate from the represented legal person.	DPP registry delegated access; proposed EBW role/authorisation model.
DR-09	Evidence objects shall be signed, sealed, hashed, timestamped or otherwise integrity-protected as appropriate, and shall retain source, issuer, method, chain of custody and verification metadata.	DPP registry proof; ETSI assessment evidence; trust-services framework.
DR-10	Claims, credentials, mandates, certificates and evidence shall expose status. Corrections and updated assessments shall supersede rather than erase prior states, preserving audit history.	ESPR lifecycle linking; registry versioning; credential/evidence status.
DR-11	Dynamic claims, especially vulnerability, support, configuration and installed-version claims, shall include assessment time, source cutoff or retrieval time, validity or review time, and freshness policy.	CRA support period and Article 31; EN 304 617 and other KEV assessments.

Source: Authors' derived requirements. Legal and standards sources are identified in the principal-basis column; the architectural formulation is the authors' design.

Table 12B: Derived design requirements: operational evidence, disclosure, safety, and reliance

ID	Design requirement	Principal basis
DR-12	The architecture shall distinguish component presence, candidate vulnerability match, exploitability, active exploitation, mitigation, VEX status, residual risk, remediation plan and fixed release. A single vulnerability Boolean is insufficient.	CRA Annex I Parts I-II; EN 304 617, 623, 624, 626.
DR-13	Security update evidence shall distinguish release, availability, delivery, package verification, installation, post-installation version, failure, authorised rollback, recovery and final verification.	CRA update duties; EN 304 620, 621, 627.
DR-14	The evidence graph shall link each applicable CRA or standards requirement to risk context, technical control, assessment method, evidence object, verdict, issuer, scope, date and open deviation.	CRA risk assessment, technical documentation and conformity assessment.
DR-15	The architecture shall represent support determination, support period, update availability, component-support dependencies, end-of-support notice, migration guidance and unsupported residual risk.	CRA Art. 13(8)-(9),(18)-(19); EN 304 619.
DR-16	Manufacture, integration, commissioning, configuration, update, repair, component substitution, substantial modification, transfer, second life and decommissioning shall be represented as attributable events linked to predecessor states.	CRA Arts. 21-23; ESPR Art. 11(d); lifecycle evidence.
DR-17	Public, customer/operator, conformity-assessment and authority views shall be separated by policy. Secrets, private keys, raw security logs, personal telemetry, exploit details and sensitive topology shall not be placed in an unrestricted passport.	CRA public/restricted information split; ESPR role-based access and data protection.
DR-18	For safety- or essential-availability-relevant products, a security change shall be able to invoke a safety-impact assessment, approval state, deployment restriction, fallback or rollback plan and post-change verification.	CRA Art. 13(2) health/safety consideration; EN 304 631-633.
DR-19	Evidence shall remain exportable, verifiable and recoverable across DPP providers and after manufacturer or provider cessation, with open formats, backups, migration, persistent identifiers and independent verification where feasible.	ESPR persistence, backup and no lock-in; CRA long retention periods.
DR-20	The system shall validate schema, vocabulary, identifiers, relationships, required fields and granularity, while clearly distinguishing structural validation from factual, technical and legal conformity.	DPP registry semantic checks and explicit non-conformity boundary.
DR-21	Relying parties shall evaluate evidence through explicit policies that consider issuer authority, evidence type, freshness, product scope, deployment context, risk and required assurance. Authentication alone shall not authorise reliance.	Zero-trust control-plane derivation; eIDAS/EBW authority and evidence context.
DR-22	The architecture shall not output an unqualified binary CRA-compliant status. Any decision shall identify the subject, version, scope, time, conformity route, evidence basis, unresolved findings and responsible decision-maker or policy.	CRA responsibility and scoped presumption; temporal and subject-specific evidence findings.

Source: Authors' derived requirements.

These requirements form the input specification for Part II. Part II defines the subject graph, evidence object model, issuer and mandate model, access-control tiers, status and supersession mechanisms, and the dual-plane reference architecture. Their legal and standards traceability must remain visible in the architecture and in every later sector evaluation.

3 Part II - Concept and Reference Architecture

Part II answers RQ2 and develops the architectural mechanisms needed to answer RQ3. It transforms the legal, standards, and derived design requirements in Part I into a technology-neutral reference architecture. The architecture is intended to support CRA implementation; it is not represented as a legal requirement, a harmonised standard, or an autonomous conformity-assessment route. Each architectural element is therefore traceable to one of three sources: an explicit legal or standards requirement, an indirect governance need derived from that requirement, or a design choice proposed by the authors.

The central design problem is not data availability alone. It is evidentiary attribution. A relying party must determine which product subject a claim concerns, who issued it, under which organisational authority, how the claim was produced, which evidence supports it, whether the evidence remains current, and whether the requested use falls within the claim's scope. A CRA-capable DPP must preserve these relations across organisational boundaries and product lifecycles without forcing all evidence into one repository.

The resulting architecture is federated by default. A registry or passport service can provide identifiers, discovery, structural validation, and controlled views. Detailed evidence can remain with manufacturers, suppliers, laboratories, operators, devices, or authorities, provided that it is persistently identifiable, integrity-protected, resolvable, status-aware, and accessible under an enforceable policy. Distribution in this paper therefore denotes distributed authority and evidence custody. It does not imply that a distributed ledger is required.

3.1 Definition, scope, and architectural boundary

3.1.1 Formal definition

A CRA-capable DPP is defined as a verifiable product evidence architecture that represents product subjects, claims, evidence artefacts, organisational authority, policies, and lifecycle relations in a form that can be resolved and evaluated by authorised parties. The architecture may extend a statutory DPP established under the ESPR or sectoral law. It may also operate as an equivalent evidence service for software or other products for which no statutory DPP is required. The term describes capability, not a legal designation.

Definition 1 - CRA-capable Digital Product Passport (authors' design)

A CRA-capable DPP is a distributed, versioned, and access-controlled evidence architecture that binds product identity, hardware and software composition, vulnerability status, security updates, conformity evidence, operational state, and lifecycle events to the relevant product model, release, instance, deployment, and remote service. It verifies issuer identity and authority, preserves evidence status and lineage, and produces only scoped, time-bound reliance decisions. It does not itself establish legal conformity.

The conceptual artefact can be represented as $\text{CRA-DPP} = \langle S, C, E, A, P, R \rangle$. S is the set of identifiable product and service subjects. C is the set of typed claims about those subjects. E

is the set of evidence artefacts and observations supporting the claims. A is the set of actor, role, and mandate assertions. P is the set of access, issuance, verification, and reliance policies. R is the set of typed and versioned relations among all elements. The notation is descriptive. It does not prescribe a graph database, credential format, identifier method, or cryptographic suite.

$$\text{CRA-DPP} = \langle S, C, E, A, P, R \rangle$$

The architecture is capable rather than compliant. Capability means that the system can represent and verify the evidence necessary for a defined CRA task. Legal conformity remains a decision made under the applicable conformity route by the responsible economic operator, conformity assessment body, market-surveillance authority, or court. The same architecture can therefore support internal control, third-party assessment, post-market surveillance, procurement, incident response, and maintenance without collapsing those functions into one status field.

3.1.2 Negative definition and excluded interpretations

Four interpretations are excluded. First, the architecture is not a static PDF, QR-code landing page, or repository of unverified manufacturer declarations. Second, it is not a substitute for the CRA technical documentation, the EU declaration of conformity, an OJEU-cited standard, a conformity certificate, or an official Article 14 reporting channel. Third, it is not a universal public dataset. Detailed composition, test, incident, configuration, and operational evidence can require restricted access. Fourth, it is not a claim that every product with digital elements must carry a statutory DPP. The legal duty to establish a DPP remains dependent on the applicable Union product legislation.

A data carrier also has limited evidentiary meaning. A printed QR code can resolve a passport identifier, but it does not establish that the code is attached to the correct device, that the resolved data concern the installed release, or that the issuer had authority to make the claim. Physical and logical binding mechanisms are therefore treated separately in Section 3.14.

3.2 Design principles and assurance objectives

The architecture is governed by ten principles. They operationalise the design requirements DR-01 through DR-22 and define the conditions under which a DPP can be relied upon as product evidence infrastructure. The principles are architectural proposals, although several are derived directly from CRA, ESPR, eIDAS, DPP registry, or candidate ETSI requirements.

Table 13: Architectural principles for CRA-capable product evidence infrastructure

Principle	Operational meaning	Failure prevented	Source status
Subject specificity	Every claim identifies whether it concerns a model, release, instance, deployment, remote service, component, actor, or event.	Model-level evidence being misapplied to an unpatched or differently configured instance.	Authors' design derived from CRA scope and ETSI version matching.
Source-status fidelity	Legal text, OJEU-cited standards, draft standards, legislative proposals, and authors' design remain machine-distinguishable.	A draft or design choice being presented as binding law or presumption of conformity.	Directly derived from CRA Article 27 and Part I source-status method.

Principle	Operational meaning	Failure prevented	Source status
Evidence before assertion	A claim references the method and artefacts that support it, including limitations and open findings.	Unsigned or unsupported statements being treated as technical proof.	Derived from CRA technical documentation and ETSI assessment models.
Temporal validity	Dynamic claims carry observation time, source cutoff, validity, review, status, and supersession.	Replay of stale VEX, support, configuration, or installed-version claims.	Derived from lifecycle duties and ETSI KEV assessments.
Accountable authority	A valid signature is combined with legal-person identity, role, mandate, scope, and status.	A technically valid signature by an unauthorised employee, service, or agent.	Derived from eIDAS, DPP registry delegation, and proposed EBW controls.
Federated custody	Evidence may remain with authoritative sources and is resolved through persistent references and verifiable presentations.	Central replication becoming a source of inconsistency, confidentiality loss, or lock-in.	Authors' design informed by ESPR interoperability and persistence.
Least disclosure	Policies expose only the evidence required for the purpose and relying party.	Publication of secrets, raw logs, exploit details, personal telemetry, or trade secrets.	Derived from ESPR role-based access, data protection, and CRA information tiers.
Append-only lineage	Corrections, new assessments, updates, and modifications supersede prior states without deleting audit history.	Loss of the state that applied at release, incident, repair, or inspection time.	Derived from DPP versioning, lifecycle evidence, and substantial-modification duties.
Risk-based reliance	Policy decisions consider authority, subject binding, method, freshness, completeness, context, and risk.	Authentication being confused with authorisation or evidence sufficiency.	Authors' zero-trust design informed by [18].
No unqualified compliance flag	Outputs state subject, version, scope, time, route, evidence basis, findings, and decision owner.	A universal green status masking partial coverage, unresolved findings, or deployment drift.	Derived from scoped presumption of conformity and DR-22.

Source: Authors' design derived from Part I, especially DR-01 through DR-22.

The architecture targets four assurance objectives. Integrity concerns whether evidence has been altered. Authenticity concerns who issued or observed it. Authority concerns whether the issuer was entitled to act for the represented organisation and subject. Sufficiency concerns whether the evidence and method are adequate for the requested decision. These properties are independent. Cryptographic authenticity does not prove authority or technical correctness; a complete test report can still concern the wrong release; and a current release claim does not prove that a particular instance has installed it.

3.3 Five-subject product model

3.3.1 Subject separation

The regulated product and its evidence cannot be represented reliably through one identifier alone. The reference architecture therefore defines five first-class subject categories: product model or type, hardware or software release or build, physical or logical instance, deployment and configuration state, and remote data processing solution. Hardware and software releases can be represented as separate nodes where they change independently, or as a controlled combined baseline where conformity is assessed against their combination. Figure 1 summarises these subject categories and their principal relations.

Source: Authors' design based on CRA product scope, remote data processing solutions, and

Five-subject model for a CRA-capable DPP




Figure 1: Five-subject model for a CRA-capable DPP.

the instance-matching patterns in EN 304 627 and EN 304 636.

Table 14: Subject classes, identifiers, authoritative state, and evidence implications

Subject	Identity and authoritative issuer	Typical state and cadence	Principal evidence questions
Product model or type	Manufacturer product identifier; statutory DPP product identifier where applicable; manufacturer as principal issuer.	Relatively stable. Changes where intended purpose, architecture, classification, or commercial type changes.	What product is placed on the market? Which intended purpose, CRA scope, classification, support policy, and conformity route apply?
Hardware or software release/build	Release identifier, build identifier, firmware revision, image digest, hardware revision; manufacturer or authorised build system.	Changes at each assessed build, patch, hardware revision, or controlled baseline.	Which composition, controls, risk assessment, tests, VEX statements, advisories, and conformity evidence apply to this exact release?
Physical or logical instance	Serial number, device certificate, workload identity, installation identifier, or equivalent; manufacturer and later operator observations.	Changes when installed release, ownership, support entitlement, or device state changes.	Which release is actually installed? Is the identifier authentic? What update, repair, or decommissioning events occurred?
Deployment and configuration state	Operator deployment identifier and configuration profile; operator or authorised integrator as principal issuer.	Potentially frequent. Event-driven or continuously observed depending on risk.	Does the operational environment match the assessed assumptions? Which deviations, mitigations, network settings, integrations, and safety constraints apply?
Remote data processing solution	Service identifier, API or endpoint identity, service release, region or tenancy where relevant; manufacturer or responsible service provider.	Can change independently from the physical product or local software.	Is the service necessary for product function and under manufacturer responsibility? Which version, dependencies, support, vulnerabilities, incidents, and compatibility constraints apply?

Subject	Identity and authoritative issuer	Typical state and cadence	Principal evidence questions
Component or supplier subject	Supplier component identifier and persistent software, firmware, hardware, cryptographic, or service identifier.	Changes with component version, supplier notice, support, or substitution.	Which final releases contain the component? What supplier evidence is authoritative, current, and applicable?
Lifecycle or assessment event	Event identifier issued by the actor that performed or observed the action.	Append-only. Each event records valid time and recorded time.	Who performed the release, test, update, repair, modification, transfer, or decommissioning action, and what state resulted?

Source: Authors' design; legal and standards basis described in Part I.

3.3.2 Subject tuple and state comparison

For a defined operational product context, the architecture uses a subject tuple

$$P = \langle m, r, i, d, s \rangle$$

. Here m denotes the product model or type, r the applicable release or controlled hardware-software baseline, i the physical or logical instance, d the deployment and configuration state, and s any relevant remote data processing solution. Elements can be absent where not applicable, but the model must be capable of representing them. The tuple is a graph projection, not a composite identifier that must be stored in every record.

$$P = \langle m, r, i, d, s \rangle$$

The architecture compares states rather than assuming equivalence. A manufacturer-declared release state, device-observed installed state, operator-configured state, and independently assessed state remain separate claim classes. A policy can determine whether they agree. A mismatch is itself an evidence object. For example, the assessed release can be R4.2, the device can self-report R4.1, the operator configuration can disable a required control, and the latest independent assessment can apply only to R4.2 in its secure-default configuration. No one of these claims should overwrite the others.

The same separation applies to remote services. A device release can remain unchanged while an API, cloud database, policy service, management plane, or model-serving endpoint changes. Compatibility and responsibility relations must therefore be explicit. The service node should identify whether it is part of the regulated product, a third-party dependency, or an operator-selected service outside manufacturer responsibility. This status is a legal and architectural scope claim, not an inference made solely from network traffic.

3.3.3 Statutory DPP granularity and cybersecurity-specific subjects

Where a statutory DPP exists, its delegated or sectoral act determines whether the passport is established at model, batch, or item level. The CRA-capable architecture reuses the legally defined identifier and does not replace that granularity. It adds linked cybersecurity-specific subjects where evidence requires greater precision. An item-level battery passport, for example, can link to multiple BMS firmware releases and site-specific deployment states. A model-level passport for a software product can link to exact build digests and logical installations. The additional subjects are proposed evidence nodes, not a redefinition of the legally mandated DPP level.

For systems of systems, the architecture uses composition and assembly relations rather than creating one monolithic record. A machine, building, BESS, or virtualisation cluster can have a system-level subject that references component passports and software releases. System integration evidence then documents architecture, configuration, interface assumptions, combined risk, and responsibility. Component evidence remains attributable to its supplier, while system-level conformity remains the responsibility of the relevant manufacturer or integrator under the applicable law.

3.4 Dual-plane reference architecture

3.4.1 Architectural separation

The reference architecture separates a control plane from an evidence plane. These are authors' architectural constructs. The control plane answers who may issue, change, request, disclose, evaluate, or rely on evidence, under which organisational mandate and policy. The evidence plane answers what is claimed about the product, which artefacts support the claim, how the claim relates to requirements and lifecycle state, and whether it remains current. The control plane cannot establish technical truth, and the evidence plane cannot grant organisational authority. Figure 2 presents the resulting dual-plane architecture.




Figure 2: Dual-plane reference architecture.

Source: Authors' design informed by the eIDAS trust framework, the proposed European Business Wallet, EPR DPP infrastructure, and zero-trust policy architecture [18–20].

Table 15: Logical components and responsibilities of the dual-plane architecture

Component	Principal function	Trust boundary	Required output
Legal-person identity and organisational wallet	Identify the accountable organisation and present identity, role, attribute, seal, or mandate evidence.	Authentic organisation identity does not by itself authorise a specific technical claim.	Verified organisation identifier, trust source, credential or seal status, and represented capacity.

Component	Principal function	Trust boundary	Required output
Role and mandate service	Represent internal role, delegation, product scope, permitted actions, time limits, and delegation constraints.	A valid principal can act only within the mandate granted by the legal person.	Authorisation decision or verifiable mandate reference with status.
Policy decision point (PDP)	Evaluate identity, authority, evidence, context, purpose, freshness, risk, and requested action.	Policy logic must be versioned and must not silently convert evidence gaps into approval.	Permit, deny, or escalate decision with reasons, policy version, and evidence snapshot.
Policy enforcement point (PEP)	Enforce issuance, update, access, presentation, and reliance decisions.	Enforcement must bind the decision to the actual API call, evidence view, or action.	Controlled response, selective presentation, audit event, and signed decision receipt where appropriate.
DPP resolver and registry	Resolve stable subject identifiers, passport location, structural metadata, operator identity, and service endpoints.	Resolution and registry checks do not prove substantive CRA conformity.	Resolved subject and endpoint, version, integrity metadata, registry status, and access path.
Versioned evidence graph	Maintain typed relations among requirements, subjects, claims, evidence, actors, events, findings, and decisions.	Graph storage does not itself make a claim true; issuer, method, and status remain external trust inputs.	Current and historical graph views with provenance and source-status metadata.
Authoritative evidence stores	Retain SBOMs, VEX, test reports, attestations, advisories, risk assessments, and technical documentation.	Evidence may be confidential and may remain under the source organisation's custody.	Content-addressed artefact, access policy, integrity proof, retention and availability status.
Operational observation sources	Report installed version, configuration, update outcome, runtime dependency, or device attestation.	Self-reported state has lower assurance than independently or hardware-rooted observation.	Observation claim with method, time, subject binding, and assurance metadata.
Status and supersession service	Publish revocation, suspension, withdrawal, expiry, correction, replacement, or end-of-support state.	Status of the issuer credential and status of the product claim are separate.	Machine-verifiable current status and link to successor or reason.

Source: Authors' reference architecture.

3.4.2 Verification and reliance sequence

A typical reliance sequence contains seven logically separate steps. The relying party submits a request with purpose, product subject, decision context, and requested evidence or action. The control plane authenticates the principal and resolves the represented legal person. It verifies role and mandate, including scope, time, delegation, and conflict constraints. The evidence plane resolves the relevant product subjects and evidence graph. Each claim is verified for integrity, issuer status, authority reference, subject binding, evidence references, freshness, and supersession. The PDP evaluates the request against a versioned policy and returns permit, deny, or escalate. The PEP then releases a purpose-limited evidence view or performs the authorised action and records a decision receipt.

Escalation is a first-class outcome. It is required where evidence is incomplete, contradictory, stale, outside the issuer's authority, or insufficient for the risk of the decision. Human hand-back is therefore part of the architecture rather than a failure path. A conformity engineer, product security incident response team, functional-safety authority, or market-surveillance official can review the evidence and issue a new attributable decision. The automated system records that decision without representing it as generated by the policy engine alone.

Reliance rule (authors' design)

A claim may be relied upon only where integrity is valid, issuer and mandate status are valid, the subject and scope match the request, the evidence method and assurance satisfy the policy, freshness requirements are met, no superseding or conflicting claim invalidates the result, access is permitted, and the requested action remains within the relying party's legal and organisational competence.

3.5 Control-plane identity, mandate, and zero-trust authority

3.5.1 From signature validation to accountable authority

The control plane distinguishes five questions that are often collapsed into signature validation: Which cryptographic key produced the proof? Which natural or legal person controls or is represented by that key? Which organisation is accountable for the claim? Which role and mandate authorised the action? Was the authority valid for the product, requirement, jurisdiction, time, and operation concerned? A positive answer to the first question is insufficient for the other four.

The European trust framework provides several possible mechanisms. Qualified electronic seals can bind a document to a legal person. Qualified and non-qualified electronic attestations of attributes can represent verified organisational attributes. European Digital Identity Wallets can support representation and mandate functions for natural persons. The proposed European Business Wallet is designed for economic operators and includes identity, signing, sealing, documents, roles, attributes, and authorisations. Because the EBW instrument remained a proposal at the evidence cut-off, the architecture treats EBW as a preferred European implementation path, not as a current legal prerequisite [19, 20].

Equivalent implementations can use regulated trust services, enterprise public-key infrastructure, organisational wallets, verifiable credentials, or controlled directories, provided that they satisfy the same authority questions. Decentralised identifiers can be used as resolvable identifiers for organisations, devices, services, or evidence subjects, but a DID does not establish legal identity or authority by itself [14]. The trust source and governance of the identifier remain part of the evidence.

Table 16: Minimum mandate object and actor-specific claim authority

Element	Required semantics	Example policy test	Reason
Principal and represented legal person	Human, workload, service, or agent identifier plus the accountable legal entity.	Does the principal validly represent manufacturer A?	Separates key holder from accountable organisation.
Role and competence	Manufacturer, supplier, laboratory, notified body, integrator, operator, repairer, authority, or other defined role.	May this laboratory issue a verdict for this assessment method?	Prevents role confusion and unsupported reliance.
Permitted action	Issue, update, revoke, observe, approve, disclose, request, or rely.	May this service issue an installed-version observation?	Authentication alone does not grant every action.
Product and requirement scope	Models, releases, instances, product groups, standards, clauses, or evidence types.	Is the mandate valid for model XR-200 and EN 304 627 clause Q?	Prevents over-broad delegation.
Temporal and geographic scope	Not-before, expiry, review date, jurisdiction, site, or market.	Was the authority valid at issuance and for the Union market?	Supports audits and time-bounded delegation.
Delegation and conflict rules	Whether sub-delegation is permitted; segregation of duties; incompatible roles.	Can the same agent issue and independently approve the test verdict?	Reduces over-delegation and self-approval risk.

Element	Required semantics	Example policy test	Reason
Status and evidence	Credential, mandate, employment, service, or registration status plus source.	Is the mandate active and not revoked or superseded?	Enables continuous verification.
Decision provenance	Policy version, inputs, result, reasons, and authorising authority.	Which rule permitted disclosure or issuance?	Makes automated authority decisions auditable.

Source: Authors' design informed by eIDAS, the DPP registry, and the proposed EBW role-authorisation model.

3.5.2 Workloads and automated agents

Manufacturing, vulnerability management, DPP maintenance, and incident response increasingly use automated workloads and AI agents. The architecture permits such systems to issue or transform evidence only through bounded organisational authority. The agent identity must be distinct from the legal person. Its mandate must specify permitted data sources, product scope, actions, approval thresholds, and conditions for human hand-back. The evidence object must state whether a claim was generated automatically, approved by a human, or independently assessed.

Automated reasoning can correlate an SBOM with vulnerability sources, detect a configuration mismatch, or prepare a draft VEX statement. It cannot acquire legal authority through technical capability. A policy engine can permit publication only where the designated authority has approved the method and the evidence satisfies the relevant threshold. High-impact decisions, such as declaring a known exploitable vulnerability non-exploitable, postponing an urgent patch, or approving a safety-relevant update, can require dual control or independent review.

3.6 Evidence object model and verifiable evidence graph

3.6.1 Separation of artefact, claim, assessment, and decision

A central contribution of the architecture is the separation of four evidence layers. An evidence artefact is the underlying object, such as an SBOM, firmware image digest, test log, architecture diagram, vulnerability-source snapshot, device attestation, or assessment report. A claim is a typed statement about a subject, such as "release R contains component C" or "instance I runs release R". An assessment links a requirement, method, artefacts, findings, and verdict. A reliance decision applies a policy to one or more verified claims in a defined context. The layers can be stored together, but their semantics must remain distinct.

This distinction prevents several common errors. A hash proves that a retrieved artefact matches a referenced byte sequence, not that the artefact is complete. A signed SBOM proves issuer attribution and integrity, not that the SBOM matches the deployed binary. A VEX statement records an issuer's product-context determination, not universal non-exploitability. A policy decision can authorise continued operation under a temporary mitigation, but it does not alter the underlying vulnerability state.

Table 17: Minimum verifiable claim envelope

Field group	Required content	Verification function	Illustrative representation
Identity and type	Claim identifier, claim type, subject identifier, predicate, value or object.	Prevents ambiguous statements and enables typed graph relations.	"installedRelease" claim for instance I.
Issuer and authority	Issuer identifier, represented legal person, role, mandate reference, delegation chain, status.	Establishes accountable authority and scope.	Operator service acting under site-maintenance mandate.
Time	Issuance, observation, valid-from, valid-until or review-by, source cut-off, recorded time.	Supports freshness, bitemporal history, and incident reconstruction.	Observed at T1; recorded at T2; review by T3.
Scope and context	Product subject, release range, deployment profile, intended use, jurisdiction, standard clause, exclusions.	Limits reliance to the context actually assessed.	Applies to R4.2 in secure-default configuration.
Method and evidence	Assessment or observation method, tool and version, parameters, evidence references, hashes, limitations.	Allows the verifier to evaluate technical sufficiency and reproducibility.	Scanner output plus manual reachability test.
Source status	Binding law, OJEU-cited standard, draft standard, proposal, guidance, contract, or authors' design.	Prevents incorrect legal inference.	Draft EN 304 627 clause reference.
Integrity and proof	Signature, seal, MAC, content hash, timestamp, certificate or credential chain, algorithm and key identifiers.	Detects alteration and attributes the proof.	Qualified seal or JWS/COSE-secured claim.
Status and lineage	Active, suspended, withdrawn, expired, contested, superseded; predecessor and successor references.	Prevents stale or corrected claims from remaining current.	VEX v2 supersedes v1 after new evidence.
Disclosure and retention	Classification, permitted purposes, authorised roles, retention, export, and redaction rules.	Enforces least disclosure and continuity.	Full SBOM available to operator and authority only.

Source: Authors' evidence object model, informed by W3C Verifiable Credentials, eIDAS trust services, CRA evidence duties, and ETSI assessment artefacts [15].

3.6.2 Graph model and typed relations

The evidence plane is modelled as a versioned directed graph $G_t = (V_t, E_t)$. Nodes represent subjects, actors, requirements, risks, controls, methods, artefacts, claims, findings, verdicts, policies, and lifecycle events. Edges have defined semantics such as `appliesTo`, `composedOf`, `assessedBy`, `produces`, `supports`, `issuedBy`, `authorisedBy`, `supersedes`, `mitigates`, `installedOn`, `dependsOn`, and `affects`. Each node and edge carries provenance and temporal metadata. The graph can be implemented in a graph database, relational system, linked-data store, object repository, or a combination. The architectural requirement concerns semantics and resolvability, not storage technology. Figure 3 shows how these relations support a scoped conformity decision.

$$G_t = (V_t, E_t)$$

Source: Authors' design based on the CRA technical-documentation chain and ETSI requirement-assessment-evidence patterns.

Requirement-to-evidence graph and scoped conformity decision




Figure 3: Requirement-to-evidence graph and scoped conformity decision.

Table 18: Core graph nodes and relations

Node or relation	Semantics	Mandatory metadata	Example
LegalRequirement	Binding obligation or essential requirement.	Instrument, article or annex, consolidated version, applicability.	CRA Annex I Part II, vulnerability handling.
StandardRequirement	Normative clause or assessment criterion.	Standard identifier, edition, clause, status, OJEU coverage.	EN 304 627 REQ-KEV-1-01 draft.
RiskContext	Intended purpose, environment, assets, misuse, safety or availability context.	Issuer, scope, date, assumptions, revision.	Public-network router used in a hospital.
Control	Technical or organisational measure addressing a requirement or risk.	Owner, implementation, configuration, affected subjects.	Signed-update verification before installation.
AssessmentMethod	Inspection, test, analysis, observation, attestation, or review procedure.	Method edition, tool, parameters, assessor, limitations.	Installed-version comparison against SBOM release.
EvidenceArtifact	Underlying content used in assessment or observation.	Persistent ID, hash, media type, location, retention, access.	CycloneDX or SPDX SBOM, test log, attestation.
Claim or Verdict	Typed statement or assessment result concerning a subject.	Claim envelope fields in Table 17.	No unresolved known exploitable vulnerability as of cutoff T.
Finding or Deviation	Open issue, exception, non-conformity, residual risk, or evidence gap.	Severity, owner, due date, affected scope, status.	Required log retention not verified for deployment D.
ConformitySnapshot	Time-bound set of accepted claims, findings, route, and evidence for a defined scope.	Subject tuple, requirements, decision owner, policy, time, status.	Internal-control snapshot for release R4.2.
supports / contradicts	Evidence relation that can strengthen or challenge a claim.	Issuer, method, time, confidence category or assurance vector.	Device attestation contradicts operator inventory.

supersedes / corrects	Lineage relation preserving prior evidence while identifying the current statement.	Reason, effective time, successor, authorising actor.	Corrected VEX after component reachability test.
-----------------------	---	---	--

Source: Authors' graph model.

3.6.3 Bitemporal state and evidence resolution

The architecture uses bitemporal semantics for material claims and events. Valid time states when a claim or state applied in the product world. Recorded time states when the evidence system received or registered it. The distinction is necessary where an incident is discovered retrospectively, a device was offline, an update receipt arrives late, or an assessment is corrected. A query such as "what was believed on 3 March?" differs from "what do we now know was true on 3 March?". Both are relevant for due diligence, incident reconstruction, liability, and market surveillance.

Evidence resolution is preferred over uncontrolled replication. The evidence graph can store metadata, hashes, status, and authorised endpoints while retrieving the authoritative artefact when needed. Cached copies are permitted where the cache records source, retrieval time, hash, and freshness policy. A verifier can therefore detect whether it is viewing current authoritative evidence, a preserved historical snapshot, or an unverified copy.

3.7 Evidence assurance vector and policy evaluation

3.7.1 Multidimensional assurance

A binary trusted-untrusted label is inadequate for heterogeneous product evidence. The architecture therefore proposes a multidimensional Evidence Assurance Vector rather than one universal score. For a claim c , decision context q , and time t , the vector $\alpha(c, q, t)$ contains authority, subject binding, method strength, freshness, completeness, independence, and reproducibility. Additional dimensions can be added for confidentiality, safety criticality, or operational availability. Values can be categorical or quantitative where a validated measurement exists.

$$\alpha(c, q, t) = \langle A, B, M, F, C, I, R \rangle$$

The vector is intentionally not reduced to an additive number by default. A high freshness score cannot compensate for an invalid mandate; an independent assessment cannot compensate for a subject mismatch; and a complete SBOM cannot compensate for an absent exploitability analysis where the decision concerns known exploitable vulnerabilities. Policies apply mandatory gates and context-specific thresholds. This approach creates a basis for evidence-graph assurance analysis without false numerical precision.

Table 19: Evidence Assurance Vector

Dimension	Question	Possible categories	Non-substitutable failure
A - Authority	Was the claim issued by an accountable and authorised actor?	Unverified; identified; role-verified; mandate-verified; independently authorised.	Invalid authority cannot be offset by strong technical evidence.

Dimension	Question	Possible categories	Non-substitutable failure
B - Subject binding	How strongly is the evidence bound to the exact model, release, instance, deployment, or service?	Name match; identifier match; signed manifest; device credential; hardware-rooted attestation.	Evidence for another build or instance is not applicable.
M - Method strength	How suitable and controlled was the assessment or observation method?	Declaration; automated scan; documented analysis; controlled test; accredited or independent assessment.	A weak method cannot support a high-risk verdict.
F - Freshness	Is the evidence current for the requested decision?	Expired; stale; within review period; event-triggered current; real-time observed.	Stale VEX or installed-state claims can invert the decision.
C - Completeness	Does the evidence cover the required components, requirements, interfaces, and states?	Unknown; partial; declared scope; validated scope; complete for defined profile.	A partial inventory cannot support a universal claim.
I - Independence	How independent is the evidence from the party benefiting from the claim?	Self-declared; segregated internal review; supplier corroborated; third-party; authority observed.	Self-declaration can be insufficient where third-party assessment is required.
R - Reproducibility	Can the method and result be reproduced or independently checked?	Opaque; documented; artefacts retained; tools and inputs retained; reproducible build or test.	An irreproducible result limits audit and dispute resolution.

Source: Authors' proposed assurance model.

3.7.2 Policy decisions and signed receipts

The PDP evaluates a request against explicit policy. Inputs include the requested action, relying party, subject tuple, intended use, risk class, applicable legal and standards profile, issuer authority, evidence assurance vector, current findings, and time. The output is permit, deny, or escalate. Permit authorises the requested reliance or action within the stated scope. It is not a declaration that the product is generally CRA compliant. Deny can result from an adverse finding or from insufficient evidence. Escalate directs the matter to a designated human or specialist function.

A decision receipt records the request, subject, policy identifier and version, evidence claim identifiers and hashes, status-check time, relevant findings, result, reason codes, and decision authority. It allows a later auditor to reconstruct why a patch was deferred, why a customer received a restricted SBOM, or why an instance was permitted to continue operating under a temporary mitigation. Sensitive evidence need not be copied into the receipt; persistent references and hashes can preserve the basis.

3.8 Composition, build provenance, and supplier evidence

3.8.1 Composition manifests as release-bound evidence

The architecture treats SBOMs, firmware manifests, HBOMs, cryptographic bills of materials, service bills of materials, and runtime dependency manifests as specialised composition evidence. They are not the passport itself. Each manifest is bound to an exact release or build

by a digest, build identifier, generation time, tool and version, authorising issuer, and scope statement. A manifest must state whether it represents source dependencies, build-time components, shipped binaries, runtime components, remote services, hardware dependencies, or a defined combination.

This distinction is necessary because product-specific standards can exceed the CRA minimum of top-level software dependencies. Virtualisation and container stacks can require runtime services, orchestrators, controllers, policy engines, secrets-management components, and hardware dependencies. Wearables or complex software can require transitive composition. A CRA-capable DPP therefore supports multiple composition profiles and records completeness relative to the applicable profile rather than labelling every object simply "SBOM".

Table 20: Composition and provenance evidence profile

Evidence element	Minimum content	Binding rule	Operational use
Release manifest	Model, release/build identifier, binary or image digests, hardware/firmware baseline.	Signed or sealed by authorised manufacturer/build service and linked to release subject.	Defines the assessed release boundary.
Software composition	Components, versions, supplier/origin, identifiers, dependency relations, direct/transitive scope.	Hash and generation metadata tied to release manifest.	Vulnerability correlation and supplier impact analysis.
Hardware and firmware composition	Hardware revisions, firmware, security processors, network interfaces, replaceable modules.	Linked to model, batch/item, and controlled hardware baseline.	Component substitution, authenticity, support and repair.
Runtime and service dependencies	Required services, APIs, orchestrators, controllers, plugins, remote processing, compatibility.	Versioned relationship to release and deployment profile.	Exposure analysis where operational dependencies differ from the shipped package.
Build provenance	Source reference, build system, toolchain, build parameters, environment, time, attestation.	Content-addressed and signed by controlled build identity.	Detects unapproved builds and supports reproducibility.
Supplier evidence	Supplier identity, component manifest, support state, VEX, advisory, certificate, update, limitations.	Federated reference preserving supplier issuer and status.	Propagates component evidence without transferring final-product responsibility.
Completeness declaration	Included and excluded classes, known blind spots, generation method, validation result.	Issued for each manifest and profile.	Prevents partial composition from being treated as complete.
Substitution event	Removed and installed component, reason, actor, compatibility tests, risk and safety impact.	Linked to predecessor and successor state of the instance/system.	Repair, maintenance, substantial-modification and second-life control.

Source: Authors' design informed by the CRA, EN 304 617, EN 304 625, EN 304 627, EN 304 634-636, SPDX, and CycloneDX [43, 44].

3.8.2 Nested and federated supplier evidence

Supplier evidence is represented as a graph of attributable statements rather than flattened into an anonymous manufacturer dataset. A component supplier can issue a signed manifest, support statement, VEX assertion, advisory, or remediation record. The final manufacturer can reference, accept, challenge, or supplement the statement and record its own product-context assessment. The graph therefore preserves both provenance and responsibility: the supplier remains accountable for its claim, while the final manufacturer remains accountable for the product placed on the market.

Nested passports can support this model where components already have statutory or voluntary DPPs. The final product passport resolves the component identifier and evidence subject, but it should retain a snapshot or hash of the supplier state used for the assessment. Otherwise, a later supplier correction could make it impossible to reconstruct the evidence basis at market release. Dynamic resolution and historical preservation are therefore complementary.

3.9 Vulnerability, exploitability, and VEX lifecycle

3.9.1 State model

Vulnerability evidence must distinguish component presence, candidate correlation, product-specific exploitability, active exploitation, mitigation, remediation, residual risk, and verification. The architecture models these as states and transitions rather than one Boolean field. A candidate match can result from a package identifier, version range, source-code match, binary analysis, or supplier notice. It becomes a product-context claim only after the matching method and release scope are recorded. Figure 4 represents the resulting vulnerability and VEX state model.

Product-context vulnerability and VEX state model

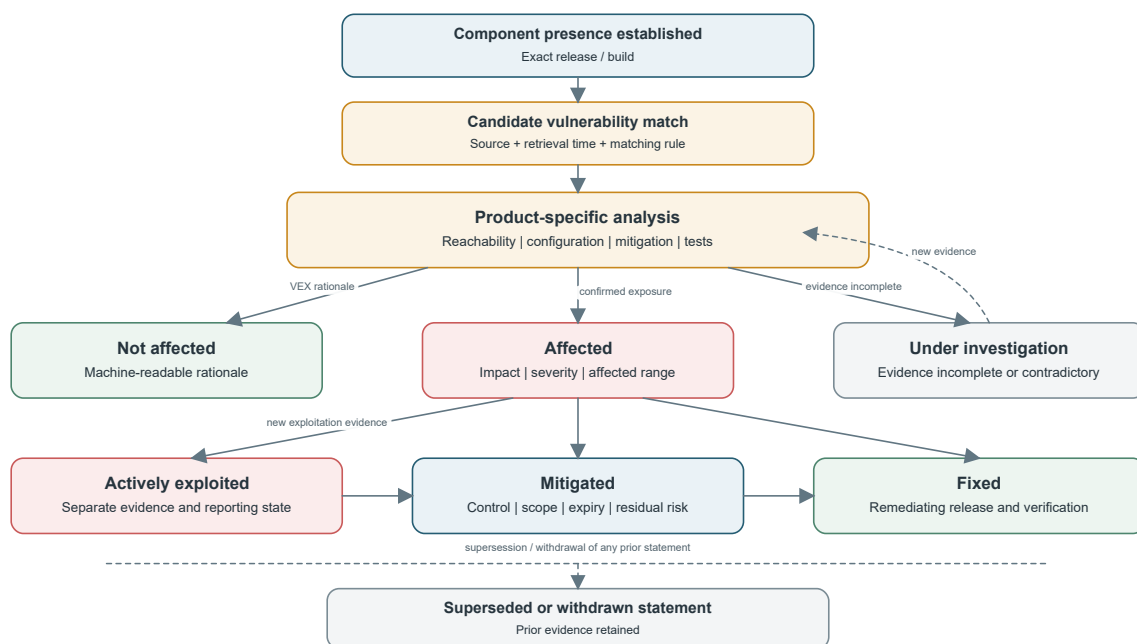


Figure 4: Product-context vulnerability and VEX state model.

Source: Authors' design based on CRA Annex I, VEX practice, CSAF, CycloneDX, and the evidence patterns in EN 304 617, EN 304 622-624, EN 304 626, and EN 304 634.

A VEX statement is treated as an attributable claim with an exact product and release scope, vulnerability identifier or source reference, status, rationale, action statement, method, evidence, issuance time, review time, and supersession relation. The status "not affected" must state why, for example because the vulnerable code is not present, the vulnerable function cannot be invoked, the execution path is unreachable under the assessed architecture, or an effective mitigating control prevents exploitation. The rationale is not interchangeable. Each basis can require different evidence.

Table 21: Vulnerability and VEX evidence object

Field	Required content	Freshness or status rule	Public/restricted treatment
Product context	Model, exact release/build, component path, deployment profile where relevant.	Superseded when composition, configuration, or service dependency changes.	Affected model/version can be public; detailed dependency path can be restricted.
Vulnerability source	Identifier, source authority, retrieval time, source snapshot or digest, matching rule.	Re-evaluate when source changes or policy review period expires.	Identifier and public advisory source normally public.
Candidate correlation	Matched component, version, confidence category, false-positive controls.	Invalid if the underlying SBOM or build binding changes.	Raw scanner output usually restricted.
Exploitability status	Under investigation, affected, not affected, mitigated, fixed, or other controlled value.	Must carry issuance, review, and supersession state.	Status can be disclosed; technical rationale may be tiered.
Rationale and method	Reachability, code presence, configuration, architecture, test, mitigation, or supplier evidence.	Re-test on relevant product or environment change.	Detailed attack-surface and test evidence restricted.
Active exploitation	Evidence source, affected geography or deployment, incident/reporting relation.	Independent of severity and VEX status; urgent update on new evidence.	Public communication separated from sensitive Article 14 reporting.
Remediation	Fixed release, mitigation, workaround, target date, verification, residual risk.	Close only after remediating release and result are verified for relevant subject.	User action and fixed version public where required.
Issuer and authority	Manufacturer, supplier, analyst, laboratory, authority, or delegated agent plus mandate.	Status checked at issuance and reliance time.	Issuer identity public where appropriate; internal delegation can be restricted.

Source: Authors' design informed by OASIS CSAF and CycloneDX VEX [44, 45].

3.9.2 Time-bounded claims and contradiction handling

A defensible vulnerability claim is time-bounded and source-bounded. It states that, for release R represented by composition hash H, no unresolved known exploitable vulnerability was identified against sources S as of cutoff T, subject to the method and limitations M. It does not state that the product is secure. When a new source record, exploit, architecture fact, or component correction appears, the earlier claim is superseded. It remains available as historical evidence of what was assessed and known at that time.

The graph can contain contradictory evidence. A supplier may declare a component not affected while the final manufacturer identifies a reachable vulnerable function in its integration. A scanner may report a candidate match while binary analysis finds that the component was removed at build time. Contradictions are represented explicitly and resolved through policy, authority, method strength, and scope. The system must not silently select the newest or most favourable statement.

3.10 Security update, installed-state, rollback, and recovery evidence

3.10.1 Update as a chain of events

An update is not one event. The architecture separates creation, testing, signing, publication, targeting, delivery, source and package verification, authorisation, installation, installed-version reporting, post-installation verification, failure, rollback, recovery, and closure. This chain follows the operational distinctions found across the CRA and ETSI drafts, especially EN 304 620, EN 304 621, and EN 304 627. It also separates manufacturer evidence from operator and device evidence. Figure 5 depicts the complete update, rollback, recovery, and closure lifecycle.




Figure 5: Security update, rollback, and recovery evidence lifecycle.

Source: Authors' design based on CRA update duties and EN 304 620, EN 304 621, EN 304 627, and EN 304 631.

Table 22: Update lifecycle events and minimum evidence

Event	Principal issuer or observer	Minimum evidence	Resulting state
Release created and tested	Manufacturer build and product-security functions.	Target release, source/build provenance, composition, tests, fixed vulnerabilities, limitations.	Remediating release candidate.
Package signed and published	Authorised release service or manufacturer.	Package hash, signer, signature, target products, prerequisites, advisory, publication time.	Authentic update available.
Affected subjects resolved	Manufacturer/operator vulnerability or fleet service.	Affected model/release/instance query, evidence snapshot, exclusions, target list version.	Defined update campaign scope.
Package received and verified	Device, management system, or authorised technician.	Source, signature, hash, compatibility, anti-rollback checks, verification result.	Eligible or rejected package.
Security-safety-availability gate	Operator, manufacturer, safety authority, or policy engine under mandate.	Impact assessment, conditions, change window, fallback, decision, approver, expiry.	Approved, conditional, deferred, or prohibited installation.

Event	Principal issuer or observer	Minimum evidence	Resulting state
Installation attempted	Device, installer, or management service.	Start/end time, package, prior version, instance, result, error code, actor.	Installed, failed, or unknown.
Installed version observed	Device, management agent, hardware attestation, or independent inspector.	Observed version, method, time, subject binding, attestation or signature.	Operational release state.
Post-install verification	Device, operator, laboratory, or service monitor.	Control tests, service health, configuration, safety/availability checks, deviations.	Verified, degraded, unsafe, or incomplete state.
Authorised rollback or recovery	Designated authority and technical executor.	Reason, approved target, anti-downgrade exception, recovery steps, result, residual risk.	Restored or rolled-back state with explicit exception.
Campaign closure	Manufacturer/operator campaign owner.	Coverage, unresolved instances, failures, exceptions, user notices, final evidence snapshot.	Closed, partially closed, or escalated campaign.

Source: Authors' update evidence model.

3.10.2 Assurance of installed state

Installed-state evidence has different assurance strengths. A manufacturer inventory can state which release should be installed. A device can self-report its version. A management agent can sign an observation. A secure boot or measured-attestation mechanism can bind the observation to a hardware-rooted state. An independent technician can inspect the device. The architecture records the method rather than treating all reports as equivalent. High-risk policies can require a stronger method or corroboration from independent sources.

Offline and air-gapped products remain supported. An authorised technician can import a signed update package, verify it locally, and issue a maintenance receipt that identifies the instance, prior and resulting release, tool, time, and verification result. The receipt can be synchronised later. Bitemporal recording then distinguishes when the update occurred from when the evidence entered the DPP system.

3.11 Continuous conformity and lifecycle lineage

3.11.1 Conformity snapshots rather than permanent status

Continuous conformity is used here as an architectural framing, not as a legal term or a claim that conformity assessment runs continuously. It denotes the ability to maintain and re-evaluate the evidence relevant to CRA obligations throughout the support period. The architecture preserves an initial conformity baseline and records events and evidence deltas that can affect it. Re-evaluation can be event-driven, periodic, risk-based, or requested by an authority, customer, or internal control function.

A conformity snapshot Σ_t is a scoped set of accepted claims, evidence, findings, route, and decision metadata for subject P and requirement set Q at time t . It identifies which requirements were assessed, which standards and editions were used, whether OJEU presumption applies, which evidence supported the result, which findings remain open, and who made the decision. A later snapshot supersedes but does not delete the earlier one.

$$\Sigma_t(P, Q) = \{\text{accepted claims, evidence basis, findings, route, decision authority}\}$$

Table 23: Events that trigger re-evaluation of a conformity snapshot

Trigger	Affected subjects	Required new evidence	Possible result
New vulnerability or active exploitation evidence	Component, release, instances, deployment, remote service.	Updated correlation, VEX, risk, advisory, mitigation, affected-instance analysis.	Snapshot remains valid with mitigation; becomes conditional; or requires remediation.
Security update or new release	Release, instances, service compatibility, deployment.	New composition, tests, package proof, installation and post-check evidence.	New release snapshot and instance-state updates.
Configuration drift or new deployment	Deployment and instance.	Observed configuration, deviation analysis, compensating controls, approval.	Permit, deny, or escalate for intended use.
Remote service change	Service, compatible releases, dependent deployments.	Service composition, interface change, tests, incident and compatibility evidence.	Updated product scope and dependency assurance.
Component substitution or repair	Instance/system, release baseline, supplier component.	Authenticity, compatibility, risk, safety impact, updated composition and test.	Updated instance state; possible substantial-modification assessment.
Support or supplier-support change	Model, release, component, service.	New support end, migration path, residual risk, contractual or supplier evidence.	Supported, extended support, unsupported, or replacement required.
Standard or legal status change	Requirement profile and conformity route.	New edition mapping, OJEU coverage, gap analysis, transition plan.	Presumption status or required assessment route changes.
Incident, authority request, or audit finding	Affected subjects and organisational processes.	Incident evidence, logs, root cause, remediation, corrective action and verification.	Corrected snapshot, escalation, notification, recall, or restriction.

Source: Authors' continuous-conformity model.


3.11.2 Append-only lifecycle and responsibility transitions

Lifecycle events are represented through an append-only lineage graph. Manufacture, placement on the market, installation, commissioning, configuration, update, repair, component replacement, ownership or operator transfer, substantial modification, second-life conversion, secure reset, decommissioning, and recycling each produce an attributable event and successor state. Prior evidence remains linked to the state for which it was valid. Figure 6 shows how lifecycle and responsibility transitions are preserved without overwriting prior states.

Source: Authors' design based on CRA responsibility for substantial modification, economic-operator traceability, and ESPR passport lineage.

A substantial modification creates a legal decision point. The architecture records the actor, modification, affected functions and components, legal analysis, risk assessment, and conclusion. Where a new manufacturer role arises, the system creates a new product lineage node and responsible legal person while preserving the original manufacturer evidence. Where the modification is not substantial, the current manufacturer or operator still records the changed state and its evidence. The graph should not infer the legal outcome automatically from the number of changed components.

Append-only lifecycle lineage and responsibility transitions



Every transition adds a new event or subject node; prior claims remain addressable for audit and responsibility analysis.

Figure 6: Append-only lifecycle lineage and responsibility transitions.

Table 24: Lifecycle event record

Event class	Minimum record	Responsibility implication	Evidence continuity rule
Manufacture or release	Product subject, baseline, actor, time, plant/build system, composition and quality evidence.	Initial manufacturer accountability.	Creates root or release node.
Installation and commissioning	Instance, release, deployment, configuration, installer, tests and acceptance.	Operator/integrator responsibility for proper installation and context.	Links assessed release to operational state.
Repair or maintenance	Fault, action, parts, firmware, actor, tests, resulting state.	Service actor accountable for event; manufacturer role depends on legal effect.	Prior state retained; successor state created.
Component substitution	Removed/added component, supplier, authenticity, compatibility, cyber and safety assessment.	Can trigger new risk assessment or substantial-modification analysis.	Composition and support dependencies updated.
Substantial modification	Change description, legal analysis, affected product scope, new risk and conformity evidence.	Potential new manufacturer for affected part or whole product.	New lineage node and responsible legal person where applicable.
Transfer or second life	New operator, intended use, location, configuration, support, residual risk, reset and credentials.	New operational duties; manufacturer role depends on modification and market action.	Original identity preserved; new deployment state linked.
Decommissioning	Secure reset, data removal, credential/key revocation, service termination, disposal or recycling.	Operator and service actors complete secure end-of-life actions.	Final state retained without retaining removed data.

Source: Authors' lifecycle event model.

3.12 Functional-safety and essential-availability interface

3.12.1 Co-assurance boundary

The CRA-capable DPP does not replace a functional-safety case, a sectoral safety management system, or certification under standards such as IEC 61508, ISO 26262, or sector-specific equivalents. It provides an interface through which cybersecurity changes can be evaluated against safety and essential-availability assumptions. The interface is required because software, firmware, network, cryptographic, and remote-service changes can alter timing, resource use, control paths, diagnostics, fail-safe behaviour, availability, and recovery.

The architecture therefore links each safety-relevant system baseline to the exact hardware, firmware, software, configuration, and remote-service subjects on which the safety argument depends. A security change creates a cyber-safety impact object. That object identifies affected safety functions, timing or resource changes, new failure modes, availability constraints, required tests, deployment conditions, approved fallback, rollback or recovery plan, and decision authority. The result can be approved, conditionally approved, deferred, or prohibited.

Table 25: Cyber-safety change gate

Gate element	Required evidence	Decision question	Outcome
Change definition	Affected release, component, configuration, interface, remote service, vulnerability and intended remediation.	What exactly changes, and which product/deployment subjects are affected?	Defined change scope.
Safety-function impact	Linked safety requirements, hazards, assumptions, diagnostic coverage, timing and failure modes.	Can the change invalidate or weaken a safety argument?	No impact, impact identified, or evidence incomplete.

Gate element	Required evidence	Decision question	Outcome
Availability impact	Essential function, downtime, update window, fallback, local operation, network/power recovery.	Can automatic or immediate installation create unacceptable availability or safety risk?	Install now, scheduled install, conditional defer, or prohibit.
Verification	Regression, integration, timing, failover, recovery, cybersecurity and safety tests.	Is the updated baseline verified for the intended deployment?	Approved or open findings.
Authority	Cybersecurity owner, product authority, functional-safety authority, operator and segregation-of-duties rule.	Who has competence to approve the change and any exception?	Valid approval, escalation, or rejection.
Fallback and rollback	Approved prior release, anti-downgrade exception, safe state, recovery instructions, residual risk.	Can the system be restored safely if installation or operation fails?	Approved fallback or unacceptable recovery risk.
Post-change observation	Installed version, configuration, operational tests, alarms, performance and safety status.	Does the deployed system match the approved baseline?	Verified, degraded, unsafe, or unknown.

Source: Authors' design informed by CRA health/safety considerations and EN 304 631-633.

3.12.2 Security exceptions and compensating controls

A deferred security update is represented as an exception with an owner, reason, affected subjects, risk assessment, compensating controls, expiry, review trigger, and approved target state. Safety or availability concerns do not erase the cybersecurity obligation. They change the remediation plan and evidence required. The same applies to products that cannot update automatically or that require planned outage windows. The DPP allows the manufacturer and operator to distinguish justified, time-bounded exceptions from unmanaged delay.

3.13 Access control, confidentiality, and selective disclosure

3.13.1 Evidence-tier model

The DPP must separate discovery from disclosure. A public resolver can identify the product, manufacturer, support end, vulnerability contact, declaration of conformity, public advisory, and update instructions. A customer or operator can require detailed composition, VEX, configuration profiles, update status, and supplier evidence. A conformity assessment body or authority can require the full risk assessment, test reports, vulnerability correlation, deviations, and incident evidence. The architecture applies policies to graph nodes, edges, artefacts, and query results rather than treating the entire passport as one public or private document.

Table 26: Role- and purpose-based evidence views

View	Typical evidence	Normally excluded	Control basis
Public and user view	Product/manufacturer identity, support period, CVD contact, public advisories, affected versions, update and decommissioning instructions, declaration of conformity.	Detailed architecture, complete SBOM, raw logs, incident evidence, personal telemetry, secrets.	Applicable CRA user/public duties and DPP access rules.
Purchaser or operator view	Detailed SBOM/HBOM, VEX, supported configurations, remote-service dependencies, update/rollback state, support entitlements, supplier evidence.	Other customers' data, private keys, unrestricted exploit detail, internal personnel data.	Contract, legitimate operational purpose, role and mandate.

View	Typical evidence	Normally excluded	Control basis
Supplier or integrator view	Component interfaces, evidence requests, affected-product feedback, integration assumptions, remediation and substitution records.	Unrelated system architecture, competitor supplier data, operator telemetry outside scope.	Supply-chain role, product scope, NDA or data-space policy.
Conformity assessment view	Risk assessment, applicable requirements, complete technical evidence, tests, deviations, quality/process evidence.	Evidence outside assessment scope or prohibited personal data.	Conformity route, assessor competence, engagement and legal powers.
Market-surveillance or reporting view	Technical documentation, current and historical evidence, economic-operator chain, incidents, Article 14 inputs, corrective actions.	Information not required for statutory purpose.	Legal authority, purpose limitation, secure channel and audit.
Incident-response view	Affected instances, versions, exploitation evidence, mitigations, update campaign, communication state, relevant operational logs.	Broad unrelated product or customer data.	Emergency policy, need-to-know, time-bound access and later review.

Source: Authors' access model derived from CRA information tiers, ESPR role-based access, and data-minimisation principles.

3.13.2 Policy enforcement and verifiable presentations

Attribute-based and policy-based access control are suitable because decisions depend on organisational role, product scope, purpose, jurisdiction, contract, incident status, and evidence sensitivity. The policy can release a claim, a redacted artefact, an aggregate result, a cryptographic proof, or a controlled reference. A verifier can receive proof that an SBOM passed a defined validation profile without receiving the full SBOM, where that level of disclosure is legally and operationally sufficient. Conversely, a market-surveillance authority can receive the complete artefact under a different policy.

Verifiable credentials and presentations can package claims with issuer, proof, validity, and status metadata [15]. JOSE or COSE-based security can protect JSON or CBOR representations [16]. Qualified seals, electronic timestamps, or attestations can provide regulated trust services. These mechanisms are implementation options. The architecture requires verifiability and policy enforcement, not one credential or cryptographic format.

Every disclosure event should record relying party, purpose, policy, requested and released evidence, time, and result. This supports accountability and helps detect over-broad collection. The audit record should contain identifiers and hashes rather than unnecessary copies of sensitive data. Raw security logs and secrets remain outside the unrestricted passport and are accessed only through controlled evidence references.

3.14 Cryptographic binding, status, and long-term verifiability

3.14.1 Binding chain

The central proposition requires cryptographic binding at multiple levels. The data carrier must resolve the intended passport. The passport identifier must resolve the intended subject. The release manifest must bind composition and assessment evidence to the exact build. An instance claim must bind the observed state to the actual device or workload. A deployment claim must bind configuration and operator context to the instance. A service claim must bind the service release and endpoint to the manufacturer-responsible remote function. Each issuer proof must bind the claim to the authorised legal person and mandate.

Table 27: Binding mechanisms by subject and assurance need

Binding target	Baseline mechanism	Higher-assurance mechanism	Residual limitation
Data carrier to passport	Persistent identifier and resolver; tamper-evident label where appropriate.	Signed data carrier payload, secure tag, registry proof or device-assisted resolution.	A cloned carrier can still point to a valid passport for another item.
Model/type to manufacturer	Manufacturer-signed or sealed product record and registry identity.	Qualified seal, authoritative registry attestation, controlled issuance process.	Identity does not prove product security or item authenticity.
Release to composition and binary	Signed manifest, SBOM hash, build identifier and binary/image digest.	Reproducible build, controlled build attestation, independent verification.	A complete manifest can still omit runtime or external dependencies outside scope.
Instance to physical/logical product	Serial or installation identifier plus manufacturer/operator record.	Device certificate, secure element, measured boot, remote attestation or independent inspection.	Compromised device roots or cloning remain possible and require lifecycle controls.
Deployment to configuration	Operator-signed configuration profile and instance relation.	Configuration digest, management-plane attestation, independent compliance scan.	Dynamic state can change after observation.
Remote service to product function	Service identifier, endpoint, release and compatibility claim.	Signed service manifest, workload identity, deployment attestation and monitored interface proof.	Cloud internals may remain opaque; contractual and audit evidence remains relevant.
Claim to issuer authority	Signature or seal plus issuer identifier.	Credential/mandate chain, status, qualified trust service, dual control.	Signature validity does not prove technical correctness.
Historical evidence to time	Trusted timestamp and retained certificate/credential status.	Archival validation data, periodic re-sealing, multi-signature or hybrid transition proof.	Long-term verification depends on preservation and algorithm migration.

Source: Authors' cryptographic binding model.

3.14.2 Status, revocation, and supersession

Status is multidimensional. An issuer credential can be revoked while a previously issued evidence object remains historically authentic. A mandate can expire while a claim issued within its validity remains attributable. A claim can be withdrawn, contested, or superseded. A product can be supported, unsupported, recalled, decommissioned, or transferred. A vulnerability claim can be under investigation, affected, mitigated, or fixed. These states should not share one status code.

Credential status mechanisms, including privacy-preserving status lists, can communicate revocation or suspension [17]. Evidence and product status can use separate registries or signed status objects. Verifiers must check the correct status dimension at the relevant time. A status endpoint also becomes an availability dependency; high-assurance deployments can cache signed status snapshots with expiry and use multiple resolvers.

3.14.3 Crypto agility and long-lived evidence

CRA evidence can require retention for periods that exceed the safe lifetime of current cryptographic algorithms, certificates, or keys. The architecture therefore records algorithm identifiers, key identifiers, certificate or credential chains, proof format, signing time, trusted timestamps, and validation data. It supports key rotation, proof renewal, archival timestamping, and migration of evidence without changing the underlying claim. A new preservation proof should reference the original evidence hash and prior proof chain.

For high-assurance or long-lived products, the architecture can support multi-signature or multi-hybrid cryptography, including parallel conventional and post-quantum proofs during transition. This is a proposed resilience option, not a CRA or DPP requirement. Semantic claim and identifier models should allow multiple verification methods and controlled phase-in or phase-out of cryptographic primitives. The reliance policy determines which combinations are acceptable at a given time and risk level.

No distributed ledger is required. A ledger can provide ordering, timestamping, or transparency proofs, but it cannot establish that an issuer was authorised, that an SBOM is complete, that a device report is genuine, or that disclosure is lawful. Those properties remain functions of the control plane, evidence methods, and governance.

3.15 Interoperability stack and deployment patterns

3.15.1 Technology-neutral logical profile

The reference architecture is intentionally compositional. No existing standard covers product identity, organisational authority, composition, vulnerability, assessment, lifecycle, access, and long-term status as one complete stack. Interoperability therefore requires profiles that combine standards while preserving their distinct scopes. The profile records the exact edition and conformance level used for every evidence object.

Table 28: Candidate interoperability stack for a CRA-capable DPP

Function	Candidate standards or mechanisms	Use in the architecture	Boundary
DPP exchange, identifiers and carrier	OJEU-cited EN 18216, EN 18219, EN 18220 and related DPP standards.	Statutory DPP discovery, exchange, product identifiers, carriers, access and lifecycle interoperability.	Presumption concerns covered ESPR DPP requirements, not CRA cybersecurity controls.
Industrial digital twin	Asset Administration Shell meta-model, APIs, submodels, and ABAC security model.	Industrial product/system representation and controlled submodels for evidence references and operational state.	AAS does not itself establish issuer authority or CRA conformity [46].
Composition and provenance	SPDX, CycloneDX, signed build attestations, product and package identifiers.	SBOM, HBOM, service, cryptographic, AI/data, and runtime composition with provenance.	Exact profile and completeness must be declared; identifiers require mapping.
Security advisories and VEX	OASIS CSAF, CycloneDX VEX/VDR, supplier advisories.	Affected products, exploitability status, remediation, action statements and machine-readable dissemination.	Issuer authority, product binding, freshness and evidence method remain additional.

Function	Candidate standards or mechanisms	Use in the architecture	Boundary
Claims and presentations	W3C Verifiable Credentials Data Model, JOSE/COSE secured credentials, signed JSON/CBOR, eIDAS seals and attestations.	Portable, tamper-evident claims and selective presentations.	Credential format does not determine legal effect or evidence sufficiency.
Identifiers and resolution	URI/URN schemes, statutory DPP identifiers, registries, DIDs where governed appropriately.	Persistent subject and evidence identifiers, service endpoints, key and verification-method resolution.	Identifier governance and legal-person mapping must be explicit.
Status and revocation	Credential status lists, certificate status, signed evidence-status objects, DPP registry versioning.	Revocation, suspension, withdrawal, expiry and supersession.	Different status dimensions must remain separate.
Policy and access	ABAC/PBAC, ODRL, XACML-like policies, policy-as-code, data-space usage controls.	Purpose, role, mandate, product scope, risk and evidence-aware decisions.	Policy correctness and legal basis require governance and testing.
Schema and semantics	JSON-LD Schema, SHACL, controlled vocabularies, RDF/JSON-LD or equivalent typed models.	Structural validation, relation validation, legal-status tagging, cross-standard mappings.	Schema validity is not factual or legal conformity.
Transport and federation	Secure APIs, registered delivery, data-space connectors, verifiable presentations, offline evidence packages.	Federated evidence resolution and controlled cross-company exchange.	Transport security does not replace data provenance or authority.

Source: Authors' interoperability analysis; standards are implementation candidates, not universally mandated choices.

3.15.2 API and resolver functions

A conforming implementation should expose functions rather than one monolithic document download. Core functions include resolving a product subject; retrieving the current and historical subject graph; requesting an authorised evidence presentation; verifying claim integrity, status, authority, and freshness; resolving composition and affected-product relations; submitting an authorised lifecycle event; recording supersession; querying support and update state; exporting an evidence package; and returning a policy decision receipt. Each response identifies schema and profile versions.

Query semantics are critical. "Show vulnerabilities for model M" differs from "identify instances in deployment D that run affected release R and lack verified mitigation X". The latter requires composition, release, installed-state, configuration, and freshness evidence. The API should return unknown or evidence-incomplete states explicitly rather than treating missing relations as negative results.

3.15.3 Deployment models

Table 29: Deployment models for CRA-capable evidence services

Model	Description	Strengths	Principal risks and controls
Central repository	Passport and most evidence are stored by one manufacturer or DPP provider.	Simple operations, consistent query interface, strong local governance.	Provider concentration, lock-in, confidentiality aggregation and outage; require export, backup, independent verification and continuity.
Federated resolution	Evidence remains with authoritative issuers; the passport stores identifiers, relations, hashes, status and access endpoints.	Preserves provenance, reduces replication, supports supplier autonomy and selective disclosure.	Endpoint availability, inconsistent policies and status; require signed snapshots, caching, common profiles and service continuity.
Hybrid registry and evidence graph	Registry provides product and actor discovery; graph metadata and critical claims are shared; detailed artefacts remain federated.	Balances interoperability, search, confidentiality, provenance and continuity.	Governance complexity; require clear source of truth, bitemporal state and conflict handling.
Offline evidence package	Signed export contains manifest, claims, artefact hashes, selected evidence, status snapshot and validation instructions.	Supports air-gapped assets, inspections, long-term archive and provider failure.	Can become stale; require expiry, source cut-off, later synchronisation and preserved status evidence.

Source: Authors' deployment model comparison. The hybrid model is the preferred reference implementation, but no topology is legally mandated.

3.16 Conformance profiles and validation rules

3.16.1 Minimum capability profiles

Not every product requires the same evidence granularity or assurance. The architecture defines three cumulative implementation profiles. They are authors' design profiles and do not correspond to CRA product classes or conformity modules. The profile selected depends on product category, intended use, operating environment, support model, safety and availability implications, and relying-party needs.

Table 30: Proposed CRA-capable DPP implementation profiles

Profile	Required subject and evidence coverage	Control-plane requirements	Typical use
Core Product Evidence Profile	Model and release subjects; manufacturer identity; scope and intended use; applicable requirements; risk and control mapping; release-bound SBOM; support period; CVD contact; public advisories; declaration and conformity route; status and lineage.	Accountable issuer, integrity proof, source-status tags, basic role access, export and continuity.	Software releases and products where instance state is not available or required for the decision.

Profile	Required subject and evidence coverage	Control-plane requirements	Typical use
Operational Assurance Profile	Core plus instance, deployment and remote-service subjects; installed version; configuration profile; VEX; update, rollback, recovery and support events; affected-instance resolution; operator evidence.	Scoped mandates, zero-trust PDP/PEP, freshness policies, selective disclosure, decision receipts and escalation.	Connected devices, managed fleets, enterprise software deployments, cloud-dependent products.
High-Assurance Cyber-Safety Profile	Operational plus stronger device/service attestation, independent assessment evidence, safety-baseline links, change gates, dual control, fine-grained audit, long-term proof preservation and crypto-agility plan.	Segregation of duties, high-assurance legal-person and mandate verification, multi-source evidence, continuity and offline operation.	Critical infrastructure, safety-relevant cyber-physical systems, long-lived industrial and energy assets.

Source: Authors' proposed implementation profiles. They are not legal CRA categories.

3.16.2 Structural, factual, and legal validation

Validation occurs at three levels. Structural validation checks schema, required fields, identifiers, value types, signatures, hashes, and graph cardinalities. Semantic validation checks controlled vocabularies, relation meaning, subject granularity, status transitions, and cross-standard mappings. Evidentiary validation evaluates method, authority, subject binding, completeness, freshness, contradictions, and limitations. Legal conformity assessment then applies the relevant law, route, presumption, and decision authority. A system must not present success at one level as success at another.

Representative validation rules include: every release-bound SBOM must identify the exact release and hash; every current vulnerability claim must identify source cutoff and review time; every installed-version observation must identify the observation method and instance binding; every draft-standard reference must be tagged as draft; every OJEU presumption claim must identify the covered clauses and essential requirements; every superseded claim must point to a successor or reason; every restricted evidence object must have an enforceable access policy; and every conformity decision must identify open findings and the decision authority.

Minimum criterion for the term "CRA-capable" (authors' design)

An implementation should not be described as CRA-capable unless it can represent source status, the five-subject model, accountable issuer authority, release-bound composition, temporal claim status, requirement-to-evidence traceability, vulnerability and update lifecycle evidence, access-controlled disclosure, append-only lineage, export and continuity, and scoped reliance decisions without an unqualified global compliance flag.

3.17 Threat model for the DPP evidence infrastructure

3.17.1 Adversarial and systemic failure modes

The DPP evidence service becomes security-relevant infrastructure. Its compromise can hide vulnerable products, misdirect updates, disclose sensitive architecture, or create false regulatory

assurance. The threat model therefore covers malicious actors, compromised issuers, insider misuse, implementation error, stale data, dependency failure, and governance failure. The service itself can constitute a product with digital elements depending on how it is supplied and used; its own CRA scope must be assessed separately.

Table 31: Threats and required architecture controls

Threat	Failure mode	Architecture controls	Residual governance issue
Issuer impersonation or key compromise	Forged SBOM, VEX, update, test, or conformity evidence.	Legal-person verification, credential and key status, timestamps, rotation, dual control, anomaly detection, revocation and re-issuance.	Timely detection and liability allocation.
Valid signer without authority	Employee, workload, or agent issues a claim outside role or product scope.	Mandate verification, product/action scope, segregation of duties, PDP/PEP enforcement and decision receipt.	Quality of enterprise role governance.
Subject substitution or cloned carrier	Valid evidence is applied to another model, release, instance, or item.	Persistent identifiers, release digests, device credentials, attestation, registry proof and cross-state comparison.	Cost and feasibility of strong physical binding.
Stale evidence replay	Expired VEX, support, status, configuration, or installed-version claim remains current.	Validity, source cutoff, review time, signed status, supersession, freshness policy and event triggers.	Availability of authoritative status sources.
Evidence omission or selective favourable view	Relevant finding, component, incident, or contradictory evidence is hidden.	Completeness declarations, required graph relations, multi-source evidence, authority access, signed query and disclosure receipts.	Audit rights and sanctions for deliberate omission.
Graph or resolver poisoning	Identifiers, endpoints, relations, or mappings redirect the verifier to false evidence.	Signed registry metadata, DNS/TLS and resolver controls, content addressing, multiple resolution sources and monitoring.	Governance of namespaces and registries.
Equivocation across recipients	Issuer presents inconsistent claims to customers, assessors and authorities.	Claim identifiers, transparency or audit logs, signed presentations, authority cross-check, historical snapshots.	Legitimate redaction must be distinguishable from factual inconsistency.
Unauthorised disclosure	SBOM, architecture, logs, personal or incident data are exposed.	ABAC/PBAC, encryption, selective disclosure, data minimisation, purpose limitation, audit and incident response.	Balancing transparency, trade secrets and security.
Provider outage or cessation	Evidence, status, or resolver becomes unavailable during support or retention period.	Export, backup, escrow or continuity provider, open formats, offline packages, multiple resolvers and recovery tests.	Funding and legal responsibility for long-term continuity.
Policy or automation error	Incorrect rule permits reliance, blocks urgent action, or outputs false assurance.	Versioned policy, testing, explainable reasons, change approval, simulation, monitoring and human escalation.	Responsibility for automated decisions.
Cryptographic obsolescence	Old proofs cannot be validated or become forgeable.	Algorithm metadata, trusted time, archival validation, re-sealing, crypto agility, multi-hybrid transition where justified.	Migration governance and long-term trust anchors.

Source: Authors' threat model.

3.17.2 Security properties of the evidence infrastructure

The architecture seeks authenticity, integrity, availability, confidentiality, accountability, non-equivocation, traceability, recoverability, and controlled transparency. It does not promise immutable truth. Evidence can be wrong, incomplete, or contested. The system makes those limitations visible, preserves provenance, and allows correction without erasure. This is a stronger and more realistic property than an immutable database containing unverified assertions.

3.18 Worked example: managed enterprise router with remote controller

3.18.1 Initial release and deployment

A fictional managed enterprise router, XR-200, illustrates the architecture. Manufacturer A creates a model subject M-XR200 and declares intended use, operating environment, support period, CRA classification, and conformity route. It releases software build R-3.4.0, signs a release manifest and SBOM, links the applicable draft EN 304 627 requirements, and retains risk, test, and assessment artefacts. A conformity snapshot records that the draft standard does not yet provide OJEU presumption.

Instance I-0147 is commissioned by Operator O at site D-BER-EDGE-07. The device reports installed release R-3.4.0 through a management agent. The operator issues a configuration claim for deployment profile DP-CRITICAL-BRANCH and records approved deviations from the manufacturer secure default. The router depends on manufacturer-controlled remote controller service S-5.8. The service is represented separately, with a compatibility relation to R-3.4.0 and a support relation to the deployment.

3.18.2 Vulnerability and remediation sequence

A new vulnerability source identifies a flaw in component libQ version 2.1. The manufacturer correlation service produces a candidate match against the SBOM for R-3.4.0. Product-specific analysis confirms that the vulnerable function is reachable through the remote-management interface in the deployed configuration. Manufacturer A issues an affected VEX claim, a public advisory, a temporary mitigation, and a remediating release R-3.4.1. The claim states source cutoff, affected range, method, evidence, residual risk, and review time.

The evidence graph resolves instances that run R-3.4.0 and deployments in which the interface is enabled. I-0147 is included. The operator policy requires a change window because the router supports an essential site connection. The cyber-safety-availability gate permits temporary operation for 48 hours under the mitigation and schedules installation. This decision is recorded as a scoped conditional permit, not a green product-compliance status.

At the change window, the device verifies the update signature, hash, target model, and anti-rollback policy. It installs R-3.4.1, reports the resulting version, and produces an attested boot measurement. The management service verifies connectivity, policy enforcement, and configuration. The update campaign record closes for I-0147. The earlier affected VEX claim remains in history; a fixed claim and new conformity snapshot supersede it for the remediated release and instance state.

Table 32: Evidence timeline for the XR-200 example

Step	New or changed evidence	Control-plane decision	Current state represented
1. Model registration	Model identity, manufacturer, intended use, support, scope and classification.	Manufacturer authority verified; release issuance permitted.	M-XR200 is the regulated model subject.
2. Release R-3.4.0	Signed manifest, SBOM, risk/control mapping, tests, draft-standard mapping, conformity snapshot.	Internal-control authority approves release within defined scope.	R-3.4.0 is assessed but draft standard provides no OJEU presumption.
3. Instance commissioning	Instance identity, installed-version observation, operator configuration and acceptance.	Operator mandate permits commissioning and configuration claim.	I-0147 runs R-3.4.0 in D-BER-EDGE-07.
4. Remote service relation	Service S-5.8 identity, release, compatibility, support and dependency.	Manufacturer service authority verified.	Deployment depends on manufacturer-responsible service S-5.8.
5. Candidate vulnerability	Source snapshot, component correlation and affected product range.	Evidence incomplete; policy escalates to product-specific analysis.	Candidate match, not yet an affected verdict.
6. Affected VEX and mitigation	Reachability test, affected claim, advisory, mitigation, residual risk and review time.	Manufacturer authority permits publication; operator receives restricted detail.	R-3.4.0 and I-0147 affected; temporary mitigation available.
7. Conditional continued operation	Deployment risk, essential-availability impact, mitigation verification and expiry.	Operator PDP returns conditional permit for 48 hours and schedules update.	Operation allowed within a narrow time and configuration scope.
8. Update R-3.4.1	Signed package, new SBOM, fixed claim, test and compatibility evidence.	Update approved for target instance and change window.	Remediating release available.
9. Installation and verification	Package verification, installation receipt, installed-version attestation and post-check.	Policy permits campaign closure for I-0147.	Instance runs R-3.4.1; fixed claim current; earlier evidence retained.

Source: Authors' illustrative example. Product and identifiers are fictional.

3.18.3 Example queries and bounded answers

The example supports distinct queries. A public user can ask whether XR-200 is supported and whether an urgent update affects release R-3.4.0. An operator can ask which managed instances are affected, which mitigation applies, and whether installation succeeded. A conformity assessor can inspect the release-bound SBOM, source snapshots, reachability test, draft-standard mapping, and open findings. An authority can reconstruct the manufacturer decision, communications, campaign coverage, and historical state. Each party receives a different evidence view, but all views refer to the same subject and claim identifiers.

A question such as "Is XR-200 CRA compliant?" is rejected as under-specified. The system requests a subject, release or instance, time, requirement scope, conformity route, and intended decision. It can then return a scoped statement, for example: "For instance I-0147 at 17:00 UTC on date T, installed release R-3.4.1 was attested, the applicable update and known-exploitable-vulnerability controls were supported by evidence set E, no open high-severity finding was recorded for those controls, and the operator policy permitted continued operation. This statement is not a general legal conformity determination."

3.19 Architecture traceability and specification for Part III

The reference architecture is complete only if it implements the design requirements derived in Part I. Table 33 groups the traceability relation. Detailed clause-level traceability will later be maintained in a machine-readable requirements register and bibliography package.

Table 33: Traceability from Part I design requirements to Part II architecture

Design requirements	Part II mechanism	Primary architecture output
DR-01, DR-14, DR-20, DR-22	Source-status metadata, requirement-to-evidence graph, validation layers, scoped conformity snapshots and bounded decisions.	Legally accurate traceability without an unqualified compliance flag.
DR-02, DR-03, DR-05, DR-06	Five-subject model, persistent identifiers, state comparison and independent remote-service subject.	Correct attribution to model, release, instance, deployment and service.
DR-04, DR-09, DR-11	Release-bound composition, verifiable claim envelope, integrity proofs, bitemporal and freshness metadata.	Integrity-protected, time-bounded evidence with provenance.
DR-07, DR-08, DR-21	Control plane, legal-person identity, roles, mandates, zero-trust PDP/PEP and decision receipts.	Accountable and policy-authorised issuance, disclosure and reliance.
DR-10, DR-12, DR-13, DR-15	Status and supersession, vulnerability/VEX state model, update lifecycle and support-state evidence.	Dynamic post-market evidence and continuous conformity.
DR-16, DR-18	Append-only lifecycle lineage and cyber-safety change gate.	Attributable modifications, second life, decommissioning and safety-aware updates.
DR-17	Evidence-tier model, selective disclosure, restricted artefacts and policy audit.	Confidentiality, data minimisation and role-appropriate transparency.
DR-19	Federated custody, open profiles, export, offline package, backup, status snapshots and crypto agility.	Provider-independent continuity and long-term verifiability.

Source: Authors' traceability analysis.

Part III evaluates this architecture through operational and sector-specific implementation patterns. The evaluation tests whether the subject model reduces ambiguity, whether the evidence graph improves exposure analysis and conformity operations, whether the control plane supports accountable cross-company automation, whether selective disclosure preserves legitimate confidentiality, and whether the cyber-safety interface supports controlled updates in systems where availability or physical harm matters. The analysis will compare software products, connected devices, remote-service products, industrial systems of systems, and repaired or second-life products.

4 Part III - Operational Value and Sector Implementation

This Part answers RQ3 and advances RQ4 by evaluating how the reference architecture developed in Part II supports operational decisions across cybersecurity, conformity, lifecycle management, and functional-safety change control. The evaluation is analytical and scenario-based. It identifies the causal mechanisms through which a CRA-capable DPP can improve a decision process, the evidence needed for that improvement, and the conditions under which the value claim fails. It does not claim measured industry-wide cost savings, risk reduction, or legal compliance.

The object of evaluation is not the passport interface or data carrier. It is an evidence-enabled

decision episode: a trigger occurs; the relevant product subject and time are identified; authorised evidence is resolved; a policy or competent person reaches a scoped decision; an action is taken; and a verifiable receipt closes or escalates the episode. The architecture has operational value only where this sequence is more precise, timely, accountable, or reusable than the counterfactual process based on static technical files, disconnected asset inventories, manually exchanged spreadsheets, and unverified supplier statements.

4.1 Evaluation design and analytical scenarios

4.1.1 Decision episode as the unit of analysis

A decision episode is represented by the tuple

$$d = \langle \text{trigger, subject, context, time, authority,} \\ \text{evidence, decision, action, receipt, residual uncertainty} \rangle$$

. This unit preserves the distinction between information availability and decision quality. A technically valid SBOM can be irrelevant if it refers to the wrong build. A current VEX statement can be insufficient if the deployment configuration is unknown. An attested installed version can be unusable if the observer lacks authority or the decision requires an independent assessment. Evaluation must therefore test the complete chain rather than the presence of isolated artefacts.

The counterfactual is a conventional fragmented process. Product identity resides in product lifecycle management; software composition in CI/CD tooling; device state in a configuration-management database or fleet service; vulnerabilities in a security platform; conformity evidence in a document repository; operator configuration in an OT or cloud system; and organisational authority in contracts, directories, or email. Such systems can be individually mature. The weakness arises where a cross-company decision depends on reliable joins among them and where the join itself is not versioned, attributable, or reproducible.

Table 34: Analytical evaluation criteria for CRA-capable DPP decision episodes

Criterion	Positive evaluation condition	Failure condition
Legal-scope fidelity	Correctly distinguishes CRA scope, exclusions, product category, conformity route, and adjacent regimes.	A technically complete passport is presented as proof of CRA applicability or compliance.
Subject precision	Binds each claim to model, release, instance, deployment, remote service, or lifecycle node.	Model-level evidence is applied to an unverified installed instance.
Temporal correctness	Records valid time, recorded time, source cut-off, freshness, expiry, and supersession.	A formerly correct vulnerability or support claim is treated as current.
Evidence traceability	Connects requirement, control, method, result, artefact, issuer, and subject.	A certificate or score cannot be traced to the tested configuration.
Authority and accountability	Verifies issuer identity, role, mandate, status, and decision authority.	A valid signature is relied on without proof that the signer could make the claim.
Operational actionability	Resolves affected assets, mitigations, update targets, exceptions, and closure evidence.	The system reports a component match but cannot identify actionable products.
Controlled transparency	Discloses the minimum evidence required for the relying purpose.	Sensitive SBOMs, logs, vulnerabilities, or topology are made public by default.
Interoperability and portability	Uses stable identifiers, explicit profiles, export, and independent verification.	Evidence is locked in one provider or loses meaning outside one platform.
Cyber-safety compatibility	Links security changes to safety and essential-availability decisions.	An urgent patch is installed or deferred without assessing its safety impact.
Auditability and reproducibility	Preserves inputs, policy, evidence versions, decision, and receipt.	A later assessor cannot reconstruct why a decision was taken.

Source: Authors' design-science evaluation framework, derived from Parts I and II.

4.1.2 Scenario set

Six recurrent scenarios are used to test the architecture. They cover the periods before market placement, during normal operation, during incident response, and after modification or support end. The scenarios are intentionally cross-sectoral. Product-specific controls and conformity routes remain governed by the CRA, applicable harmonised standards, and sectoral legislation.

Table 35: Analytical scenarios used in Part III

Scenario	Decision question	Minimum evidence scope
S1 - New exploitable vulnerability	Which releases and deployed instances are affected, and which are not?	Component, release, instance, deployment, remote service; SBOM, VEX, configuration and freshness.
S2 - Urgent security update	Can the update be authorised, delivered, installed, verified, rolled back, and closed?	Release, instance, deployment; package proof, mandate, installation, attestation, rollback and recovery.
S3 - Remote-service change or outage	Does a cloud, API, controller, or model service change the assessed product state?	Remote-service version, compatibility, dependency, incident, support, and deployment relation.
S4 - Support end in a mixed fleet	Which products remain supported, exposed, or subject to migration and compensating controls?	Model, release, instance, component support, operator inventory, migration and exception evidence.
S5 - Repair, integration, or substantial modification	Who is responsible for the changed product and what evidence must be renewed?	Lifecycle lineage, changed components and functions, actor mandate, legal assessment, risk and conformity evidence.
S6 - Conformity or market-surveillance inquiry	Can the responsible actor produce the correct technical evidence for a defined product and time?	Requirement mapping, standards status, tests, risk assessment, declaration, post-market events and access policy.

Source: Authors' scenario set. The scenarios do not alter statutory reporting channels or decision competences.

4.1.3 Interpretation limits

A positive analytical result means that the architecture can represent the evidence and decision path without internal contradiction and with a plausible implementation route. It does not establish that an organisation has accurate source data, that suppliers will participate, that operators will permit device observation, or that a regulator will accept a particular evidence format. Those questions are treated as governance and external-validity limits in Part IV.

The analysis also rejects a common but weak value claim: more product data do not automatically produce better cybersecurity or compliance. Value depends on subject binding, freshness, provenance, authority, policy, and action closure. A DPP that increases the volume of unverified claims can increase false assurance and investigation cost.

4.2 Cybersecurity operations value

4.2.1 From component match to affected deployed product

Security operations commonly begin with a broad signal: a supplier advisory, vulnerability database entry, incident report, exploit notification, or threat-intelligence observation. The first correlation against an SBOM produces a candidate set, not an affected-product set. The

product context must establish whether the component and vulnerable function are present in the assessed build, whether the affected code path is reachable, whether the deployed configuration enables the path, whether a remote service changes exposure, and whether an effective mitigation or fixed release has been verified.

For a vulnerability v , decision context q , and time t , the architecture resolves an affected-instance set rather than equating component presence with exploitability:

$$A(v, q, t) = \{ i \mid \text{installed}(i, t) = r; \text{contains}(r, c); \text{applies}(v, c, r); \\ \text{reachable}(v, i, q, t); \text{no effective mitigation or fixed state is verified} \}$$

The expression is a logical model, not a legal formula. It makes the required joins explicit. Missing installed-state or configuration evidence produces an unknown state, not a negative finding. Candidate correlation remains useful for triage, but high-impact actions should be based on the strongest available product-context evidence. Figure 7 translates this logic into a verified affected-instance resolution process.




Figure 7: From component signal to verified affected-instance set.

Source: Authors' operationalisation of the evidence graph defined in Part II.

Table 36: Product-context vulnerability states and permitted operational meaning

State	Meaning	Required evidence	Permitted operational use
Candidate match	Affected component identifier or version appears in composition evidence.	SBOM/HBOM and source correlation with cut-off time.	Prioritise analysis; do not yet assert product impact.
Under assessment	Product-specific applicability or reachability is unresolved.	Build, architecture, interface, configuration and test evidence.	Escalate; preserve uncertainty and deadline.
Affected	Vulnerable function is present and relevant in the product context.	Product-specific analysis, VEX affected state, deployment relation.	Mitigate, update, notify, report where legally required.

Table 36: Product-context vulnerability states and permitted operational meaning

State	Meaning	Required evidence	Permitted operational use
Mitigated	A control reduces exposure for a defined configuration and time.	Mitigation instructions, operator implementation and verification.	Permit conditionally; monitor expiry and configuration drift.
Fixed	A remediating release is installed and verified.	Fixed-release claim, package proof, installed-state observation, post-check.	Close the instance campaign while retaining history.
Not affected	Evidence demonstrates absence, non-inclusion, non-reachability, or another valid basis.	VEX rationale, build provenance, architecture or test evidence.	Exclude from campaign subject to evidence freshness.
Unknown	Evidence is absent, stale, contradictory, inaccessible, or not attributable.	Explicit evidence gap and failed resolution record.	Treat according to risk policy; never translate unknown into safe.

Source: Authors' state model, aligned with the VEX and evidence-resolution architecture in Part II.

4.2.2 Coordinated vulnerability response and Article 14 readiness

Article 14 creates short reporting sequences for actively exploited vulnerabilities and severe incidents. The CRA-capable DPP does not replace the single reporting platform, the CSIRT designated as coordinator, ENISA, or manufacturer judgement about whether the legal threshold is met. It can reduce the time required to assemble consistent product, version, impact, mitigation, and communication evidence and can preserve the evidence package used for each reporting stage [1, Art. 14].

The strongest contribution is affected-set resolution. A manufacturer can link a vulnerability to affected release ranges and then resolve known customers, operators, products, services, or instances under the applicable access and data-protection policy. A later notification can supersede the early warning without erasing it. The graph records which facts were known at each stage and which remained estimates. This supports bitemporal reconstruction when the full scope becomes clear only after investigation.

Table 37: Evidence package for CRA vulnerability and incident response readiness

Evidence group	Content	CRA-capable DPP contribution	Boundary
Product scope	Product identity, affected models and releases, relevant remote services.	Supplies structured product identifiers and version relations.	The manufacturer remains responsible for legal scope and completeness.
Vulnerability or incident	Identifier, exploitation or incident state, discovery channel and time.	Links source snapshot, analysis and later correction.	Source presence alone does not prove active exploitation.
Impact and severity	Technical impact, users or functions affected, safety or availability implications.	Resolves impact evidence and controlled sector context.	Severity scoring does not replace legal reporting assessment.
Affected population	Known affected products, releases, operators, geographies or instances.	Produces an evidence-bounded affected set and unknown set.	Coverage depends on registration, customer and installed-state evidence.

Table 37: Evidence package for CRA vulnerability and incident response readiness

Evidence group	Content	CRA-capable DPP contribution	Boundary
Mitigation and corrective measure	Temporary control, fixed release, deployment instructions and limitations.	Distributes version-specific advice and records implementation evidence.	A published patch is not proof that deployed products are remediated.
Timeline and communications	Early warning, notification, final report, user notice and advisory versions.	Preserves time, recipients, content hashes and supersession.	The official reporting and user-notification duties remain outside the DPP service.
Closure evidence	Installed state, verification, rollback, residual risk and unresolved exceptions.	Supports campaign and post-incident closure.	Closure is scoped to observed products and evidence freshness.

Source: Authors' operational mapping based on [1, Art. 14] and the Part II evidence model.

4.2.3 Patch campaign, rollback, and recovery

A patch campaign is a chain of distinct states. Release publication, delivery, signature verification, authorisation, installation, restart, installed-version observation, post-change validation, rollback, and recovery cannot be collapsed into one update status. The distinction is explicit in the ETSI drafts for VPN products, network management systems, routers, and firewalls, and it is necessary for operational closure [29, 30, 36, 42].

The DPP graph permits manufacturer and operator evidence to coexist without conflating responsibility. The manufacturer can prove the identity, integrity, target scope, and security rationale of the release. The operator can prove scheduling, change approval, installation result, deployment-specific validation, and any time-bounded exception. A device or trusted management service can provide an observed or attested installed version. Each source remains attributable.

Table 38: Patch-campaign stages and closure evidence

Stage	Minimum decision evidence	Primary issuer or observer
Release authorised	Security review, target products, compatibility, safety/availability decision, signer authority.	Manufacturer release and mandate evidence.
Package delivered	Package identifier, hash, source, recipient or target group, delivery time.	Distribution service receipt.
Authenticity verified	Signature, certificate/status, hash and target-model validation.	Device or update client verification record.
Installation authorised	Operator mandate, policy, change window, exception conditions.	Control-plane decision receipt.
Installation executed	Start/end time, result, error, restart and rollback readiness.	Device, agent or service event.
Installed state observed	Resulting release, observation method, device binding and freshness.	Device attestation or management observation.
Post-change validation	Connectivity, security control, application, safety and availability checks.	Operator or independent assessment evidence.
Rollback or recovery	Authorised target, reason, anti-rollback exception, outcome and residual risk.	Auditable lifecycle event; not a silent downgrade.
Campaign closed or escalated	Known affected denominator, verified remediated set, unknowns and exceptions.	Scoped campaign receipt and open-finding list.

Source: Authors' synthesis of CRA update duties and [29, 30, 36, 42].

4.2.4 Supplier propagation, remote services, and fleet coordination

The final manufacturer remains responsible for the product, but component suppliers can issue evidence that reduces response latency and ambiguity. A supplier advisory or VEX statement is linked to the component version and supplier authority. The final manufacturer evaluates its use in each build, issues a product-specific claim, and resolves affected releases. Operators then resolve deployed instances and configurations. This staged propagation preserves the legal and technical distinction between component evidence and final-product responsibility.

Remote data processing solutions require the same discipline. A cloud controller, API, database, policy service, or model-serving endpoint may change independently from the device. A service incident or incompatible service release can make a device function unavailable or alter its risk state without changing local firmware. Representing the service as a separate subject permits compatibility, support, incident, and dependency evidence to be evaluated without treating the entire cloud estate as public passport data.

Fleet coordination benefits from a shared subject graph but not necessarily from a shared database. Manufacturers, suppliers, service providers, and operators can retain authoritative evidence in their own systems. The resolver obtains current, authorised claims, verifies status and mandate, and records the evidence versions used. This federated pattern is especially important where operators cannot disclose asset locations or configurations to all suppliers and where manufacturers cannot expose confidential build information to the public.

4.2.5 Operational metrics and residual uncertainty

The architecture supports operational indicators, but the indicators must not be presented as legal compliance scores. A metric is valid only for a defined population, evidence cut-off, and observation method. The unknown population must remain visible. For example, a remediation rate calculated only over online devices can materially overstate fleet closure.

Table 39: Proposed operational indicators for evidence-enabled cybersecurity management

Indicator	Definition	Required qualification
Evidence coverage	Subjects with current minimum evidence / subjects in the declared population.	Population definition and required evidence profile must be explicit.
Installed-state freshness	Age distribution of the latest trusted version observation.	Offline or unmanaged assets remain unknown, not current.
Affected-set resolution	Affected, not affected, mitigated, fixed and unknown counts for a vulnerability.	State definitions and source cut-off must be retained.
Verified remediation completion	Verified fixed instances / known affected instances.	Exclude unverified delivery; disclose unknown affected population.
Contradiction backlog	Open conflicts among supplier, manufacturer, device, operator or assessor claims.	Conflict count measures evidence quality, not product risk by itself.
Unsupported exposure	In-use instances or releases past support end, grouped by consequence and mitigation.	Support status does not establish exploitability but changes residual risk.
Exception discipline	Open exceptions, expired exceptions, and exceptions without verified compensating controls.	An approved exception is not equivalent to remediation.
Decision latency	Time from signal to scoped decision and from decision to verified closure.	Useful for process improvement; not a statutory threshold unless law or contract provides one.

Source: Authors' proposed metrics. They are management indicators, not CRA conformity measures.

4.3 Compliance operations value

4.3.1 Dynamic technical-documentation index

The CRA technical documentation remains the legally relevant evidence set. A CRA-capable DPP can serve as a dynamic, access-controlled index that binds each requirement to its applicability rationale, control, assessment method, evidence artefact, issuer, subject, result, and validity. The architecture does not require every confidential artefact to be copied into the passport service. It requires that the artefact can be resolved, its integrity and provenance can be verified, and the relationship to the assessed product can be reconstructed.

This model addresses a frequent documentation failure: the technical file contains a correct test report, but the report is not linked to the software build later placed on the market. Release-bound evidence and conformity snapshots make the boundary explicit. Where a component, build process, remote service, or intended-use profile changes, the system can identify which requirement mappings and tests require review. It cannot determine sufficiency without the competent technical and legal assessment.

4.3.2 Conformity snapshots and standards status

A conformity snapshot records what was assessed, against which legal requirements and standards, under which route, on what evidence, by whom, and with which open findings. It also records source status. A draft ETSI requirement, a final European standard, an OJEU-cited harmonised standard, a common specification, and a cybersecurity certificate have different legal effects. The DPP must preserve these distinctions and must not silently upgrade a draft mapping when a standard is later cited.

The use of snapshots supports continuous documentation without turning conformity into a permanent runtime flag. A product can retain a valid pre-market declaration while an installed instance becomes operationally exposed because it was not updated. Conversely, an instance can be technically remediated while a revised product release still requires formal documentation or assessment before market placement. The evidence graph allows both facts to coexist.

4.3.3 Market surveillance and economic-operator accountability

A market-surveillance request is a controlled evidence query. The authority may require the risk assessment, technical file, test evidence, standards mapping, vulnerability analysis, update process, support-period basis, and post-market records. Role- and purpose-based access permits the authority to receive complete evidence while the public resolver remains limited to user-facing information. The disclosure receipt records the request, legal basis, released evidence, version, and time.

Economic-operator identity and lineage are equally material. Manufacturer, authorised representative, importer, distributor, integrator, repairer, and modifying actor claims must remain distinguishable. The control plane verifies the legal person and mandate; the evidence plane records the product event. This supports the ten-year traceability duty without implying that every actor is legally responsible for every claim [1, Arts. 21-23].

4.3.4 Substantial modification and responsibility transition

Substantial modification is a legal and technical decision, not a database event that can be inferred from a component-count threshold. The DPP can assemble the facts needed for the

decision: original intended purpose, changed functions, altered interfaces, new components, security and safety impact, updated risk assessment, actor identity, and whether the changed product is made available on the market. The competent organisation then determines whether a new manufacturer role and conformity process arise.

Where responsibility changes, the graph creates a new lineage node and does not overwrite the original product history. Original manufacturer evidence remains attributable and can be reused only within its valid scope. The new responsible actor issues the modified-product identity, composition, support commitment, risk assessment, and conformity evidence. This pattern is central to repair, remanufacturing, integration, and second-life markets.

Table 40: Compliance workflows supported by the CRA-capable DPP

Workflow	Evidence resolved	Automatable contribution	Reserved decision
Scope and classification	Product scope, exclusions, core functionality, important/critical category and intended use.	Validated source data and legal analysis.	Scope determination and legal classification.
Risk assessment	Assets, threats, foreseeable use, controls, residual risk and support assumptions.	Requirement-control-evidence graph and versioning.	Adequacy and proportionality of the assessment.
Standards selection	Edition, clauses, OJEU status, partial application and coverage.	Automated status and traceability checks.	Whether the chosen standards sufficiently cover the product.
Technical documentation	Architecture, components, SBOM, tests, processes, update solution and declaration.	Integrity, completeness, subject and freshness validation.	Technical truth and sufficiency of evidence.
Post-market maintenance	New vulnerabilities, VEX, updates, support changes, incidents and corrections.	Event lineage, supersession and open-finding detection.	Regulatory response and risk acceptance.
Article 14 reporting readiness	Affected products, timeline, impact, mitigation, corrective measure and users.	Structured evidence package and communication targets.	Whether the threshold is met and what must be reported.
Market surveillance	Authority request, evidence release, audit trail and corrective action.	Purpose-bound disclosure and reproducible evidence package.	Authority findings and enforcement decision.
Modification and second life	Changed product, actor, purpose, composition, risk and conformity evidence.	Lineage and responsibility-transition package.	Whether modification is substantial and who becomes manufacturer.

Source: Authors' compliance-process mapping based on [1].

4.3.5 Boundary of compliance automation

The architecture can automate verification of syntax, identifiers, signatures, status, clause references, evidence freshness, required relationships, and policy rules. It can detect missing, stale, contradictory, or wrongly scoped evidence. It can calculate affected sets and prepare structured documentation packages. These functions reduce clerical work and improve auditability.

The architecture must not autonomously assert legal compliance where the decision depends on interpretation, technical sufficiency, proportionality, independence, or sector-specific judgement. A policy engine can deny a transaction because mandatory evidence is missing. It cannot transform complete evidence into a legally binding conformity decision unless the applicable law assigns that decision to the authorised process and actor represented in the control plane.

Table 41: Boundary between automated evidence processing and reserved compliance decisions

Decision class	Examples
Suitable for deterministic automation	Schema and identifier validation; signature and status checks; evidence-hash verification; source-status tagging; mandatory-field and graph-cardinality rules; freshness and expiry alerts; affected-version resolution; campaign coverage; access enforcement.
Suitable for decision support	Risk prioritisation; evidence assurance vector; conflict detection; test-coverage gaps; substantial-modification fact package; safety-impact screening; draft report preparation; supplier evidence comparison.
Requires competent human or legally authorised process	CRA scope; product core functionality; risk acceptability; sufficiency of controls and tests; substantial-modification conclusion; conformity assessment; Article 14 threshold; market-surveillance finding; functional-safety acceptance; liability allocation.
Prohibited presentation	A universal green compliance badge; a permanent "CRA compliant" state; an OJEU presumption based on a draft; an unknown evidence state presented as not affected; a platform registration presented as product conformity.

Source: Authors' automation-boundary model.

4.4 Business and lifecycle value

4.4.1 Procurement and supplier assurance

A buyer can express evidence requirements as machine-readable procurement conditions rather than as a general security questionnaire. Examples include support until a defined date, a release-bound SBOM, current exploitability evidence, a secure-update mechanism, an identified remote-service dependency, a remediation commitment, and conformity evidence for the intended deployment. The control plane can verify that the supplier or assessment body is authorised to issue the requested claims. The evidence plane can disclose a complete artefact or a bounded proof, depending on purpose and confidentiality.

This supports comparative supplier assurance but does not create a universal product score. A supplier may provide strong evidence for one release and weak evidence for another. An independent test can be strong but stale. A complete SBOM can coexist with poor remediation performance. Procurement rules should therefore apply mandatory gates and explicit dimensions rather than compress all evidence into one ranking.

4.4.2 Maintenance, service, warranty, and support

Maintenance decisions require the correct product state. A service provider can resolve the approved replacement component, compatible firmware, security advisory, update package, configuration profile, warranty or support entitlement, and required post-maintenance checks. The resulting repair or update event is attributed to the service organisation and linked to the affected instance. This reduces the risk that a functionally compatible component silently changes cybersecurity or safety assumptions.

Support entitlement and regulatory support period must remain distinct. A commercial service tier can provide faster response, monitoring, or extended support. It cannot reduce the statutory duties that apply to the manufacturer. The passport should expose the actual support state, responsible provider, covered product and release, and end date without presenting paid evidence access as a condition for legally required user information.

4.4.3 Targeted campaigns and customer communication

Product recalls, security campaigns, and service notices are often over-broad because manufacturers cannot reliably distinguish affected releases and instances. Subject-specific evidence permits more precise targeting. A campaign can be limited to a component batch, software release, configuration, remote-service dependency, or deployed instance set. The manufacturer can record which recipients received which advisory and which products reached verified closure.

Targeting must not be confused with narrowing legal duties. Where the evidence population is incomplete, the manufacturer may need to communicate broadly. The DPP makes the limitation visible by separating confirmed affected, confirmed not affected, and unknown populations.

4.4.4 Residual value, refurbishment, and second life

A used or refurbished product has value only if a buyer can understand its identity, installed state, remaining support, modifications, reset status, unresolved vulnerabilities, and responsible economic operator. The DPP lineage can preserve original evidence while adding repair, ownership transfer, secure data removal, component replacement, and new support commitments. This is especially material for long-lived industrial assets, energy systems, building equipment, and batteries.

The mechanism can support residual-value and insurance decisions, but the evidence should be treated as an input rather than a valuation or underwriting outcome. A passport cannot prove physical condition, future exploitability, or absence of undisclosed modifications unless relevant observation and assessment evidence is provided. Strong business claims require later empirical validation.

4.4.5 Value conditions and implementation cost

Business value is expected to be highest where product state changes frequently, installed fleets are large or distributed, multiple organisations contribute evidence, support periods are long, and the same evidence is reused across procurement, cybersecurity, compliance, maintenance, and second life. Value is lower for short-lived products with no update path, no remote dependency, one responsible actor, and few downstream decisions. This is a qualitative design proposition, not a quantified return-on-investment model.

Table 42: Business-value mechanisms, enabling conditions, and principal risks

Process	Value mechanism	Enabling condition	Principal risk
Procurement and onboarding	Reusable proof of identity, support, composition, vulnerability process, conformity route and deployment fit.	Evidence profiles, supplier mandates, selective disclosure and current status.	Automating acceptance of complete but weak evidence.
Supplier management	Component provenance, support state, advisory quality, update responsiveness and unresolved findings.	Federated supplier evidence and product-specific final-manufacturer analysis.	Treating supplier evidence as transfer of manufacturer responsibility.
Maintenance and field service	Correct parts, releases, configuration, work instructions, update and post-check evidence.	Instance identity, authorised service role and lifecycle event capture.	Unverified observations or unauthorised modification.
Targeted recall or security campaign	Narrower affected set, better communication and verifiable closure.	Current installed-state and customer/operator relations.	False precision where fleet evidence is incomplete.

Table 42: Business-value mechanisms, enabling conditions, and principal risks

Process	Value mechanism	Enabling condition	Principal risk
Warranty and support service	Transparent entitlement, response level, support end and responsible provider.	Contract and statutory support states kept separate.	Using commercial terms to obscure legal obligations.
Resale, refurbishment and second life	Remaining support, reset, modification, component and vulnerability history.	Append-only lineage and new responsibility evidence.	Assuming digital history proves physical condition.
Risk transfer and financing	Evidence input for due diligence, residual risk and control conditions.	Independent assessment and explicit model limitations.	Presenting a DPP score as an actuarial or safety guarantee.

Source: Authors' analytical business-value model. Effects require empirical validation.

Table 43: Principal implementation costs and governance responses

Cost or risk area	Why it arises	Required response
Source-system integration	PLM/ALM, CI/CD, SBOM, vulnerability, fleet, service, quality and conformity systems must expose stable, versioned evidence.	Begin with high-value evidence and preserve authoritative ownership rather than replicate everything.
Evidence quality	Poor identifiers, incomplete SBOMs, stale configurations and ambiguous ownership can undermine the graph.	Define validation profiles, data-quality owners, evidence gaps and supplier escalation.
Identity and mandate governance	Cross-company claims require legal-person identity, roles, delegation, status and separation of duties.	Use verifiable organisational authority and record reliance decisions.
Confidentiality and trade secrets	Composition, tests, topology and incidents may be sensitive.	Separate discovery from disclosure; release claims, proofs or controlled artefacts by purpose.
Operational continuity	A DPP provider or supplier may cease operation or withdraw evidence.	Provide export, backups, persistent identifiers, independent verification and continuity plans.
Long-term cryptography	Long-lived assets outlast algorithms, certificates, services and key owners.	Maintain crypto agility, timestamping, re-sealing and evidence-preservation procedures.
Governance and liability	More visible evidence can reveal contradictions, delays and responsibility gaps.	Define correction, dispute, reliance and escalation rules; do not suppress adverse evidence.

Source: Authors' implementation analysis.

4.5 Functional safety and essential availability

4.5.1 Cybersecurity evidence as a safety-case dependency

The CRA targets cybersecurity. Functional safety remains governed by sectoral legislation, safety standards, and competent safety processes. The interaction is nonetheless operationally unavoidable. Regulation (EU) 2023/1230 requires machinery within both regimes to satisfy the CRA cybersecurity requirements and the Machinery Regulation essential health and safety requirements. It also addresses protection against corruption and the safety and reliability of control systems [47].

A CRA-capable DPP supports co-assurance by binding the safety baseline to the exact hardware, firmware, software, configuration, network component, cryptographic state, and remote

service on which the safety argument depends. A cybersecurity event can then trigger a safety-impact review when it changes a referenced subject or assumption. The safety case remains authoritative within its own lifecycle and access regime.

4.5.2 Cyber-safety update gate

An urgent security update can reduce cyber risk while creating new timing, resource, availability, diagnostic, or control-path risk. Conversely, deferring an update can preserve a validated safety baseline while leaving an exploitable interface exposed. The architecture therefore requires an authorised change gate that evaluates both dimensions and records the decision as approved, conditionally approved, deferred, or prohibited. EN 304 631 provides a direct standards example by allowing product-specific limits on automatic updates where essential-function availability or safety would face unacceptable risk [37]. Figure 8 depicts the resulting cyber-safety decision and evidence-closure workflow.




Figure 8: Cyber-safety update decision and evidence closure.

Source: Authors' co-assurance workflow informed by [37, 38, 47].

Table 44: Cybersecurity changes that can invalidate functional-safety or availability assumptions

Cyber event	Potential safety or availability effect	Required linked evidence
Unauthorised parameter or configuration change	Unsafe setpoint, disabled protection, changed operating envelope.	Approved baseline, actor authority, parameter diff, safety-function impact and restoration proof.
Security update or library replacement	Changed timing, memory, interfaces, diagnostics, toolchain or failure behaviour.	Release diff, safety-impact analysis, regression tests, approved deployment and rollback state.
Deferred patch	Known exposure remains because immediate change threatens availability or validated safety state.	Exception owner, threat and safety analysis, compensating controls, expiry and target release.
Rollback	Restores prior function but may reintroduce a known vulnerability or incompatible service state.	Authorised target, anti-rollback decision, vulnerability status, safety baseline and post-check.

Table 44: Cybersecurity changes that can invalidate functional-safety or availability assumptions

Cyber event	Potential safety or availability effect	Required linked evidence
Remote-service outage or change	Loss of control, monitoring, authentication, update, alarm or decision support.	Dependency relation, local fallback, recovery test, service compatibility and incident state.
Cryptographic migration	Changed boot, communication, key size, latency, hardware load or trust anchors.	Crypto profile, performance and timing tests, dual-operation plan, key and certificate lineage.
Sensor, log, or time-source compromise	Incorrect safety response, missing event sequence, unreliable forensic evidence.	Authenticated source, time assurance, integrity, plausibility checks and contradiction handling.

Source: Authors' co-assurance analysis; the table does not replace sector-specific hazard analysis.

4.5.3 Time-bounded exceptions and compensating controls

A deferred update is represented as a controlled exception, not as a successful security state. The exception identifies the responsible authority, affected products and deployments, decision rationale, threat exposure, safety or availability constraint, compensating controls, monitoring, expiry, review triggers, and approved target state. The policy engine can prevent the exception from being silently renewed and can escalate where evidence of the compensating control becomes stale.

Examples include isolating an industrial controller while a safety-certified release is prepared, disabling a vulnerable remote interface, restricting operation to a local mode, adding network monitoring, or scheduling a grid asset update during an approved outage window. The effectiveness of each control must be verified in the actual deployment. A generic mitigation instruction is not an operator implementation receipt.

4.5.4 Runtime validation and limitation of the DPP

Post-change evidence should include the installed release, configuration, boot or runtime measurement where appropriate, connectivity, safety function, alarm, recovery, and essential-service checks. High-assurance deployments may require independent or dual-control approval. Where the product cannot provide attestation, the observation method and its limitations must be stated.

The DPP cannot determine that a system is safe merely because the approved versions are present. Physical degradation, environmental conditions, hidden common-cause failures, sensor drift, and process hazards remain outside a purely digital evidence graph unless relevant observations are supplied. The architecture provides traceable configuration and change evidence to the safety process; it does not replace hazard analysis, validation, certification, or operational safety management.

4.6 Reusable implementation patterns

The architecture is implemented through recurring patterns rather than one universal passport. The patterns differ in subject coverage, evidence sources, access needs, and operational decisions. They can coexist for one product ecosystem. A BESS, for example, may use component passports for security controllers, release passports for firmware, instance passports for devices,

a system-of-systems passport for the site, a remote-service passport for the energy-management platform, and modification lineage for second-life conversion.

4.6.1 Software-release passport

The software-release pattern is the minimum implementation for browsers, VPN software, SIEM, operating systems, security tools, and other software products. The primary subject is a signed release or build. Core evidence includes product identity, intended use, support period, release manifest, SBOM, build provenance, known-exploitable-vulnerability analysis, update channel, advisories, conformity mapping, and supersession. A package identifier, repository reference, API, or signed manifest can replace a physical data carrier.

A release passport cannot establish the state of a customer deployment. Where deployment assurance is required, the operational pattern adds installed-version and configuration evidence. The distinction prevents a vendor release claim from being interpreted as proof that every customer is patched.

4.6.2 Connected-device and fleet passport

This pattern adds a persistent physical or logical instance, installed release, device or management observation, configuration profile, update receipt, rollback and recovery evidence, and operator relation. It is suitable for routers, gateways, smart-home devices, wearables, industrial controllers, scanners, charging equipment, and other managed fleets. Stronger device binding can use certificates, secure elements, measured boot, or remote attestation where proportionate.

4.6.3 Remote-service passport

A manufacturer-controlled remote service is represented separately where it is necessary for a product function or materially changes cybersecurity, availability, or compatibility. Evidence includes service identity, responsible legal person, release or change state, API or protocol version, product compatibility, support, incidents, vulnerabilities, and service-to-deployment dependencies. The public view should disclose the dependency and support consequences without revealing internal cloud topology.

4.6.4 Industrial system-of-systems passport

A system integrator composes component and release evidence into an approved architecture and deployment. The system passport references controllers, operating systems, network interfaces, security products, applications, sensors, gateways, remote-maintenance services, and safety systems. It adds the system-level risk assessment, integration tests, approved configuration, operator responsibilities, change process, and combined conformity evidence. Component passports remain authoritative for their own subjects.

4.6.5 Component and supplier passport

The component pattern supports software libraries, firmware, hardware modules, network interfaces, security chips, operating-system components, and other supplier inputs. It provides stable component identity, version or revision, supplier, origin, support period, composition or subcomponent evidence, vulnerabilities, advisories, updates, certification, and status. The final

manufacturer evaluates the component in product context and does not outsource final-product responsibility to the supplier credential.

4.6.6 Substantial-modification and second-life passport

This pattern creates a successor product state and, where the legal analysis requires it, a new responsible manufacturer. It records the original product, changed purpose, components, software, configuration, safety and cybersecurity impact, new support commitment, and renewed conformity evidence. It is suited to remanufacturing, customised machinery, second-life batteries, refurbished connected devices, and major system upgrades. Append-only lineage prevents the changed product from inheriting claims beyond their valid scope. Figure 9 compares the evidence subjects required by the principal implementation patterns.




Figure 9: Implementation patterns and required evidence subjects.

Source: Authors' implementation-pattern synthesis.

Table 45: Comparison of CRA-capable DPP implementation patterns

Pattern	Subject coverage	Dominant evidence	Part II profile	Typical domains
Software release	Model, release; optional deployment and service.	SBOM, build, VEX, update, support, conformity mapping.	Core Product Evidence Profile.	Enterprise software, browsers, VPN, OS, SIEM.
Connected device / fleet	Model, release, instance, deployment, service, lineage.	Installed state, configuration, update, attestation, operator events.	Operational Assurance Profile.	Routers, gateways, smart home, wearables, charging.
Remote-service product	Model, release, service, deployment and compatibility.	Service change, API, dependency, incident, support and availability.	Operational Assurance Profile.	Cloud-managed products, SaaS-dependent devices, AI services.
Industrial system of systems	All five subject layers plus integration and safety baseline.	Component graph, architecture, configuration, integration tests, change gate.	High-Assurance Cyber-Safety Profile.	Machinery, robotics, BESS, buildings, OT.

Table 45: Comparison of CRA-capable DPP implementation patterns

Pattern	Subject coverage	Dominant evidence	Part II profile	Typical domains
Component / supplier	Component model/revision, release, supplier and lifecycle.	Composition, provenance, support, vulnerability, update, certification.	Core or federated supplier profile.	Semiconductors, firmware, libraries, network modules.
Modification / second life	Original and successor subjects, instance, deployment and responsibility lineage.	Modification, new purpose, risk, support, reset and conformity evidence.	Operational or High-Assurance Profile.	Refurbishment, remanufacturing, second-life batteries and machines.

Source: Authors' pattern comparison. Profiles are proposed architecture profiles, not CRA product classes.

4.6.7 Public, operator, assessor, and authority views

The implementation pattern does not determine disclosure. A consumer-facing view should answer practical questions: whether the product remains supported, the current supported release, whether an urgent advisory applies, how to update or reset the product, how to report a vulnerability, and how to decommission or remove data. Operator views can add composition, configuration, installed state, VEX, service dependencies, and campaign evidence. Assessors and authorities can receive the full restricted technical evidence under purpose-bound access.

This multi-view model is necessary for software products with no physical DPP carrier and for industrial products where the public cannot receive the detailed architecture. The stable product identifier and resolver connect the views while policies prevent the passport from becoming either an empty public label or an uncontrolled technical-file publication channel.

4.7 Sector implementation and applicability boundaries

4.7.1 Enterprise software, cloud platforms, and cybersecurity products

This sector has the most direct relationship to the 17 ETSI Public Enquiry drafts. Browsers, password managers, antimalware, VPN products, network management systems, SIEM, boot managers, PKI software, operating systems, virtualisation stacks, and firewalls are covered by product-specific drafts [26–42]. The dominant passport subject is the software release. Composition, build provenance, exploitability analysis, update, support, and conformity evidence therefore form the minimum profile.

Operational value increases when the software is deployed as a managed service or on customer infrastructure. The deployment profile can identify the installed release, enabled features, plugins, runtime services, orchestration components, remote dependencies, and configuration. For virtualisation and container stacks, EN 304 635 demonstrates why runtime and orchestration dependencies are part of the evidence graph rather than an optional appendix [41].

Sensitive content requires strict access control. Password vault content, private keys, detailed detection logic, exploit instructions, tenant topology, and raw logs remain outside unrestricted passport data. The DPP exposes claims, capability statements, hashes, and controlled evidence references.

4.7.2 Network and telecommunications equipment

Routers, internet modems, switches, network management systems, network interfaces, VPN products, and general firewalls map directly to the ETSI corpus. The highest operational value lies in fleet evidence: installed firmware, update checking, configuration profile, remote controller dependency, support end, and verified remediation. EN 304 627 directly supports the release-to-installed-version binding [36].

Telecommunications network functions require additional sector mapping. EN 304 642 remained an early work item at the evidence cut-off and was not included in the clause-level corpus. Network operators may also be subject to NIS2 risk-management and incident duties. The CRA-capable DPP can provide product evidence to those operator processes but does not replace network-level risk assessment, service resilience, or NIS2 reporting [25, 48].

4.7.3 Smart homes, connected buildings, toys, and wearables

Smart-home assistants, smart-home security products, internet-connected toys, and specified personal wearables are important CRA product categories and are covered by EN 304 631-634. Their implementation requires a dual audience. Consumers and parents need clear support, update, cloud-dependency, reset, data-removal, and urgent-advisory information. Manufacturers and authorities need restricted composition, vulnerability, logging, harm, recovery, and test evidence [37–40].

Connected-building systems add an integration layer. A camera, access controller, alarm panel, gateway, building-management application, and cloud service may be individually assessed but create new dependencies when integrated. A building-level system passport can record approved topology, local fallback, service dependencies, recovery tests, and maintenance events while retaining component passports and responsibilities.

4.7.4 Industrial machinery, robotics, and operational technology

Connected machinery, industrial controllers, engineering software, robot control, gateways, remote-maintenance products, operating systems, and network interfaces can fall within CRA scope. Machinery that is also a product with digital elements must meet both the CRA and Machinery Regulation requirements [1], recital 53; [47]. This makes the industrial system-of-systems and High-Assurance Cyber-Safety profiles the relevant patterns.

The system integrator needs evidence for component compatibility, approved firmware, network architecture, remote access, safety baseline, change control, and operator configuration. The product evidence graph can reduce ambiguity between an OEM release, an integrator configuration, and the state installed at a plant. It also supports maintenance and substantial-modification analysis when a controller, safety function, or remote service is replaced.

Standards scope must remain explicit. EN 304 636 excludes industrial operational-technology products and points to prEN 50770-1. A passport must not cite the general firewall draft as evidence for an OT appliance outside its scope [42]. Sector-specific industrial standards and assessment methods must be recorded as separate profiles.

4.7.5 Energy systems, batteries, charging, and grid equipment

Potential CRA-relevant products include smart meter gateways, grid and plant gateways, connected inverters, battery-management systems, energy-management systems, charging

equipment, heat-pump controllers, flexibility-control platforms, and remote monitoring services. Smart meter gateways and certain secure cryptoprocessing devices are critical CRA product categories [1, Annex IV]. Long asset lives, mixed supplier stacks, remote control, and high availability make support and installed-state evidence material.

The Battery Regulation requires a battery passport from 18 February 2027 for each LMT battery, each electric-vehicle battery, and each industrial battery with a capacity greater than 2 kWh. The passport contains information concerning the battery model and information specific to the individual battery [49, Art. 77]. It can share identifiers, lifecycle events, and access mechanisms with CRA evidence, but it does not by itself demonstrate the cybersecurity of the battery-management system, inverter, energy-management system, gateway, or remote service.

A BESS therefore requires a hierarchy. Pack, BMS release, inverter, container or system, site deployment, gateway, energy-management service, and remote-maintenance service remain distinct subjects. Swapping a pack, replacing a controller, changing grid-service configuration, or converting a battery to second life creates a new system state and can trigger cybersecurity, safety, and conformity review. A container- or system-level DPP should reference component passports rather than erase pack identity and lineage.

4.7.6 Automotive and mobility ecosystems

Products covered by the motor-vehicle type-approval cybersecurity regime under Regulation (EU) 2019/2144 are excluded from the CRA to the extent specified by Article 2. This exclusion must be recorded explicitly and must not be inferred for every product used in an automotive context [1, Art. 2].

EV charging equipment, general fleet software, workshop and diagnostic tools, aftermarket devices, connected accessories, energy equipment, and enterprise ICT may remain within CRA scope depending on their own product function and market placement. Battery-passport duties also continue for qualifying EV batteries. A mobility evidence architecture therefore needs a regulatory-scope node for each product and must keep vehicle type-approval, battery, CRA, and operator evidence separate but linkable.

4.7.7 Healthcare, medical, and pharmaceutical operations

Medical devices and in vitro diagnostic medical devices covered by Regulations (EU) 2017/745 and 2017/746 are excluded from the CRA under the specified conditions. General-purpose operating systems, network products, identity systems, gateways, security software, and building or logistics equipment used in healthcare can still fall within CRA scope. A hospital deployment should therefore avoid treating the sector of use as the legal classification of every product [1, Art. 2].

In pharmaceutical manufacturing and distribution, CRA-relevant products can include packaging and inspection machinery, warehouse scanners, cold-chain sensors, industrial gateways, remote-maintenance software, operating systems, and network-security products. The DPP can separate regulated process or product-quality evidence from the CRA evidence of the digital equipment. Shared identifiers can support incident analysis where a cyber event threatens product integrity, cold-chain continuity, or manufacturing availability.

4.7.8 Semiconductors, secure elements, and security components

The CRA important and critical categories include security-related microprocessors, microcontrollers, ASICs and FPGAs, tamper-resistant processors and controllers, hardware security devices, secure cryptoprocessing products, smartcards, and secure elements. Commission Implementing Regulation (EU) 2025/2392 further defines the technical descriptions of important and critical categories [50]. Not all categories were represented by a final ETSI Public Enquiry draft at the evidence cut-off.

A component passport can record hardware revision, firmware, supplier and manufacturing identity, security functionality, supported cryptographic primitives, vulnerability and advisory state, update capability, certification, support period, and lifecycle status. The final manufacturer can reference this evidence in its product risk assessment and composition graph. Device authenticity and hardware-rooted binding can also strengthen instance evidence, but the presence of a secure element does not prove the security of the complete product.

4.7.9 AI, edge AI, and physical AI

An AI-enabled product can combine an operating system, container runtime, model-serving software, drivers, accelerator firmware, network interfaces, tools, remote APIs, and a connected physical device. Each subject can change independently. The architecture therefore keeps the AI model or model version separate from the product release and deployment. Model cards, AI-system documentation, and AI Act registration serve different purposes and should not be relabelled as CRA evidence.

CRA Article 12 creates a regulatory bridge for products with digital elements that are also high-risk AI systems. Where the CRA conditions are met, the product can be deemed to meet the AI Act Article 15 cybersecurity requirements to the extent covered by the CRA declaration. The risk assessment must also consider AI-specific attack paths such as data poisoning and adversarial manipulation [1], Art. 12; [51], Art. 15.

A CRA-capable DPP can bind the model-serving release, container image, tool connectors, drivers, accelerator firmware, remote model service, update state, and device attestation to the deployed AI function. For physical AI, the safety interface also requires authenticated sensors, secure input/output paths, approved control policies, human hand-back conditions, and evidence that a software or model change remains within the validated operating envelope.

4.7.10 Logistics, food, agriculture, and retail

Connected scanners, warehouse robots, tracking devices, refrigeration controls, processing and packaging machinery, edge gateways, point-of-sale systems, farm automation, and remote-monitoring products can fall within CRA scope according to their own characteristics. These sectors are not represented by one dedicated ETSI vertical draft, so the manufacturer must use the horizontal CRA requirements and the applicable product and sector standards.

The main value lies in fleet and system evidence. An operator can identify affected devices, verify updates, maintain local fallback, and link cybersecurity state to cold-chain, processing, or logistics continuity. The DPP should not absorb transactional traceability or food-safety data unless those data are needed for a defined product-security or assurance decision.

4.7.11 Aerospace, marine, defence, and other scope boundaries

Products certified under the civil-aviation framework and equipment within the Marine Equipment Directive are excluded under CRA Article 2. Products developed or modified exclusively for national-security or defence purposes, and products specifically designed to process classified information, are also outside the general regime under the conditions stated by the CRA. General enterprise ICT, ground systems, logistics products, and non-certified equipment used by these sectors require separate scope analysis [1, Art. 2].

The evidence architecture remains technically useful in excluded sectors, but the claim must change. It can be described as a product evidence or cyber-assurance passport, not as proof of CRA duties that do not apply. Evidence access may also be subject to security classification, export control, or sector-specific authority. The passport must record the applicable legal basis and disclosure restriction.

Table 46A: Sector implementation matrix: software, connected products, industry, and energy

Sector	CRA applicability boundary	Dominant pattern	High-value evidence	Primary decisions
Enterprise software and cybersecurity	Generally CRA in scope where the product definition is met; many important/Class II categories.	Software-release plus deployment and service.	Release SBOM/VEX, build, update, support, conformity mapping.	Vendor, enterprise SecOps, assessor.
Network and telecom equipment	Routers, switches, interfaces, management, VPN and firewalls directly covered; network functions need further sector mapping.	Connected fleet and remote service.	Installed firmware, configuration, controller dependency, patch and support.	Manufacturer and network operator.
Smart home and connected buildings	Important categories for specified assistants and security products; integrated building system requires system analysis.	Connected device and system of systems.	Consumer support/advisory, cloud dependency, fallback, recovery, restricted tests.	Consumer, facility operator, authority.
Toys and wearables	Specified internet-connected toys and non-medical wearables are important categories.	Connected device with consumer view.	Support, update, reset, data removal, harm-sensitive evidence and cloud dependency.	Parent/user, manufacturer, authority.
Industrial machinery and robotics	CRA can apply alongside Machinery Regulation; product and integration scope is case-specific.	High-assurance system of systems.	Component graph, approved configuration, safety baseline, remote access, change gate.	OEM, integrator, plant operator.
Energy and BESS	Digital controllers, gateways and services can be in scope; smart meter gateways are critical category.	System of systems plus battery and component passports.	BMS/inverter/EMS releases, installed state, site configuration, support and second-life lineage.	OEM, integrator, asset operator, grid actor.

Source: Authors' sector mapping based on [1, 26–42, 47, 49].

Table 46B: Sector implementation matrix: mobility, health, components, AI, logistics, and excluded regimes

Sector	CRA applicability boundary	Dominant pattern	High-value evidence	Primary decisions
Automotive and mobility	Type-approved vehicle products under Regulation (EU) 2019/2144 excluded; adjacent chargers, tools, aftermarket and ICT assessed separately.	Scope-aware connected product and battery linkage.	Regulatory scope, charging or software release, battery identity, support and modification.	OEM, charging operator, workshop, fleet.
Healthcare and pharma	Covered medical devices/IVDs excluded; general ICT and manufacturing/logistics equipment can remain in scope.	Release, connected fleet and industrial system.	Scope boundary, updates, network dependencies, cold-chain or production availability.	Provider, manufacturer, operator.
Semiconductors and secure components	Important/critical CRA categories; standards coverage remains incomplete.	Component and supplier passport.	Revision, firmware, security function, certification, vulnerability, lifecycle and authenticity.	Component supplier and final manufacturer.
AI, edge and physical AI	CRA applies to qualifying product with digital elements; Article 12 bridges defined high-risk AI cybersecurity requirements.	Deployment plus service and cyber-safety profile.	Runtime, model-serving release, drivers, tools, remote service, attestation and AI-specific risks.	AI provider, product manufacturer, operator.
Logistics, food, agriculture and retail	General CRA scope based on each connected product; no single vertical ETSI draft.	Fleet and industrial system.	Installed state, gateway/service dependency, update, fallback and continuity evidence.	Equipment maker and operator.
Aerospace, marine and defence	Specified certified aviation, marine and defence/security products excluded; adjacent general ICT analysed separately.	Scope-bound assurance passport.	Legal basis, classification, product release, restricted evidence and authority.	Sector authority, manufacturer, operator.

Source: Authors' sector mapping based on [1, 50, 51].

4.8 Cross-sector scenario evaluation

4.8.1 Scenario A: exploitable library in a virtualisation stack

A supplier identifies an actively exploited vulnerability in a runtime library used by a container execution stack. A static process can correlate the library against vendor SBOMs but may not know which product builds include the vulnerable function, which clusters run those builds, or whether an admission controller, plugin, or mitigation changes exposure. The CRA-capable DPP links the supplier advisory to the release-bound runtime composition required by EN 304 635, records product-specific VEX, and resolves deployed clusters through installed-version and configuration evidence [41].

The resulting decision episode separates candidate clusters, confirmed affected clusters, mitigated clusters, fixed clusters, and unknown clusters. The vendor prepares Article 14 evidence where the legal threshold is met; operators schedule updates according to service criticality; installation and post-change evidence close each deployment. The analytical improvement is decisional specificity. It depends on current cluster inventory, configuration evidence, and supplier participation.

4.8.2 Scenario B: BESS controller update with safety and grid-availability constraints

A BMS or energy-management controller requires an urgent security update. The BESS provides grid services and cannot be taken offline without an approved operational window. The update also changes communication and memory behaviour relevant to protection logic. A static battery passport identifies the battery and sustainability data but cannot establish the controller release, remote service, site configuration, or safety impact.

The CRA-capable system passport resolves affected BMS instances, related packs and containers, inverter and gateway dependencies, remote-service compatibility, and the approved safety baseline. The cyber-safety gate permits a temporary mitigation for a defined period, records compensating controls, schedules the update, verifies package authenticity and installed state, runs protection and grid-interface checks, and issues a new evidence snapshot. The improvement is controlled coordination across manufacturer, integrator, operator, and safety authority. It depends on clear system hierarchy and change authority.

4.8.3 Scenario C: second-life industrial robot cell

A robot cell is removed from one plant, refurbished, fitted with a new controller and gateway, integrated into a different safety environment, and placed on the market for a second use. Original certificates and software records remain available, but they do not cover the new controller, network architecture, operating environment, or remote-maintenance service. A passport that overwrites the original record would obscure responsibility and invalidate traceability.

The modification pattern creates a successor product and deployment state. It retains original OEM evidence, identifies the refurbisher and integrator, records changed components and intended use, supports the substantial-modification analysis, links new SBOM and support commitments, and requires renewed cyber-safety assessment. The improvement is accountable lineage and evidence reuse without inappropriate inheritance. It depends on access to original evidence and a competent legal conclusion about manufacturer responsibility.

Table 47: Cross-sector scenario evaluation

Scenario	Core problem	Evidence resolved	Analytical value mechanism	Critical dependency
A - Virtualisation vulnerability	Release and deployment ambiguity across runtime, plugins and clusters.	Supplier advisory, runtime SBOM, VEX, installed cluster state, configuration and update receipts.	Affected-cluster resolution and verified campaign closure.	Stale inventory or incomplete transitive composition.
B - BESS controller update	Cyber urgency conflicts with safety and grid availability.	BMS/EMS releases, site configuration, remote services, safety baseline, exception and post-checks.	Time-bounded mitigation, authorised change window and co-assurance closure.	Unclear system hierarchy or decision authority.
C - Second-life robot cell	Original evidence no longer covers modified product and deployment.	Original lineage, modifications, new composition, intended use, support, risk and conformity evidence.	Responsibility transition and scoped evidence reuse.	Unavailable OEM evidence or incorrect substantial-modification conclusion.

Source: Authors' scenario evaluation. Products and scenarios are illustrative.

4.9 Findings and implementation priorities for Part IV

4.9.1 Cross-sector design propositions

The scenario and sector analyses support a set of design propositions. They are propositions for evaluation, not empirical findings about deployed markets. Part IV tests their legal accuracy, technical coherence, interoperability, security, scalability, governance requirements, and failure modes.

Table 48: Design propositions carried into Part IV

Proposition	Claim	Part IV evaluation focus
P1 - Subject separation	Operational assurance improves when model, release, instance, deployment and remote-service claims cannot be silently substituted for one another.	Test query correctness and false-assurance cases.
P2 - Temporal evidence	Vulnerability, support and conformity decisions require valid-time, recorded-time, source cut-off, expiry and supersession.	Test retrospective incident and stale-evidence scenarios.
P3 - Authority before reliance	Cryptographic integrity is insufficient without legal-person identity, mandate, role and status.	Test forged, expired and overbroad authority claims.
P4 - Contextual exploitability	SBOM correlation becomes actionable only when combined with product-specific VEX, configuration and installed state.	Test affected-set precision and unknown-state handling.
P5 - Verified installation	Publishing an update is not equivalent to remediating a fleet.	Test delivery, installation, attestation, rollback and closure states.
P6 - Cyber-safety gating	Safety and essential-availability constraints require authorised, time-bounded exceptions and post-change evidence.	Test urgent patch and deferred-update scenarios.
P7 - Federated disclosure	Evidence resolution and selective disclosure can preserve authority and confidentiality better than universal replication.	Test provider failure, access policy and independent verification.
P8 - Conditional business value	Value increases with state volatility, fleet scale, actor distribution, decision criticality, lifecycle duration and evidence reuse.	Test sector patterns and implementation cost assumptions.
P9 - Bounded automation	Automated validation can support but must not replace reserved legal, conformity and safety decisions.	Test policy outputs, human escalation and prohibited claims.

Source: Authors' propositions derived from Parts I-III.

4.9.2 Phased implementation sequence

The architecture should be implemented incrementally. A manufacturer does not need high-assurance device attestation, full EBW integration, and sector-wide federation before it can create value. It does need a stable subject model and governance boundary from the beginning. Otherwise early identifiers and evidence objects become difficult to reconcile later.

Table 49: Phased implementation sequence for CRA-capable DPP capability

Phase	Primary work	Minimum output
Phase 0 - Scope and governance	Define products, legal regimes, subject granularity, actors, evidence owners, disclosure classes and prohibited claims.	Approved scope model and responsibility matrix.
Phase 1 - Core release evidence	Establish model/release identifiers, manufacturer identity, intended use, support, SBOM, advisories, standards status and conformity mapping.	Core Product Evidence Profile.
Phase 2 - Operational state	Add instance, deployment, installed version, configuration, update, rollback, recovery and remote-service relations.	Operational Assurance Profile and fleet queries.

Table 49: Phased implementation sequence for CRA-capable DPP capability

Phase	Primary work	Minimum output
Phase 3 - Cross-company control plane	Add supplier credentials, legal-person verification, mandates, status, selective disclosure and decision receipts.	Accountable federated evidence exchange.
Phase 4 - High assurance and safety	Add attestation, independent assessment, safety-baseline links, dual control, exception gates and offline continuity.	High-Assurance Cyber-Safety Profile.
Phase 5 - Continuous evaluation	Measure evidence coverage, freshness, contradictions, campaign closure, unsupported assets and policy performance.	Auditable improvement loop and Part IV validation data.

Source: Authors' implementation sequence.

4.9.3 Boundary for the governance evaluation

Part III demonstrates where the architecture can create operational value and where that value is conditional. The highest-value cases combine dynamic software state, distributed responsibility, long support periods, large or safety-relevant fleets, and repeated cross-company decisions. The same features also create the strongest governance risks: confidential evidence, uncertain authority, provider dependence, stale observations, inconsistent sector rules, and pressure to replace scoped judgement with a simple compliance label.

Part IV evaluates those risks directly. It examines legal accuracy, interoperability, scalability, security of the evidence infrastructure, transparency and trade-secret protection, long-term continuity, provider dependency, dispute and correction mechanisms, and the limits of automated conformity claims. The evaluation retains the central boundary established in Part I: the CRA-capable DPP is a proposed product evidence architecture. It is not a DPP mandate created by the CRA and is not an autonomous source of legal conformity.

5 Part IV - Evaluation, Governance, and Research Implications

Part IV evaluates the CRA-capable DPP as a legal-technical design artefact. The purpose is not to certify a product, a DPP platform, or a conformity assessment procedure. The purpose is to test whether the architecture developed in Part II and applied in Part III is internally coherent, traceable to the legal and standards corpus, secure against foreseeable failure modes, interoperable across organisations, and governable over the product lifecycle. The evaluation therefore answers Research Question 4: under which legal, technical, organisational, and evidentiary conditions can the proposed architecture support reliable decisions without overstating the legal effect of the underlying evidence?

The evaluation is *ex ante* and analytical. It combines design-science evaluation, legal consistency analysis, architecture threat analysis, conformance-test design, and falsification-oriented scenario reasoning [21–23, 52]. It does not substitute for field trials, formal security proofs, notified-body assessment, market-surveillance practice, or judicial interpretation. Four evaluation lenses are used throughout: legal fidelity, technical assurance, governance durability, and bounded automation. The evidence cut-off is 15 August 2026. At that date, the CRA is binding law, six DPP standards have been cited in the Official Journal for the ESPR requirements they cover, 17 ETSI CRA draft standards were available as Final Draft European Standards under Public Enquiry, and the European Business Wallet remains a legislative proposal [1, 3, 5, 20, 24].

5.1 Evaluation method and validity conditions

5.1.1 Design-science claim ladder

A recurrent weakness in digital-compliance architecture is the compression of distinct claims into one assertion. The fact that an architecture can represent a legal requirement does not establish that it improves operational performance. The fact that evidence can be signed does not establish that it is complete or correct. The fact that a policy engine can return a decision does not give that decision legal effect. This article therefore evaluates the artefact through a claim ladder. Each higher claim requires evidence that cannot be inferred from the lower level.

Representation validity is the lowest level. It asks whether the subject model and evidence graph can express the relevant legal and operational distinctions. Decision-support coherence asks whether those representations can be resolved into bounded, explainable answers. Security and resilience potential asks whether the infrastructure can preserve integrity, availability, confidentiality, provenance, and continuity under realistic threats. Process performance asks whether organisations achieve faster, more accurate, or less costly outcomes. Legal effect asks whether an authorised actor or competent body may rely on the evidence for a specific statutory decision. The present study provides strong analytical support for the first two levels, conditional support for the third, and testable propositions for the fourth. It expressly does not establish the fifth.

Table 50: Design-science claim ladder and strength of evaluation

Claim level	Evaluation question	Evidence available in this study	Assessment
1. Representation validity	Can the architecture express the required subjects, claims, time states, authorities, and lifecycle relations without semantic conflation?	Normative mapping, standards analysis, design requirements DR-01 to DR-22, subject graph, evidence object model.	Analytically supported.
2. Decision-support coherence	Can a verifier derive bounded and explainable answers from the represented evidence?	Resolution sequence, policy model, evidence assurance vector, worked examples, cross-sector scenarios.	Supported within the stated assumptions.
3. Security and resilience potential	Can integrity, authority, confidentiality, status, availability, and continuity be protected under foreseeable threats?	Threat model, control design, provider-neutral continuity model, negative test cases.	Conditionally supported; implementation assurance remains necessary.
4. Operational and economic performance	Does the architecture improve exposure precision, patch closure, audit effort, service delivery, or residual value?	Causal mechanisms and proposed measures, but no field dataset or controlled deployment study.	Plausible and testable; not empirically established.
5. Legal effect and conformity	Does the evidence establish CRA conformity or create a presumption of conformity?	Legal boundary analysis only. Legal effect depends on applicable law, assessment route, OJEU status, and accountable decision-maker.	Not established by the architecture.

Source: Authors' evaluation framework, informed by [21–23, 52].

5.1.2 Hard admissibility gates and gradable assurance

The evaluation distinguishes admissibility from assurance quality. Certain conditions are binary for a given decision. Evidence about the wrong product release is not made useful by a high-

quality signature. A claim issued outside the issuer's authority is not repaired by semantic completeness. A revoked or expired mandate cannot be offset by a favourable vulnerability score. These conditions form hard gates. For a decision episode d , the minimum admissibility condition is represented as:

$$G(d) = L(\text{scope}) \wedge S(\text{subject}) \wedge A(\text{authority}) \wedge I(\text{integrity}) \wedge T(\text{time}) \wedge P(\text{purpose}) \wedge M(\text{minimum})$$

The terms denote legal-scope compatibility, subject match, issuer authority, cryptographic and record integrity, temporal validity and status, lawful purpose and access, and the presence of the minimum evidence required by the applicable policy. Where a mandatory gate fails, the safe result is denial, isolation, or escalation. A favourable average score must never override a failed gate.

After admissibility is established, assurance remains multidimensional. The evidence assurance vector introduced in Part II is retained rather than collapsed into a universal score:

$$Q(d) = \langle \text{provenance, freshness, completeness, consistency,} \\ \text{independence, observability, continuity, residual uncertainty} \rangle$$

A relying party may apply purpose-specific thresholds or weights to this vector for triage. The policy version, weights, thresholds, missing evidence, and residual uncertainty must remain visible in the decision receipt. This model prevents the term "CRA-compliant score" from being used as a substitute for legal analysis or conformity assessment.

Table 51: Hard admissibility gates and gradable evidence-assurance dimensions

Gate or dimension	Question	Failure result	Gradable after gate?
Legal scope	Is the product, actor, claim, and intended decision within the relevant legal regime and assessment route?	Reject the claimed legal effect or refer for legal determination.	No.
Subject match	Does the evidence identify the exact model, release, instance, deployment, or remote service at issue?	Reject or resolve the subject mismatch.	No.
Issuer authority	Was the issuer authorised, competent, and within mandate scope when the claim was issued?	Reject, suspend reliance, or escalate.	No.
Integrity and provenance	Is the artefact intact and cryptographically or procedurally bound to its issuer and subject?	Reject as inadmissible.	No.
Time and status	Was the claim valid at the relevant event time and is it not revoked, withdrawn, or superseded?	Reject current reliance; retain only as historic evidence.	No.
Purpose and access	Is the requested processing lawful, necessary, and permitted for the requester and purpose?	Deny or reduce disclosure.	No.
Minimum evidence	Does the policy-defined minimum set exist, including mandatory negative or exception evidence?	Deny or escalate as incomplete.	No.
Provenance quality	How directly does the evidence originate from an accountable and authoritative source?	Lower assurance or require corroboration.	Yes.
Freshness	How recent is the observation relative to the volatility of the claim?	Shorten validity, flag stale state, or reacquire.	Yes.

Table 51: Hard admissibility gates and gradable evidence-assurance dimensions

Gate or dimension	Question	Failure result	Gradable after gate?
Completeness and consistency	Are required fields and relations present, and are contradictions resolved or exposed?	Reduce assurance or escalate.	Yes.
Independence and observability	Is the claim corroborated by an independent assessor, device observation, or operational receipt?	Reduce assurance where self-declaration is insufficient.	Yes.
Continuity and uncertainty	Can the evidence be verified after provider failure, and what uncertainty remains?	Apply fallback, archive resolution, or explicit uncertainty.	Yes.

Source: Authors' evaluation model. The gate logic operationalises DR-01, DR-02, DR-07 to DR-11, DR-17, DR-20, and DR-21.

5.1.3 Failure-oriented and falsifiable evaluation

The proposed architecture is evaluated by its safe behaviour under adverse conditions, not only by its nominal functionality. It fails its design objective where model-level evidence is silently applied to a different installed release; an expired mandate is accepted because the signature remains valid; stale vulnerability evidence is presented as current; a contradiction is suppressed; provider failure makes previously issued evidence unverifiable; restricted composition data is disclosed without a lawful purpose; or an automated policy emits an unqualified statement of CRA conformity. These are observable failure conditions. They provide a basis for conformance tests and later empirical evaluation.

This failure-oriented method also constrains claims of novelty. The evidence graph is valuable only if it preserves differences that conventional document stores or asset inventories tend to erase. The control plane is valuable only if it can establish accountable authority across organisational boundaries. Continuous conformity is valuable only if new evidence can change the operational decision while preserving the historic record. The evaluation therefore treats ambiguity, stale state, unqualified authority, and hidden uncertainty as first-order defects.

5.2 Legal accuracy and regulatory coherence

5.2.1 CRA-capable does not mean CRA-mandated

The first legal test is terminological. Regulation (EU) 2024/2847 does not require a Digital Product Passport. It requires manufacturers and other economic operators to meet defined cybersecurity, documentation, conformity, support, vulnerability-handling, corrective-action, and reporting duties [1]. A DPP may be required under the ESPR, the Battery Regulation, or future delegated acts for specific product groups. These legal bases remain distinct. The term CRA-capable DPP is therefore an architectural capability label proposed by this article. It means that a DPP-compatible evidence service can represent and resolve information relevant to CRA processes. It does not mean that the CRA mandates that service.

Three statements must remain independent. First, a DPP exists and is registered or accessible in accordance with the applicable product-passport regime. Second, the DPP infrastructure satisfies the technical and governance requirements applicable to its own operation. Third, the product and the manufacturer's processes conform to the CRA. Evidence for the first statement cannot prove the third. Evidence for the second may improve trust in transmission and per-

sistence but cannot establish the technical correctness of the product claim. The architecture therefore prohibits a generic status field that merges these statements.

5.2.2 Presumption of conformity and source-status control

Article 27 CRA links a presumption of conformity to harmonised standards, common specifications, or European cybersecurity certification schemes under the conditions set by the Regulation. A candidate standard, a Final Draft European Standard under Public Enquiry, an adopted EN not yet cited in the Official Journal, and an OJEU-cited harmonised standard have different legal effects. The 17 ETSI vertical drafts analysed in Part I remain under Public Enquiry at the evidence cut-off. They are relevant engineering and assessment inputs, but they do not yet create a CRA presumption of conformity [1,3].

The six DPP standards EN 18216:2026, EN 18219:2026, EN 18220:2026, EN 18221:2026, EN 18222:2026, and EN 18223:2026 have been cited in the Official Journal. Their presumption is limited to the ESPR Articles 10 and 11 requirements covered by those standards [24]. Their citation does not confer a CRA presumption on a product whose evidence is exchanged through a conforming DPP. DR-01 is therefore legally necessary: every requirement-to-evidence relation must preserve the source, edition, clause or requirement identifier, legal status, OJEU reference status, date, and scope of the claimed presumption.

5.2.3 Scope, classification, remote services, and substantial modification

The architecture can organise the facts relevant to legal scope, but it cannot autonomously settle every scope question. Determining whether software is supplied in the course of a commercial activity, whether a remote data processing solution is part of the product, whether a function constitutes the core functionality of an important or critical product category, whether a sectoral exclusion applies, or whether a modification is substantial may require legal interpretation and fact-sensitive technical assessment [1]. The DPP can provide product purpose, contractual responsibility, functional dependency, release lineage, modification records, and technical change evidence. It must not convert those inputs into a legally binding classification unless an actor authorised for that determination issues the claim.

A safe implementation represents scope and classification as versioned, attributable assertions. It records who made the determination, the legal basis, the facts considered, the date, the products and functions covered, the assumptions, and any open issue. A later interpretation or delegated act may supersede the assertion without rewriting the historic state. The same approach applies to substantial modification. A lifecycle event can trigger a mandatory review, but the triggering event and the legal conclusion remain separate objects.

5.2.4 Economic-operator responsibility and evidence reuse

Federated supplier evidence does not transfer the final manufacturer's responsibility for the product placed on the market. Component suppliers, open-source maintainers, laboratories, integrators, importers, distributors, operators, and repairers may issue useful evidence. The manufacturer must still exercise the due diligence and perform the assessment required by its role [1]. A CRA-capable DPP improves attribution by preserving the issuer, subject, mandate, method, and evidence boundary of every claim. It should therefore reduce the common error of treating a supplier certificate or VEX statement as a final-product conformity conclusion.

Hosting or resolving evidence also does not, by itself, make a DPP service provider the man-

ufacturer, conformity assessment body, or source of the substantive claim. The provider is accountable for the security, availability, integrity, and contractual performance of its own service. The legal person named as issuer remains accountable for the claim, subject to applicable law and contract. The architecture must preserve this separation even where the provider supplies user interfaces, validation services, or policy tooling.

5.2.5 Cross-regime evidence without compliance inheritance

Persistent product identifiers and common evidence objects create legitimate opportunities for reuse across the CRA, ESPR, the Battery Regulation, the Machinery Regulation, the AI Act, sectoral safety law, procurement, and private assurance schemes. Reuse should occur at the level of verified facts: product identity, manufacturer identity, component composition, update history, lifecycle events, test results, or support dates. Legal conclusions must preserve their own legal basis, scope, assessment route, validity period, decision-maker, and access policy. Conformity under one regime is inherited by another only where the law explicitly provides that effect.

Table 52: Legal-facing claims and the boundary of architectural proof

Legal-facing claim	What the architecture can establish	What it cannot establish by itself	Reserved or accountable actor
A DPP is registered	A registry record, product identifier, operator identity, version, timestamp, and proof of registration exist.	That the underlying product data are true or that the product conforms to CRA or ESPR requirements.	Economic operator and competent registry authority under the applicable DPP regime.
A DPP package is structurally valid	Schema, identifier, signature, link, and mandatory-field checks pass for a named profile version.	Factual correctness, completeness beyond the profile, or legal sufficiency.	Profile governor and relying party; substantive issuer remains accountable.
An issuer signature is valid	The artefact was signed by the holder of a corresponding key and has not been altered.	Issuer competence, mandate scope, or truth of the signed statement.	Issuer-recognition and mandate authorities; relying party.
An issuer was authorised	A valid role or mandate credential covers the action, subject, purpose, and time.	That the technical assessment was correct or that no conflict of interest exists.	Mandating legal person, accreditation or notification authority where applicable.
An SBOM is bound to a release	Manifest hash, build identifier, release identifier, and provenance evidence form a verifiable relation.	That the SBOM is complete, the build process was uncompromised, or the deployed instance runs the release.	Manufacturer/build authority; independent assessor where required.
Requirement evidence is sufficient	All policy-required artefacts are present, current, attributable, and internally consistent.	The final legal or technical judgement where expert assessment is reserved.	Manufacturer, conformity assessment body, authority, or other competent decision-maker.

Table 52: Legal-facing claims and the boundary of architectural proof

Legal-facing claim	What the architecture can establish	What it cannot establish by itself	Reserved or accountable actor
A standard creates a presumption	The exact edition and OJEU reference can be resolved and related to covered requirements.	That the product applied the standard correctly or satisfies uncovered requirements.	Manufacturer and conformity assessment route under Article 27 CRA.
An EU declaration of conformity is authentic	The declaration, issuer, product scope, version, and integrity can be verified.	That its substantive conclusion is legally correct in every circumstance.	Manufacturer and competent market-surveillance or judicial authority.
A deployed instance has an acceptable cyber state	Installed release, configuration, VEX, update receipt, observation time, and policy result are available.	A permanent or context-free security guarantee.	Operator and manufacturer under the applicable operational policy.
A change is a substantial modification	The change, lineage, affected functions, components, and risk evidence are recorded and a review is triggered.	The binding legal characterisation without an authorised determination.	Responsible economic operator, competent assessor, authority, or court.

Source: Authors' legal-technical boundary analysis based on [1, 4, 5, 24].

5.3 Interoperability and semantic governance

5.3.1 Interoperability layers

Interoperability is not achieved by selecting a shared JSON schema. A CRA-capable DPP crosses at least seven layers: legal scope, organisational governance, semantic meaning, syntactic representation, cryptographic trust, transport and resolution, and operational decision policy. Two systems may exchange the same format while disagreeing about the product subject, the authority of the issuer, the meaning of an exploitability state, the applicable standard edition, or the validity time. Conversely, different serialisations can remain interoperable where they map to a common semantic and trust profile.

The evaluation therefore treats interoperability as a governed relation among legal and technical layers. The upper layers constrain permitted claims and their legal meaning. The middle layers establish issuer authority and evidence semantics. The lower layers provide exchange, resolution, status, archival access, and decision receipts. Corrective feedback must flow in the opposite direction so that disputed, revoked, or superseded evidence changes later decisions. Figure 10 illustrates this layered interoperability model.

Source: Authors' architecture. Legal effect, semantic meaning, cryptographic trust, and operational use are governed at different layers.

5.3.2 Canonical subject identifiers and equivalence

Subject identity is the first interoperability dependency. The architecture requires stable identifiers for the product model, hardware and software release, physical or logical instance, deploy-

Layered interoperability and governance for CRA-capable product evidence

Legal effect, semantic meaning, cryptographic trust, and operational use are governed at different layers.

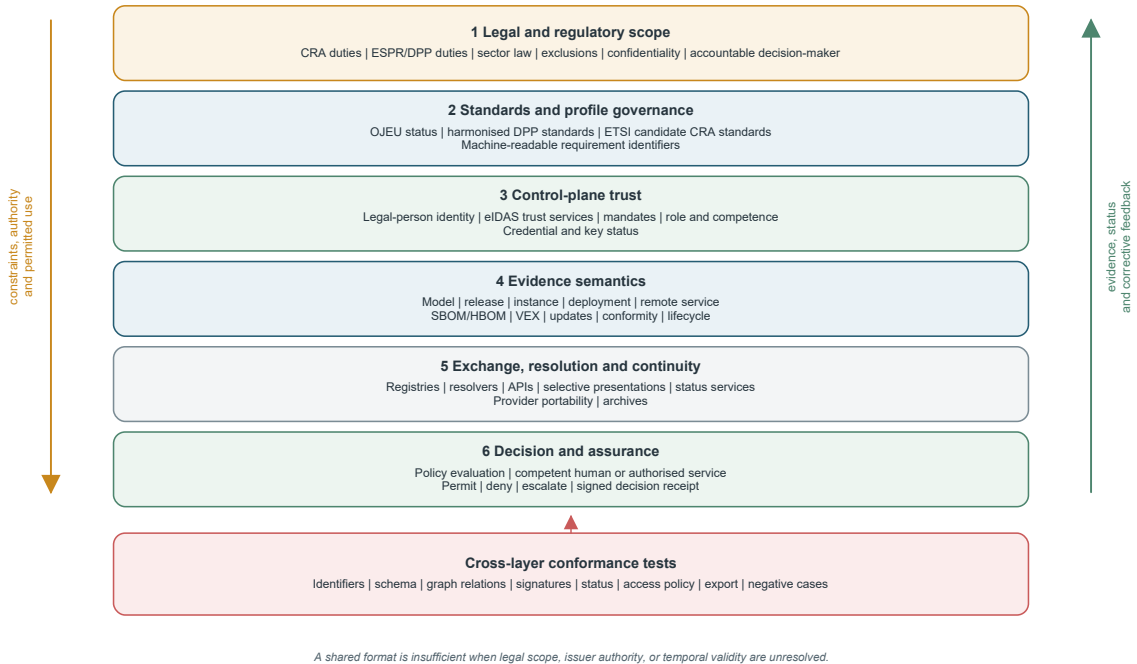


Figure 10: Layered interoperability and governance for CRA-capable product evidence.

ment or configuration state, and remote service. These identifiers may originate from sector schemes, manufacturer namespaces, software package systems, device identities, DPP identifiers, or service registries. Interoperability does not require one universal identifier scheme. It requires explicit typing, governance of the namespace, resolvability or reference data, and controlled equivalence relations.

An equivalence assertion is itself evidence. A mapping between a manufacturer part number and an ESRP product identifier, or between a software build digest and a marketed release, must identify its issuer, scope, method, time, and confidence. Alias resolution must not collapse a product family into one release or a release into one deployed instance. Decentralised identifiers may be used for some actors, services, or devices, but the architecture does not depend on DIDs and does not treat a DID as evidence of legal-person status or authority.

5.3.3 Evidence profiles and format pluralism

The evidence plane should permit format pluralism within governed profiles. SPDX and CycloneDX can represent software composition; CSAF and VEX-compatible profiles can express advisories and product-context exploitability; W3C Verifiable Credentials secured with JOSE or COSE mechanisms can protect claims; Asset Administration Shells and JSON-LD can provide industrial and graph-oriented semantics; and the harmonised DPP standards provide protocol, identifier, carrier, persistence, API, and system-interoperability building blocks [15,16,24,43–45]. None of these formats alone establishes issuer authority or legal effect.

A conformance profile must therefore specify more than a schema. It should define mandatory subject identifiers, claim types, cardinalities, controlled vocabularies, time semantics, status mechanisms, signature and key metadata, issuer-role requirements, minimum evidence links, disclosure rules, validation behaviour, and permitted extension points. Profiles should also

define how unknown, unavailable, disputed, and contradictory states are represented. Interoperability is reduced where absence is silently interpreted as false or where the same label is used for both candidate vulnerability and product-specific exploitability.

5.3.4 Cross-layer conformance tests

Conformance testing should include positive and negative cases. A positive test demonstrates that two implementations exchange and verify a valid evidence package. A negative test demonstrates that they fail safely when a mandatory identifier, authority chain, status record, purpose restriction, or time condition is invalid. Negative cases are critical because permissive parsing and silent fallback can create false trust while producing apparently successful interoperability.

Table 53: Cross-layer conformance-test catalogue

Layer	Test object	Positive test	Negative test	Governance owner
Legal status	Law, standard edition, clause, OJEU citation, and scope metadata.	Covered requirements and legal status resolve to the correct date and jurisdiction.	Draft or uncited standard is rejected as a source of presumption.	Regulatory source-status service and relying-party policy owner.
Identifier	Typed model, release, instance, deployment, and service identifiers.	Exact subject and permitted alias resolve consistently.	Family-level or ambiguous identifier is not accepted for an instance-level decision.	Identifier-scheme owner and profile governor.
Schema	Evidence envelope and mandatory fields.	Valid object passes the named schema and profile version.	Missing mandatory field, unknown critical extension, or invalid cardinality fails.	Profile and schema maintainer.
Semantic graph	Typed relations and controlled vocabularies.	Claim, artefact, requirement, and lifecycle relations preserve intended meaning.	Invalid relation or semantic category error is exposed.	Semantic profile governance body.
Cryptographic integrity	Signature, hash, timestamp, and verification material.	Artefact integrity and signer key are verified for the relevant time.	Altered artefact, unsupported algorithm, or invalid historic key path fails.	Trust-service and cryptographic-profile governance.
Authority and status	Legal-person identity, issuer role, mandate, credential status.	Issuer is recognised and mandated for the claim and subject at issuance time.	Valid signature with expired, revoked, or out-of-scope mandate fails.	Mandating organisation, accreditation or notification authority, trust-list governor.
Access and disclosure	Purpose, requester role, data tier, selective presentation.	Only the minimum permitted claims are disclosed and logged.	Enumeration, overbroad request, or unauthorised onward use is denied.	Data controller, evidence issuer, and ecosystem policy owner.
Portability and continuity	Export package, resolver mapping, status history, validation data.	Independent verifier reproduces the decision after provider migration.	Provider-specific export without schemas or historic validation data fails.	DPP provider, manufacturer, and continuity auditor.
Temporal behaviour	Valid time, transaction time, freshness, supersession.	Historic and current queries return the correct state.	Stale or superseded evidence is not presented as current.	Issuer and status-service operator.
Decision receipt	Policy version, evidence snapshot, result, actor, uncertainty.	Decision can be reproduced or its change explained.	Unversioned policy or missing evidence snapshot blocks high-impact reliance.	Relying-party policy owner and audit function.

Source: Authors' conformance-test design. The tests operationalise DR-01 to DR-03, DR-07 to DR-11, DR-17, DR-19 to DR-21.

5.3.5 Profile versioning and change governance

Profiles will change as CRA guidance, ETSI standards, DPP delegated acts, sector vocabularies, cryptographic recommendations, and implementation experience mature. A profile update must not silently reinterpret old evidence. Every claim and decision receipt should reference the profile and policy version used. Governance should publish release notes, deprecation dates, migration rules, test vectors, semantic crosswalks, and compatibility statements. Breaking changes should create new versions rather than modify the historic meaning of an identifier or status value.

The DPP registry provides important infrastructure for identifiers, registration, semantic resources, and interoperability. It does not by itself govern the complete semantics of CRA evidence, issuer competence, exploitability assessment, cyber-safety exceptions, or relying-party decisions [5]. A coordinated profile-governance function is therefore required. It may be sectoral or cross-sectoral, but it must remain transparent, versioned, contestable, and aligned with authoritative legal sources.

5.4 Scalability, performance, and proportionality

5.4.1 Workload model

The scale problem is not equivalent to the number of DPP records. A single product model may have many releases, millions of deployed instances, multiple configuration profiles, and several remote services. A new vulnerability can generate a burst of resolution, correlation, notification, and update events across the fleet. The operational workload can be represented at a high level as:

$$\Lambda = N(r)\lambda(r) + N(i)\lambda(i) + N(d)\lambda(d) + N(s)\lambda(s) + N(v)\lambda(v)$$

$N(r)$, $N(i)$, $N(d)$, and $N(s)$ denote the numbers of releases, instances, deployments, and remote services, respectively. $N(v)$ denotes vulnerability, incident, update, and status events. The corresponding λ terms represent average query or event rates. The model is intentionally simple. Its main implication is that instance and event workloads can dominate model-level registration. A design optimised only for publishing one passport per marketed model will not support large patch campaigns or high-frequency cloud-service state.

5.4.2 Federation instead of universal replication

The architecture should avoid a universal evidence lake. Authoritative artefacts can remain with the manufacturer, supplier, assessor, operator, or regulated archive. The graph stores typed references, hashes, status metadata, and the minimum information needed for discovery and policy evaluation. Evidence is retrieved or selectively presented when needed. Frequently used public data may be cached. Restricted evidence should be disclosed by policy and purpose. This federated pattern reduces unnecessary duplication and limits the consequences of one repository breach.

Federation introduces its own costs: resolver latency, inconsistent service levels, dependency on status endpoints, and the risk that evidence disappears. The architecture therefore requires

explicit availability classes, signed snapshots for defined decisions, provider-independent validation material, and continuity paths. It is not sufficient to keep a URL in the evidence graph. The system must know who is responsible for the referenced artefact and what happens when that source cannot be reached.

5.4.3 Operational service levels and measures

Scalability should be evaluated against decision-specific service levels rather than aggregate transaction counts. A market-surveillance request, an actively exploited vulnerability, a routine procurement check, and a ten-year archive verification have different latency, availability, and completeness needs. The architecture should expose measurable service characteristics but should not claim universal numerical thresholds without sector evidence.

Table 54: Proposed performance and assurance measures for operational evaluation

Function	Proposed metric	Measurement rule	Risk if missed
Subject resolution	Identifier resolution latency and ambiguity rate.	Measure end-to-end resolution and the share of requests returning more than one unresolved subject.	Evidence may be applied to the wrong product or arrive too late for response.
Evidence retrieval	Required-evidence completeness at decision time.	Compare retrieved objects with the minimum set in the versioned policy.	False assurance, manual rework, or unsafe automatic action.
Verification	Cryptographic and semantic verification throughput.	Measure valid, invalid, and indeterminate results separately under normal and burst load.	Backlogs can convert urgent exposure analysis into stale evidence.
Freshness	Source-to-decision evidence lag.	Measure from authoritative observation or issuance time to policy evaluation.	New vulnerabilities, revocations, or updates are missed.
Status propagation	Revocation and supersession propagation delay.	Measure time until independent verifiers reject or downgrade the prior claim.	Compromised or withdrawn evidence remains relied upon.
Availability	Resolver and status-service availability by evidence class.	Measure primary and alternate paths, including degraded and offline operation.	Evidence becomes a single point of operational failure.
Contradiction control	Unresolved contradiction rate and age.	Count material conflicting claims and time to disposition.	Policy output may hide uncertainty or choose evidence arbitrarily.
Patch closure	Verified affected-instance closure time.	Measure from advisory to verified install, compensating control, retirement, or accepted exception.	Published fixes are mistaken for fleet remediation.
Audit reproducibility	Decision reproduction success rate.	Re-evaluate archived evidence snapshot and policy version with retained validation data.	Historic decisions cannot be defended or corrected.
Archive recoverability	Independent restoration and verification success.	Periodic recovery exercise from export and archive packages without the primary provider.	Long-term legal and operational evidence is lost.

Source: Authors' proposed empirical measurement framework. No numerical service-level objective is asserted without sector-specific validation.

5.4.4 Offline and partitioned operation

Industrial sites, vehicles, remote infrastructure, crisis environments, and regulated networks may be intermittently connected. A verifier should be able to use a signed evidence snapshot, retained trust and status material, and a defined policy for offline operation. The output must state the snapshot time, the maximum permitted age, the unavailable sources, and the residual

uncertainty. Once connectivity returns, the decision should be reconciled against current status and any newly available evidence.

Offline evidence must not be displayed as live evidence. A policy may permit a bounded action under crisis conditions where availability is essential, but the exception must be explicit, time-limited, attributable, and followed by review. In this way, resilience is not confused with permanent bypass of the control plane or evidence freshness requirements.

5.4.5 Proportionality and adoption by smaller economic operators

A high-assurance instance-attestation architecture is not proportionate for every product. The phased profiles defined in Part II remain important. A core profile can provide typed product and release identity, a version-bound SBOM, support information, public advisories, source-status metadata, and evidence portability. An operational profile adds installed state, configuration, update receipts, and remote-service relations. A high-assurance profile adds hardware-rooted evidence, independent assessment, cyber-safety gates, and stronger continuity controls.

Managed services may lower implementation costs for small and medium-sized enterprises, but accountability cannot be outsourced by interface design. The manufacturer should be able to export, inspect, correct, and independently verify its evidence. Contracts should define service levels, subcontractors, security controls, continuity, exit rights, and incident responsibilities. Proportionality concerns the depth and automation of assurance, not the removal of statutory duties.

5.5 Security and resilience of the evidence infrastructure

5.5.1 The DPP evidence service as a protected system

A CRA-capable DPP becomes security-relevant infrastructure. Attackers could manipulate product identity, suppress an advisory, substitute an SBOM, forge an update receipt, revoke a valid mandate, expose confidential architecture, or deny evidence during an incident. Depending on how it is supplied and whether it performs product functions, the DPP software or a connected remote service may itself require a CRA scope assessment. Irrespective of that classification, the service should apply secure development, vulnerability handling, secure updates, access control, tenant isolation, monitored administration, backup, and incident response [1, 53].

The security objective is not merely tamper-evident storage. It includes end-to-end correctness from issuance to decision. A valid signature on a poisoned build manifest remains a false claim. A secure registry can still resolve an identifier to a compromised issuer endpoint. A policy engine can correctly evaluate a maliciously changed policy. Defence must therefore cover evidence production, subject binding, issuer authority, transport, resolution, status, policy administration, user interfaces, automation, and archival validation.

5.5.2 Threats, controls, and residual risk

Table 55: Threat model and required controls for the CRA-capable DPP evidence infrastructure

Threat	Failure mode	Required controls	Residual risk and monitoring
Issuer-key compromise	Attacker issues apparently valid SBOM, VEX, conformity, or update claims.	Hardware-backed or strongly protected keys; scoped keys; rapid suspension; historic validation data; reissuance and compromise notices.	Compromise before detection; monitor anomalous issuance and key use.
Mandate escalation	Employee, service, or agent signs outside its organisational authority.	Machine-readable mandate scope; least privilege; separation of duties; status checks; approval thresholds.	Incorrect mandate issuance; audit cross-role anomalies.
Replay or stale evidence	Old favourable state is reused after vulnerability, revocation, or modification.	Valid time, transaction time, freshness rules, nonce or audience binding where relevant, supersession status.	Offline windows and delayed sources; expose age and uncertainty.
Resolver or registry hijack	Identifier is redirected to an attacker-controlled evidence source.	Authenticated resolution; signed mappings; DNS and transport protection; alternate resolvers; consistency monitoring.	Governance compromise; compare independent resolution paths.
Evidence-graph poisoning	False nodes or relations bind a component, release, or assessment to the wrong subject.	Authoritative relation issuers; schema and graph constraints; provenance; independent corroboration for high-impact edges.	Trusted insider or upstream false statement; contradiction detection.
Equivocation	Issuer presents inconsistent evidence to different relying parties.	Signed transparency receipts, version identifiers, witness or log mechanisms, challenge process.	Confidential claims limit public comparison; monitor disclosed hashes and status.
Schema or semantic confusion	Parser accepts unknown critical fields or maps one state to a different meaning.	Version-pinned profiles; strict validation; controlled vocabularies; critical-extension rules; test vectors.	Implementation defects; continuous interoperability testing.
Policy manipulation	Attacker changes thresholds, weights, permitted actions, or escalation rules.	Signed and versioned policy; change approval; segregation of duties; policy-as-code tests; immutable decision receipt.	Authorised but unsafe change; independent governance review.
Confidentiality or inference attack	SBOM, supplier, fleet, vulnerability, or operational details are exposed.	Purpose limitation; selective disclosure; query controls; encryption; anti-enumeration; aggregation thresholds.	Inference from repeated lawful queries; monitor query patterns.
Denial of service or ransomware	Evidence, status, or policy services are unavailable during an incident.	Federation; offline snapshots; immutable backup; alternate providers; recovery exercises; minimum local controls.	Correlated provider and customer attack; measure degraded operation.
Time-source compromise	False timestamps change validity, order, or freshness decisions.	Trusted timestamping; multiple time sources; monotonic device evidence; tolerance policy; anomaly detection.	Long partition or historic clock error; retain uncertainty.
Cryptographic obsolescence	Algorithm, key length, implementation, or trust anchor becomes unsafe before evidence retention ends.	Algorithm metadata; crypto agility; resealing; archived validation data; transition plans; hybrid methods where justified.	Unknown future attacks and stranded devices; periodic cryptographic review.
Automated-agent abuse	Agent gathers excess evidence or performs an unsafe update or compliance action.	Bounded mandate; tool and data limits; human escalation; action receipts; runtime monitoring; revocable workload identity.	Prompt or tool-chain compromise; monitor anomalous action patterns.

Source: Authors' threat analysis, extending Table 31 and the zero-trust principles in [18].

5.5.3 Zero-trust control and segregation of duties

Zero trust requires every decision to verify identity, device or workload state where relevant, mandate, purpose, policy, evidence status, and context. Membership in the manufacturer's domain, a trusted provider network, or a sector data space is not sufficient [18]. Administrative access to a DPP platform should not confer the authority to issue product evidence. The roles that create a build, approve an SBOM, assess exploitability, approve an update, sign a conformity statement, change a policy, and administer keys should be separated in proportion to risk.

High-impact actions can require dual control or independent approval. Examples include changing the release-to-SBOM binding, withdrawing a critical advisory, accepting a safety-relevant update exception, changing the issuer recognition policy, or issuing a final conformity assertion. The evidence graph should record both the substantive decision and the authorisation chain that permitted it.

5.5.4 Auditability, transparency, and no distributed-ledger dependency

Append-only histories, signed receipts, cryptographic commitments, and witnessed logs can make equivocation and unauthorised change easier to detect. A distributed ledger is one possible implementation for selected shared events. It is not an architectural requirement. Many evidence objects contain confidential or personal data that should not be replicated on a public or consortium ledger. The required property is independently verifiable history with governed correction and status, not a specific consensus technology.

Audit evidence should be proportionate and purposeful. Every read event need not become a permanent cross-sector record. High-impact issuance, status, policy, disclosure, update, and decision events should have attributable receipts. Audit data itself requires retention, access control, integrity protection, and minimisation.

5.5.5 Long-term cryptographic verifiability

CRA evidence may need to remain verifiable for at least ten years after placement on the market or for the support period, whichever is longer. Long-lived industrial products may require substantially longer operational evidence. A signature that is valid today may be unverifiable or insecure at the end of that period if trust anchors, status services, algorithms, or software disappear. Each secured claim should therefore carry algorithm and key identifiers, issuance and validation time, certificate or credential path where applicable, status evidence, and links to archival validation material.

Crypto agility should permit algorithm phase-in and phase-out without changing the semantic identity of the claim. Periodic resealing or preservation signatures can extend archival assurance. Hybrid or multi-hybrid signatures may be justified during post-quantum transitions for high-value and long-lived evidence, but the architecture does not mandate one primitive. Governance must define selection criteria, overlap periods, downgrade resistance, key recovery limits, and evidence for the algorithm transition.

5.5.6 Incident response for the evidence infrastructure

An incident affecting the evidence infrastructure requires its own response playbook. The organisation must be able to suspend an issuer, revoke or qualify claims, disable a resolver,

invoke an alternate status service, notify relying parties, preserve forensic evidence, restore from a known state, and reissue affected records. Recovery should preserve the distinction between claims that were false, claims whose integrity is uncertain, and claims that remain valid but temporarily unavailable.

Evidence-infrastructure incidents can also affect product incident response. A compromised update-status service, for example, may prevent reliable fleet closure. The decision policy must then represent the infrastructure failure as evidence uncertainty rather than infer that no affected product exists. Incident exercises should combine product security, DPP service, identity and key services, conformity evidence, and provider continuity.

5.6 Transparency, confidentiality, data protection, and trade secrets

5.6.1 Lawful transparency and protected evidence

Transparency under the CRA and DPP legislation is purpose-specific. Users need accessible support information, security advisories, update instructions, manufacturer contact data, and conformity information. Market-surveillance authorities may require technical documentation. Operators and integrators may need detailed composition and configuration evidence. The public is not automatically entitled to the complete SBOM, internal risk assessment, attack surface, raw security logs, or confidential supplier data. Article 63 CRA protects confidential information and trade secrets obtained in the application of the Regulation, subject to the applicable legal conditions [1]. The Trade Secrets Directive and the GDPR remain relevant [54, 55].

The architectural consequence is a layered disclosure model rather than a public or private binary. Public evidence should communicate what users need to act. Operator evidence should support secure deployment and maintenance. Assessor and authority evidence may include the complete technical file. Supplier evidence may remain federated and selectively disclosed. Every access path should preserve purpose, requester, scope, retention, onward-use rules, and an auditable decision.

5.6.2 Purpose-bound disclosure test

Table 56: Purpose-bound disclosure test for CRA-capable product evidence

Decision question	Required control	Safe failure state
Is there a legal or contractual basis for the disclosure?	Reference the applicable duty, permission, consent, or contract and its jurisdictional scope.	Deny or seek a valid basis.
Who is requesting the evidence?	Verify legal-person, natural-person, device, or workload identity as appropriate.	Deny anonymous or unresolved high-risk access.
What role and mandate does the requester hold?	Validate role, organisation, competence where relevant, subject scope, purpose, and validity time.	Reduce disclosure or escalate.
What decision purpose is declared?	Bind the request and presentation to a defined purpose and relying party.	Deny broad, reusable, or unspecified collection.
What is the minimum necessary evidence?	Apply data-tier and claim-level minimisation; disclose hashes or derived claims where sufficient.	Return an insufficient-evidence result rather than excess data.

Table 56: Purpose-bound disclosure test for CRA-capable product evidence

Decision question	Required control	Safe failure state
Does the evidence contain trade secrets or security-sensitive information?	Apply sensitivity labels, contractual controls, secure channel, and onward-use restrictions.	Withhold or provide a restricted alternative.
Does the evidence contain personal data?	Apply GDPR role, purpose, minimisation, retention, and rights analysis.	Remove or pseudonymise personal data; deny where no lawful basis exists.
How long may access remain valid?	Use expiry, one-time presentation, time-limited token, or re-authorisation.	Expire access and require a new policy decision.
May the recipient redistribute or train systems on the evidence?	Express machine-readable onward-use and model-training restrictions where lawful and enforceable.	Prohibit onward use or provide an aggregate result.
How is emergency access governed?	Require a defined emergency role, reason, scope, logging, expiry, and post-event review.	Deny unqualified break-glass access.
How is the disclosure audited and contested?	Record requester, purpose, evidence set, policy version, outcome, and correction path.	Treat the disclosure as non-auditable and block high-impact use.

Source: Authors' access-control evaluation, informed by [1, 4, 54, 55].

5.6.3 Selective disclosure and proof techniques

Verifiable presentations can disclose a subset of signed claims, and some cryptographic techniques can prove a predicate without disclosing all source attributes [15, 16]. A supplier might prove that a component has a valid support date beyond a threshold without revealing a full contract. An operator might prove that an installed release is within an approved range without disclosing its entire configuration. These mechanisms can reduce disclosure, but they do not remove the need for semantic and policy governance.

Zero-knowledge proofs are not a universal solution. They can increase implementation complexity, constrain later audit, and obscure the source evidence where the predicate is poorly defined. A proof that a numeric score exceeds a threshold is only as meaningful as the score model. Selective disclosure should therefore be used for narrow, well-governed predicates. High-impact conformity or safety decisions may still require access to the underlying technical evidence by an authorised assessor or authority.

5.6.4 Personal data and operational telemetry

Product evidence can become personal data where device identifiers, operator actions, location, maintenance history, or usage patterns relate to an identifiable person. The control plane should prefer organisational roles and workload identifiers where a named individual is not necessary. Human identity should be disclosed only to the extent required for accountability, competence, approval, or legal process. Device identifiers require contextual analysis rather than an assumption that they are always non-personal.

Raw operational telemetry should not be placed in a general passport. The DPP can store an observation result, signed attestation, test outcome, or hash and controlled pointer to the source record. This reduces privacy and security exposure while preserving verifiability. Retention rules should distinguish long-term conformity evidence from short-lived operational logs.

5.6.5 Inference and aggregation risk

Even correctly authorised individual disclosures can create risk when combined. Repeated SBOM queries may reveal supplier strategy. Instance-resolution queries may expose fleet size or location. Vulnerability-status requests can identify an unpatched population. The access-control system should therefore consider query rate, aggregation, enumeration, correlated purposes, and unusual access patterns. Controls can include throttling, aggregation thresholds, delayed public release, query budgets, secure data rooms, and independent oversight for sensitive ecosystem analytics.

5.7 Provider dependency, portability, and long-term continuity

5.7.1 Provider-neutral source of truth

A DPP provider may store, index, resolve, validate, and present evidence. It should not become the only source of the product claim. The accountable issuer signs or otherwise attests the evidence, and the subject identifier remains stable across providers. A manufacturer, authority, or successor provider should be able to verify the evidence without trusting the former provider's database. This is the practical meaning of provider neutrality.

Provider neutrality does not imply that all evidence is public or copied to every participant. It requires portability of identifiers, signed artefacts, schemas, status history, access policies, and validation material. Restricted evidence can remain encrypted and governed during migration. The new provider must not be able to reinterpret the old claims or silently discard negative and superseded states.

5.7.2 Minimum exit and continuity package

Table 57: Minimum exit and continuity package for provider-neutral evidence

Package element	Required content	Why it is needed	Must not contain or imply
Subject identifiers and mappings	Typed model, release, instance, deployment, service identifiers; namespaces; governed aliases.	Preserves subject resolution and prevents provider-specific relabelling.	Unverified equivalence or silent merging of subject layers.
Signed claims and core artefacts	Claims, SBOM/HBOM, VEX, advisories, update records, conformity evidence, hashes, and attachments as permitted.	Allows independent integrity and provenance verification.	Provider-generated claims presented as issuer claims.
Schema and vocabulary package	Profile versions, schemas, controlled vocabularies, semantic contexts, extension definitions, migration notes.	Preserves the meaning of historic evidence.	A current schema applied retroactively without version control.

Table 57: Minimum exit and continuity package for provider-neutral evidence

Package element	Required content	Why it is needed	Must not contain or imply
Issuer and mandate chain	Legal-person identifiers, issuer credentials, mandates, recognition rules, accreditation or notification references.	Allows historic verification of authority and scope.	Private signing keys or broad reusable administrator secrets.
Status and supersession history	Revocation, suspension, withdrawal, correction, challenge, and supersession records with event time.	Prevents invalid claims from appearing current after migration.	Only the latest favourable state.
Historic validation material	Certificates or credentials, trust anchors, timestamps, algorithm metadata, validation receipts, retained status evidence.	Supports verification after keys, providers, and algorithms change.	An assumption that live endpoints will exist for the full retention period.
Access and disclosure policy	Evidence sensitivity, permitted roles and purposes, retention, onward-use rules, emergency access, policy versions.	Preserves confidentiality and lawful access after migration.	Plaintext export of restricted evidence without key and policy migration.
Audit and decision receipts	Material issuance, access, policy, correction, and reliance receipts with evidence snapshots.	Supports accountability, reconstruction, and dispute handling.	Raw personal or security logs beyond the defined purpose.
Storage and replication manifest	Object inventory, locations, checksums, encryption state, backup copies, restore procedure, missing items.	Makes completeness and recoverability testable.	A database dump with unknown semantics or unverifiable completeness.
Key and secret migration plan	Re-encryption, key rotation, escrow or recovery roles where lawful, destruction evidence, continuity of trust.	Prevents data loss and unauthorised access during transfer.	Export of issuer private keys as a normal portability requirement.

Source: Authors' continuity specification, implementing DR-19 and extending the persistence principles in [4, 5].

Figure 11 illustrates the provider-neutral continuity and independent-verification model.

Source: Authors' architecture. The provider stores and resolves evidence; the accountable issuer remains the source of the claim.

5.7.3 Failure scenarios and recovery

Continuity planning should cover provider outage, insolvency, acquisition, contractual termination, manufacturer cessation, domain-name loss, key compromise, schema deprecation, network partition, and loss of a component supplier. Different failures require different recovery paths. A temporary outage may be handled by cached status and signed snapshots. Insolvency requires export, alternate resolution, and possibly an escrow or regulated continuity arrangement. Key compromise requires suspension, historic validation analysis, reissuance, and targeted notification. Manufacturer cessation requires preservation of public support and safety information and a defined responsibility path for remaining obligations.

Recovery exercises should demonstrate that an independent verifier can resolve the subject, verify the issuer and historic mandate, validate the artefact and status at the relevant time, apply the correct schema and policy, and reproduce a prior decision. A successful database import is not sufficient. The test must include loss of the primary provider.




Figure 11: Provider-neutral evidence continuity and independent verification.

5.7.4 Retention beyond support and market entry

The CRA requires technical documentation, the EU declaration of conformity, and user information to remain available for at least ten years after the product is placed on the market or for the support period, whichever is longer. Security updates made available during the support period must remain available for at least ten years after issuance or for the remainder of the support period, whichever is longer [1]. These periods create a minimum continuity requirement for the evidence needed to identify the product, explain its support state, verify relevant updates, and reconstruct conformity evidence.

Operational needs may continue beyond the statutory support period. Industrial machinery, building systems, energy assets, and second-life products may remain in service. The DPP should distinguish legal retention, active support, extended commercial support, archive-only availability, and unsupported operation. Historic evidence must remain verifiable without implying that the manufacturer continues to provide security support.

5.7.5 Provider qualification and concentration risk

Provider selection should assess security engineering, key management, tenant separation, service continuity, subcontractors, data location, interoperability, export quality, incident response, financial stability, and independent audit. Concentration risk arises where many manufacturers depend on one resolver, cloud service, semantic repository, credential-status service, or policy provider. Federation reduces but does not eliminate this risk. Shared dependencies should be identified and exercised under failure.

Registration in a DPP registry, successful API conformance, or use of a qualified trust service establishes only the properties covered by the relevant legal and technical scheme. It does not prove that the provider's entire evidence service is secure, that the product claims are true, or that every continuity requirement is satisfied. Qualification claims must remain scoped.

5.8 Governance and accountability model

5.8.1 Multi-layer governance

The architecture requires governance at several levels. Legal and regulatory governance defines duties, competent authorities, confidentiality, conformity routes, and legal effect. Standards and profile governance defines requirements, identifiers, semantics, test suites, and versions. Trust-infrastructure governance defines recognised issuers, trust services, credentials, status, and cryptographic policy. Sector and ecosystem governance defines roles, evidence profiles, access purposes, service levels, and dispute processes. Enterprise governance defines mandates, segregation of duties, product ownership, risk acceptance, and evidence quality. Runtime governance defines the policy that permits, denies, or escalates a concrete action.

These layers should be connected but not merged. A sector consortium cannot change the legal effect of an OJEU citation. A cryptographic trust service cannot determine whether a vulnerability is exploitable. A manufacturer cannot self-assign the status of a notified body. A policy engine cannot expand its own mandate. The evidence graph must identify which governance layer authorised each assertion and which actor remains accountable.

5.8.2 Governance functions and owners

Table 58: Governance functions, accountable owners, and required outputs

Governance function	Primary owner	Required output	Oversight or control
Legal source and status	EU or national legislator, Commission, competent authority, official publication service.	Authoritative law, implementing act, guidance, OJEU citation, effective date, scope.	Legal review and source-status change control.
Standard and profile	European standards organisation, sector profile body, semantic governance group.	Versioned standard, profile, vocabulary, mappings, test suite, migration policy.	Open review, appeal, interoperability testing, conflict-of-interest controls.
Identifier scheme	DPP registry, sector identifier authority, manufacturer namespace owner.	Typed identifier rules, namespace governance, resolution, alias and deprecation policy.	Uniqueness tests, misuse handling, continuity plan.
Issuer recognition	Regulator, accreditation or notification authority, trust-list or ecosystem governor.	Recognised role, competence or trust status, scope, validity, suspension and revocation.	Independent assessment and status publication.
Organisational mandate	Accountable legal person.	Authorised actor or workload, actions, subjects, purpose, limits, expiry, approval chain.	Least privilege, segregation of duties, periodic review.
Evidence issuance	Manufacturer, supplier, assessor, operator, device, service, or authority within role.	Signed and scoped claim with subject, method, evidence, time, status, and confidentiality.	Quality management, technical review, audit, correction process.
DPP service operation	DPP provider and contracting economic operator.	Secure storage, resolution, access enforcement, portability, continuity, logs, incident response.	Contract, audit, testing, concentration-risk review.
Reliance policy	Relying organisation and accountable process owner.	Versioned minimum evidence, gates, thresholds, actions, exceptions, escalation, receipt.	Risk and legal approval; policy-as-code tests.
Dispute and correction	Issuer, relying party, ecosystem dispute body, competent authority as applicable.	Challenge status, evidence preservation, correction, supersession, notification, appeal.	Due process, non-retaliation, independent escalation.

Table 58: Governance functions, accountable owners, and required outputs

Governance function	Primary owner	Required output	Oversight or control
Cyber-safety change	Manufacturer, operator, safety authority, independent assessor as required.	Impact assessment, approved baseline, restrictions, rollback, compensating controls, closure.	Functional-safety process and change authority.
Incident and vulnerability response	Manufacturer, CSIRT/SOC, ENISA or CSIRT reporting channel, operator.	Affected subjects, timeline, mitigation, reporting package, user communication, closure evidence.	Statutory deadlines, exercises, post-incident review.
Audit and assurance	Internal audit, conformity assessment body, authority, independent assurance provider.	Scope, evidence sample, findings, limitations, corrective action, retained report.	Independence, competence, records, regulatory oversight.

Source: Authors' governance model. Legal roles remain determined by applicable Union and sector law.

5.8.3 eIDAS trust services and the prospective European Business Wallet control plane

The European Digital Identity Framework and eIDAS trust services provide legally recognised mechanisms for electronic signatures, seals, timestamps, registered delivery, website authentication, and electronic attestations of attributes [19]. These mechanisms can support legal-person identity, claim integrity, time evidence, and status. They do not automatically prove technical competence, product responsibility, or mandate scope. Those properties require appropriate attribute issuers, recognition rules, and organisational governance.

The proposed European Business Wallet would provide a more direct control-plane instrument for economic operators, organisational representatives, mandates, documents, and machine interaction [20]. At the evidence cut-off it remains a proposal and must not be treated as applicable law. The architecture is therefore technology-neutral: an EBW may carry or present legal-person and mandate evidence once the legal and technical framework is adopted, while other conformant identity and authorisation mechanisms may provide the same logical functions.

A wallet proves only the claims and trust properties within its scheme. A valid business identity does not establish that the company manufactures the product. A valid employee credential does not establish that the employee may approve a safety-relevant update. A mandate must bind the legal person, actor or workload, action, product subjects, purpose, limits, and validity. The verifier must check the mandate and credential status at decision time.

5.8.4 Correction, contradiction, and dispute

A product evidence system must support disagreement. A supplier and manufacturer may issue different exploitability assessments. An operator observation may conflict with a manufacturer-declared installed version. A laboratory may challenge a test result. An authority may require correction. The graph should preserve competing claims, their provenance, and their status rather than overwrite the less favourable statement. A challenged claim can be marked as disputed while remaining available for audit.

Correction should occur through a new version or superseding claim. The correction record should identify the original claim, reason, authorising actor, affected relying parties, effective

time, and whether prior decisions must be reviewed. Material corrections should trigger notification. Where a decision affects market access, safety, or legal rights, the process should provide escalation and appeal consistent with the applicable governance scheme.

5.8.5 Accountability and liability boundary

The decision receipt supports accountability by recording the evidence snapshot, policy version, actor or workload, outcome, exception, and residual uncertainty. It does not transfer the underlying statutory duty to the platform. The manufacturer remains responsible for manufacturer obligations. An operator remains responsible for its operational decisions. A conformity assessment body remains responsible for the assessment within its notified scope. A DPP provider remains responsible for its service. Contract may allocate operational tasks and liability, but the interface should not obscure legal roles.

The DPP provider is not a conformity assessment body by default. A validation service that checks signatures, schema, required fields, and status should describe its output as structural or evidentiary validation. It should not label that output as a certification unless the provider is legally authorised and has performed the required assessment. This boundary is central to DR-20 and DR-22.

5.9 Limits of automated conformity and proof-graph evidence scoring

5.9.1 Validation ladder

Automation is useful when its output is precisely named. The architecture distinguishes seven validation levels. Syntax validation checks representation. Cryptographic validation checks integrity and signer control. Status and authority validation checks recognition and mandate. Semantic validation checks subject, relation, time, and controlled meaning. Evidence-sufficiency validation checks whether a policy-defined set is present. Technical assessment judges whether controls and evidence satisfy a technical requirement. Legal or safety determination assigns the effect reserved by law, conformity procedure, or safety governance. Success at one level does not imply success at the next.

Table 59: Validation ladder and limits of automated output

Validation level	Degree of automation	Human or competent-body role	Permitted output	Prohibited output
1. Syntax and schema	High for machine-readable profiles.	Defines and governs schema and critical extensions.	Valid or invalid representation for profile version.	Statement that the data are true or sufficient.
2. Cryptographic integrity	High where algorithms and validation data are supported.	Governs trust anchors, algorithms, historic validation, and exceptions.	Signature, hash, timestamp, or proof valid, invalid, or indeterminate.	Statement that the signer was authorised or truthful.
3. Status and authority	High where recognised roles, mandates, and status are machine-readable.	Establishes recognition and mandate rules; handles exceptional cases.	Issuer recognised and action within mandate at a stated time.	Technical correctness or legal conformity.
4. Semantic and temporal match	Medium to high for governed identifiers and vocabularies.	Resolves ambiguous identity, scope, and semantic disputes.	Evidence applies, does not apply, or is indeterminate for the named subject and time.	Assumption that a model claim applies to every instance.

Table 59: Validation ladder and limits of automated output

Validation level	Degree of automation	Human or competent-body role	Permitted output	Prohibited output
5. Evidence sufficiency	High for an explicit versioned policy.	Defines minimum evidence and acceptable uncertainty.	Complete, incomplete, contradictory, stale, or escalated evidence package.	Substantive assessment where judgement is reserved.
6. Technical requirement assessment	Mixed; some tests and rules can be automated.	Competent engineer, laboratory, manufacturer, or assessor interprets risk, test quality, and residual findings.	Requirement-specific assessment with scope, method, limits, and result.	Unqualified legal presumption or CE decision.
7. Legal, conformity, or safety determination	Automation may support but should remain bounded.	Manufacturer, notified body, authority, safety authority, or court as legally assigned.	Scoped decision, declaration, certificate, approval, restriction, or enforcement action.	Self-authorising platform decision outside legal competence.

Source: Authors' boundary model, extending Table 40 and 41. The permitted legal effect depends on applicable law and actor competence.

5.9.2 Proof-graph evidence scoring

This article proposes proof-graph evidence scoring as a discipline for evaluating the readiness and assurance of a specific evidence path. The term proof graph does not imply mathematical proof of legal conformity. It denotes a graph in which each material decision can be traced from requirement and context through claims, artefacts, provenance, issuer authority, temporal status, validation results, and policy to an outcome. The score concerns the quality of that path, not the abstract security of the product.

Hard admissibility gates are evaluated first. After the gates pass, the evidence assurance vector may support ranking, triage, or escalation. A purpose-specific scalar may be computed for operational convenience, but its weights, thresholds, profile, and limitations must be disclosed. Different purposes require different models. Freshness may dominate an incident-response decision, while independence and completeness may dominate a conformity review. No scalar should be labelled a universal CRA-compliance score.

5.9.3 AI agents and machine-to-machine action

AI agents can resolve evidence, correlate vulnerabilities, prepare Article 14 reporting packages, propose an update campaign, or monitor support obligations. Their authority should be bounded through the control plane. The mandate should identify the legal person, agent or workload, permitted tools, data sources, product scope, actions, financial or safety limits, validity period, and required escalation. The agent must not create or expand its own authority.

For a high-impact action, the evidence plane should provide a signed and reproducible decision context. The policy engine may permit, deny, or escalate. The action service should verify the decision receipt before acting and issue an action receipt afterwards. Human approval remains necessary where legal, conformity, safety, or organisational policy reserves the decision. This separates technical authentication from authority and authority from sufficient evidence.

5.9.4 Contestability and explanation

An automated result must be contestable. The relying party should be able to inspect the subject resolved, evidence snapshot, issuer and mandate chain, policy version, failed or passed gates, assurance dimensions, contradictions, missing evidence, and reason for the outcome. Where evidence cannot be disclosed, the system should provide the minimum explanation consistent with confidentiality and identify a review path. A later correction should identify which prior decisions were affected.

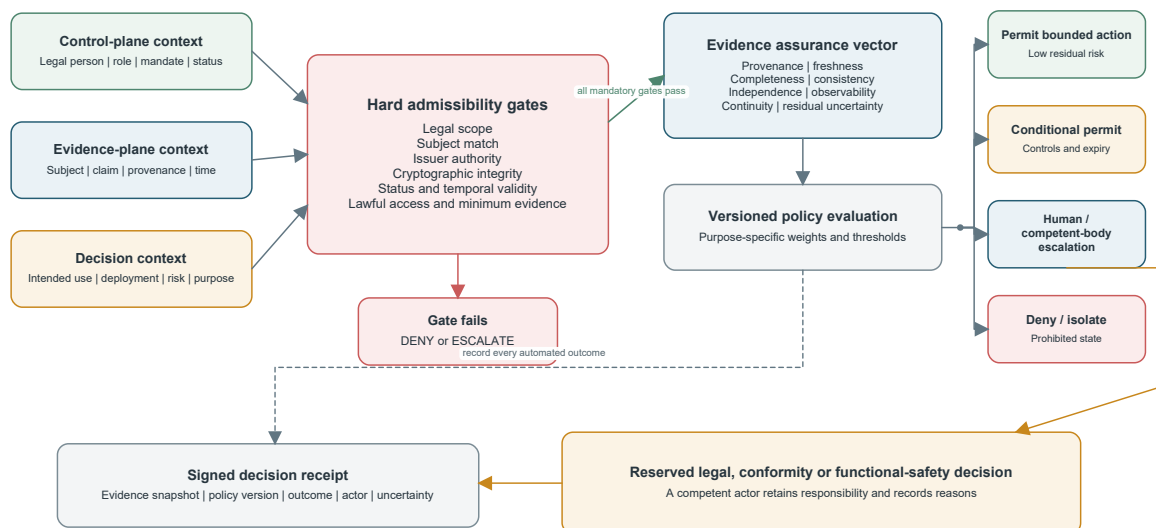
Explanation is not generated prose alone. It is a structured relation between the decision and its evidence. A natural-language summary may aid users, but the authoritative record is the signed decision receipt and linked evidence. This protects against persuasive explanations that are not grounded in the actual policy evaluation.

5.9.5 Automation bias and interface controls

User interfaces should make unknown, stale, disputed, and inapplicable states visible. A green badge derived from a partial SBOM, an old VEX statement, or a structurally valid DPP can create more risk than no badge. The interface should state the subject and time of the result, the decision purpose, the evidence scope, the legal status of standards, and residual uncertainty. High-impact users should be able to inspect the evidence path rather than rely on colour alone. Figure 12 summarises the boundary between automated validation, proof-graph scoring, and reserved human decisions.

Bounded automation and proof-graph evidence scoring

Automation may validate and prioritise evidence. It does not create an unqualified CRA-conformity decision.



No scalar "CRA-compliant" score. Evidence readiness may be ranked only after mandatory gates are satisfied.

Figure 12: Bounded automation and proof-graph evidence scoring.

Source: Authors' architecture. Automation may validate and prioritise evidence but does not create an unqualified CRA-conformity decision.

5.10 Integrated evaluation of propositions and design requirements

5.10.1 Evaluation of the Part III propositions

Part III formulated nine propositions for evaluation. The analytic result is stated below with its strongest support and the empirical test still required. "Supported" denotes coherence with the legal-technical analysis and the worked scenarios. It does not denote statistically validated operational performance.

Table 60: Evaluation of the nine design propositions from Part III

Proposition	Analytic result	Strongest support	Remaining empirical test
P1. Subject separation	Supported.	The five-subject model prevents model, release, instance, deployment, and service evidence from being silently conflated.	Measure affected-instance precision and false decisions against conventional asset and document systems.
P2. Temporal evidence	Supported.	Bitemporal validity, status, supersession, and snapshots preserve historic and current state.	Test status propagation, stale-evidence detection, and audit reconstruction at scale.
P3. Authority before reliance	Supported.	The control plane distinguishes signature validity from legal-person identity, competence, and mandate.	Pilot cross-company mandate issuance, revocation, and automated verification, including prospective EBW integration.
P4. Contextual exploitability	Conditionally supported.	SBOM, dependency path, configuration, VEX, and observation time reduce component-match ambiguity.	Benchmark exploitability classification accuracy and false-positive reduction in real product fleets.
P5. Verified installation	Supported.	Release publication, delivery, installation, observation, rollback, and recovery are separate events.	Measure patch-closure time and the share of apparent closures not supported by instance evidence.
P6. Cyber-safety gating	Conditionally supported.	The architecture binds security change evidence to a sector safety and availability decision process.	Evaluate with certified safety processes and quantify update delay, exception quality, and unsafe change prevention.
P7. Federated disclosure	Conditionally supported.	Purpose-bound access and selective presentation can preserve distributed authoritative sources and confidentiality.	Test interoperability, availability, inference risk, and user burden across independent providers.
P8. Conditional business value	Plausible but unresolved.	Causal mechanisms exist for procurement, maintenance, targeted campaigns, support, and residual value.	Measure costs, avoided losses, adoption incentives, and value distribution across supply-chain roles.
P9. Bounded automation	Supported as a normative safeguard.	The validation ladder, hard gates, decision receipts, and reserved decisions prevent semantic overreach.	Test human-agent interaction, automation bias, escalation quality, and policy error recovery.

Source: Authors' integrated analytic evaluation. Results remain subject to the limitations in Section 5.11.

5.10.2 Evaluation of DR-01 to DR-22

The design requirements derived in Part I remain the primary traceability baseline. The evaluation confirms that all 22 requirements are represented in the architecture. The result is not uniform. Some requirements are directly satisfied by data and trust structures. Others depend

on external governance, operational services, or sector processes that the DPP can support but cannot supply alone.

Table 61: Integrated evaluation of design requirements DR-01 to DR-22

Requirement group	Evaluation result	Principal architecture evidence	Residual issue
DR-01 to DR-03: legal status, subject layers, persistent identifiers	Satisfied at design level.	Source-status object; five-subject model; typed identifiers and mappings.	Authoritative source updates, sector identifier governance, and alias disputes require external governance.
DR-04 to DR-06: composition binding, state separation, remote services	Satisfied at design level.	Release-bound SBOM/HBOM; declared/observed/configured/assessed evidence classes; service subject.	Completeness and trustworthy device or service observation require implementation assurance.
DR-07 to DR-11: legal person, mandate, provenance, status, time	Satisfied logically; governance-dependent.	Control plane; claim envelope; status and supersession; bitemporal model.	Cross-border role recognition, competence, historic validation, and mandate adoption need operational trust frameworks.
DR-12 to DR-16: vulnerability, updates, conformity, support, lifecycle	Satisfied at design level.	VEX state model; update event chain; conformity snapshots; support state; append-only lineage.	Supplier evidence quality, fleet observability, legal classification, and timely update practices remain external.
DR-17 to DR-19: disclosure, cyber-safety gate, continuity	Conditionally satisfied.	Evidence tiers; purpose policy; cyber-safety change gate; exit and continuity package.	Sector safety authority, inference risk, provider service levels, and long-term archival operations need empirical validation.
DR-20 to DR-22: validation layers, explicit reliance policy, no binary compliance flag	Satisfied and reinforced by Part IV.	Validation ladder; hard gates; versioned policy and receipt; prohibited unqualified status.	User-interface discipline and organisational resistance to oversimplified badges require governance and testing.

Source: Authors' traceability evaluation against Table 12A, 12B, and 33.

5.10.3 Falsification cases

The following cases define observable behaviour that would falsify the safety and assurance claims of the architecture. They should be implemented as negative conformance tests and exercised in pilots.

Table 62: Falsification cases for the CRA-capable DPP architecture

Failure test	Expected safe behaviour	Unsafe behaviour that falsifies the design
Wrong-release evidence	Reject or mark inapplicable; require a verified release or instance relation.	Accept model or family evidence as proof for a different installed release.
Expired or revoked mandate	Reject the action or escalate, even where the signature remains cryptographically valid.	Treat signature validity as sufficient authority.
Stale VEX or vulnerability source	Expose cutoff time and freshness; reacquire or downgrade before reliance.	Display "not affected" without time and source scope.
Conflicting assertions	Preserve both claims, identify provenance, apply conflict policy, and escalate where material.	Select the favourable claim silently or overwrite the adverse claim.

Table 62: Falsification cases for the CRA-capable DPP architecture

Failure test	Expected safe behaviour	Unsafe behaviour that falsifies the design
Primary provider outage	Use alternate resolver, signed snapshot, or archive; state degraded assurance and reconcile later.	Return no evidence as proof of no vulnerability or make historic evidence unverifiable.
Safety-relevant emergency update	Apply cyber-safety gate, documented exception, bounded compensating control, and post-event review.	Deploy automatically on cybersecurity severity alone without safety or availability assessment.
Request for confidential SBOM	Verify requester, mandate, purpose, minimum need, and disclosure tier; provide a restricted result where sufficient.	Publish the full composition because a DPP is presumed public.
Automated conformity badge	Return scoped evidence readiness, requirement-specific result, or escalation with legal status visible.	Emit an unqualified "CRA compliant" status from schema, signature, or evidence-score success.
Provider migration	Independent verifier reproduces subject, status, authority, and decision from the continuity package.	Only current records migrate, or verification requires the former provider's proprietary service.

Source: Authors' negative-test specification.

5.10.4 Design-science assessment

The architecture is internally coherent with the normative and standards analysis. It preserves legal status, separates product subjects, distinguishes declared and observed state, binds evidence to issuers and mandates, supports time and supersession, and provides a bounded path from evidence to decision. The cross-sector scenarios show that the same logical model can be instantiated for software releases, connected devices, remote services, industrial systems, and modified or second-life products without asserting that one assurance profile fits every domain.

The evaluation also identifies conditions that the architecture cannot satisfy alone. It depends on accurate issuer evidence, product and fleet observability, recognised roles, robust status services, secure implementation, profile governance, provider continuity, and competent legal and safety decisions. Operational efficacy, cost, regulatory acceptance, and usability therefore remain empirical questions. The principal design result is not a claim that the DPP automates conformity. It is a specification for making the evidence and authority behind a decision explicit, verifiable, and contestable.

5.11 Limitations

5.11.1 Regulatory and standards volatility

The legal and standards environment is changing. The analysis uses an evidence cut-off of 15 August 2026. The 17 ETSI product documents are Final Draft European Standards under Public Enquiry and may change before adoption and OJEU citation. Additional horizontal standards, common specifications, delegated acts, guidance, product-category descriptions, sector rules, or judicial interpretations may alter the mapping. The European Business Wallet remains a proposal. Every implementation must therefore resolve authoritative and current sources rather than embed the paper's cut-off state as permanent law.

5.11.2 Analytical rather than empirical evaluation

The study does not report a controlled field deployment, statistically evaluated operational dataset, economic impact study, or user trial. The operational and business mechanisms are analytically plausible but unquantified. No claim is made about actual false-positive reduction, patch-closure acceleration, audit-cost savings, or provider performance. The proposed metrics and falsification cases define how such claims should be tested.

5.11.3 Scope and sector heterogeneity

The paper analyses the CRA, selected adjacent legislation, 17 ETSI vertical drafts, DPP infrastructure, and illustrative sectors. It is not an exhaustive analysis of every standard requested under M/606, every sector exclusion, national market-surveillance practice, procurement rule, safety standard, certification scheme, or product configuration. Sector-specific legal advice and engineering assessment remain necessary. The architecture is a reference model, not a ready-made conformity profile.

5.11.4 Evidence quality and incentive assumptions

Cryptographic verification cannot correct false source data, an incomplete SBOM, poor testing, or a deliberately narrow VEX assessment. The architecture assumes that accountable actors can be identified and that governance creates incentives to issue, correct, and maintain evidence. Commercial confidentiality, liability concerns, supplier power, and unequal digital capability may reduce evidence availability. These political-economic factors require empirical study and contractual or regulatory responses.

5.11.5 Cryptographic, privacy, and safety abstraction

The paper specifies cryptographic and privacy properties at architecture level. It does not provide formal protocol verification, algorithm profiles, key ceremonies, side-channel analysis, or a complete data-protection impact assessment. It identifies the interface between cybersecurity evidence and functional-safety change control but does not perform a safety assessment under a sector standard. The proposed evidence assurance vector and proof-graph scoring model require calibration and governance before operational use.

5.11.6 Legal character of the article

This article is an academic legal-technical design study. It is not legal advice, a conformity assessment, a harmonised standard, a common specification, or an official interpretation by the European Commission, ENISA, a market-surveillance authority, or a court. The term CRA-capable DPP has no independent legal status.

5.12 Research, standards, and implementation implications

5.12.1 Standards and regulatory implementation

CRA and DPP implementation would benefit from machine-readable requirement identifiers, explicit standard and OJEU status metadata, normative evidence templates, temporal validity rules, and conformance tests for negative states. Product standards should identify the subject

level to which an assessment applies and the evidence required to bind a tested release to a marketed product. DPP profiles should support release, instance, deployment, and remote service relations even where the statutory passport is registered at model, batch, or item level.

Standards should also distinguish structural validation, cryptographic verification, issuer authority, technical assessment, and legal effect. This would reduce the risk that a successful registry check or schema validation is presented as substantive product conformity. Public test suites should include stale status, wrong-subject, out-of-scope mandate, conflicting claim, provider failure, and confidentiality cases.

5.12.2 European trust infrastructure

European trust infrastructure should support legal-person identity, organisational representation, scoped mandates, service and workload identity, revocation, historic validation, and machine-readable recognition. eIDAS trust services provide important building blocks. A future European Business Wallet could make organisational authority portable across B2B and B2G processes, provided that the final framework defines machine-verifiable mandates and does not reduce business identity to a generic company credential.

The control plane should be interoperable with product evidence without becoming product-specific. One legal person may issue evidence for many products, and one product graph may contain claims from many legal persons. Recognition rules should express which actor may issue which claim for which product subject and purpose. Workload credentials and action receipts are required where agents and automated services act on behalf of organisations.

5.12.3 Industrial implementation sequence

Early implementations should prioritise evidence that improves both compliance readiness and operational security: typed product and release identity, version-bound SBOMs, support period, security contact, standard-edition and OJEU status, advisories, VEX, and update-package integrity. The next step is to connect deployed instances, configuration profiles, remote services, and installation receipts. Cross-company control-plane mandates, selective disclosure, cyber-safety gates, and high-assurance attestation can then be added where the risk and value justify them.

Implementation should begin with a bounded decision episode rather than an attempt to digitise the entire technical file. Examples include identifying affected router instances, verifying BESS controller update closure, or assessing a second-life robot-cell modification. The pilot should specify its legal boundary, subjects, evidence minimum, authority model, failure cases, metrics, and exit package before technology selection.

5.12.4 Empirical research agenda

The next research phase should test the architecture across technical, legal, organisational, and economic dimensions. Multi-organisation pilots are required because the claimed value arises at supply-chain and lifecycle boundaries rather than within one database.

Table 63: Empirical research agenda for CRA-capable Digital Product Passports

Research question	Suggested method	Primary measures	Suitable domain
Does subject separation improve exposure precision?	Compare conventional inventory and document workflows with evidence-graph resolution for known vulnerability events.	Precision, recall, unresolved identity rate, time to affected-instance set, analyst effort.	Routers, operating systems, virtualisation, industrial gateways.
Does verified installation improve patch closure?	Prospective patch campaign with release, delivery, installation, observation, rollback, and recovery receipts.	Verified closure time, false closure rate, exception age, recovery success.	Enterprise fleets, smart-home hubs, industrial controllers.
Does the dynamic technical-file index reduce audit effort?	Parallel conformity or market-surveillance exercise using conventional and graph-indexed evidence.	Evidence retrieval time, missing evidence, assessor rework, reproducibility, dispute rate.	CRA important products and machinery integrations.
Can evidence remain portable across providers?	Planned provider migration and destructive recovery exercise without the primary service.	Export completeness, independent verification rate, resolution continuity, historic decision reproduction.	Cross-sector DPP service providers.
Does the cyber-safety gate improve update decisions?	Controlled tabletop and field exercises with cybersecurity urgency, safety impact, availability constraints, and rollback.	Decision time, unsafe action prevention, exception quality, service interruption, post-event closure.	BESS, robotics, building controls, transport infrastructure.
Can proof-graph evidence scoring be calibrated?	Expert-labelled case set and prospective triage study using hard gates and multidimensional assurance.	Calibration, inter-rater agreement, escalation quality, automation bias, error severity.	Vulnerability response and compliance evidence review.
Can organisational mandates support safe automation?	Cross-company pilot using current eIDAS mechanisms and, when available, EBW-compatible credentials and workload identities.	Authorisation success, revocation propagation, least-privilege violations, manual fallback, auditability.	Data spaces, supplier assurance, agentic procurement and maintenance.
Is the architecture proportionate for SMEs?	Longitudinal implementation and cost study across core, operational, and high-assurance profiles.	Implementation cost, recurring effort, provider dependency, compliance readiness, business benefit distribution.	Component suppliers, software SMEs, connected-product manufacturers.
How should long-term crypto agility be governed?	Migration experiment with archived claims, resealing, algorithm transition, compromised-key scenario, and offline verification.	Verification survival, downgrade resistance, migration effort, stranded evidence, recovery time.	Industrial and energy assets with long lifetimes.
How do users understand bounded assurance outputs?	Human-factors study comparing binary badges, evidence vectors, conflict displays, and structured explanations.	Decision accuracy, over-reliance, comprehension, time, trust calibration, escalation behaviour.	Procurement, operators, assessors, market surveillance.

Source: Authors' proposed research programme.

This agenda establishes a validation path from analytic design to deployed evidence infrastructure. The key test is not whether organisations can publish more product data. It is whether independent actors can resolve the correct product subject, verify who had authority to make a claim, determine the evidence state at the relevant time, preserve legitimate confidentiality, survive provider and cryptographic change, and make a bounded decision whose basis can be reproduced and contested.

6 Conclusion

6.1 Synthesis of findings

This article has examined how Digital Product Passport mechanisms can be extended into a CRA-capable product evidence architecture. The central conclusion is that the Cyber Resilience Act does not mandate a Digital Product Passport, yet many of its lifecycle obligations generate information, traceability, and evidence needs that are difficult to satisfy efficiently through static files, isolated portals, or disconnected product records alone [1, 2]. A CRA-capable DPP is therefore best understood not as a new legal artefact, but as a technical and governance architecture through which regulatory evidence can be structured, cryptographically bound, disclosed in a controlled manner, and maintained over time.

The normative analysis showed that the CRA is, in substance, a lifecycle evidence regime. Its obligations reach beyond the pre-market moment of initial conformity and continue through support, vulnerability handling, security updates, incident and vulnerability reporting, user communication, and product change. The analysis also showed that candidate ETSI standards, while still under Public Enquiry at the evidence cut-off of this study, already point toward a richer evidence model than is usually assumed in DPP discussions. In particular, version-specific SBOMs, product-context vulnerability evaluation, update and rollback evidence, component-to-instance traceability, and support-state signalling recur across several draft standards. These requirements do not merely call for more product data; they call for structured and verifiable product evidence.

The design-science contribution of the paper is the definition of the CRA-capable DPP as a distributed, versioned, and access-controlled evidence graph. Its decisive architectural move is the separation of five evidence subjects: product model, release or build, product instance, deployment or configuration state, and remote service. This separation prevents one of the most common compliance errors in digital-product governance: the implicit assumption that a model-level conformity claim automatically describes the security state of every installed instance and all service dependencies. By introducing a dual-plane architecture, the paper further separates organisational authority, mandate, and policy control from product evidence, status, and lifecycle claims. This control-plane/evidence-plane distinction is necessary for accountable machine-readable trust.

The cross-sector evaluation demonstrated that the value of such an architecture is not confined to one product class. It is directly relevant to software-only products, connected devices, network equipment, industrial control environments, energy systems, batteries and battery-management components, smart-home products, wearables, and second-life or substantially modified products. Across these domains, the strongest benefits arise where the security state changes over time, where products depend on upstream or remote components, where vulnerability status must be assessed in context, and where several accountable organisations contribute evidence across the supply chain.

At the same time, the article identified clear boundaries. A CRA-capable DPP cannot create legal conformity by itself. It cannot substitute secure development, vulnerability testing, incident response, formal conformity assessment, or safety engineering. It cannot compensate for incomplete SBOMs, poor source evidence, weak governance, or false issuer claims. Its value lies elsewhere: in binding claims to the correct subjects, preserving evidence provenance and time context, making disclosure role-sensitive, supporting continuous conformity operations, and reducing the gap between legal obligations and operational product reality.

Taken together, these findings support a precise final proposition. For CRA-relevant products, the most promising role of a Digital Product Passport is to function as reusable evidence in-

frastructure: a structured and governable layer that links legal requirements, technical controls, assessed artefacts, operational state, lifecycle changes, and accountable issuers across the product support period. In this sense, the CRA-capable DPP is not a compliance label. It is a persistent assurance substrate for cybersecurity, compliance, functional safety coordination, and product lifecycle governance.

6.2 CRA requirements heat map and final interpretation

Figure 13 synthesises the article's central evaluative insight in a two-dimensional heat map. The vertical axis groups major CRA requirement clusters. The horizontal axis groups the main evidence and governance capabilities of a CRA-capable DPP. The chart does not show whether the CRA requires a DPP, nor does it indicate automatic conformity. It shows where a CRA-capable DPP contributes most strongly as an evidence architecture and where its role is necessarily indirect or limited.




Figure 13: CRA requirements heat map: where a CRA-capable DPP provides the strongest support.

Source: Authors' evaluative synthesis based on CRA requirements, ETSI draft standards, and the reference architecture developed in this article. The heat map shows support intensity as an evidence architecture, not legal sufficiency or presumption of conformity.

The heat map highlights four areas of particularly strong contribution. First, the DPP provides high support for product and manufacturer identification once identity, scope binding, and issuer authority are treated as first-class evidence objects. Secondly, the DPP strongly supports the CRA objective of addressing known exploitable vulnerabilities when version-bound SBOM or HBOM evidence is combined with product-context vulnerability analysis and VEX statements. Thirdly, it contributes strongly to security-update governance, support-state management, and user communication when update status, rollback, recovery, and end-of-support events are represented as versioned lifecycle evidence. Fourthly, it provides very strong support for technical

documentation, conformity traceability, and substantial-modification lineage, because these domains depend heavily on structured provenance, subject separation, and append-only change history.

The heat map also makes visible the limits of DPP support. Requirements such as secure design, secure-by-default implementation, access control, confidentiality, integrity, or incident-preventive technical measures are only partially assisted by a DPP. In these areas, the passport can document assessed configuration baselines, declared controls, test results, and deployed state, but it does not itself implement the control. The architecture is therefore most valuable where CRA duties depend on evidence continuity, cross-organisational accountability, and version-aware traceability. It is less determinative where security quality depends primarily on product-internal engineering properties.

The final implication is therefore clear. A CRA-capable DPP should be pursued as a European product assurance infrastructure, especially where products are software-intensive, connected, remotely dependent, safety-relevant, long-lived, or subject to supply-chain modification. Its strategic value lies in enabling continuous conformity rather than static documentation, machine-readable accountability rather than isolated declarations, and evidence-based lifecycle governance rather than episodic compliance exercises. Future work should now move from design to deployment, including sector pilots, provider-neutral interoperability testing, and empirical measurement of compliance and security effects under real operational conditions [21, 22].

A Requirements-to-evidence matrix and versioned source register

This appendix clarifies the research-control artefacts referred to in Sections 1.7 and 1.9. The requirements-to-evidence matrix is a structured mapping from a legal or standards requirement to the evidence needed to interpret, assess, or operate that requirement. The source register records the legal status and exact version of each source used in the mapping. Together, these artefacts prevent a later draft, harmonised standard, implementing act, or legislative amendment from silently replacing the basis of an earlier finding. The complete clause-level extraction register is a working research dataset; the schema and representative records are provided here to make the method transparent and reproducible.

A.1 Matrix and source-register schema

Table A1: Fields used in the requirements-to-evidence matrix and source register

Field	Required content	Change-control function
Record identifier	Stable identifier for the mapping record.	Preserves references to the record across revisions.
Source category and legal status	Binding Union law, OJEU-cited harmonised standard, draft standard, legislative proposal, guidance, contract, or authors' design.	Prevents sources with different legal effects from being conflated.
Source identifier	Instrument, standard, work item, or document identifier.	Enables unambiguous retrieval of the authoritative source.
Edition and date	Consolidated legal version, standards edition, publication date, and evidence cut-off status.	Permits later editions to supersede, rather than overwrite, earlier mappings.
Article, annex, clause, or requirement ID	Precise normative or assessment reference.	Supports clause-level traceability and partial application.

Field	Required content	Change-control function
OJEU and conformity status	OJEU citation status, covered requirements, common-specification status, or absence of legal presumption.	Prevents a draft or unrelated harmonised standard from being presented as CRA presumption of conformity.
Product scope and exclusions	Product category, intended use, deployment profile, jurisdiction, exclusions, and assumptions.	Constrains the mapping to the context actually assessed.
Requirement summary	Concise interpretation of the duty or assessment expectation.	Provides a readable statement while preserving the exact source reference.
Evidence object	Document, manifest, test result, observation, attestation, status record, or lifecycle event.	Connects the requirement to evidence that can be represented or resolved.
Product subject	Model, release or build, instance, deployment or configuration, remote service, actor, component, or event.	Prevents evidence for one subject from being applied to another.
Issuer and responsible actor	Issuer role, accountable legal person, competence, and mandate.	Preserves claim provenance without transferring the manufacturer's legal responsibility.
Assessment or observation method	Inspection, test, analysis, scan, attestation, operational observation, or review procedure.	Makes evidentiary sufficiency and reproducibility assessable.
Temporal and status metadata	Issuance, observation, validity, review, source cutoff, revocation, withdrawal, and support state.	Prevents stale or superseded evidence from remaining current.
Access and retention class	Public, customer or operator, conformity-assessment, authority, or restricted; retention and export rules.	Separates necessary evidence retention from unrestricted disclosure.
Predecessor, successor, and supersession	Links to corrected, replaced, or later records and the reason for change.	Creates an auditable version history rather than silent replacement.
Verification note	Known limitations, unresolved questions, independent checks, and validation status.	Makes uncertainty and incomplete verification visible.

Source: Authors' methodological schema. The fields operationalise the source-status, evidence, and subject tests defined in Part I.

A.2 Representative versioned records

Table A2 illustrates how the schema is applied. The records are representative rather than exhaustive. A production register would include each relevant legal provision and each applicable standard requirement at clause level.

Table A2: Illustrative requirements-to-evidence records

Record and source status	Requirement and scope	Evidence, subject, and issuer	Versioning and verification
RTE-CRA-0138 Binding Union law: Regulation (EU) 2024/2847, Article 13(8)-(9).	Determine and document the support period; make issued security updates available for at least ten years after issuance or for the remainder of the support period, whichever is longer.	Support-period determination, update-availability record, and end-of-support notice. Primary subjects: product model and release family. Principal issuer: manufacturer.	Versioned on a change to expected use, component support, operating environment, or legal requirements. Public support status may reference a restricted rationale.

Record and source status	Requirement and scope	Evidence, subject, and issuer	Versioning and verification
RTE-ETSI-0627-KEV-01 Candidate harmonised standard: ETSI EN 304 627 V1.0.1, Public Enquiry draft at 15 August 2026.	Selected draft requirement REQ-KEV-1-01: maintain a machine-readable SBOM for each product version and assess the relation between the declared and installed versions.	Release-bound SBOM, SBOM hash, installed-version observation, assessment result, and mismatch finding. Subjects: release and device instance. Issuers: manufacturer and authorised observation source.	No CRA presumption of conformity at the evidence cut-off. A later ETSI edition or OJEU citation creates a successor record; the draft record remains available for reconstruction.

Source: Authors' illustrative application of the schema. The second record reproduces the draft identifier verbatim; its legal status is explicitly limited to the evidence cut-off.

References

- [1] European Union (2024a) Regulation (EU) 2024/2847 of the European Parliament and of the Council of 23 October 2024 on horizontal cybersecurity requirements for products with digital elements (Cyber Resilience Act). Official Journal of the European Union, L 2024/2847, 20 November 2024. ELI: <http://data.europa.eu/eli/reg/2024/2847/oj>.
- [2] European Commission (2025a) Commission Implementing Decision C(2025) 618 final of 3 February 2025 on a standardisation request to the European standardisation organisations as regards products with digital elements in support of Regulation (EU) 2024/2847 (M/606). Brussels: European Commission.
- [3] ETSI (2026r) 'ETSI launches approval process for 17 European Standards supporting the Cyber Resilience Act', press release, 13 August 2026. Available at: <https://www.etsi.org/newsroom/press-releases/etsi-launches-approval-process-for-17-european-standards-supporting-the-cyber-resilience-act/> (Accessed: 15 August 2026).
- [4] European Union (2024b) Regulation (EU) 2024/1781 of the European Parliament and of the Council of 13 June 2024 establishing a framework for the setting of ecodesign requirements for sustainable products (ESPR). Official Journal of the European Union, L 2024/1781, 28 June 2024. ELI: <http://data.europa.eu/eli/reg/2024/1781/oj>.
- [5] European Commission (2026b) Commission Implementing Regulation (EU) 2026/1778 of 16 July 2026 laying down the implementation arrangements for the digital product passport registry set up under Regulation (EU) 2024/1781. Official Journal of the European Union, L 2026/1778, 17 July 2026. ELI: http://data.europa.eu/eli/reg_impl/2026/1778/oj.
- [6] NTIA (2021) Vulnerability-Exploitability eXchange (VEX): An Overview. Washington, DC: National Telecommunications and Information Administration, U.S. Department of Commerce, 27 September 2021.
- [7] Plociennik, C., Pourjafarian, M., Saleh, S., Hagedorn, T., do Carmo Precci Lopes, A., Vogelgesang, M., Baehr, J., Kellerer, B., Jansen, M., Berg, H., Ruskowski, M., Schebek, L. and Ciroth, A. (2022) 'Requirements for a Digital Product Passport to Boost the Circular Economy', in INFORMATIK 2022. Bonn: Gesellschaft für Informatik. doi:10.18420/inf2022__127.
- [8] Zhang, A. and Seuring, S. (2024) 'Digital product passport for sustainable and circular supply chain management: a structured review of use cases', International Journal of Logistics Research and Applications, 27(12), pp. 2513-2540. doi:10.1080/13675567.2024.2374256.
- [9] Hulea, M., Miron, R. and Muresan, V. (2024) 'Digital Product Passport Implementation Based on Multi-Blockchain Approach with Decentralized Identifier Provider', Applied Sciences, 14(11), 4874. doi:10.3390/app14114874.
- [10] De Diego, S. and Gutiérrez-Aguero, I. (2025) 'Decentralized Digital Product Passport Building Blocks for Enhancing Supply Chain Sovereignty and Circular Economy Practices', IEEE Access, 13, pp. 137973-137985. doi:10.1109/ACCESS.2025.3594826.
- [11] Psarommatis, F. and May, G. (2024) 'Digital Product Passport: A Pathway to Circularity and Sustainability in Modern Manufacturing', Sustainability, 16(1), 396. doi:10.3390/su16010396.
- [12] Xia, B., Bi, T., Xing, Z., Lu, Q. and Zhu, L. (2023) 'An Empirical Study on Software Bill of Materials: Where We Stand and the Road Ahead', in 2023

- IEEE/ACM 45th International Conference on Software Engineering (ICSE), pp. 2630-2642. doi:10.1109/ICSE48619.2023.00219.
- [13] Kloeg, B., Ding, A.Y., Pellegrom, S. and Zhauniarovich, Y. (2024) 'Charting the Path to SBOM Adoption: A Business Stakeholder-Centric Approach', in Proceedings of the 19th ACM Asia Conference on Computer and Communications Security, pp. 1770-1783. doi:10.1145/3634737.3637659.
 - [14] W3C (2022) Decentralized Identifiers (DIDs) v1.0. W3C Recommendation, 19 July 2022. Available at: <https://www.w3.org/TR/did-core/> (Accessed: 15 August 2026).
 - [15] W3C (2025a) Verifiable Credentials Data Model v2.0. W3C Recommendation, 15 May 2025. Available at: <https://www.w3.org/TR/vc-data-model-2.0/> (Accessed: 15 August 2026).
 - [16] W3C (2025b) Securing Verifiable Credentials using JOSE and COSE. W3C Recommendation, 15 May 2025. Available at: <https://www.w3.org/TR/vc-jose-cose/> (Accessed: 15 August 2026).
 - [17] W3C (2025c) Bitstring Status List v1.0. W3C Recommendation, 15 May 2025. Available at: <https://www.w3.org/TR/vc-bitstring-status-list/> (Accessed: 15 August 2026).
 - [18] NIST (2020) Rose, S., Borchert, O., Mitchell, S. and Connelly, S., Zero Trust Architecture. NIST Special Publication 800-207. Gaithersburg, MD: National Institute of Standards and Technology. doi:10.6028/NIST.SP.800-207.
 - [19] European Union (2024c) Regulation (EU) 2024/1183 of the European Parliament and of the Council of 11 April 2024 amending Regulation (EU) No 910/2014 as regards establishing the European Digital Identity Framework. Official Journal of the European Union, L 2024/1183, 30 April 2024. ELI: <http://data.europa.eu/eli/reg/2024/1183/oj>.
 - [20] European Commission (2025b) Proposal for a Regulation of the European Parliament and of the Council on the establishment of European Business Wallets, COM(2025) 838 final, 19 November 2025. Brussels: European Commission.
 - [21] Hevner, A.R., March, S.T., Park, J. and Ram, S. (2004) 'Design science in information systems research', MIS Quarterly, 28(1), pp. 75-105.
 - [22] Peffers, K., Tuunanen, T., Rothenberger, M.A. and Chatterjee, S. (2007) 'A Design Science Research Methodology for Information Systems Research', Journal of Management Information Systems, 24(3), pp. 45-77. doi:10.2753/MIS0742-1222240302.
 - [23] Gregor, S. and Hevner, A.R. (2013) 'Positioning and presenting design science research for maximum impact', MIS Quarterly, 37(2), pp. 337-355. doi:10.25300/MISQ/2013/37.2.01.
 - [24] European Commission (2026a) Commission Implementing Decision (EU) 2026/1736 of 14 July 2026 on harmonised standards for digital product passports drafted in support of Regulation (EU) 2024/1781. Official Journal of the European Union, L 2026/1736, 15 July 2026. ELI: http://data.europa.eu/eli/dec_impl/2026/1736/oj.
 - [25] ETSI (2026s) EN 304 642 V0.0.15, Cybersecurity requirements for network functions of telecommunications systems. ETSI work item, early draft status at 31 July 2026. ETSI Work Programme (Accessed: 15 August 2026).
 - [26] ETSI (2026a) EN 304 617 V1.0.0, Cyber Resilience Act: cybersecurity requirements for web browsers. Final Draft European Standard under Public Enquiry, 12 August 2026. Available at: <https://docbox.etsi.org/CYBER/EUSR/Open/> (Accessed: 15 August 2026).

- [27] ETSI (2026b) EN 304 618 V0.2.2, Cyber Resilience Act: cybersecurity requirements for password managers. Final Draft European Standard under Public Enquiry, 22 June 2026. Available at: <https://docbox.etsi.org/CYBER/EUSR/Open/> (Accessed: 15 August 2026).
- [28] ETSI (2026c) EN 304 619 V1.0.0, Cyber Resilience Act: cybersecurity requirements for software that searches for, removes, or quarantines malicious software. Final Draft European Standard under Public Enquiry, 15 July 2026. Available at: <https://docbox.etsi.org/CYBER/EUSR/Open/> (Accessed: 15 August 2026).
- [29] ETSI (2026d) EN 304 620 V1.0.0, Cyber Resilience Act: cybersecurity requirements for virtual private network products. Final Draft European Standard under Public Enquiry, 12 August 2026. Available at: <https://docbox.etsi.org/CYBER/EUSR/Open/> (Accessed: 15 August 2026).
- [30] ETSI (2026e) EN 304 621 V1.0.5, Cyber Resilience Act: cybersecurity requirements for network-management systems. Final Draft European Standard under Public Enquiry, 12 August 2026. Available at: <https://docbox.etsi.org/CYBER/EUSR/Open/> (Accessed: 15 August 2026).
- [31] ETSI (2026f) EN 304 622 V1.0.0, Cyber Resilience Act: cybersecurity requirements for security information and event management systems. Final Draft European Standard under Public Enquiry, 12 August 2026. Available at: <https://docbox.etsi.org/CYBER/EUSR/Open/> (Accessed: 15 August 2026).
- [32] ETSI (2026g) EN 304 623 V0.1.3, Cyber Resilience Act: cybersecurity requirements for boot managers. Final Draft European Standard under Public Enquiry, 17 June 2026. Available at: <https://docbox.etsi.org/CYBER/EUSR/Open/> (Accessed: 15 August 2026).
- [33] ETSI (2026h) EN 304 624 V1.0.0, Cyber Resilience Act: cybersecurity requirements for public-key-infrastructure and digital-certificate-issuance software. Final Draft European Standard under Public Enquiry, 5 August 2026. Available at: <https://docbox.etsi.org/CYBER/EUSR/Open/> (Accessed: 15 August 2026).
- [34] ETSI (2026i) EN 304 625 V1.0.0, Cyber Resilience Act: cybersecurity requirements for physical and virtual network interfaces. Final Draft European Standard under Public Enquiry, 10 August 2026. Available at: <https://docbox.etsi.org/CYBER/EUSR/Open/> (Accessed: 15 August 2026).
- [35] ETSI (2026j) EN 304 626 V1.0.1, Cyber Resilience Act: cybersecurity requirements for operating systems. Final Draft European Standard under Public Enquiry, 12 August 2026. Available at: <https://docbox.etsi.org/CYBER/EUSR/Open/> (Accessed: 15 August 2026).
- [36] ETSI (2026k) EN 304 627 V1.0.1, Cyber Resilience Act: cybersecurity requirements for routers, internet modems, and switches. Final Draft European Standard under Public Enquiry, 6 July 2026. Available at: <https://docbox.etsi.org/CYBER/EUSR/Open/> (Accessed: 15 August 2026).
- [37] ETSI (2026l) EN 304 631 V1.0.0, Cyber Resilience Act: cybersecurity requirements for smart-home general-purpose virtual assistants. Final Draft European Standard under Public Enquiry, 20 July 2026. Available at: <https://docbox.etsi.org/CYBER/EUSR/Open/> (Accessed: 15 August 2026).
- [38] ETSI (2026m) EN 304 632 V1.0.0, Cyber Resilience Act: cybersecurity requirements for smart-home products with security functions. Final Draft European Standard under Public Enquiry, 20 July 2026. Available at: <https://docbox.etsi.org/CYBER/EUSR/Open/> (Accessed: 15 August 2026).

- [39] ETSI (2026n) EN 304 633 V1.0.0, Cyber Resilience Act: cybersecurity requirements for internet-connected toys. Final Draft European Standard under Public Enquiry, 20 July 2026. Available at: <https://docbox.etsi.org/CYBER/EUSR/Open/> (Accessed: 15 August 2026).
- [40] ETSI (2026o) EN 304 634 V1.0.0, Cyber Resilience Act: cybersecurity requirements for personal wearable products. Final Draft European Standard under Public Enquiry, 24 July 2026. Available at: <https://docbox.etsi.org/CYBER/EUSR/Open/> (Accessed: 15 August 2026).
- [41] ETSI (2026p) EN 304 635 V1.0.1, Cyber Resilience Act: cybersecurity requirements for virtualisation and container execution stacks. Final Draft European Standard under Public Enquiry, 24 June 2026. Available at: <https://docbox.etsi.org/CYBER/EUSR/Open/> (Accessed: 15 August 2026).
- [42] ETSI (2026q) EN 304 636 V1.0.0, Cyber Resilience Act: cybersecurity requirements for firewalls, intrusion-detection systems, and intrusion-prevention systems. Final Draft European Standard under Public Enquiry, 9 July 2026. Available at: <https://docbox.etsi.org/CYBER/EUSR/Open/> (Accessed: 15 August 2026).
- [43] SPDX (2024) The System Package Data Exchange Specification, Version 3.0.1. The Linux Foundation. Available at: <https://spdx.github.io/spdx-spec/v3.0.1/> (Accessed: 15 August 2026).
- [44] CycloneDX (2025) CycloneDX Bill of Materials Standard, Version 1.7, released 21 October 2025. OWASP Foundation and Ecma International. Available at: <https://cyclonedx.org/specification/overview/> (Accessed: 15 August 2026).
- [45] OASIS (2026) Hagen, S. and Schmidt, T. (eds), Common Security Advisory Framework Version 2.1. Committee Specification Draft 02, 25 February 2026. OASIS Open. Available at: <https://docs.oasis-open.org/csaf/csaf/v2.1/csd02/csaf-v2.1-csd02.html> (Accessed: 15 August 2026).
- [46] IDTA (2025) Specification of the Asset Administration Shell, Part 1: Metamodel, IDTA-01001, Version 3.1.1. Industrial Digital Twin Association. doi:10.62628/IDTA.01001-3-1-1.
- [47] European Union (2023b) Regulation (EU) 2023/1230 of the European Parliament and of the Council of 14 June 2023 on machinery. Official Journal of the European Union, L 165, 29 June 2023, pp. 1-102. ELI: <http://data.europa.eu/eli/reg/2023/1230/oj>.
- [48] European Union (2022) Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on measures for a high common level of cybersecurity across the Union (NIS 2 Directive). Official Journal of the European Union, L 333, 27 December 2022, pp. 80-152.
- [49] European Union (2023a) Regulation (EU) 2023/1542 of the European Parliament and of the Council of 12 July 2023 concerning batteries and waste batteries. Official Journal of the European Union, L 191, 28 July 2023, pp. 1-117. ELI: <http://data.europa.eu/eli/reg/2023/1542/oj>.
- [50] European Commission (2025c) Commission Implementing Regulation (EU) 2025/2392 of 28 November 2025 on the technical description of the categories of important and critical products with digital elements pursuant to Regulation (EU) 2024/2847. Official Journal of the European Union, L 2025/2392, 1 December 2025. ELI: http://data.europa.eu/eli/reg_impl/2025/2392/oj.

- [51] European Union (2024d) Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union, L 2024/1689, 12 July 2024. ELI: <http://data.europa.eu/eli/reg/2024/1689/oj>.
- [52] Venable, J., Pries-Heje, J. and Baskerville, R. (2016) 'FEDS: a Framework for Evaluation in Design Science Research', *European Journal of Information Systems*, 25(1), pp. 77-89. doi:10.1057/ejis.2014.36.
- [53] NIST (2022) Souppaya, M., Scarfone, K. and Dodson, D., *Secure Software Development Framework (SSDF) Version 1.1: Recommendations for Mitigating the Risk of Software Vulnerabilities*. NIST Special Publication 800-218. Gaithersburg, MD: National Institute of Standards and Technology. doi:10.6028/NIST.SP.800-218.
- [54] European Union (2016a) Directive (EU) 2016/943 of the European Parliament and of the Council of 8 June 2016 on the protection of undisclosed know-how and business information (trade secrets) against their unlawful acquisition, use, and disclosure. Official Journal of the European Union, L 157, 15 June 2016, pp. 1-18.
- [55] European Union (2016b) Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation). Official Journal of the European Union, L 119, 4 May 2016, pp. 1-88.