PACE: From Physical AI Trials to Industrial Operation

Verifiable safety and cyber evidence for faster commissioning, controlled change and new validation services

Executive Brief

Dr Carsten Stöcker

Spherity GmbH

9 September 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.

Executive summary

Physical AI creates value when machines enter productive service and remain fit for their tasks. A capable arm still needs evidence for its installed tool, payload, software, workspace and interaction with people. Suppliers, integrators and operators must establish which evidence applies and what changes require validation. Repeated document searches can delay commissioning, consume engineering capacity and postpone revenue.

Verifiable Safety Assurance Evidence Graphs (VSAEG) offer a common method for connecting that evidence. Component passports contribute identity, design, production and test records. The arm passport combines these records with installation, application and operating evidence. Explicit links connect each claim to its supporting tests, assumptions, configuration and current status. This creates a basis for automating repeatable checks and directing expert effort towards the decisions that require engineering judgement. [1]

PACE — Physical AI Compliance Evidence — is presented here as a best-practice concept and industry leadership initiative for applying the VSAEG across company borders. It connects functional safety and cyber resilience with practical evidence-sharing, validation and monitoring services. This executive brief translates the physical AI compliance architecture into priorities for component manufacturers, machine and robot OEMs, operators, assessors and trust providers.

The commercial case builds upon fewer repeated evidence searches, faster justified return to service and reusable validation services. Gains require measurement against a document-based process. Start with component identity, a verified evidence index and configuration matching, then extend to automated test coordination and lifecycle monitoring.

Best practices: bind evidence to the installed configuration, assign accountable issuers, monitor change and drift, and revalidate affected claims. **Management recommendation:** nominate

technical and commercial owners for one cross-company infrastructure, with evidence-sharing terms, measurable benefits and investment gates.

- **Component and robot suppliers:** contribute one module or arm configuration and the evidence customers repeatedly request.
- **Integrators, operators and assessors:** provide a representative commissioning case, a qualified review method and a measurable comparison process.
- **Service and trust providers:** connect test, identity, credential and status services through portable interfaces with clear responsibilities.

The deployment principle is simple: an authorised task needs an accepted assurance basis and active local protection. The graph supports evidence-based decisions; qualified parties remain responsible for assessment and the local safety controller enforces the validated operating limits.

1 The challenge of Physical AI

1.1 A capable component is only the starting point

Physical AI combines perception and planning software with equipment that moves, grips, lifts or applies force. Relevant applications include machine tending, intralogistics, collaborative assembly, inspection and maintenance. A manufacturer may supply the same arm platform into several of these markets before the eventual tool, operator exposure or future requirements are known. A later care application introduces another assurance boundary, especially where personal support or a medical intended purpose is involved.

The unit of deployment is therefore the installed application. It includes the component set, firmware, AI model, tool, payload, workspace, safeguards and task. Evidence produced for one combination has to be checked before reuse in another. The VSAEG reference model uses seven arm axes and treats motors, gears, brakes, sensors, controllers and mechanical structures as traceable modules. Its assurance method can also accommodate other joint counts and machine types through an explicit configuration model. [1]

1.2 Three changes that expose the business problem

A different gripper or workpiece. An arm moves rounded parts in a guarded cell. The integrator fits a sharper gripper and introduces a heavier workpiece for a second customer. Component identities may remain valid, while collision, retention, stopping and workspace evidence require review. A useful service identifies the affected claims and the justified tests, so the customer receives a decision for the actual application.

Torque-sensor drift after commissioning. A force estimate becomes biased during use. The arm could apply more contact force than expected while its original test report remains unchanged. A signed observation tied to the sensor and baseline can trigger an assessed restriction, maintenance and targeted revalidation. The response depends on the hazard: a suspended load may require controlled holding and braking.

A vulnerability in deployed firmware. A supplier advisory affects software used in several controller versions. The operator needs to identify the installed builds, affected products and

relevant safety assumptions, then assess the update and any resulting revalidation. A software inventory linked to product and configuration records makes this investigation traceable across organisational boundaries.

These are proposed engineering scenarios, rather than reported incidents. They explain why searchable documents alone provide limited support for a changing assurance case: the service also needs the dependencies between evidence, claims and decisions.

1.3 Compliance becomes a lifecycle operation

The Cyber Resilience Act (CRA) adds urgency. Its reporting obligations apply from 11 September 2026; its main product obligations apply from 11 December 2027. Reporting covers actively exploited vulnerabilities and severe security incidents, with an early warning within 24 hours of awareness and a notification within 72 hours. Timely reporting depends on staged findings and a controlled investigation; complete fleet analysis may continue after the early warning. [2,3]

Requirements also change over time. A maintained applicability profile must record the relevant legal version, product scope and transition rules. When a new requirement applies, the service can identify affected claims and prepare a revalidation plan for responsible review. This prevents every regulatory update from becoming a fresh manual search through the entire product history.

2 The VSAEG solution initiative

2.1 From component passports to an application assurance graph

A component Digital Product Passport (DPP) is proposed here as an access-controlled evidence subgraph. It identifies the component and links its capabilities, firmware, calibration, tests, assumptions and responsible issuers. The robot arm passport composes the passports of the installed components and adds assembly, application and runtime evidence. Their linked claims and assurance relationships form the arm-level VSAEG.

Composition preserves each source's identity, scope and status. A brake test remains a laboratory result for a particular brake and configuration. An integration claim explains how that brake contributes to holding a specified load. An application assessment considers the complete hazard and control strategy. System safety depends on these relationships and the supporting argument; adding component certificates cannot establish it by arithmetic.

The passport supplies discovery and controlled access. The graph supplies meaning and dependencies. Verifiable Credentials (VCs) protect the origin and integrity of scoped claims when used with a defined securing mechanism, trusted issuer policy and status checks. Technical acceptance additionally requires valid methods, calibrated equipment, uncertainty information and competent assessment. [4]

2.2 PACE as an industry best-practice concept

PACE puts the VSAEG architecture into practice through shared terminology, governed validation profiles and interoperable evidence services. Testing, evaluation, verification and validation (TEVV) combines digital twins, simulation, bench work and physical tests. Each result records its subject, configuration, method, calibration, uncertainty, acceptance criterion and issuer. Best practices include agreed confidentiality rules, explicit treatment of conflicting claims and

incomplete dependencies, and accountable ownership of evidence updates. A modular robot arm provides the first implementation case; reusable profiles support adoption across other machine modules.

The cross-company infrastructure connects a component supplier, a robot OEM or integrator, an industrial user, a test or assessment organisation and the evidence/trust services. Each contributes records under agreed access rules. One organisation's evidence can then support another's decision while the original issuer remains identifiable. Existing product-lifecycle, engineering, maintenance and cybersecurity systems should connect through adapters, with responsibilities for data quality and updates assigned at each interface.

2.3 Start with a commercially useful foundation

The first stage can provide a verified evidence index for one component family, exact configuration matching and a record of missing, expired or mismatched prerequisites. This already addresses a recurring customer task: finding the correct safety manual, calibration result or firmware baseline before engineering work can proceed. Suppliers can package these records with their products; integrators can reduce evidence collection effort; assessors can receive a more consistent review package.

Later stages add change-impact analysis, test coordination and runtime monitoring. Adoption follows demonstrated customer value and evidence quality. An operational release, a conformity decision and a credential's technical status retain their respective meanings throughout this progression.

Proposed service outcome: "For this installed arm and task, these claims have accepted evidence; these prerequisites remain unresolved; this change affects these tests; this authorised party owns the next decision." This is the management information that turns a graph into a useful industrial service.

3 Technical architecture

3.1 A shared meaning for distributed evidence

The Resource Description Framework (RDF) provides the graph data model; JSON-LD represents linked relationships in an interoperable data format. Open-world semantics allow further facts to be contributed as suppliers, integrators and operators learn more about the product. A missing calibration record remains unknown. A versioned validation profile defines the finite evidence required for a particular deployment decision, and validators check those requirements against a bounded evidence snapshot. The Shapes Constraint Language (SHACL) supports structural checks; additional services verify proofs, time, status and authority. [5, 6]

For managers, this means the evidence model can grow with new tools, tasks and requirements while the release decision still uses explicit acceptance criteria. Shared identifiers and governed terminology are essential: two records labelled "joint torque" may refer to different measurement locations, methods or uncertainty assumptions. The graph must retain these distinctions. Optional compact encodings such as CBOR-LD are a later transport choice, independent of the core business case. [1]

3.2 Zero-trust Edge Wallets connect machines to responsible organisations

A robot-resident Edge Wallet protects device and workload keys and holds narrowly scoped credentials. The operator's organisational identity service provisions its authority: which task the robot may perform, in which area, for which principal and for how long. Access is evaluated for the requested resource and current context, consistent with zero-trust principles. Hardware identity and installation evidence bind the service to the actual arm; a protected key alone identifies only its holder. [1, 7]

The **control plane** manages identity, mandates, access rules, renewal and revocation. The **data plane** signs bounded observations, interaction records and evidence manifests under those mandates. Raw high-rate telemetry stays in controlled repositories. The graph holds signed summaries, trace references and integrity digests. This controls cost and access while preserving a route to the underlying evidence.

The eIDAS 2.0 trust framework connects evidence to recognised trust services and accountable persons. Natural persons sign; legal persons can apply organisational seals. A qualified seal supports presumptions concerning origin and integrity, and a qualified timestamp supports evidence of time. Qualification depends on the applicable service and certificate conditions. It does not follow from using a VC format. [8]

Prospective European Business Wallets provide an organisational interface for identity, representation, signing, sealing and exchange. The proposal remains under legislative consideration. A practical PACE profile can use existing identity and eIDAS services now, with an upgrade path to the final wallet specifications. Robot Edge Wallets remain a proposed technical profile. [9]

For an organisational evidence package, the wallet verifies authority and invokes the permitted seal service over a manifest of secured credentials. Each contributing credential retains its own issuer and proof. This gives the customer both attributable organisational origin and inspectable engineering evidence.

Monitoring covers configuration, safety-function performance, AI behaviour, cybersecurity and operational context. A drift event identifies affected claims and can suspend reliance on a scoped assurance credential. Revocation, mandate withdrawal and device-key compromise have distinct consequences. Accepted revalidation creates a new baseline and credential; historical evidence remains auditable.

3.3 How the architecture connects to the arm

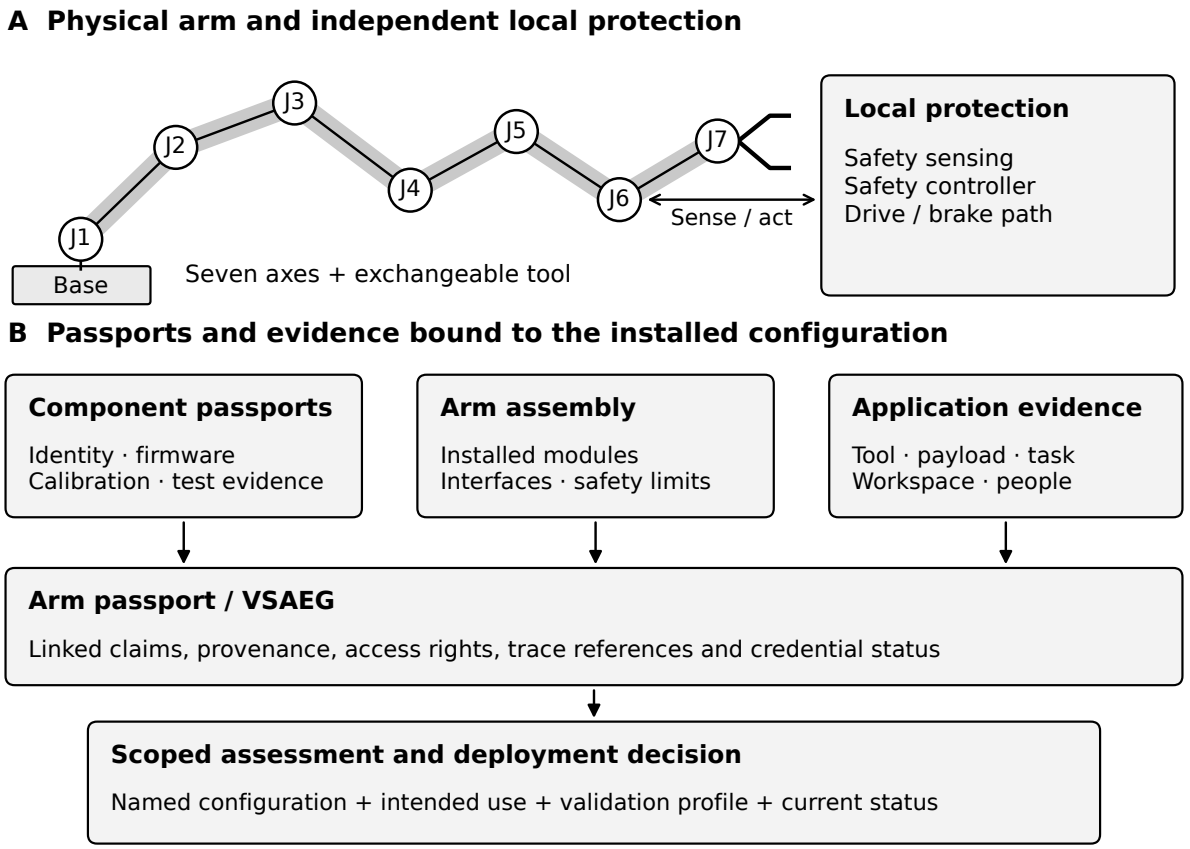

Figure 1: *From the physical arm to a scoped assurance decision. Component passports, arm-assembly records and application evidence contribute to an access-controlled VSAEG. Assessment uses the installed configuration, intended use and current status. Immediate sensing and protective action remain in the local safety path. Adapted from Figure 1 of the VSAEG manuscript; proposed architecture. [1]*

The reference arm contains three shoulder axes, elbow flexion, forearm rotation and two wrist axes. Figure 1 uses an expanded chain to make all seven identifiers visible; its geometry is schematic. Arrows from the component, assembly and application records represent evidence contributions. The lower arrow represents evidence supplied to assessment, rather than a motion command. [1]

Inside-out sensing measures joint position, torque-related signals, temperature and brake state. AI compute proposes motion. An independent local safety path checks the validated limits and enforces the required drive or holding response. Outside-in infrastructure adds scene observations and authenticated requests; any credited safety function needs its own validated sensing, communication and response-time basis.

NVIDIA Halos provides a concrete example of outside-in observation and perception monitoring. Component ecosystems such as Infineon and STMicroelectronics can contribute implementations of the required control, sensing and hardware-trust roles. Product and software selection remains tied to the exact configuration and supporting documentation. These are non-exclusive implementation examples; selection follows the application's evidence and integration requirements. [1, 10]

Graph queries, wallet services and remote status retrieval sit outside the deterministic motion loop. Their outcomes enter through validated supervisory transitions. A cloud outage or revoked mandate therefore invokes a predefined operating policy, while immediate protection remains local.

4 Ecosystem roles and monetisation

4.1 Allocate responsibility where evidence originates

The proposed service depends on complementary expertise. Component suppliers know design assumptions and diagnostic limits. OEMs and integrators know the installed system. Operators observe real use. Laboratories and assessors establish results and conclusions within their competence. Trust providers make origin, authority, access and status verifiable. Table 1 links these roles to potential commercial value. [1]

Table 1: *Ecosystem contributions, potential value and infrastructure deployment measures. Potential benefits require validation against an agreed baseline; assessment responsibilities remain with authorised parties.*

Participant	Evidence or service contribution	Potential value / commercial model	Deployment measure
Component manufacturers	Identity, safety manuals, test and calibration records, firmware lineage	Reusable evidence packs; reduced customer-support effort; lifecycle services	Evidence retrieval time; repeat enquiries
Robot and machine OEMs	Installed baseline, interfaces, system limits and change dependencies	Faster application engineering; supported updates and upgrade packages	Time and effort per validated variant
Integrators and engineering providers	Application risk assessment, configuration binding and validation workflow	Repeatable commissioning and usage-based validation services	Engineering hours; justified tests reused
Industrial operators	Task context, maintenance history, runtime events and deployment decisions	Earlier productive use; targeted maintenance and revalidation	Elapsed return-to-service time
Laboratories and assessors	Calibrated measurements, test methods and scoped independent reviews	Digital test and review services; recurring evidence maintenance	Review effort; invalid evidence detected
Evidence and trust providers	Semantics, credentials, access, mandates, seals, status and portable interfaces	Subscription or usage fees for managed assurance infrastructure	Availability; verification cost; portability
Associations and regulators	Shared expectations, profile feedback, transfer and oversight requirements	Sector adoption and more consistent evidence submissions; public oversight	Cross-company reuse; review usability

Source: management interpretation of the VSAEG architecture. Each contract must preserve the authority and independence of the relevant assessor. [1]

An industry implementation brings together complementary capabilities: component manufacturers such as Infineon and STMicroelectronics; machine and robot OEMs; integrators and operators; assessment organisations such as TÜV bodies; and trust infrastructure providers such

as Spherity. VDMA or comparable associations can help align requirements across machine builders. Collaboration centres on specified deliverables, evidence rights and accountable service roles. These examples illustrate the ecosystem, without implying agreed participation.

4.2 Build a business case around the cost of each change

The economic unit is a justified deployment or change decision. A customer needs to know how much engineering time it consumes, how long the asset waits and how much evidence can be reused. Initial services can charge for onboarding a component family, establishing a configuration baseline or preparing an assessment package. Recurring services can cover evidence freshness, vulnerability-to-configuration mapping, status management and maintenance-triggered revalidation.

Management teams can evaluate usage-based validation fees and contracts linked to evidenced savings. A sustainable model needs a clear payer, a defined service unit, sufficient event volume and a credible comparison baseline. Any savings-based fee should depend on independently observable process improvements. Assessment fees should preserve professional independence and avoid incentives linked to a favourable safety outcome.

Illustrative calculation — assumptions, not a forecast. An operator processes 40 changes per year. If a pilot shows a reduction of 20 engineering hours per change, at a loaded cost of €100 per hour, the annual capacity value is €80,000. With €25,000 in recurring service and governance costs, that leaves €55,000 before onboarding costs and any additional production benefit. These assumptions must be replaced by partner data. Released staff capacity becomes a cash saving only when it changes expenditure; otherwise it is capacity available for other work.

Earlier production creates additional value only where demand, material availability and the rest of the commissioning process support it. Use contribution margin for recovered productive hours, and avoid counting the same delay reduction twice. The full case must include adapters, data remediation, protected hardware, test integration, profile maintenance, cybersecurity, external review and supplier participation costs.

5 Regulatory alignment

VSAEG organises evidence around applicable duties; each regime retains its own scope, assessment route and responsible actors. For industrial robots, ISO 10218-1 addresses the robot and ISO 10218-2 its application and cell. ISO 13849-1 addresses safety-related control-system design. Targets for individual safety functions follow applicable requirements and risk assessment. Service, care and medical uses require separate applicability analysis. [11]

Table 2: *Regulatory and security alignment. The graph organises evidence for the applicable decision; it does not confer conformity, certification or legal authority.*

Framework / scope	Evidence linked in VSAEG	Accountable decision	Management priority
Machinery safety and industrial robotics [11, 12]	Hazards, specified safety functions, integration assumptions and validation	Manufacturer and relevant integrator roles; notified body where the route requires it	Agree the application boundary and validated safety limits
CRA: products with digital elements in scope [2, 3]	Software inventory, vulnerabilities, update history, support and reporting records	Responsible manufacturer; conformity-assessment route as applicable	Connect affected versions to products and corrective action
AI Act: classification by function and use [13]	Intended purpose, model version, evaluation, oversight and monitoring	Provider and deployer duties; applicable conformity route	Assess product/safety-component and sensitive-use criteria
eIDAS and prospective Business Wallets [8, 9]	Identity, representation, mandate, proof, seal, time and status	Authorised principal and issuer; trust-service provider within its role	Select the required legal trust level and access policy
DPP / ESPR product framework [14]	Product identifiers, discovery, access rights and evidence manifests	Responsible economic operator under applicable product rules	Check product-specific duties and protect sensitive records

Sources: official instruments and implementation guidance cited in each row. An assurance credential records a scoped statement; each regulatory regime retains its own decision and documentation requirements.

The Machinery Regulation generally applies from 20 January 2027. The Commission's current AI Act timeline distinguishes high-risk Annex III uses, with rules applying from 2 December 2027, from high-risk AI embedded in Annex I regulated products, from 2 August 2028, following the 2026 amendment. Product and AI-function classification should precede any programme date or assessment commitment. [12, 13]

For machinery, that amendment also provides for relevant AI requirements to be incorporated into the Machinery Regulation through delegated acts. The applicability profile must track this sector-specific route as its implementing requirements are finalised. [13]

The first PACE profile should record the intended purpose, jurisdiction, legal versions, applicable standards and evidence owners. A named owner should maintain it when requirements or guidance change. Publication of a new standard calls for an applicability and impact review; revalidation follows the obligations and engineering changes relevant to the actual system.

5.1 Connect CRA evidence to the product and its safety case

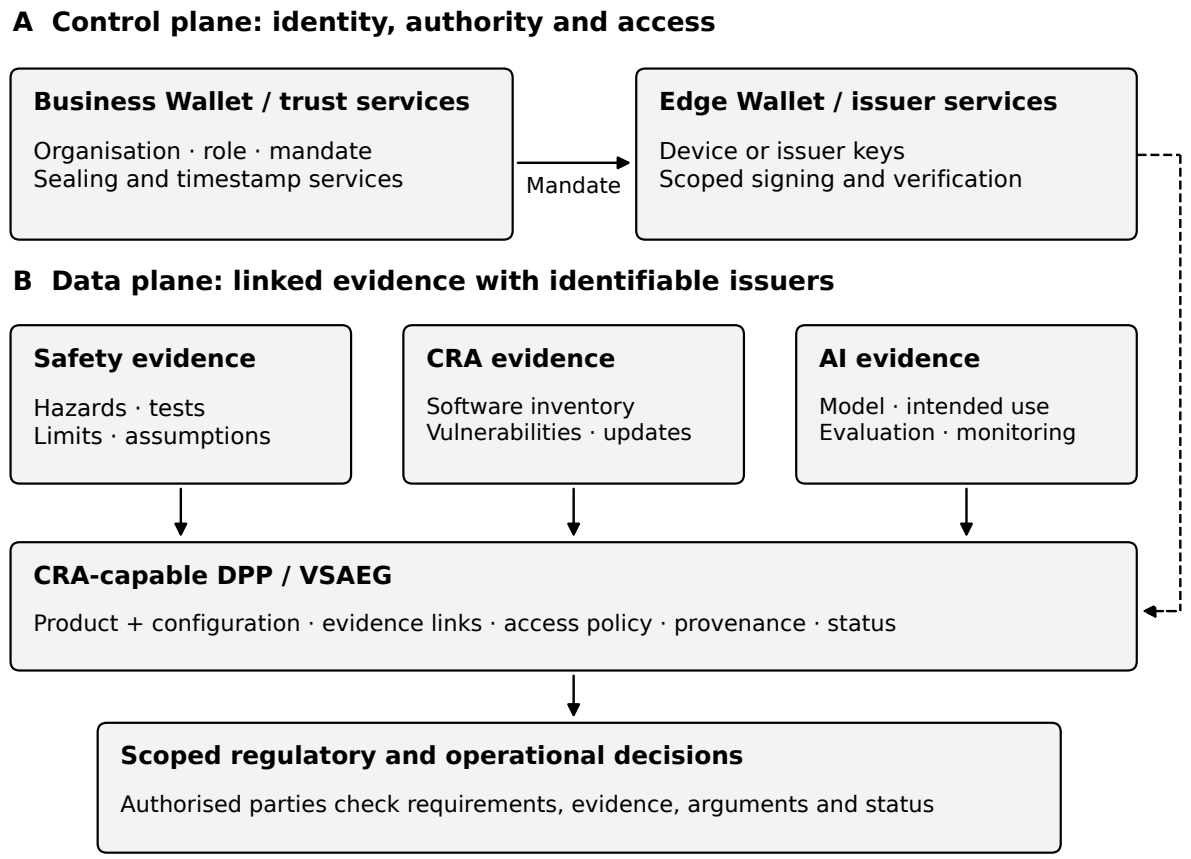

Figure 2: *Shared evidence with distinct regulatory responsibilities. Safety, CRA and AI records are linked through the product and configuration in a CRA-capable DPP/VSAEG. Organisational trust and Edge Wallet services govern authority and access; authorised parties evaluate the evidence under the applicable profile. Adapted from Figure 9 of the VSAEG manuscript; Business Wallet integration is prospective. [1, 9]*

Figure 2 adapts the manuscript's CRA-capable DPP architecture. Safety, cybersecurity and AI evidence feed a common discovery and assurance graph through product and configuration identifiers. The upper control plane governs issuer authority and access. The lower decision stage represents separate conformity, assessment and operational responsibilities. A public passport can reveal discovery information while protected engineering and security records remain access-controlled. [1]

A practical incident workflow starts with a supplier vulnerability notice, identifies the affected software component and deployed builds, and follows dependencies to exposed functions and operating assumptions. The responsible manufacturer coordinates the applicable CRA notification and corrective action. An update record then identifies any required safety regression tests before return to service. The passport is a means of organising this work; the CRA itself does not mandate a DPP. [2, 3]

Product passports under the Ecodesign for Sustainable Products Regulation depend on the

relevant product rules. PACE uses compatible discovery and evidence concepts without assuming that every robot already has a statutory DPP duty. A shared evidence package supports several reviews, while each legal conclusion remains tied to its proper basis. [14]

6 Call to action: establish a pilot implementation with investment gates

6.1 Define the minimum working ecosystem

The next decision is to resource a joint scoping exercise. Spherity and interested industry partners nominate accountable technical and commercial leads. The group should select one arm, one industrial task and three change events: a gripper change, a software update and sensor drift. A component supplier, OEM or integrator, industrial user, assessment partner and trust/evidence providers form the minimum practical chain.

Three adoption stages provide investment gates for the PACE industry initiative. Partners agree timing and commitments against customer priorities, evidence readiness and the business case. Each stage produces a usable capability and a decision on further investment.

- **Stage 1 — foundation and customer value.** Agree the use case, risk and threat boundaries, shared identifiers, access terms and comparison process. Produce one component passport and an installed-arm baseline. Gate: a second organisation can retrieve and verify the required evidence, identify missing prerequisites and demonstrate a reduction in repeat evidence-collection effort.
- **Stage 2 — operational pilot and change assurance.** Connect simulation, bench tests and hardware-in-the-loop evidence. Exercise the three changes plus tamper, status-loss, stale-attestation and contradictory-evidence cases. Gate: technical feasibility and assurance utility pass separately against pre-agreed measures; qualified review supports the proposed revalidation boundary.
- **Stage 3 — industrial evaluation and scaling.** Run repeated change cases with a user and assess the service economics. Test portability with a second provider or module. Gate: agreed performance and quality criteria, a sustainable operating model and a transfer package support a decision on subsequent industrial rollout.

6.2 Make the proof of value as explicit as the proof of technology

Proof of Technology should demonstrate physical/configuration binding, proof and status verification, controlled evidence access, dependency analysis and isolated evidence processing. The demonstrator should show how a failure changes the applicable credential status and how accepted revalidation establishes a new baseline.

Proof of Value should establish whether another company can make a better-supported decision with less elapsed time or effort. It should include assessor usability, confidentiality, willingness to pay, onboarding cost and an auditable service-level definition. This prevents technical integration success from being mistaken for customer adoption.

Partners should enter the scoping exercise with one evidence package, one named application owner, one representative change history and a baseline estimate of current effort. Regulators and standardisation representatives can help identify evidence expectations and transfer questions

while retaining their independent public roles. The concrete output is a jointly reviewable pilot charter with responsibilities, data rights, costs, success criteria and a decision date for proceeding.

References

Selected source material and primary references. Online sources verified on 8 September 2026. Figures and tables are proposed adaptations for this brief.

[1] Stöcker, C. Verifiable Safety Assurance Evidence Graphs for Modular Humanoid Robot Arms. Spherity research manuscript; reviewed Word source and revised LaTeX edition, 8 September 2026. Technical architecture and illustrative scenarios.

[2] European Commission. Cyber Resilience Act: scope and implementation timetable for Regulation (EU) 2024/2847. Updated 7 September 2026. CRA scope and timetable

[3] European Commission. Cyber Resilience Act — Reporting obligations. Article 14 reporting triggers and staged notifications. Updated 31 July 2026. CRA reporting obligations

[4] W3C. Verifiable Credentials Data Model v2.0. Recommendation, 15 May 2025. Claims, issuers, verification and status. Verifiable Credentials Data Model v2.0

[5] W3C. RDF 1.1 Concepts and Abstract Syntax (2014); JSON-LD 1.1 (2020); OWL 2 Web Ontology Language Primer, second edition (2012). Graph representation, linked identifiers and open-world interpretation. RDF 1.1 · JSON-LD 1.1 · OWL 2 Primer

[6] W3C. Shapes Constraint Language (SHACL). Recommendation, 20 July 2017. Validation of graph structures against declared shapes. SHACL

[7] Rose, S., Borchert, O., Mitchell, S. and Connelly, S. Zero Trust Architecture. NIST Special Publication 800-207, August 2020. NIST SP 800-207

[8] European Union. Regulation (EU) No 910/2014 on electronic identification and trust services, as amended by Regulation (EU) 2024/1183. Consolidated text, 18 October 2024; in particular Articles 25, 35 and 41. Consolidated eIDAS framework

[9] European Parliament. European Business Wallets, procedure 2025/0358(COD), concerning COM(2025) 838. Legislative Train record updated 1 August 2026; status: tabled. European Business Wallet legislative status

[10] NVIDIA. Inside NVIDIA Halos for Robotics: A Full-Stack Functional Safety System for Physical AI. Technical Blog, 22 June 2026. Vendor architecture example; application-level validation remains necessary. NVIDIA Halos architecture

[11] ISO. ISO 10218-1:2025, Industrial robots; ISO 10218-2:2025, Industrial robot applications and robot cells; ISO 13849-1:2023, Safety-related parts of control systems — General principles for design. Official scope records; clause-level work requires the applicable standard texts. ISO 10218-1 · ISO 10218-2 · ISO 13849-1

[12] European Union. Regulation (EU) 2023/1230 on machinery. Product safety, conformity responsibilities and protection against corruption; general application from 20 January 2027. Machinery Regulation

[13] European Union. Regulation (EU) 2024/1689 (AI Act), as amended by Regulation (EU) 2026/1744. European Commission implementation overview, updated 3 August 2026, explains the revised high-risk application timetable. AI Act implementation overview · 2026 amending Regulation

[14] European Union. Regulation (EU) 2024/1781 establishing a framework for ecodesign requirements for sustainable products. Articles 9–15 address Digital Product Passports within the applicable product framework. Ecodesign for Sustainable Products Regulation