This paper was presented at the XXXVII ISPIM Innovation Conference, held in Granada, Spain on 07 June to 10 June 2026. ISBN 978-952-65771-2-8.

Full-Stack Digital Product Passport Architecture Supporting Interoperability and AI-Enabled Use

Kristiina Valtanen*

VTT Technical Research Centre of Finland Ltd, PO Box 1100, 90590 Oulu, Finland
E-mail: Kristiina.Valtanen@vtt.fi

Markus Ylikerälä

VTT Technical Research Centre of Finland Ltd, PO Box 1000, FI-02044 VTT, Finland
E-mail: Markus.Ylikerala@vtt.fi

Liisa Hakola

VTT Technical Research Centre of Finland Ltd, PO Box 1000, FI-02044 VTT, Finland
E-mail: Liisa.Hakola@vtt.fi

* Corresponding author

Abstract: Digital Product Passports (DPPs) are advancing under the EU Ecodesign for Sustainable Products Regulation, but implementations often remain fragmented and difficult to integrate across systems. Using Design Science Research Methodology, we investigate a modular full-stack architecture and the infrastructure-level boundary resources it provides, aiming to improve interoperability and governed access in ways that could later enable reliable AI-based and agentic use. The solution integrates a minimum viable DPP based on Asset Administration Shell (AAS) models and a dataspace exchange pattern, Smart Tags for item-level identification and condition signaling, and a Progressive Web App reader for low-friction access. Scenario-based demonstrations in pilot environments show that cross-context reuse is achievable when variability is confined to AAS submodels, interfaces, and transformation services while the core remains stable. The evaluation synthesizes infrastructure-level challenges (semantics, identifier consistency, governance, onboarding, scalability) and discusses innovation management implications for building reusable, interoperable DPP ecosystems.

Keywords: Digital Product Passport; DPP; Asset Administration Shell; AAS; interoperability; Smart Tags; Progressive Web Application; circular economy; circular innovation; AI agents.

1 Introduction and motivation

Despite growing regulatory momentum and extensive discourse (Voulgaridis et al., 2024; Zhang & Seuring, 2024) on the strategic potential of Digital Product Passports (DPPs) to enhance transparency, circularity and data-driven value creation across supply chains, practical implementations often remain fragmented and technically immature. Current DPP solutions are in many cases developed in isolation, lack semantic or technical interoperability and rely on proprietary or tightly coupled system architectures that hinder reuse, extension and cross-value-chain collaboration. (Zhang & Seuring, 2024) From an innovation management (IM) perspective, this limits companies' ability to leverage DPPs as platforms for ecosystem-level innovation, constraining opportunities for new services, business models and digitally enabled circular practices that depend on shared data and modular digital infrastructures.

At the same time, evolving DPP and data related regulation and standardization create a moving target for industry adoption, making it difficult to design solutions that are both compliant and futureproof. As a result, there is a pressing need for a modular full-stack DPP architectures that can function not only as a compliance mechanism but as a generative digital layer (Yoo et al., 2010) capable of supporting services, third-party extensions, ecosystem-level interactions as well as future AI-enabled automation. Rather than proposing a sector-specific DPP solution or evaluating AI applications, this study focuses on infrastructural design choices that condition reuse, interoperability, and future innovation across DPP ecosystems.

2 Foundations

The development of DPPs is driven by the Ecodesign for Sustainable Products Regulation (ESPR)(European Union, 2024), which entered into force in July 2024 and introduces DPPs as a core instrument within the EU's broader 'twin transition' agenda that couples climate-neutrality goals with digital leadership. While ESPR sets the overarching legal framework, the detailed functional and data requirements for different product groups will only be defined through delegated acts, which are expected to be released progressively over the coming years and create anticipation and uncertainty among industries and solution providers. In parallel, DPP standardisation efforts, especially CEN/CENELEC JTC 24, are actively shaping the technical foundations for future deployments, but the standards remain under development and are not yet fully aligned.

Alongside policy and standardization, multiple industrial initiatives and research projects have explored early DPP concepts, pilots and sector-specific solutions. Structured literature reviews identify growing experimentation across supply chains (e.g. (Zhang & Seuring, 2024) while empirical studies in electronics and manufacturing highlight both the benefits and persistent barriers related to interoperability, data availability and integration complexity (Chaudhuri et al., 2025).

Interestingly, within digital innovation research, the concept of layered modular architecture (Yoo et al., 2010) illustrates how loosely coupled physical–digital product layers can enable ecosystem-wide innovation—a principle directly applicable to DPPs, where modular digital layers could support new services, circular business models and distributed value-chain collaboration. However, existing DPP implementations seldom operationalize these principles in practice. Building on Yoo et al.'s theoretical foundation, our research develops and demonstrates a full-stack, modular DPP architecture that embodies generativity and extensibility, positioning DPP as a scalable digital platform capable of evolving with emerging requirements, including future AI-enabled capabilities.

3 Design Science Research Methodology (DSRM)

This study follows a Design Science Research Methodology (DSRM) to design and assess a full-stack DPP architecture, addressing the following research question:

> *RQ1: How can a modular, scalable, and interoperable full-stack DPP architecture be designed and incrementally extended to meet current industrial needs and remain extensible for emerging AI-mediated use?*

Following Peffers et al.'s DSRM framework (Peffers et al., 2007), the research process comprises six activities: (1) problem identification and motivation, (2) definition of the objectives for a solution, (3) design and development, (4) demonstration, (5) evaluation, and (6) communication. The goal of this research is to understand design challenges of modular DPP frameworks through experimentation under still-evolving reference implementations and requirements.

The work was conducted within two EU research projects (DigInTraCE, DaCapo) and implemented as proof-of-concepts (PoCs). The PoCs were deployed in didactic factory research environments and tailored to companies' circular economy (CE) use cases. The study does not claim production readiness, but instead, it emphasizes the use of open-source and standardized components to support transparency, reproducibility and extensibility. Next, we provide a concise summary of the DSRM activities:

1. *Problem identification and motivation:* The study was motivated by an observed gap between the strategic expectations placed on DPPs and the practical realities of implementing them across organisational and technical boundaries, where solution approaches often remain largely solution-specific. In such contexts, integration work tends to become repetitive and costly, because data structures, identifiers, access conditions, and lifecycle events are handled inconsistently across components and stakeholders, which in turn constrains interoperability and reduces the potential for reuse. Concurrently, the regulatory and standardisation landscape introduces new targets for product-group requirements and data expectations, which makes it difficult for organisations to commit to a single implementation approach without risking redesign, lock-in or misalignment as requirements mature.
2. *Definition of the objectives for a solution:* Based on the identified problem, we defined design objectives emphasizing modularity, interoperability, and extensibility as core requirements for a reusable DPP architecture adaptable to different product categories and evolving needs. The solution should support incremental deployment, allowing organisations to start with a minimum viable

capability and extend the architecture over time through clearly specified interfaces and shared model elements.

3. *Design and development:* Guided by these objectives, we designed a modular full-stack architecture and implemented three complementary artefacts that together operationalise the proposed DPP framework:
   - a minimum-viable DPP (MV-DPP) based on Asset Administration Shell (AAS) and a dataspace exchange pattern
   - dynamic-ink Smart Tags for item-level identification and condition signaling
   - a Progressive Web App (PWA) reader enabling low-friction access
   - In addition, shared AAS-based data models for DPPs and Smart Tags were defined to support reuse across product categories and considering potential links between AAS semantic identifiers and FAIR-aligned ontologies.

   A detailed description of the proposed architecture and the developed artefacts is provided in Chapter 4. In the Figure 1 is depicted the general operation of Smart Tag-enabled DPP. Upon scanning a Smart Tag, the DPP application links the tag identifier to its DPP, processes the tag's visual condition indicators and displays integrated product, condition and lifecycle information.

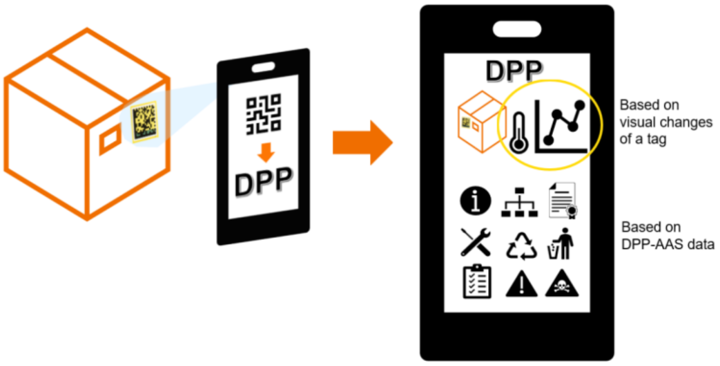

Figure 1. Basic functionality of Smart Tag-enabled DPP.

4. *Demonstration:* We demonstrated the artefacts and the data exchange between DPP instances through structured end-to-end scenarios executed in VTT's Data Innovation Lab (DIL) research server environment and in cloud installations (Microsoft Azure), with the aim of validating that the components operate together as a coherent full-stack solution and that the architecture can be tailored to different product-category needs without breaking core interfaces. The demonstrations covered identifier-based DPP access and retrieval, component data retrieval over the data space, lifecycle event capture and propagation, Smart

Tag reading via the PWA with linkage of observations to the DPP, and targeted model extensions in which new attributes or submodel elements were introduced to accommodate product-category requirements. Where applicable, we also demonstrated dataset publication and retrieval using the dataspace pattern, while keeping the focus on feasibility and integration behaviour rather than on production-grade contractual governance or marketplace-level operations.

5. *Evaluation:* The evaluation was structured to validate the study's main findings: namely, that the architecture's core components (MV-DPP, Smart Tags, PWA reader and shared data models) can operate together as a coherent full-stack solution and can be adapted across pilot contexts. We therefore conducted a scenario-based feasibility and interoperability assessment in heterogeneous pilot environments, focusing on technical interoperability (protocol/API compatibility and interface reuse) and data interoperability (shared model elements and cross-system exchange effort). The evaluation combined formative ex ante reflection during iterative design and development with ex post assessment following the demonstrations. The results are synthesized as design challenges and derived design principles in Section 5. In addition, we assessed at a conceptual and architectural level how the proposed DPP architecture supports emerging structural AI and agent-based opportunities, drawing on a targeted literature review and limited exploratory demonstration.
6. *Communication:* The communication activity is realised primarily through this article, which reports the research process, artefacts and findings. In addition, recorded demonstration videos published on the DaCapo and DigInTraCE projects' dissemination channels illustrate the operation of the developed artefacts.

4 DPP architecture and artefacts

Figure 2 illustrates the proposed layered, full-stack DPP architecture, which is designed to support incremental deployment, interoperability across organizational boundaries and future extensibility, while avoiding tight coupling between components. The implementation builds on widely adopted open-source frameworks and reference implementations, avoiding proprietary dependencies. There, AAS acts as the established information layer widely supported by manufacturing industries. It is decoupled from user interfaces, application logic and cross-organizational data exchange mechanisms. At the interaction layer, DPP data can be accessed, depending on the use case requirements, through multiple interchangeable interfaces, including lightweight Progressive Web Apps (PWA) for mobile scanning, programmatic APIs for system-to-system integration, and domain specific UI applications such as the Electronic Maintenance Book implemented in the DaCapo project.

At the application layer, a dedicated backend orchestrates access, identity management and relationships between DPPs and physical identifiers. The information layer is anchored in standardized AAS representations, which encapsulate product, lifecycle, and Smart Tag-related data in a modular and machine-interpretable form. Cross-organizational data sharing is realized using emerging dataspace protocols aligned with European standardization efforts. Data spaces separate data governance and access control from DPP application logic and avoid centralized data pooling.

Smart Tags are shown as optional edge components linking physical products to the digital stack. They provide item-level identification and condition information that can be associated with DPP instances, thereby grounding the architecture in real-world product contexts while remaining loosely coupled to the DPP core.

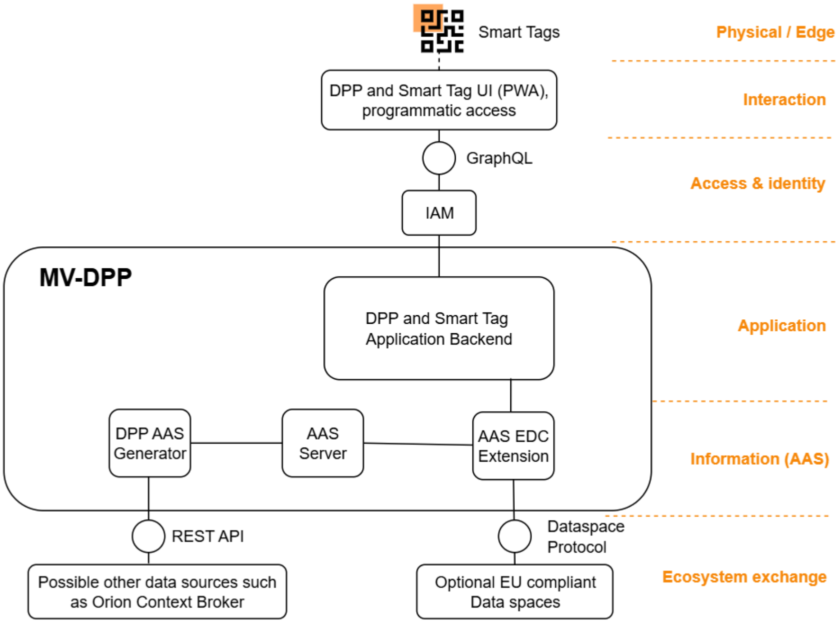

Figure 2: Layered, modular full-stack DPP architecture.

MV-DPP on AAS and Dataspaces

The MV-DPP constitutes the core artefact of the architecture, grounded on standardized AASs combined with Eclipse Dataspace Components (EDC). EDC operationalizes policy-driven, cross-organizational data exchange, which well reflects the EU DPP principle of decentralized data management. The module called *DPP and Smart Tag application backend*, implemented using the Spring framework, orchestrates access to DPP content, manages associations between Smart Tags and DPP instances as well as enforces role-based access control (RBAC) through integrated identity and access management (*IAM*) based on Keycloak. Also, it exposes its functionality via a GraphQL API to user-facing interfaces and third-party systems requiring programmatic access, enabling both human interaction and system-to-system integration without a graphical user interface.

DPP- and Smart Tag-related data are represented and persisted as AAS instances managed by an *AAS server* (Eclipse BaSyx), which serves as the repository of structured product and lifecycle information. These AAS representations are generated through a *DPP AAS Generator*, which exposes a REST-based API to transform heterogeneous source data, such as enterprise systems or external context brokers, into standardized AAS submodels, thereby avoiding source-specific assumptions in the DPP core. Cross-organizational data exchange is handled through an *AAS EDC extension*, which implements the dataspace protocol and enables controlled publication and retrieval of AAS artefacts between participants. The extension supports fetching AAS data from the AAS server, creating and consuming dataspace artefacts, and delivering them to the application backend under defined usage policies.

Together, these components allow the MV-DPP to function as a reusable infrastructure element rather than a product-specific application, easing integration effort and supporting gradual extension as regulatory requirements and standards evolve.

Smart Tags

Smart tags are data carriers that combine item level identification with environmental or item condition monitoring (Gligoric et al., 2019). The use of monitoring functionalities enables access to dynamic environmental or condition data that can be used to make more informed decisions when combined with DPP data, e.g. decisions on product repairability based on exposure to inappropriate or hazardous conditions. As an example, electronic devices are sensitive to extreme temperatures or contact with water. Therefore, it would be beneficial to know about exposure to these conditions when deciding if recycle or component replacement is the most feasible option, for example.

In the context of DPPs smart tags are usually considered as a combination of 2D bar codes and visual sensors as illustrated in Figure 3. Usually, the visual sensors are printable inks that are sensitive to conditions, such as temperature, humidity, UV light or presence of oxygen or volatiles, among others. When exposed to a pre-defined threshold value the sensor changes its colour either immediately or gradually over a short period of time. The change can be reversible or irreversible. Usually, the visual sensors give a simple on-off response i.e. if the threshold value has been reached or not. However, the gradual colour change can be used also for semi-quantitative analysis to assess the duration and level of exposure (Hakola et al., 2025).

Building on these general principles, the Smart Tags developed and evaluated in this study, can complement the MV-DPP architecture by operating as DPP data carriers; however, in the DigInTraCE project their role is emphasized as distributed sensors, contributing humidity and condition data that augment DPP lifecycle information. Changes in the visual state of a Smart Tag provide a simple but robust signal of product history, which can be interpreted at access time rather than continuously streamed.

Figure 3: A principle of a smart tag for DPPs: (a) a standard QR code is combined with a sensor area (yellow area in the top left corner) that (b) changes its colour (from yellow to transparent) when exposed to a threshold value of a condition it is sensitive to (in this case > +37 °C).

To support interoperability and reuse, an AAS model for Smart Tags was designed in parallel with the DPP models. This model captures identifier mappings, sensor observations, and contextual metadata in a way that allows Smart Tags to be flexibly associated with different DPP instances or used independently in non-DPP scenarios. By treating Smart Tags as modular data carriers rather than tightly coupled sensors, the architecture supports heterogeneous use cases without imposing unnecessary dependencies on the DPP core.

Progressive Web App Reader (PWA)

The Progressive Web App Reader (*PWA*) provides a low-friction access layer for interacting with DPP and Smart Tag data. Implemented as an installable web application, it can be used on standard devices without dedicated hardware or platform-specific deployment, supporting practical adoption across organizational contexts. Further, this approach is consistent with ESPR requirements, which mandate that DPP data be machine readable and accessible via standardized interfaces using persistent identifiers, implying access through general purpose devices.

Each Smart Tag encodes a unique identifier as a URL, which is read by the PWA (based on React front-end library) using standard barcode-scanning libraries (e.g. react-zxing). When a Smart Tag is scanned, the scanning library returns both the decoded identifier and the associated result points, which are used by the application to interpret the visual state of the Smart Tag's sensor area. Based on the resolved identifier, the PWA links the observed Smart Tag data to the corresponding Smart Tag and DPP instance and retrieves the associated DPP information through the backend interfaces. In this way, the PWA functions as a controlled access point that combines physical product cues with structured digital lifecycle data, while remaining decoupled from the underlying AAS-based data models and cross-organizational data exchange mechanisms.

Importantly, the PWA is not treated as a primary locus of business logic or intelligence. Instead, it illustrates how diverse user interfaces can be layered on top of the same modular DPP infrastructure without affecting the integrity or reuse of the core components. Figure 4 shows the PWA frontend detecting the Smart Tag.

This product has been stored in suitable conditions and is ready for use

Scanned QR Code: https://digintrace.vtt.fi/?smarttag=urn:DigInTraCE:aas:SmartTag:1

AI Sensor (TL): Hot (29.6°C)

Multi-QR detections: 0

Top-Left: Hot (29.6°C)

Top-Right: Unknown

Bottom-Left: Yellow (0.7)

Bottom-Right: Unknown

Figure 4. PWA detecting the Smart Tag.

Integration conventions and shared model elements

Beyond the individual artefacts, several integration conventions and shared design decisions proved critical for achieving interoperability and portability across product categories. First, variability was deliberately confined to AAS submodels and external interfaces, while the MV-DPP core structure was kept fixed. This reduced the need for bespoke mappings when adapting the architecture to different use cases.

Second, shared AAS-based data models were defined for both DPPs and SmartTags to support reuse across contexts. In doing so, exploratory work was also conducted to align selected AAS semantic identifiers with FAIR-aligned ontologies, with the aim of reducing semantic ambiguity and supporting future machine-interpretable use. While not evaluated exhaustively, this alignment effort illustrates how semantic interoperability can be strengthened without imposing rigid or premature standardization choices.

Finally, interfaces were treated as first-class architectural elements. By exposing DPP data through well-defined APIs and dataspace protocols, the architecture enables new services, analytics components and future AI-assisted tools to consume DPP information without requiring changes to the underlying data structures. These integration conventions collectively position the proposed architecture as a generative foundation that can evolve alongside regulatory, technological and ecosystem developments.

5 Evaluation

Architectural feasibility of the proposed full-stack DPP architecture was evaluated by instantiating and examining selected, partially overlapping configurations of DPP modules across multiple use-case scenarios and product categories. During the examination, the internal MV DPP structure remained largely unchanged, while variations were introduced through differing AAS content, external data sources, integration interfaces, and deployment environments, ranging from isolated intranet setups to hybrid intranet–cloud configurations. The MV-DPP approach was shown to support scalability and extensibility using dataspaces for sharing composition-based and hierarchically modelled data among supply-chain stakeholders. In practice, this was enabled by leveraging e.g. AAS Hierarchical Structures, including the Bill of Materials (BoM) submodel, to represent product assemblies and component relationships. The examined variants indicate that reuse across contexts was feasible only when early architectural decisions confined variability to AAS content and external interfaces, rather than embedding context-specific assumptions into the MV-DPP core. However, effective data-driven collaboration between stakeholders depends critically on semantic interoperability. Although AAS provides mechanisms to support semantic identifiers and established classification systems such as ECLASS are available, semantic alignment across organizational boundaries remained largely manual and requires further maturation, including ongoing EU-level vocabulary and ontology standardization efforts.

In addition, while the MV-DPP architecture is designed to scale, the evaluation was conducted in controlled research environments using simulated or limited datasets and a small number of dataspace connectors operating concurrently. As a result, throughput, caching strategies, and discovery and lookup services could not be assessed under higher ecosystem load. Furthermore, organizational adoption dynamics and governance enforcement mechanisms were outside the scope of the pilot studies. These limitations point to the need for future work focusing on large-scale validation, operational governance enforcement, and more explicit alignment with emerging DPP standards and specifications as they evolve.

As a summary, the Table 1 synthesizes the evaluation outcomes as a set of recurring design challenges that emerged during integration and cross-context adaptation, with emphasis on issues revealed at the infrastructure level rather than within isolated components. These challenges provide the analytical foundation for the derived design principles also presented in the table and innovation management implications discussed in the next section.

Table 1 Design challenges and derived design principles

Observed design challenges	*Derived design principle*
DC-1 Constraining variability without losing adaptability: A central challenge was supporting multiple product categories, use cases and heterogeneous incoming data sources while preserving a non-volatile and reusable DPP core. At the same time, the limited availability of standardized AAS submodels increased pressure to implement ad-hoc extensions.	Localize variability to submodels, interfaces, and dedicated transformation layers (e.g. DPP AAS generator), and keep the DPP core steady and reusable.
DC-2 Semantic interoperability across organizational boundaries: Established classification systems (e.g. ECLASS) provided partial support but did not fully resolve cross domain interpretation. The absence of mature, shared EU level DPP ontologies further complicated semantic alignment and the lack of explicit semantic representation in API interfaces (e.g. GraphQL) limited machine interpretability	Make semantics explicit and machine-interpretable across AAS models and APIs by aligning semantic identifiers with shared vocabularies and FAIR-oriented ontologies, rather than relying on classifications alone.[1]
DC-3 Governance layering between dataspace and DPP requirements: Questions emerged about where DPP specific rules should reside: within dataspace governance, within application logic, or as a separate governance layer. In addition, dataspace negotiations and policy enforcement introduced latency into DPP data queries, indicating a need for caching or architectural mitigation strategies.	Separate generic dataspace governance from DPP-specific regulatory and lifecycle rules, allowing governance responsibilities to be layered rather than hard-coded.
DC-4 Identifier consistency across physical and digital layers: Maintaining persistent links across multiple identifier schemes, lifecycle states and product granularity levels was identified as crucial for reliable DPP integration but was not extensively studied in this work. However, the SmartTag data model was designed to support flexible identity schemes.	Design identifiers as persistent, resolvable, and lifecycle-aware boundary objects, supporting multiple schemes, explicit mapping between identifiers (e.g. SmartTag ID ↔ product DPP ID), and product granularity.[2]
DC-5 Supporting incremental adoption across heterogeneous maturity levels: Even with containerized deployments (e.g. Docker packages), practical adoption revealed dependencies on local infrastructure knowledge, such as firewall configuration or connector version compatibility.	Enable minimum-viable participation with clear upgrade paths and minimize infrastructure dependencies to support onboarding across varying digital maturity levels.
DC-6 Scaling dataspace interactions beyond pilot conditions: Throughput, caching strategies, discovery mechanisms, and lookup services could not be meaningfully assessed at scale.	Design dataspace mechanisms for scalability and robustness, so performance and reliability can be validated as the ecosystem grows.

---

[1] In practice, this requires closely tracking emerging DPP-related semantic standardization efforts, including CEN/CENELEC JTC 24 work on DPP semantics, delegated acts under ESPR, and wider European semantic interoperability work conducted e.g. in ETSI TC DATA committee.

[2] For identifier standardization, see CEN/CENELEC JTC 24's work on DPP unique identifiers and data carriers, which introduces harmonized frameworks for cross-system interoperability and product granularity management.

This paper was presented at the XXXVII ISPIM Innovation Conference, held in Granada, Spain on 07 June to 10 June 2026. ISBN 978-952-65771-2-8.

6 Discussion and implications

Towards Structural AI for DPP

Within circular economy, AI has been characteristically applied inter alia to predictive maintenance, circular decision support, generative design, operational optimization and assistance systems (Celik et al., 2025; Handoyo & Sueb, 2026), many of which have also been demonstrated in the DaCapo project. In contrast to these AI applications that embed decision logic into task-specific systems, structural AI in DPP contexts operates at the infrastructure level, shaping how data is governed, interpreted, and exchanged and thereby enabling a wide range of downstream decision-support and service applications. Accordingly, this subsection shifts attention from evaluating application-level AI systems towards synthesizing structural AI opportunities that are enabled by the architectural choices presented earlier.

In this regard, recent DPP research also increasingly positions AI not merely as an add-on feature, but as an infrastructural mechanism for transforming regulatory DPP and other digital traceability systems into *scalable service platforms*. As an illustration, empirical and conceptual studies highlight concrete AI opportunities in areas of:

- **facilitating data quality assurance and validation at scale**, where machine learning models detect anomalies and validate heterogeneous sustainability indicator submissions prior to integration into a DPP platform, particularly relevant in fragmented, multi-tier value chains such as textiles (Cruz et al., 2025) Enhanced data accuracy and interoperability support DPP data governance (Tamm et al., 2025).
- **AI-mediated DPP interpretation and interaction**, in which Retrieval-Augmented Generation (RAG) and conversational interfaces provide a governed access layer over structured DPP data by retrieving authoritative lifecycle information (e.g. AAS-encoded DPP elements at runtime) and synthesizing grounded explanations and comparisons (Hilgarth & Schicker, 2026; Stjepandić et al., 2025). As highlighted by (Tamm et al., 2025), such capabilities are particularly relevant in e-commerce settings, where AI-based explanations, comparisons and sustainability narratives can be generated dynamically without duplicating underlying DPP data.
- **AI-supported interoperability and compliance monitoring**, where machine-readable representations of regulatory requirements (e.g., digital regulation documents) are linked to DPP data to enable proactive conformity checks and auditable compliance processes (Hilgarth & Schicker, 2026)
- **agentic and machine-to-machine DPP consumption**, where AI agents query DPPs as machine-accessible API infrastructure rather than human-readable documents, enabling automated comparison and recommendation workflows based on evidence quality and trust boundaries (Tamm et al., 2025; Teikari & Fuenmayor, 2026)

Across these use cases, the underlying promise is that AI can reduce administrative burdens, increase the usability of DPP data in operational decisions and enable new circular-economy services, provided that the DPP ecosystem can deliver sufficiently structured, interoperable and trustworthy data. While only a subset of the reviewed work explicitly adopts an AI-agent perspective, the structural AI opportunities identified in the literature remain consistent, as they primarily concern infrastructural capabilities that are prerequisites for both agent-based and non-agentic AI implementations.

In parallel with the structural opportunities, literature consistently identifies also structural barriers that constrain AI adoption in DPP ecosystems. Data fragmentation as well as inconsistent identifiers and metadata across tiers limit the feasibility of AI-driven analytics and automation. Multiple case evidence shows interoperability and standardization gaps along with knowledge barriers and IP concerns as recurring obstacles (Chowdhury, 2025; Wan et al., 2025). Furthermore, data quality and verification challenges, including missing, noisy or unverifiable sustainability claims, undermine the reliability of AI outputs, motivating machine learning (ML) based validation approaches and highlighting the dependence of RAG systems on retrieval quality (Cruz et al., 2025; Stjepandić et al., 2025). Heterogeneous digital maturity and cost constraints, especially for SMEs and lower-tier suppliers may create uneven participation and limit access to the sensing, digitization and integration capabilities needed for AI-enabled DPPs (Chowdhury, 2025; Wan et al., 2025). Semantic ambiguity and "conceptual silos" in DPP ontologies can impede reliable machine reasoning. Ontology work (Tuzun et al., 2025) suggests that aligning DPP concepts with an upper ontology such as Basic Formal Ontology (BFO) helps preserve critical distinctions (e.g. between physical processes and digital instructions), thereby strengthening semantic interoperability and downstream agent-based reasoning. Alrdaan (Alrdaan, 2026) also brings up relevant aspect about legal and governance constraints becoming more prominent as automation increases. This highlights robust auditability, safeguards for contestability and careful treatment of data rights and liability as critical constraints on AI-supported DPP infrastructures.

As a first step toward agentic utilization, VTT conducted a limited RAG experiment to explore AI-assisted DPP generation and interaction. These initial results suggest that, consistent with Teikari's perspective, scaling toward agent-driven DPP use will require higher degrees of automation and orchestration at the infrastructure level, beyond interactive AI assistance. Thus, as future work, focusing on structural AI opportunities is a natural continuation of the proposed full-stack modular DPP architecture perspective, as it examines how intelligence can be embedded at the infrastructure level to enable scalable, trustworthy downstream applications without hard-coding decision logic into the DPP itself. The existing architecture supports AI uptake beyond its role as a compliance implementation by treating DPPs as machine-interpretable information infrastructures rather than static disclosure artefacts. Particularly, AAS's modular information structure and standardized REST interfaces can reduce integration friction and make DPP data more directly tool-accessible for AI components and future agent tool calls (Mateo-Casali et al., 2025; Stjepandić et al., 2025). Further, our experiments of aligning AAS semantic identifiers with FAIR-aligned ontologies point toward improved semantic interoperability, which is also consistent with work of (Kuiper et al., 2025) showing that RDF/OWL-based representations allow AAS models to interoperate with broader semantic standards. At the governance layer, emerging dataspace concepts provide a baseline for sharing DPP data under clearly defined access rights and usage policies, but at the same time, the regulatory, lifecycle and accountability requirements associated with DPPs suggest that additional, DPP-specific governance mechanisms may be required on top of general dataspace governance to support trustworthy AI-mediated use. Given that DPPs may contain business-critical and confidential information, including trade secrets, proprietary product data, and sensitive lifecycle records, robust governance frameworks are essential to ensure appropriate protection.

To conclude, advancing DPP infrastructures toward more reliable machine-to-machine use by strengthening semantic alignment and clarifying governance responsibilities as automation will be needed.

Implications for innovation management

The modular, full-stack DPP architecture developed in this study provides a structured pathway for organizations to move beyond compliance-driven pilots and toward generative, innovation-ready DPP ecosystems. While empirical use was limited to research pilots and industrial proof-of-concepts conducted in controlled environments, these settings allowed early adopters to explore architectural choices and organizational readiness without requiring full ecosystem deployment. This work also operationalizes Yoo et al's layered modular architecture theory, demonstrating how loosely coupled digital layers and boundary resources can enable ecosystem-wide innovation and incremental adoption.

A central implication for IM is the strategic value of boundary resources. The architecture exposes standardized interfaces, such as REST and GraphQL APIs, that serve as tangible boundary resources for ecosystem actors. Modular data objects, including AAS submodels (e.g., Digital Nameplate), Smart Tag observations, and DPP instance records, function as reusable building blocks. These resources allow complementors to develop e.g. comparison tools, compliance checkers or analytics services independently. This boundary resource logic directly reflects Yoo et al's generativity principle, where platform openness and modularity foster third-party innovation without compromising architectural stability.

Incremental adoption is both feasible and necessary, as observed in the staged deployments carried out in research and pilot environments. Architecture enables a staged pathway for organizations, especially SMEs, to join DPP ecosystems without prohibitive upfront investment. Firms can start with minimum viable compliance (basic DPP instance and access), progress to operational traceability (event capture and Smart Tag linkage), achieve ecosystem interoperability (dataspace sharing and shared semantics), and in later stages reach the service innovation layer (e.g. validation, analytics, AI-mediated use). Each stage is supported by specific artefacts: MV-DPP, Smart Tags, PWA, and the dataspace pattern. This staged adoption pathway makes DPP implementation manageable and scalable, while preserving incentives for ongoing capability development. At the same time, scalability and standards alignment are essential for ecosystem success.

Semantic interoperability has shifted from a primarily technical concern to being also a managerial challenge. Aligning AAS semantic identifiers with shared vocabularies and FAIR-oriented ontologies enables automation, AI-supported services, and agentic use of DPP data across organizational boundaries. IM leaders should invest in semantic governance, track emerging standards, and ensure that APIs and data models expose machine-interpretable meaning. This is not just about compliance, but about future-proofing the organization for AI-driven innovation.

Further, governance emerges as a potentially important design variable influencing participation, innovation, and value capture in DPP ecosystems. Although separating generic dataspace governance from DPP-specific regulatory and lifecycle requirements offers a useful architectural framing, the practical realization of such layered governance was not a central focus of this study. The increasing role of automation and AI nevertheless highlights auditability and contestability as critical open questions for scaling DPP infrastructures.

AI-readiness should also be considered as an infrastructural investment, not just an application choice. The architecture's modular information structure and standardized interfaces make DPP data more directly tool-accessible for AI components and future agent workflows. Structural AI opportunities, such as data quality assurance, semantic mediation, compliance monitoring, and machine-to-machine usability, depend on the architectural foundations established here. Thus, for IM, preparing for AI-driven innovation starts with infrastructure, not with isolated applications.

As a summary, the study demonstrates how interoperable, reusable, and cross-value-chain DPP solutions can be practically constructed and incrementally adopted in heterogeneous industrial contexts. By articulating and evaluating a set of concrete architectural artefacts, the work provides the IM field with a practice-grounded reference model for implementing digital traceability infrastructures that support circularity-oriented innovation.

7 Conclusions

This study operationalizes a layered, modular view of DPPs by implementing a reusable infrastructure stack (AAS + dataspaces) with optional Smart Tags and a PWA reader. The primary contribution is not a finished DPP solution, but an empirically grounded architectural reference that exposes recurring infrastructure-level design challenges and corresponding design principles with direct innovation-management relevance. The pilots show how well-defined boundary resources (established information models, APIs, and exchange patterns) support incremental adoption while also creating preconditions for future AI-enabled services. For innovation management, the key implication is that design choices at the infrastructure level condition ecosystem evolution by shaping participation, innovation pathways, and value capture. Remaining challenges related to semantics, identifiers, and governance layering highlight areas where coordination and continuous standardization tracking are essential.

Acknowledgement

This research builds on work conducted within the DaCapo (GA 101091780) and DigInTraCE (GA 101091801) projects, funded under the EU Horizon Europe research and innovation programme. Views and opinions expressed are those of the authors only and do not necessarily reflect those of the European Union or HADEA. Neither the European Union nor the granting authority can be held responsible for them.

References

Alrdaan, R. (2026). *Automating Accountability: Smart Contracts and the Legal Future of Digital Product Passports in the EU Circular Economy*. In Review. https://doi.org/10.21203/rs.3.rs-8397708/v1

Celik, S. E., Chen, Z., Yldzbasi, A., & Sarkis, J. (2025). Generative Artificial Intelligence for Adopting Circular Economy in Supply Chains: Opportunities, Challenges, and Future Pathways. *International Journal of Generative Artificial Intelligence in Business*, *1*(1), 10074843. https://doi.org/10.1504/IJGAIB.2025.10074843

Chaudhuri, A., Wæhrens, B. V., Treiblmaier, H., & Jensen, S. F. (2025). Impact pathways: Digital product passport for embedding circularity in electronics supply chains. *International Journal of Operations & Production Management*, *45*(6), 1213–1226. https://doi.org/10.1108/IJOPM-01-2024-0012

Chowdhury, N. (2025). *Impact of AI-Driven Digital Product Passport (DPP) for Enhancing the Circular Economy of the Green RMG Industry*. https://dx.doi.org/10.2139/ssrn.5666330

Cruz, E. F., Silva, P., Serra, S., Rodrigues, R., Alves, M., Oliveira, J., & Rosado Da Cruz, A. M. (2025). Machine Learning-Based Data Quality Assessment for the Textile and Clothing Digital Product Passport. *Applied Sciences*, *15*(18), 10259. https://doi.org/10.3390/app151810259

European Union. (2024). *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*. https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=OJ:L_202401781

Gligoric, N., Krco, S., Hakola, L., Vehmas, K., De, S., Moessner, K., Jansson, K., Polenz, I., & van Kranenburg, R. (2019). SmartTags: IoT Product Passport for Circular Economy Based on Printed Sensors and Unique Item-Level Identifiers. *Sensors*, *19*(3), 586. https://doi.org/10.3390/s19030586

Hakola, L., Abedi, F., Nordman, S., Smolander, M., & Paltakari, J. (2025). Smart Tags as Enablers for Digital Product Passports in Circular Electronics Value Chains. *Circular Economy and Sustainability*. https://doi.org/10.1007/s43615-025-00500-y

Handoyo, S., & Sueb, M. (2026). Integrating Digital Technologies Into the Circular Economy: A Systematic Literature Review of Trends, Challenges, and Opportunities. *Circular Economy and Sustainability*, *6*(2), 111. https://doi.org/10.1007/s43615-026-00876-5

Hilgarth, B., & Schicker, G. (2026). Data Interoperability in European Product Conformity Management for Circular Supply Chains: Digital Twins and a Legal Knowledge & Data Model for linking law making and product manufacturing. *Procedia Computer Science*, *277*, 2723–2732.

Kuiper, J., Chirvasuta, T., & Breteler, J. (2025). *A general-purpose tool to convert Asset Administration Shell Templates to RDFS/OWL*.

Mateo-Casali, M. A., Boza, A., Fraile, F., & Moya-Ruiz, L. (2025). Reference Architecture for Digital Product Passports: Leveraging AI in Industry 4.0 Systems. *2025 IEEE International Conference on Engineering, Technology, and*

Innovation (ICE/ITMC), 1–9. https://doi.org/10.1109/ICE/ITMC65658.2025.11106524

Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A Design Science Research Methodology for Information Systems Research. *Journal of Management Information Systems*, *24*(3), 45–77. https://doi.org/10.2753/MIS0742-1222240302

Stjepandić, J., Holland, M., Pfirmann, F., & Zeifang, R. (2025). Predictive Generation of Digital Product Passport as Decision Support in Circular Networks. In F. Trigos, C.-J. Chou, & J. Stjepandić (Eds), *Advances in Transdisciplinary Engineering*. IOS Press. https://doi.org/10.3233/ATDE251099

Tamm, M., Draheim, D., Tammet, T., & Pappel, I. (2025). Integrating Digital Product Passports in E-Commerce: An Architectural Framework for Sustainability and AI-Driven Value. *Information Integration and Web Intelligence: 27th International Conference, iiWAS 2025, Matsue, Japan, December 8–10, 2025, Proceedings, Lecture Notes in Computer Science*, *16330*, 315–325. https://doi.org/10.1007/978-3-032-11976-6

Teikari, P., & Fuenmayor, N. (2026). *Digital Product Passports as Agentic Supply Chain Infrastructure: A Strategic Framework for Fashion*.

Tuzun, A., Warnau, J., Tripathi, S., Bachmann, N., Brunner, M., & Jodlbauer, H. (2025). *From Conceptual Silos to a Circular Economy: The BFO-Based Framework for the Digital Product Passport*.

Voulgaridis, K., Lagkas, T., Angelopoulos, C. M., Boulogeorgos, A.-A. A., Argyriou, V., & Sarigiannidis, P. (2024). Digital product passports as enablers of digital circular economy: A framework based on technological perspective. *Telecommunication Systems*, *85*(4), 699–715. https://doi.org/10.1007/s11235-024-01104-x

Wan, P. K. F., Linder, E., Arica, E., & Fragapane, G. (2025). A Multiple Case Study Providing Insights Into Applying Digital Product Passport in Several Industries. *IFAC-PapersOnLine*, *59*(10), 637–642. https://doi.org/10.1016/j.ifacol.2025.09.109

Yoo, Y., Henfridsson, O., & Lyytinen, K. (2010). The New Organizing Logic of Digital Innovation: An Agenda for Information Systems Research. *Information Systems Research*, *21*(4), 724–735. https://doi.org/10.1287/isre.1100.0322

Zhang, A., & 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*, 1–28. https://doi.org/10.1080/13675567.2024.2374256