

crunch^{IS}

The Data Leader's Guide to Digital Product Passport Architecture

How to Build One Data Layer for Every ESPR Category



Contents

1. What the Regulation Settles, and What It Leaves Open
2. Where Product Information Management Stops
3. Where the Fields Actually Live
4. Six Structural Gaps
5. A Category-Agnostic Data Layer
6. The Twelve Months Before Your Delegated Act
7. Conclusion

In brief

The Ecodesign for Sustainable Products Regulation — Regulation (EU) 2024/1781 — has been in force since 18 July 2024. It reaches products one group at a time, through delegated acts that each carry their own field list, verification rules, and date. Only one of those field lists exists today: the battery passport, mandatory from 18 February 2027 under the separate Battery Regulation (EU) 2023/1542.

Most manufacturers read that and wait, because scoping a project against an unpublished field list is not possible. The waiting is the expensive part, and it rests on a misreading: the field list is the only part that is genuinely open. How a passport must be built — the carrier, the identifiers, the access rules, the registry, the back-up copy — is already settled law, and the menu a delegated act may draw its fields from is fixed in Annex III.

Underneath all of it sits work that no delegated act will shorten: product identity that resolves to an item, supplier data arriving in a structure you can validate, provenance you can audit, and disclosure that separates public from confidential. The field list is a week of reading. That is a year.

This guide maps where passport data actually sits in a manufacturer's estate today, names the gaps that most often set the critical path, and sets out an architecture that treats the category-specific part as configuration rather than as a rebuild.



1. What the Regulation Settles, and What It Leaves Open

The ESPR is a framework. It establishes the Digital Product Passport as an instrument, sets out what such a passport is for, and then delegates part of the detail — which products, which fields, which verification, and by when — to acts adopted product group by product group. The Commission adopted the first ESPR and Energy Labelling Working Plan on 16 April 2025, naming the groups it will address first.

That structure matters more than it first appears, because it splits your planning problem in three rather than in two.

What is settled. The regulation is in force. The instrument exists. The direction is fixed, and it has been fixed since July 2024. The [digital battery passport](#) becomes mandatory in February 2027 under its own regulation, and its field list is public, detailed, and instructive.

What is not settled. For every product group other than batteries, the field-level obligation does not yet exist in adopted law. Anyone publishing a definitive field list for textiles or electronics today is reading a draft, a consultation, or a competitor's guess.

What is settled and widely missed. Between the two sits a body of law that specifies how a passport must be built, and it is where most of the architecture actually lives.

Question	Status today
Does the ESPR create a Digital Product Passport?	Settled — Regulation (EU) 2024/1781, in force 18 July 2024
Which product groups come first?	Named in the working plan adopted 16 April 2025
What may a passport be asked to carry?	Settled — Annex III fixes the menu a delegated act draws from
How must it be built — carrier, identifiers, access, back-up?	Settled — Articles 10 and 11
Where are identifiers registered, and who checks them?	Settled — Article 13 registry; Article 15 customs
At what level does the passport sit — model, batch, or item?	Set per category, from the three named in Article 10(1)(f)
What exact fields must a given category publish?	Only batteries, and under a separate regulation
When does a given non-battery category's obligation begin?	Not yet fixed in an adopted delegated act

○ What Annex III already fixes

Annex III lists the elements a delegated act may require, and a delegated act cannot invent items outside it. That list is more instructive than it is usually given credit for. Alongside the product identifier and compliance documentation, it names the Global Trade Identification Number, relevant commodity codes such as TARIC, the importer's EORI number, **unique operator identifiers** for the manufacturer and for other actors, **unique facility identifiers** for the sites involved, and the reference of the service provider hosting the back-up copy of the passport.

Three of those are worth pausing on, because they rarely appear in a manufacturer's readiness plan.

Operator and facility identifiers are not product data. They identify companies and sites across your supply chain, in a scheme that must be globally unique, and Article 12 puts the burden on the economic operator creating the passport: where a supplier or a site has no identifier yet, you request one on their behalf and pass them the details. That is a master-data workstream with an external dependency, and it does not appear anywhere in a PIM.

The back-up copy is an obligation, not a precaution. Article 10(4) requires the operator placing the product on the market to make a back-up copy available through a digital product passport service provider, and Article 11(e) requires the passport to remain available for the period the delegated act specifies — including after insolvency, liquidation, or the operator ceasing activity in the Union. A passport has to outlive the company that issued it.

○ The registry deadline has already passed

Article 13 required the Commission to set up a digital registry storing at least the unique identifiers **by 19 July 2026**. For products entering the customs procedure 'release for free circulation' it also stores the commodity code, and it stores the battery identifiers required under Article 77(3) of the Battery Regulation.

The operator's duty is concrete. You upload the identifiers, the registry returns a unique registration identifier for that product, and under Article 15 anyone placing a covered product under release for free circulation must give that registration identifier to customs. The Commission, national authorities, and customs have access.

None of that is in a delegated act. It is an outbound integration to a Commission system, an identifier you have to store back against every item, and a customs process to fit into — and it is settled today.

○ Why waiting costs more than it saves

The delegated act tells you which fields to publish. It does not tell you where those fields will come from, and it will not shorten the time it takes to make them available.

Consider what a passport requires of a manufacturer regardless of category. Every item, or every batch, needs an identifier that survives its whole life and resolves to a record. Data has to arrive from suppliers in a form a machine can validate rather than a form a person has to retype. Every value needs a provenance trail showing who submitted it and when, because Article 11(g) requires data authentication, reliability, and integrity — and a regulator asking "how do you know" is asking about lineage, not about the number. Public disclosure and commercial confidentiality have to coexist in the same record. Someone has to own the whole thing.

Not one of those depends on the field list. All of them are integration work across systems your organization already runs, and integration work has a floor on how fast it can be done honestly. Suppliers move at their own pace. Serialization touches the shop floor. Governance requires agreement between functions that do not currently meet.

When the delegated act arrives, the field list is a week of reading. The plumbing is the year.

The battery category is the exception, and it is already running. If you make LMT batteries, industrial batteries above 2 kWh, or electric-vehicle batteries, the field list, the deadline, and the verification requirements are all published — Article 77 of Regulation (EU) 2023/1542 applies from 18 February 2027. Our [EU Battery Regulation compliance checklist](#) sets out the steps in order, and the [digital battery passport guide](#) goes deeper on the data architecture that category demands.



2. Where Product Information Management Stops

Ask the market how to prepare for a product passport and a clear consensus comes back: get your product information in order, consolidate it, clean it, and publish from one place. Product information management is the named answer almost everywhere the question is asked.

It is good advice, and it is roughly a third of the problem.

A PIM is built to do a specific job extremely well. It holds the commercial and descriptive attributes of a product — names, marketing copy, imagery, specifications, category taxonomies, translations — and syndicates them, consistently, to every channel that sells the product. It is modeled at the level of the product or the SKU, because that is the level at which a product is sold.

The passport asks a different set of questions, and it asks most of them below the level a PIM is designed to describe.

It asks below the model. Article 10(1)(f) lets a delegated act set the passport at model, batch, or item level, and the battery passport already asks for both — information about the battery model, and information specific to the individual battery. Two units of the same SKU, made in different plants from different material lots, carry different footprints and different histories. Wherever a category lands on batch or item, a PIM holds one record for that SKU by design, and that design is correct for its purpose.

It asks about origin. Material composition, sourcing declarations, and emissions figures originate outside your organization, at suppliers who have no access to your PIM and often no structured data to send.

It asks about time. Repair events, ownership transfers, and end-of-life records accumulate after the product has shipped, from systems your commercial teams do not operate.

It asks who said so. Regulators inspect the trail behind a value. A PIM answers "what is the current value" — exactly what a syndication system should answer — and it overwrites the previous value when that value changes.

It asks who may see what. Composition data is frequently both a disclosure requirement and the most commercially sensitive thing a manufacturer holds. Article 11(b) names the parties who must be served, from customers and professional repairers to recyclers, market surveillance authorities, and customs — each to their own access rights.

A note on the PIM framing

None of this makes product information management the wrong investment. A clean PIM is a genuine asset, and the descriptive attributes it already governs will populate part of the passport directly. The risk is narrower and more specific: a program scoped as a PIM project will deliver the fields the PIM already held, on schedule, and discover the remaining fields late — at the point where closing them depends on supplier cooperation and shop-floor change, which are the two slowest things in the plan.

Scope the PIM work as what it is. Then scope the rest separately, and start it earlier.

The framing that holds up is not "how do we extend our product information system to cover the passport." It is "what data layer can assemble a per-item record from every system that holds a piece of it, prove where each piece came from, and serve it selectively for the next ten years." That is an integration and [data governance](#) program, and it belongs to the data organization rather than to the commercial one.

Whether you buy that layer or build it is a separate question, and a real one — [off-the-shelf and custom digital product passport platforms](#) fail in different places, and the deciding factor is usually how deep the integration into your existing systems has to go.



3. Where the Fields Actually Live

The passport reads as one document. It is assembled from six or seven systems, plus a set of fields that exist nowhere in the estate today.

The map below sets out the pattern we see most consistently across manufacturers. Readiness varies by industry and by company, so treat the readiness column as a starting hypothesis to test against your own landscape rather than as a finding.

Question	Usual system of record	Typical readiness	What is usually missing
Product identity, model, and specification	PIM, PLM	Strong	One key that reconciles across PLM, ERP, and PIM
Bill of materials, chemical composition	PLM, ERP	Moderate	Export-ready at field level; substance detail below assembly
Manufacturing record, plant, date, lot	MES, ERP	Strong at batch level	Resolution to the individual item
Test and quality results	QMS, LIMS	Moderate	Linkage to the item rather than to the order
Material origin, supplier declarations, certificates	Supplier correspondence	Weak	Any structured pipeline at all
Carbon and environmental footprint	LCA tooling, spreadsheets	Weak	Calculation at lot and site level, verified
Repairability, spare parts, service history	Service systems, connected products, dealers	Variable	A route back from whoever holds it
Operator and facility identifiers	Nowhere	Absent	A scheme, and a way to obtain them for suppliers and sites
Registry upload and registration identifier	Nowhere	Absent	An outbound integration, and somewhere to store what returns
Ownership transfer, second life, disposal	Nowhere	Absent	The data does not exist yet
Provenance — who submitted each value, and when	Nowhere	Absent	No system in the estate logs this

Three rows in that table deserve more attention than the rest, because they are the ones that cannot be solved by integration alone.

The last two rows describe data that does not exist anywhere today. No system logs the identity of whoever supplied a value and the moment they supplied it, because no business process has ever required it. No manufacturer holds a record of what happened to an item after the warranty expired. These are not integrations waiting to be built. They are new capabilities, and they need a design decision before they need an engineer.

The supplier row is the one that most often sets the critical path — not because the engineering is hard, but because it runs on other organizations' calendars.

Material origin and emissions data has to arrive from companies you do not control, in a structure a machine can validate, at a cadence that matches production. Today it usually arrives as a PDF attached to an email, once a year, in response to a request from procurement. Moving that to a structured pipeline is partly an engineering task and mostly a commercial one.



Wondering whether this is theory?

We built an auditable Scope 2 carbon forecasting platform that turns unstructured supplier invoices into figures that satisfy CSRD and ISO 14064-1 audit requirements — the same shape of problem the passport creates, one regulation earlier.

[SEE OUR WORK ↗](#)

4. Six Structural Gaps

The gaps below recur across categories and company sizes. They share a property worth naming: each one takes longer to close than to describe, and none of them closes faster because a delegated act has been published.

#	Gap	The symptom you would recognize	What closing it takes
01	Identity stops at the batch	You can name the lot but not the unit, and PLM, ERP, and PIM disagree on the part number	Identity stops at the batch
02	Supplier data has no pipeline	Certificates arrive as PDFs by email and get retyped	A schema, a validating portal, and enough commercial weight
03	Footprint is modeled, not measured	One annual figure per model, built on supplier averages	A data-collection program, and third-party verification
04	Records overwrite instead of accumulate	You cannot show what a value said last year, or who changed it	Append-only storage, decided before the build
05	Disclosure and confidentiality collide	Composition is both a disclosure duty and your trade secret	A data-collection program, and third-party verification
06	Nobody owns it	The program moves whenever some function has spare capacity	A named owner, with cross-functional decision rights

If four or more of the symptoms in that middle column are familiar, the constraint on your program is not the delegated act.

01

Identity stops at the batch

Manufacturing systems track production by order and by lot, which is what production planning needs. The passport attaches to an item, or at minimum to a lot small enough to carry a meaningful footprint. Closing the distance means serialization — marking, scanning, and a schema keyed to the instance — and it reaches the shop floor, which makes it the change most likely to meet resistance and the one worth starting first.

There is a quieter half to this gap. Before an item-level key can mean anything, the model-level keys have to agree: the PLM part number, the ERP material number, and the PIM SKU are three identifiers for one product in most estates, reconciled by convention rather than by system. A passport that assembles data across all three cannot rely on a convention.

02 Supplier data has no pipeline

Origin records, certificates, and emissions figures arrive as documents and get retyped. There is no validation at the point of entry, no rejection of an incomplete submission, and no audit trail. Building the pipeline requires a portal, a schema, and enough commercial weight to get suppliers to use them. The engineering half is a data ingestion problem with a well-understood shape; the commercial half runs on your suppliers' timelines.

03 Footprint is modeled, not measured

Most environmental figures held today are annual estimates at model level, built on assumed supplier averages. The passport expects calculation at a granularity that ties to actual production, and, in the battery case, third-party verification. The distance between an estimate and a verified figure is a data-collection program, not a calculation change.

04 Records overwrite instead of accumulate

Enterprise systems keep the current value and discard its predecessor, because that is what an operational system should do. A passport has to answer what a record said three years ago and who changed it, and Article 11(g) makes authentication and integrity an essential requirement rather than a nice-to-have. Append-only storage is straightforward to design at the start and expensive to retrofit, which makes it one of the few decisions genuinely worth making before the requirements are final.

05 Disclosure and confidentiality collide

Composition, sourcing, and performance data sit at exactly the point where a disclosure obligation meets a commercial secret. Resolving it passport by passport, by hand, does not scale past a pilot. It has to be resolved in the data model, with classification attached to the field and access decided by the identity of whoever is asking.

06 Nobody owns it

The passport spans sustainability, procurement, engineering, IT, quality, and legal. In most organizations no single role holds the whole thing, so it advances at the pace of whichever function last had capacity. This is the least technical gap on the list and reliably the most expensive, because every week of unclear ownership is a week the slower work is not running. Naming an owner — and the decision rights that come with the role — is the first data governance act of the program, not a formality that follows it.

5. The Crunch-IS DPP Data Core: A Category-Agnostic Data Layer

The architecture that survives the next decade separates two things that most programs merge: the machinery that assembles, proves, and serves a per-item record, and the definition of what a given category has to publish. We call the first the Crunch-IS DPP Data Core.

The Data Core is invariant. It is the same for a battery, a jacket, a laptop, or a chair. What sits on top of it — the category profile — changes with every delegated act, and it should change by editing configuration rather than by rebuilding a platform.

Gap	Purpose	What closing it takes
1 — Ingestion	Collect from every system and party holding a piece of the record	ERP, PLM, and MES connectors · supplier portal and API · quality and lab feeds · LCA tooling · telemetry from connected products · verification bodies
2 — Repository	Hold the assembled record as the single source of truth	Instance-keyed schema · append-only event store · classification on every field · operator and facility identifier registers · encryption in transit and at rest
3 — Processing	Turn collected data into a publishable passport	Category profile engine · completeness validation before shipment · footprint aggregation and allocation to the item · lifecycle state transitions
4 — Sharing	Serve the right view to the right party, and to the authorities	Resolver behind the data carrier · role-based views · partner APIs · signed credentials for downstream verification · registry upload and the registration identifier it returns · back-up copy with a passport service provider

Layer 2 is the one most often underestimated. An instance-keyed, append-only store with field-level classification is not a schema change to an existing warehouse — it is a data platform decision, and it is the part of this architecture that is genuinely hard to retrofit. It is also the line item that most often surprises a budget holder, because an append-only record per item, retained for the life of the product, is a sizing question as much as a schema question. Answer it at design time and it is a choice. Answer it after the first category ships and it is a migration.

○ The category profile is the part that changes

Layer 3 carries the idea that makes this architecture worth building before your delegated act exists.

A category profile is a machine-readable definition of one product group's obligation: which fields are required, what each one means, which are public and which are restricted, what validation each must pass, and which need third-party verification. It is data, not code. Adding textiles alongside batteries means adding a profile and mapping its fields to sources already connected, rather than commissioning a second platform.

Two consequences follow, and both are worth stating plainly to a budget holder.


The first is that Layers 1, 2, and 4 can be built and operated in full today, for a category whose profile has not yet been written, because none of them depend on the field list — and because Annex III already bounds what that field list can ask for.

The second is that when the act is published, the work in front of you is a mapping exercise against systems that are already integrated, already serialized, and already logging provenance.

○ Start from the one category that is finished

The [battery passport](#) is the only DPP obligation with a published, detailed, enforceable field list. Whatever you make, it is the best available specification of what a mature passport asks for: identity, composition, footprint at lot and site level, verified sourcing, dynamic performance, lifecycle events, and end-of-life.

SEE THE ACCELERATOR IN ACTION ↗



A note on standards

Openness is not a preference here. Article 10(1)(d) requires passport data to be based on open standards, machine-readable and transferable through an open interoperable exchange network without vendor lock-in, and Article 10(1)(c) requires the data carrier and the product identifier to follow the standards named in Annex III until harmonised standards are published.


Within that, the technical detail is still being written, and several bodies are working the problem from different directions — GS1 on identification and resolution, the Industrie 4.0 community on the Asset Administration Shell, Catena-X on data exchange across automotive supply chains, and the European standardization organizations on the ESPR's own requirements.

The practical response is to stay legible to all of them rather than to bet on one. Open schema formats, identifiers that resolve over the web, and API-first integration keep the cost of a standards decision low.

Proprietary structures and a data carrier that resolves only inside one vendor's platform raise it — and in that case the bill is not only commercial, because Article 10(1)(d) is a requirement, not a recommendation.

Want the architecture costed against your systems?

Every estimate turns on the same three questions: how many source systems have to be connected, whether your manufacturing data already resolves to an item, and how many suppliers have to be onboarded. We answer those against your landscape before anyone talks about a build.

An abstract graphic featuring flowing, metallic-looking liquid in shades of silver and grey, with a prominent red highlight. Overlaid on the right side is a red button with white text and an icon.

ESTIMATE MY PROJECT ↗

6. The Twelve Months Before Your Delegated Act

The sequence below is ordered by how long each item takes and how badly it blocks the next, rather than by how urgent it feels.

○ Months 1 to 3 — Settle identity and ownership.

Name a single accountable owner with a mandate across functions. Decide the granularity your category is likely to demand — Article 10(1)(f) allows model, batch, or item, and the delegated act, not you, will fix it — then confirm your manufacturing systems can carry it. If they cannot, serialization starts now, because it is the only item on this list that touches the production line. Reconcile the PLM, ERP, and PIM keys in the same pass.

○ Months 2 to 6 — Open the supplier channel.

Define the schema for what you need from suppliers, stand up a portal that validates at the point of entry, and start with the tier-one suppliers covering most of your volume. Collect operator and facility identifiers while you are there, because Article 12 makes obtaining them your problem where a supplier has none. Expect this to be a commercial negotiation as much as an integration, and expect it to run longer than the engineering estimate.

○ Months 3 to 9 — Build the repository, provenance included.

Instance-keyed, append-only, classification on every field. This is the decision that is cheap now and expensive later, and it is the one part of the platform that is genuinely hard to retrofit.

○ Months 6 to 12 — Load the battery profile as a test.

Run the only complete field list that exists against your assembled data and read the result as a gap report. It will tell you what is missing before a regulator does.

○ Months 9 to 12 — Close the two obligations outside your estate.

Decide how identifiers reach the Commission registry and where the registration identifier it returns is stored against each item, and choose the passport service provider that will hold your back-up copy. Neither is a large piece of work. Both are easy to discover late, because neither lives in a system you already run.

○ When the act arrives — Map, validate, publish.

Write your category profile against integrations that already run. This is the short part, and it is short only because everything above it happened first.

What not to build yet

Do not commission a public passport interface, a data carrier rollout, or a consumer-facing experience against a field list that has not been adopted. Those are the visible parts, which makes them the tempting ones, and they are the cheapest to change once the requirement is fixed. Build the parts nobody will see.



7. Conclusion

The Ecodesign for Sustainable Products Regulation has been in force since July 2024, and the Commission has named the product groups it will reach first. What most manufacturers are still waiting for is the delegated act that turns a framework into a field list.

The wait is rational, and the conclusion drawn from it usually is not. The field list is the short part of this program, and it is a smaller part than it looks: Annex III already bounds what can be asked, Articles 10 to 13 already fix how a passport is built and registered, and only the category's own field list is genuinely open.

The long parts — an identifier that resolves to an item, supplier data arriving in a structure you can validate, a record that keeps its own history, and a model that can disclose one thing publicly while protecting another — are the same regardless of what the field list eventually says, and they are the parts that cannot be compressed once a date is fixed.

The difference between them is mostly a matter of when the work starts, and that is a decision available today.

Find out which of the six gaps you actually have

Chapter 4 describes the gaps we see most often, and the symptom table is where most readers recognize their own estate. Which of them are yours is a question about your systems, and it takes a look at your systems to answer.

A [digital product passport](#) readiness assessment maps your product data against the passport requirements and gives you four things:

- The fields you already hold, and which system each one lives in
- The fields with no source yet — the ones that set your critical path
- What a build involves against your landscape, not a reference architecture
- A sequence for the months before your category's act lands, and what can safely wait

You finish with a plan you can take to a budget holder. No commitment to build follows from it.

About Crunch-IS

Crunch-IS is an AI-enabled custom software engineering company. Founded in 2018, we build, modernize, and integrate the systems businesses run on across the US, UK, and DACH. Our AI Pods pair senior engineers with AI agents across the lifecycle, shortening delivery without reducing engineering depth.

[DIGITAL PRODUCT PASSPORT SOLUTIONS ↗](#)[DIGITAL BATTERY PASSPORT SOLUTIONS ↗](#)[CASE STUDIES ↗](#)[CONTACT ↗](#)