



Trusted AI Advisory Alliance

Data Truth in the AI Era

Problem. Alliance. Solution. The path to the new network.

How companies control their original data, improve AI advisory, and jointly build an open truth layer — the facts, the risks, the value for every market participant, and the invitation to the brands, manufacturers, Resellers, media, associations, governments, and investors of the world.

For brands, OEMs and manufacturers, retail and distribution, publishers and media, associations, investors, and public institutions, worldwide. For internal and external use.

Michael Mohr

Founder & CEO, plugilo / DCI · Initiator of the Trusted AI Advisory Alliance

Version 24.30 (Global Edition, English) · September 2026

This Master Paper is accompanied by the Onepager (one page) and the Quick Guide (six pages) for all target groups, both on corporate.plugilo.com, plus the Business Model (prices and terms) and the Partner Operations Guide (setup, timelines, roadmap). This Master Paper is the shared overview and reference of the Alliance.

Contents

Short on time? Read the executive summary, then “The building blocks in brief”, Chapter 14 (your value), Chapter 16 (the business model), Chapter 20 (next steps), and Appendix B (key arguments per target group).

Foreword · Executive summary · The building blocks in brief

PART 1 — THE PROBLEM

1. AI becomes the decision layer
2. Where AI gets its knowledge
3. Knowledge islands: why no system holds the original
4. The consequences and the open window
5. What a viable solution must deliver

PART 2 — THE ALLIANCE

6. The Trusted AI Advisory Alliance
7. The open standard for all AI systems

PART 3 — THE SOLUTION

8. One original per object: the Master Data Card
9. The architecture: three layers
10. The services: Cards & Stacks, Wallet, pluggy, plug-it, Partner Portal
11. Real time, governance and rights: what “always current” really means
12. How it all works together: the loop and the overall comparison
13. The trust architecture: why trust becomes scalable

PART 4 — THE VALUE

14. What every group gains
15. Market and position
16. The business model: growing in the network

PART 5 — THE EXECUTION

17. Availability and expansion stages
18. Alignment with regulation and standards
19. The growth plan
20. The call to action and next steps

PART 6 — THE PEOPLE BEHIND IT

21. Stewardship and team

Appendix A: The theses of the Alliance

Appendix B: Key arguments per target group

Glossary · Sources and evidence · Closing note

Foreword

For more than thirty years I have worked with product, company, and trade data. No wave of technology has changed the relationship between companies and their customers as fundamentally as the current one. **AI is becoming the first place people turn to when they search for information, get advice, and prepare decisions.**

We founded the Trusted AI Advisory Alliance because these answers already help decide purchases, reputation, and liability, while the data they rest on consists of unverified copies. Our goal is easy to state. The truth of the data owners should reach every answer, controlled by those who created it, open to every serious system, and carried by a worldwide alliance in which every participant gains. This paper is meant to convince you to become part of it.

This Master Paper first explains how AI systems arrive at their answers and why data errors amplify there instead of fading. It shows why every AI system builds its own truth, and why no existing approach, from PIM systems, meaning a manufacturer's central product data management, through data pools to interfaces and provenance standards, lets the original reach the End User as the original. From this analysis it derives five requirements for a robust solution. Next comes the Alliance that organizes this path. Then the solution itself is described: the original for each object, the architecture of three layers, the services, the realistic limits of real time and control, the interplay in comparison with today's approaches, and the trust architecture. Only those who know the building blocks can judge the value, and that is why the benefits for each group, the market and the business model follow after that. The state of implementation, the people behind it and, as an appendix, the theses and the core arguments for each audience close the paper. Anyone who only wants an advance overview of the products will find them explained in brief right after the executive summary. One disclosure belongs here: plugilo is the initiator of the Alliance and at the same time a commercial provider. The problem analysis in the first part holds regardless of that, and anyone can verify it.

The language of this paper needs an explanation. We speak pointedly of "truth" and of the "source of truth," and there is a reason for that. The internet consists largely of copies, unchecked reuse, fake news, and increasingly of machine-generated arbitrariness. In this situation, the authorized, provenance-secured, and current version from the responsible party needs a clear name, and we call it the truth of the owner. What is meant is responsibility, provenance, version, and currency, not infallibility. External findings, measurable quality, and conflicts always remain visible (Chapters 8.1 and 8.3); the glossary defines the term bindingly.

The ambition is global. Every brand, every manufacturer, Reseller, publisher, association, and institution should be able to control its own data truth. The build-out starts in Europe and is organized internationally in parallel, because the question of whose truth AI carries arises in every country on earth. The answer is being decided now.

What follows is the case for it: the problem, the solution, the benefits, the Alliance, and the way forward.

Michael Mohr, September 2026

Executive summary

What this is about, in one sentence: The original of every data owner is to be recognized as the original everywhere, in original quality, in real time, under the owner's control, in every system, and the owner stays connected to the customer. No system delivers this today; the Alliance builds it together. And the value starts immediately: manufacturers deploy pluggy as a digital employee for advice and sales, strengthen their distributors and Resellers, and reach additional target groups through new channels, with no technical integration at the start. Sell more. Strengthen partners. Win new customers. Control original knowledge.

1. **AI is becoming the first advisory and decision layer.** More and more information paths end before customers reach a company website, a Reseller, or a manufacturer; the industry calls it Zero-Click. Decisions rest on a knowledge base of copies, abridgments, and translations; reliable proof of the original, its version date, and responsibility is usually missing.
2. **Errors are not corrected there, they multiply.** Frequently copied false claims appear as consensus, AI increasingly copies from AI, and the corrected original loses out statistically. The consequences are operational and strategic: poor advice, safety and liability risks, lost demand, damaged brands, higher service and returns costs. Plus a loss of trust that affects everyone.
3. **Every AI system builds its own truth.** Resellers, marketplaces, publishers, and brands are building their own AI applications; the same knowledge is imported, shortened, and interpreted N times. PIM, data pools, feeds, interfaces, RAG (looking something up in your own sources at the moment of the answer), MCP (an open connection between AI systems and data sources), and provenance standards are strong building blocks, but none of them guarantees that the original reaches the End User as the original: an interface guarantees no use, structure guarantees no presentation, grounding guarantees no verbatim answer, provenance, meaning proof of origin, guarantees no truth.
4. **AI takes over execution.** Agents operate tools, complete multi-step tasks and are already changing software decisions; at the same time, reliability, data access and cost remain the bottlenecks. The more a company delegates to AI, the more important the one question becomes: which original do its own agents and those of its partners access? Master data and metadata thus become a leadership task (Chapter 4).
5. **A robust solution needs five properties:** control at the owner, traceable provenance, structured and versioned data, presence at the point of decision, plus openness and neutrality. If one of them is missing, the whole does not hold.
6. **A certified real-time original for every object:** the Master Data Card, machine-readable for AI, visible for people, connected to its source. Trust is stated in three separate answers instead of one blanket grade: the trust level (Level 1 to Level 5; the blue seal from Level 2 upward, Level 5 "Managed Source Truth" additionally with a star), the measurable Data Quality Profile for data quality, the Evidence Passport for external evidence.
7. **Four services form the productive core.** Cards and Stacks hold the original, the plugilo Wallet is the End User's personal place, pluggy advises, and plug-it brings everything to every point of decision; in the Partner Portal the owner steers his data and receives the

demand back. Commercially there are two products, plugilo Cards and pluggy, and business value is paid in three units: Touchpoints, meaning active openings of a Card at the point of sale or point of interest (free on your own pages), Sessions, meaning real advisory conversations, and Matches, meaning fitting hits between customer need and released original knowledge. Operators who show the Cards of others receive Rewards. Getting started is free.

8. **The Trusted AI Advisory Alliance organizes this path** as an open, worldwide alliance of data owners; the unchangeable Charter (Chapter 6.1) secures data sovereignty, provenance and fair value creation. The data stays open to all serious AI systems; the first Chapters start in Europe. Above the platforms: an owner-controlled net, network not noise, quality not spam.
9. **All market participants gain:** brands, mid-sized companies, distribution, retail, media, associations, End Users, the public sector, and AI providers. Existing standards are integrated, not replaced; GS1 Germany and plugilo are working on a joint Proof of Concept (Chapter 18). Containers, GTIN, internet protocols, financial networks: open standards beat isolated solutions. Entry begins with the Mirror, the analysis of how AI systems talk about your own brand today.
10. **Trust is preserved across partner boundaries.** The release granted by the brand makes Resellers and publishers authorized contact points, the blue seal shows the original source, the connection to the owner, and traceable updating, and the End User takes the original along into the Wallet. Trust in a brand thus becomes a trust network across all partners and contact points, and the comparison with all of today's approaches shows that exactly this connection has been missing.

The building blocks in brief

This paper speaks of Cards, the Wallet, pluggy and the plug early on. So that nobody has to guess, here is each building block in a single sentence; Part 3 describes them in detail, Chapter 12 shows them working together.

- **The Master Data Card** is the certified original of an object, such as a product, a company or a piece of content: the owner's data, visible to people, machine-readable for AI, always current and connected to its source. Its technical core is the Master Data Record; **Stacks** bundle Cards into topics, comparisons, projects or assortments.
- **The plugilo Wallet** is the End User's personal place, as a Mobile App and as a Web Portal: here he collects originals, comparisons and projects, follows brands and partners, and keeps his demand profile under his own control.
- **pluggy** is the AI advisor of the network, a digital employee who advises from the approved original data on the principle of "Certified first" and backs every answer with its source.
- **plug-it, the plug**, is the connector at the point of decision: a button in shops, on websites, in media and, in future, at the shelf in the store that makes Cards portable and brings the originals to where customers read, compare and buy.
- **The Partner Portal** is the control center of the data owners: this is where profiles are created, where releases are granted and revoked, and where demand comes back as a signal.
- **A Skill** is the approved knowledge base from which pluggy advises; your own Skill leads as the master Skill, and approved Skills from partners add to it as sub-Skills.

- **The Trusted AI Advisory Alliance**, the Alliance for short, is the open association of the data owners with its Charter; the Trusted AI Advisory Network is the network that emerges from it. plugilo is the initiator of the Alliance and the operator of the infrastructure.

Rule of thumb: the Card is the original, the Wallet is the End User's place, pluggy is the advisor, the plug is the connection, the Partner Portal is the control center. The Alliance sets the rules.

PART 1 · THE PROBLEM

Why something has to happen at all: how AI systems arrive at their answers, why every system builds its own truth, what no existing approach delivers, what wrong data costs, why the window is open now, and what a robust solution must therefore deliver.

Part 1 at a glance: AI answers the questions before customers reach the supplier. Its data foundation consists of copies without provenance, and there the most frequent claim wins instead of the correct one. Every AI system also builds its own knowledge island, and no existing approach guarantees that the original arrives as the original. That costs trust, demand, and margin, and it gets more expensive every month. A robust answer needs owner control, verifiable provenance, structure, presence at the point of decision, plus openness and neutrality.

1. The starting point: AI becomes the decision layer

It begins with a finding: AI is taking over the first contact between suppliers and customers, and that is why this paper starts here and not with the technology.

AI assistants have spread at record speed; in early 2023 ChatGPT was considered the fastest-growing End User application to date. They now answer billions of questions every day, increasingly questions about buying, choosing, and comparing. Search engines also deliver direct answers more and more often: the majority of many search sessions already end without a click to an external website.¹ The numbers point in one direction: in the United States, according to Adobe, 39 percent of respondents had already used AI for online shopping in spring 2026, and traffic from AI sources to US retail sites grew by 393 percent in the first quarter of 2026 compared with the previous year; among those who use AI for research, 42 percent rely on it often or always as their main source for advice and purchases.² With protocols such as Google's Universal Commerce Protocol, search and checkout move into a single conversation (Chapter 15.2). Zero-Click and GEO (Generative Engine Optimization) mark the shift from a web of pages to a web of objects. **What will be visible is no longer who runs the best page, but whose structured data feeds the AI answer.** Anything that is not available as a structured object with verified provenance simply does not appear in the answer.

This shifts the first point of contact. Website, catalog, and sales conversation are often pre-empted by an answer that an AI system assembles from third-party sources: convincing in its language, regardless of whether the content is correct. Answers turn into advice and pre-selection; AI agents prepare decisions or carry them out. The first salesperson is an AI: the customer asks the AI, and the AI asks everyone. How little of this reaches the source is shown

by Pew: in Google searches with an AI summary, only one percent of visits led to a click on a cited source.³ And the news test by the Tow Center shows how uncertain attribution is: more than 60 percent of the source attributions of eight AI search services were wrong, and the systems also cited copies instead of originals; that is not an error rate for product advice, but it is evidence that source attribution does not work by itself.⁴ Anyone who is misrepresented, incompletely represented, or absent from this chain loses demand before any direct contact occurs. In many cases they never find out.

How much is at stake is shown by the customers themselves. In the Edelman Trust Barometer 2026, 88 percent of respondents named trust in a brand as an important or decisive purchase criterion, on a par with quality (89 percent) and value for money (88 percent); 41 percent attributed the greatest influence on their brand trust to unpaid third parties, only 20 percent to the brand's own statements.⁵ Salesforce adds the AI side: 61 percent of customers say AI makes the trustworthiness of companies even more important, 72 percent want to recognize when they are talking to an AI agent, and 71 percent consider human review of AI results important.⁶ And the path to a decision is getting broader: according to PwC, the average consumer uses about 3.6 shopping channels, and 51 percent count producers and manufacturers among the three main parties responsible for healthy food; responsibility thus reaches far beyond the company's own website.⁷ **Trust is therefore no longer a reputational value but a sales factor, and it increasingly arises where the brand does not speak for itself.**

In advice-intensive industries, in the kitchen studio as much as in mechanical engineering, AI does not replace the personal expert conversation. It does shape the knowledge, the expectations, and the pre-selection customers bring with them when they arrive. The strategic task is therefore not to prevent AI, but to make sure that it works with the official truth of the providers wherever people inform themselves, compare, and decide. Brands have lived through this loss of control once before: years ago, review portals took opinion sovereignty away from them, because it was no longer the manufacturer who told people what a product was like, but the reviews, genuine and invented alike. Now it is about product sovereignty. It is not the manufacturer who tells the AI what is inside his product, but the sum of the copies, as long as no one marks the original. With a search engine a brand could hold its ground with well-maintained pages; in the AI answer, what decides is whose data is recognizable as the original. How arbitrary that verdict is today is shown by an experiment at Bosch Rexroth. Asked for the five hydraulics exhibitors worth seeing at the Hannover Messe, an AI system placed the company first. Asked for the top five in factory automation for the SPS trade fair in Nuremberg, it did not appear at all. Same company, same competence, two verdicts. Whoever ranks poorly in a search engine can analyze the problem and correct it; whoever is misrepresented in a language model often does not even know what is causing it. That a market is emerging from this is shown by the funding of Profound with 180 million dollars at a valuation of 1.8 billion dollars, purely for keeping brands visible in AI answers.⁸

The question is not whether AI talks about a company. It already does. Being in the AI is not enough. Being recognized as the original is what counts.

2. The structural problem: where AI gets its knowledge

When AI gives the answers, its sources decide everything that follows, and this chapter shows what those sources actually consist of.


2.1 The open web: a body of copies without provenance

AI systems learn and research mainly on the open web. This web was never built as a reliable data foundation. Product information is copied, shortened, translated, and recombined thousands of times over by portals, marketplaces, affiliate sites, and content systems, usually without a source, a timestamp, or any marking of the original. **No authority exists that tells the original from the transcript.** For a human reader that is a nuisance. For a machine that weighs millions of sources statistically it is fatal, because it lacks exactly the information that matters: who is responsible, and which version is binding.

2.2 The copy cascade: how one error becomes a “truth”

This mechanism runs millions of times over today, fully automatically. Five steps describe it, and Figure 1 shows them at a glance.

The copy cascade: how an error becomes a “truth”



The mechanism runs millions of times, fully automatically. In today's system, the most frequently copied information wins.

Figure 1: The copy cascade: one source becomes hundreds of copies, and in the end the most frequent claim wins, not the correct one.

- 1. An error occurs.** A transposed digit in a data sheet is enough, or an outdated product version, a specification copied incorrectly, a machine mistranslation. Such errors are everyday and, at first, harmless.
- 2. The error is copied.** Reseller pages, comparison portals, affiliate networks, and content systems adopt the claim unchecked. One wrong source becomes twenty, fifty, three hundred. An example makes it tangible: a manufacturer changes an additive in a packet soup and corrects its ingredient list. The original is up to date, but three hundred old copies of that list sit on the web, at Resellers, on marketplaces, in recipe portals and comparison services. An AI asked about the ingredients finds the old entry three hundred times and the new one once. For someone with an allergy that is not a detail, and ingredient and allergen information is regulated by law in the EU.⁹ The ratio of one to three hundred is an illustration, not a measurement; that duplicates measurably change the behavior of language models, by contrast, is documented.¹⁰ Three hundred copies are not three hundred confirmations (Figure 2).
- 3. Frequency turns into certainty.** AI systems have no knowledge of the world of their own; they weigh what they encounter, how often, and how consistently. If a system meets three hundred identical false claims and only one correct manufacturer original, it treats the false claim as a consensus confirmed many times over. Statistically, the original is a single opinion.
- 4. AI copies from AI.** A substantial share of new web content is already machine-generated or machine-assisted, and there is no sound method for measuring that share exactly. These texts adopt the dominant claim and become sources for further systems in turn. The error circulates between machines without a human ever checking it again. Research shows under laboratory conditions that models can lose measurable accuracy and

diversity when synthetic data progressively replaces real data ("Model Collapse"). Continued access to authentic original data limits this risk.¹¹

5. **The correction loses.** Even a correction made by the owner stands against hundreds of uncorrected copies and against models that are already trained. Without machine-readable proof of provenance, the truth is structurally at a disadvantage.


AI believes what is described often. Not what is actually true. In today's system it is not the correct information that wins, but the most frequently copied one.

An example: the ingredient list of a packet soup, without and with the original

Today: the copy wins



With Master Data Card: the original wins



The ratio of one to three hundred is an illustration, not a measurement. Three hundred copies are not three hundred confirmations.

Figure 2: The same change, two outcomes: without the original the copy wins, with the Master Data Card the current ingredient list wins.

2.3 Why the problem intensifies on its own

- **Scale:** AI produces content orders of magnitude faster than people can check it; the body of unverified information keeps growing.
- **Persuasiveness:** Correct and false AI answers sound equally assured. The natural warning signal of contradictory sources disappears, and users can no longer tell differences in quality.
- **Loss of feedback:** When research ends inside the AI system, the data owner sees neither the error nor the drop-off. There is no complaint and no signal, only demand that fails to arrive.
- **Automated decisions:** The more selection and procurement are delegated to AI agents, the more directly data errors become revenue, contract, and supply chain risks, with no human stepping in to correct them.

Larger models do not solve the underlying problem either: they phrase things better, but they draw on the same data foundation without proof of provenance. What is missing is not better statistics about copies, but a structured, source-bound knowledge layer that every model can rely on in real time (Chapters 5.2 and 7).


3. Knowledge islands: why no system holds the original

The copy cascade explains why errors win. This chapter explains why they cannot be switched off as long as every AI system builds its own knowledge, and why the brands' current methods do not carry the original all the way to the End User either.

Alongside the cascade in the open web, a second problem arises, and it arises precisely where companies mean well. Every organization today builds its own AI applications: the Reseller a product advisor, the marketplace a shopping agent, the publisher a knowledge assistant, the distributor a procurement AI, the brand itself a support bot. Each of these systems chooses its sources itself and processes them by its own rules: its own crawlers, meaning programs that collect web pages automatically, its own data pipelines, its own search indexes, its own rules for what is found, shortened, and displayed. The same product knowledge is therefore imported, transformed, shortened, assessed, and presented not once but N times (Figure 3). **One original turns into many knowledge islands, and on every island a different version applies.**


Knowledge islands: N times today, once in the target picture

Today: every system builds its own truth



N systems × (mapping + integration + testing + operations). The brand never sees what arrives.

Target picture: one original, all systems connected



Maintain once, recognized as the original everywhere; changes reach every connected surface.

Today the same knowledge is imported, shortened and interpreted N times. The Card turns it into one original with N connections.

Figure 3: The N-times problem: the same original knowledge is imported, shortened, and interpreted anew in every AI system, and the brand sees none of these versions.

3.1 Six levels of fragmentation

Fragmentation is not a marginal phenomenon of individual providers; it arises on six levels at once. Sources sit apart: websites, PDFs, PIM data, product feeds, support knowledge, editorial

content, and internet copies. Semantics are modeled differently: attributes, taxonomies, variants, units, and relationships. Technology is built several times over: crawlers, interfaces, vector databases for AI search, chunking and update rules. Governance is managed anew for each application: releases, rights, country rules, audiences, versions. Presentation is decided by each surface itself: which fields, which source references, which actions the End User sees. And trust is lost at every step, because provenance and responsibility rarely survive copies, summaries, and generated answers. For the brand this means: even if its own data is perfect, on every island someone else decides what of it gets through.

3.2 Training is not knowledge: why “teaching it once” does not work

In everyday language, people say an AI is “trained” on company data. For product and brand knowledge that is the wrong mechanism, because training anchors knowledge in model weights, and changes do not reach it automatically. In practice, Google, Microsoft, and OpenAI therefore make company knowledge available at runtime: through retrieval, meaning looking things up in the sources at the moment of the question, through file search, vector stores, search indexes, and tools, each with application-specific ingestion, chunking, indexing, and ranking.¹² That shifts the brand's task fundamentally. Delivering good data is not enough. The brand would have to make sure that every single application indexes the right source, processes changes, uses the appropriate version, and traces the answer back to the source, and to do so for each system separately. **A living knowledge source is not trained once. It must stay permanently connected, versioned, and verified at the moment of use.**

3.3 The price: integration N times, control zero

How high the price is becomes clear from a brand integration today, in eight steps. The brand consolidates data from ERP, PIM, DAM, meaning the management of images and documents, from websites and support systems. For each recipient it analyzes that recipient's data model and creates a mapping. Mandatory fields, media formats, taxonomies, and country rules are adjusted, then a feed, an interface, or a connector is built. The target system imports or indexes by its own rules. The brand tests visibility, completeness, and currency, as far as the target system allows. Changes to sources, schemas, interfaces, and platform rules are monitored permanently, and errors are analyzed and corrected platform by platform. The Google Merchant Center shows the principle: product data has to meet platform-specific formats and quality rules, faulty or missing attributes lead to rejections or wrong presentation, and price and availability are kept current by the brand itself.¹³ The arithmetic is therefore: total effort equals the number of target systems times the recurring integration and operating work, plus central data management. With every new AI system the first factor grows, and what finally reaches the End User still stays outside the brand's control.

Every AI builds its own truth. The brand delivers the data N times and in the end controls not a single version.

Brands are not idle. They run PIM systems, use GDSN, the global exchange network for product master data, send feeds to marketplaces, publish structured data, offer interfaces, provide MCP servers, and set up RAG databases for individual applications. Each of these routes makes sense, and none of them is wrong. They were simply built for other tasks. None of them obliges a receiving system to retrieve the current original live, to adopt it completely, to prefer it over copies, and to make it visible to the End User as the original of the brand. Four gaps show where control ends in each case.

3.4 Four guarantee gaps

- **An interface guarantees no use.** An API can deliver current data. It cannot force the recipient to use that data on every request. A system can cache it, pick out individual fields, mix the source with other data, or fall back on an older copy. The recipient decides whether, when, and how.
- **Structured data guarantees no presentation.** Schema.org defines a widely used vocabulary for products and properties, and Google uses such markup to understand pages better. At the same time, Google makes clear that presentation is at the discretion of the respective application and can change.¹⁴ A vocabulary is not a control system.
- **Grounding guarantees no verbatim answer.** Retrieval methods reduce the risk of invented answers, but they do not replace governance. The search selects part of the material; ranking, context window, system rules, and the generating model determine what appears in the answer. A source can be correct while the answer shortens, combines, or misunderstands it.¹²
- **Provenance guarantees no truth.** Provenance standards such as C2PA make the origin, history, and integrity of a piece of content verifiable, and Verifiable Credentials confirm who made a statement. Both expressly do not say whether the content is factually correct; that is decided by the verifying recipient under its own rules.¹⁵ Provenance is necessary. It is not sufficient.

Trust needs content, provenance, identity, quality checking, currency, responsibility, and visible presentation at the same time. No single existing standard connects all of these elements. How this plays out across all of today's methods is shown in Chapter 12.2, once the solution has been described.

3.5 The market environment: AI-ready Product Data is not Trusted Product Truth

The market is moving too. PIM and syndication providers no longer position themselves as data management alone; they combine data storage, enrichment, distribution, quality, and AI functions under terms such as AI-ready Product Data and Agentic Commerce, meaning purchasing by AI agents. That confirms the relevance of the task and at the same time calls for a precise distinction. Based on the reviewed public descriptions of the leading providers, none of them documents the end-to-end combination this paper is about: a Card object visible to the End User and controlled by the owner, with recognizable provenance, the release of authorized partners, the End User taking it along into a Wallet of his own, and the direct return channel to the owner.¹⁶ "Not documented" does not mean that a provider has ruled the function out, but that there is no public evidence of the identical overall logic; this assessment is reviewed regularly. For the Alliance, a clear consequence follows from this. Anyone who only exports data competes in a strong market for management and distribution. The difference arises only when the original is recognized as the original: in original quality, in real time, under the control of the owner, in every place and with a connection back to the owner. The full map of the ecosystem is shown in Chapter 15.1, and the functional comparison of the target architecture with these approaches in Section 12.2.

An interface delivers data. A Master Data Card preserves the original. No system delivers that today, and that is exactly why the Alliance is needed.

4. The consequences and the open window

Here the structural finding turns into costs and risks, and because both grow with every month, the open window of time belongs in the same chapter.

4.1 What wrong data actually causes

Poor data quality was already expensive before AI.¹⁷ There is also a problem of access: Adobe determined an average machine-readability score of 66 percent for product pages in US retail; part of the knowledge that exists is not readily accessible to AI systems at all.² **What is new is the nature of the risk: errors no longer stay inside internal systems, they are actively delivered to customers and agents, in the tone of reliable advice.**

Affected	Concrete consequences of wrong or uncontrolled data
Manufacturers and brands	Products are recommended with wrong properties, specifications, or availability, or not at all; purchase decisions go against the company without it ever knowing. Brand and liability risks loom (safety and compatibility statements, recalls), service and returns costs rise. Investments in data, catalogs, and training dissipate in uncontrolled copies.
Mid-sized and smaller vendors	Few own sources drown in the noise; third parties shape the picture. Strong niche vendors are underrepresented or misrepresented, not because of their performance but because of their data situation.
Retail, specialist trade, and distribution	Advice competes with wrong AI statements the customer brings along; unmaintained copies cause wrong orders, wrong statements, and margin losses. The skilled-labor shortage sharpens the situation: ever more complex product questions meet ever fewer specialists.
Publishers and media	Content becomes raw material for third-party answers without consent or compensation, while traffic for refinancing declines. Quality loses twice: it is copied and becomes invisible at the same time.
End Users	Bad purchases through wrong statements; in electrics, child safety, health, nutrition, or chemicals, wrong specification, dosage, and compatibility statements can trigger real dangers, delivered in the tone of reliable advice.
Economies and the public sector	If the AI data basis lies outside the control of domestic owners, new dependencies arise. Documented product and information quality, whether "made in Germany", "Swiss made", "made in Japan", or any other mark of origin, loses effect exactly where decisions will be made. The question of data sovereignty is universal.

Table 1: Wrong data hits every group differently, but every group concretely.

On top of that comes a shared damage that weighs more heavily than any single one: loss of trust. An information system in which no one can tell original, copy, and responsibility apart harms providers, customers, and AI systems alike, and it does so permanently. An authorized source layer is therefore also in the interest of serious AI providers (Section 14.7). Early licensing and access models show that the market already recognizes reliable sources and puts a price on them.¹⁸

4.2 Why now: three developments at once

Three developments are running at the same time, and all three argue for acting quickly:

- **Habits are forming now.** Hundreds of millions of people are establishing the habit of preparing decisions with AI. Habits that form in months last for years. Whoever is represented correctly in this formative phase becomes part of the new normal; whoever is missing has to fight established patterns later.
- **Models and standards are taking shape now.** Which sources count as reliable, and which interfaces and labels become established, is being decided in the current phase. Early, well-structured sources shape these rules disproportionately; late corrections have little effect (Section 2.2).
- **The copy cascade is accelerating now.** Every month enlarges the body of machine-generated and unverified content that any later correction has to work against. The effort needed to establish your own data truth is as low today as it will ever be.

There is a fourth point that ties the three developments together: the gateway through which owner-controlled data flows into AI systems is still unclaimed. Standards organizations, measurement partners from the market research industry (Chapter 15), and plugilo can claim it together, and with the brands of the Alliance behind them, this combination gains an influence over the large AI models that no single party would have alone. The alternative is that someone else does it. Every brand that commits to a siloed solution today is harder to win tomorrow. **Starting together now means setting the standards. Waiting means leaving them to others.**

4.3 The wake-up call: AI takes over execution, and original data becomes a leadership task

Imagine this: an AI agent advises your customer, compares offers and prepares an order. A second agent, in your trade partner's purchasing department, checks your terms and delivery times. A third, in your own company, answers a new employee's question about the responsible contact and the valid policy. Which knowledge about your company do these three use? The answer decides revenue, liability and trust, and in most companies today it is decided by chance: from copies, old data sheets, orphaned intranet pages and whatever the web yields about them.

AI is moving from answering to executing, and the pace is measurable. In the OSWorld benchmark for computer tasks, task success rose within a short time from around 12 to 66 percent; about a third of tasks still fail, but the direction is unambiguous.¹⁹ Large companies are already scaling: 40 percent of corporations with more than one billion in revenue report scaled AI agents, a year earlier it was 27 percent, and almost every third company surveyed did not buy at least one software product because it built the function itself with coding agents.²⁰ Sales, customer retention, partner management, purchasing, service, communication: there is no business area in which agents are not already taking over or preparing tasks. What used to be software is increasingly becoming a capability you describe rather than a product you buy and operate.

And that is exactly where the trap lies. The same studies show what automation fails on: not the models, but the data. 63 percent of the data leaders surveyed consider their data practice not, or not certainly, AI-ready, and Gartner expects that through 2026 six out of ten AI projects lacking AI-ready data will be abandoned.²¹ More than 40 percent of agent projects are expected to be canceled by the end of 2027 because of cost, unclear value or missing risk controls.²² Only 37 percent of companies report any positive earnings contribution from AI at all.²⁰ Whoever builds AI on outdated, contradictory information automates his mistakes: a

wrong value in the data sheet no longer travels into a copy, but into a thousand decisions that no human checks individually any more. The copy cascade from Chapter 2 becomes an action cascade.

From answering to executing: agents improve faster than the data foundation

Capability grows on the left, the bottleneck remains on the right. The gap between them is the risk of every AI project.

Capability grows

12% to 66%

Task success on computer tasks (OSWorld), Stanford AI Index 2026

27% to 40%

Corporations (> USD 1 bn revenue) with scaled AI agents, one year later; McKinsey 2026

32%

of companies did not buy software because coding agents built the function themselves; McKinsey 2026

The data foundation slows it down

63%

of data leaders consider their data practice not, or not certainly, AI-ready; Gartner 2025

60%

of AI projects lacking AI-ready data will be abandoned through 2026; Gartner forecast 2025

> 40%

of agentic AI projects will be canceled by end of 2027 (cost, value, risk); Gartner forecast 2025

Benchmarks and surveys, not success rates for business processes; forecasts are forecasts. Sources in the appendix. The consequence stands: whoever builds AI on copies automates his mistakes.

Figure 4: From answering to executing: agent capabilities are growing faster than companies' data foundations; the gap between them is the risk, and it closes only at the original.

From this follows the guiding thesis of the Alliance: **reliable automation needs a reliable knowledge foundation.** Goals, responsibility, releases, security and the checking of results remain tasks of the company; AI can support data maintenance too. But the one task that no model and no agent can take over is the decision which version applies: which product, which terms, which contact, which policy, which campaign, which statement is the original, who is accountable for it, and where it may be used. Companies that organize their authoritative knowledge as identifiable, structured, versioned and accountable original sources can build every AI application on it, internally and externally, today and at the next model change. Companies that do not build one more knowledge island with every agent (Chapter 3) and pay for it several times over. **Content remains king, but only if it is maintained, marked and networked as the original. Copies must lose their claim to validity and be replaced by the original.**

4.4 What companies must secure now: master data, metadata, connectivity

Three things make knowledge reliable for AI, and all three are in the owner's hands. **Master data** determines unambiguously what it is about: which product, which service, which company, which person with which responsibility. **Metadata** explains what applies: who is accountable for the information, which version is current, how long it is valid and what it may be used for. **Connectivity** makes this knowledge effective: in your own applications, at authorized partners and in connected AI agents, with changes that arrive everywhere and derived content that keeps its reference to the source. This is not only about products. Company and partner data, employees and responsibilities, campaigns and terms, news and publications, policies and process rules: each of these objects needs an authoritative body that is accountable for it and an original source that people and agents can recognize. The Master Data Card is the form for this, the Master Data Record its machine-readable core (Chapter 8); existing PIM, ERP, CRM, HR and document systems remain leading sources, the Card connects identity, release and use. What AI expressly does not need is one more tool with one more copy of the data.

Provenance and verification are not a plugilo invention but established requirements: the NIST profile for generative AI treats confabulation, information integrity and content provenance as central risks and recommends documented provenance with checks before and during operation; the W3C has described provenance since 2013 through entities, activities and responsible agents; and the large platforms retrieve enterprise knowledge at run time and check permissions per query instead of retraining every change into model weights.²³ Open protocols such as MCP and A2A connect agents with data and with one another (Chapter 7). What they do not determine is which source is authoritative. That is exactly the gap the Alliance closes. And the rules themselves are becoming machine-readable: Gartner expects that by 2030 half of organizations will use agents to translate governance rules into checkable data contracts.²⁴ A data package therefore needs its rules of use alongside its content, and it needs them now, not once the agents are already running.

4.5 Re-justifying investments: five questions before every AI project

The consequence for company leadership is uncomfortable but clear: every new expense must be measured by whether it contributes permanently to automated value creation or whether AI will take it over in a short time anyway. Isolated chatbots with their own data copies, repeated manual preparation of the same knowledge for every channel, rigid custom developments for standard tasks and multi-year commitments without data export or usable interfaces are not automatically wrong, but from now on they need a justification against the alternative of organizing the knowledge once as the original and reusing it everywhere. Five questions before every project:

1. **Which business result is supposed to improve**, and how high is today's effort?
2. **Which original source applies**, and who keeps it current?
3. **Which actions may the system execute itself**, and when must it stop?
4. **What do integration, operation, tokens, controls and rework cost in total?**
5. **Do knowledge, rules and evidence remain usable when the model or provider changes?**

The second question is the one this paper is about, and it is the only one that cannot be delegated. Strengthening investments are therefore always the same: name those accountable for the most important knowledge objects, clarify leading sources and contradictory versions, establish unambiguous identifiers and definitions, provide rights, versions and change events, secure open interfaces and exportable data. That is exactly what the Trusted AI Advisory Alliance exists for: to educate companies, to focus them on their core in the new AI world, and to connect them effectively with partners and AI applications through their controlled original data. **AI accelerates action. Together we make sure that the authoritative knowledge shows the way.**

The truth layer of AI is being organized. The only open questions are by whom, under which rules, and with whose participation. The AI advises your customers, with you or without you.

5. What a robust solution has to deliver

Before any answer is presented, this chapter names the standards every solution has to meet, no matter who ends up building it.

5.1 The lesson: standards beat isolated solutions

New infrastructure waves follow a pattern that repeats with remarkable reliability throughout economic history. First every company builds its own solution, because everyone wants control and differentiation. Then a viable, open standard emerges. Finally the market tips: customers, partners, and systems move to the standard, because it offers what no island can. **Learn once, use everywhere; connect once, be connected to everyone.** The isolated solutions disappear, along with the investments and the years locked inside them.

The examples are unambiguous. Before the freight container, every shipping line and every port worked to its own procedures; the container made all of them obsolete. Introduced commercially in 1956, it became the shared foundation of world trade only through international ISO standardization from 1968 onward. Before the GTIN, every retail chain maintained its own article codes; the one shared number prevailed, and it is scanned billions of times a day today. The closed online services of the 1990s lost out to the open internet protocols, which belong to no one alone. In payments, plugs, document formats, and radio standards, the same game repeated itself, and Section 6.5 shows in detail how it played out in finance.²⁵ In none of these cases did the best individual system become the world system, but the shared, connectable foundation.

The economic logic behind this is simple. An isolated solution carries its development and maintenance costs alone, and still delivers the customer only a fraction of the value, because the customer lives in the whole market. A standard spreads the costs across everyone and connects each participant to the value created by all. Given the choice, the customer picks the standard, not out of loyalty but out of self-interest. Once a standard passes a critical threshold, its network effects can no longer be stopped.

Whoever builds the base layer alone pays alone and stays alone. Standards beat island solutions, not out of idealism but out of the self-interest of everyone involved.

For the AI truth layer, that question arises now. Companies should keep developing their distinctive data, processes, and advisory expertise. The base layer of structured, provenance-secured original data that connects everywhere is not something they should reinvent on their own. Anyone who sees that an open standard is emerging in parallel saves time and capital by using it and helping shape it. That is why the Trusted AI Advisory Alliance is opening its founding circles (Section 19.3). The knowledge islands from Chapter 3 are these isolated solutions in their latest form. How well the counter-pattern holds is shown by the precedents in Section 6.5.

5.2 Five requirements for the solution

Regardless of who implements it, a viable answer to this situation has to meet five requirements; the ten test criteria of the full functional comparison (Chapter 12.2) spell them out in detail later:

1. **Control at the owner.** The binding version of every piece of information is maintained, released, changed, and withdrawn by the data owner, at any time. Aggregation without owner control only reproduces the existing problem.
2. **Traceable provenance.** Every claim carries its source, version date, and release in machine-readable form. Only then can AI systems and people distinguish original from copy and devalue the apparent consensus of the cascade.

3. **Structure instead of prose.** Data must be uniformly structured, versioned, multilingual, and machine-readable. A beautiful website is not enough, because it is built for people, not for the systems that give the answers today.
4. **Presence at the point of decision.** Authorized data has to reach the places where advice is given and decisions are made: shops, portals, media, advisory conversations in the store at the point of sale, and AI assistants. Open interfaces make sure that every serious system can access it.
5. **Openness and neutrality.** The infrastructure must be open to all owners, regardless of size and industry, and it must not become a new gatekeeper, because then it would only replace one dependency with the next. Binding, verifiable rules matter more than any technology and more than the legal form of the operator. What counts is published rules on access and allocation, the same conditions for everyone, and the owner's approval, revocable at any time.

Part 2 shows who organizes this path. Part 3 shows what the solution looks like in detail, and Part 4 what every group gains from it.

PART 2 · THE ALLIANCE

How we address the problem: the open alliance of the data owners, its rules and stages, the evidence from history, and the open standard for all AI systems.

Part 2 at a glance: the Alliance organizes the data sovereignty of the owners as an open alliance with an unchangeable Charter. Every level of participation counts, from quietly signing to helping shape it. Shared trust infrastructure has long proven itself in retail, payments, and product identification. What is emerging here is therefore not a closed service but an open standard for every serious system. Maintain it reliably once, advise trustworthily worldwide.

Companies do not have to choose between short-term business success and a long-term AI strategy, because the two belong together. The same original knowledge that pluggies use today in sales, at partners, and in new Channels also forms the controlled data basis for tomorrow's AI advice. **The immediate sales advantage thus finances and accelerates the long-term transformation.** The Alliance is the roof above it, and this part describes how it is built.

6. The Trusted AI Advisory Alliance

Now comes the answer to the finding of the first part: the open alliance of the data owners, with its rules, its stages, and its precedents.

6.1 The alliance of the data owners

The Trusted AI Advisory Alliance is the open alliance of the data owners: manufacturers, brands, Resellers, publishers, associations, self-employed professionals, and public institutions. Together they put the five requirements from Chapter 5.2 into practice. Their task can be stated in one sentence: the data owners bring their original into all systems as a Master Data Card, keep their hand on it, and stay connected with their customers, so that the original

is recognized as the original everywhere, in original quality, in real time, and under their control. No system delivers that today (Chapter 3); together it becomes the standard. The Alliance has been designed as a worldwide movement from the start: one shared, unchangeable Charter as the global core, with national and industry Chapters beneath it. Every country and every industry helps found its own movement instead of joining someone else's. The first Chapters start in Germany and Europe; further ones in North America and Asia are forming in parallel.

The goal is global, the leverage is concentrated. Our thesis: the 1,000 consumer brands with the widest reach in the world are likely to reach nearly all people in everyday life together, through their products and the connected retail networks. The assessment rests on published reach figures of large brand portfolios: Procter & Gamble names five billion people who use its products, Unilever 3.7 billion daily users in 190 countries, and Hindustan Unilever alone reaches nine out of ten Indian households through more than nine million points of sale; the Coca-Cola system counts 2.2 billion beverage servings a day in more than 200 countries.²⁶ A global measurement of the top 1,000 adjusted for overlaps does not underlie this thesis. As an illustration, around 50 of these brands as active supporters could be enough to carry the Alliance into the wider market (a scenario, not a calculated threshold). A few right decisions will make this undertaking succeed. Behind it stands a simple conviction: **reliable data is not a competitive advantage for a few, it is the business foundation for everyone.**

The Charter is unchangeable and applies to everyone involved, including the initiator. The data belongs to the owner, every approval is revocable, and every piece of information carries its provenance. Advice follows the principle "Certified first": for each statement, the authoritative certified source takes precedence, the owner for the core of his object, the competent bodies for their own findings (Chapter 8.1). The network is open and neutral in a double sense, because it competes with no data owner and is open to every serious AI system. Access is independent of size, industry, and technology partner (on stewardship, see Chapter 21), and European data protection and transparency standards apply. The business model is designed to give data owners and participating sales partners a share of the value created, according to transparent, contractually agreed rules. Users get control over their data and an immediate functional benefit.

6.2 The shared goals: what participation stands for

The attitude behind it is clear. **The Alliance does not sell the idea of data sovereignty, it organizes it and provides the tools and the network for it, open to all.** So that everyone knows what he is signing up for, the shared goals are stated; the Alliance and its initiator must be measured against them:


1. **Return data sovereignty.** The owners again control what counts as binding about their products, services, and content, not platforms, aggregators, or intermediaries. That is the core; everything else follows from it.
2. **A real-time original for every object.** A disclosed, standardized data package from the owner: retrievable worldwide, updatable at any time, permanently connected to its source.
3. **Mark the truth visibly.** Users and systems must be able to recognize which version is authorized. Correct data must not merge invisibly with copies.
4. **Open to all systems.** Open interfaces connect every reputable AI model and every application. Existing norms, classifications, and taxonomies are mapped onto each other together with associations and expert bodies, for maximum compatibility with the industry standards.

5. **Preserve the direct channel.** Anyone who uses a Card stays connected to the owner for updates, questions, and contact, even after an AI system has answered. Control means that the owner decides which knowledge he releases, where, for whom, and on what terms it may be used, and how it may be displayed, and he can revoke any release at any time. What the End User does with an answer, whether he buys, saves, or asks on, he decides himself; that freedom is exactly what makes the advice credible.
6. **Create collective leverage.** Only when owners, associations, retail, media, and policymakers pull together does the weight emerge that makes open standards effective toward global systems. That is how the influence of the actual market participants on the next internet stays secured.

The more owners the Alliance represents, the better it can assert their influence worldwide in every system. Every participation counts.

6.3 Three levels of participation: each one counts

Three levels of participation: every one counts



You do not have to become an active partner to be part of the Alliance. But no one who owns their data should be missing.

Figure 5: Three levels of participation: observing costs nothing, and helping to shape is open to every data owner.

- **Observe and sign:** Support the shared goals free of charge and without obligation, follow developments firsthand, and move to a more active level at any time. Anyone who received this Master Paper through the registration is usually an observer already; leaving is possible at any time.
- **Secure your presence:** take over the prepared official data profile, review it, and approve it. That creates the binding version of your own information, under your own control and with a risk-free start (Chapter 16).
- **Help shape it:** connect your own sources and partner networks, use Demand Intelligence (beta), and help shape the standards, the status, and the start-phase terms in limited industry founding circles. Founding partners keep their founding status permanently.

You do not have to become an active partner to be part of the Alliance. But no one who owns data should be missing.

Onboarding, timelines, communication guides, and template texts are collected in the Partner Operations Guide.

6.4 It is time for a new network

The established networks have used up their promise: spam, misleading AI mass production, opaque algorithms, and distraction at any price. Some governments now restrict access for minors by law.²⁷ Anyone who wants to reach customers, members, or readers in a serious way no longer finds a reliable path there.

The Trusted AI Advisory Network is the answer to that, and it is built differently: owner-controlled instead of algorithm-driven, serious and high in quality instead of loud and arbitrary, with marked original sources instead of anonymous mass content. Spam, misleading AI content, and distraction have no structural place here, because only owner-released sources with a marked sender enter the network, and there is no open posting of content without someone accountable. The network emerges above the existing platforms and stays independent in its data sovereignty, because no gatekeeper decides on access, order, or the rules of the truth layer. Independence does not mean isolation. Complementary live services such as search or maps are used where they serve the user (Chapter 10.3); the source of product truth stays with the owner. What emerges is a network in which the owners stand together, use their combined reach and voice, and connect directly with the users of the world, platform-independent and by their own verifiable rules.

The owners of the data connect the End Users of the world. On an open meta level above all systems, independent of today's gatekeepers.

6.5 The evidence: shared trust infrastructure has worked for decades

The pattern from Chapter 5.1 has concrete, proven precedents. In 1973, 239 banks from 15 countries founded the SWIFT cooperative, operational since 1977 and today the global standard for financial messaging. In the same year, industry and retail settled on the GTIN, the basis of the GS1 standards that more than two million companies use today. GS1 barcodes are scanned around ten billion times a day worldwide, carried by around 120 national GS1 organizations in more than 25 sectors. Visa began as a cooperative of competing banks and is a listed company today, without the standard losing any of its authority in the process. Catena-X and Manufacturing-X currently show that industry and government are willing to invest together in owner-controlled data spaces.²⁵ The Alliance carries this proven pattern over to the data foundation of AI advisory: **the structured, provenance-secured Card is, figuratively speaking, the GTIN of the truth layer.**

What is carried over is the pattern, not the legal form of the time. Those systems were mostly user-owned; plugilo remains a commercially backed operator, disclosed in Section 21.3. The core of the success is the same in both cases: a shared standard, open connections, and the same rules for everyone.

A second type of precedent, often overlooked, shows that this combination holds: testing and certification bodies. TÜV, DEKRA, and the certificate authorities of the internet are private, commercially operating companies, and their judgment still counts worldwide. It is not their legal form that does this, but the fact that their criteria are published, their procedures documented, and their work externally verifiable. This principle underlies the trust model in Chapter 8: verifiable rules instead of promises about ownership structure.

7. The open standard for all AI systems

Because an alliance carries only as far as its connections, this chapter describes the open standard that is available to every serious AI system.

None of what this chapter describes is a precondition for a company to get started; the question here is how the original can later be read in every system in the world. The Alliance is not building an island for a single AI advisor. Its stated goal is that the certified original data of the owners is used and referenced in every reputable AI system, model, and application in the world. **Our own advisory layer is the first application, not the boundary.**

The positioning is precise: plugilo is developing an openly specified and interoperable data and connection layer. The standards and interfaces are intended to be openly usable; plugilo operates the platform itself and the services offered on it commercially. plugilo is therefore not an open-source project: the specification and the interfaces are open, the source code of the operation is not.

The target state is easy to describe. For every object, meaning every product, company, service, and piece of content, there is exactly one machine-readable real-time data record from the owner: retrievable worldwide, updatable at any time, and with a direct line to the responsible party. Systems then reference the original instead of the copy. Where wanted, original data can be fed back into existing databases and replaces outdated duplicates there. Today the internet is full of copies that no one is responsible for anymore. The Alliance dissolves this absence of responsibility: every data record has an owner, a date, and an address again.

The standard has two sides, and Chapter 8.1 describes them in detail: the machine-readable Master Data Record for the systems and the visible, marked Master Data Card for people. What matters here is what follows from that. Users may take this data package with them, save it, share it, and organize it, never as a copy but as a living reference. When the owner changes something, the change is re-indexed automatically and carried through on all connected surfaces, without manual intervention. Currency, provenance, and the link to the responsible party are preserved. The owner keeps control, and the user works with the current version on every connected surface (the guarantee levels for this are named in Chapter 11).

One authoritative original, referenced everywhere, copies with an expiry date: this principle is no utopia, it is the success pattern of the most reliable systems on the internet. DNS, the address book of the internet, works with one authoritative source per zone, served redundantly through authoritative name servers, and all copies are caches with an expiry date. The DOI system points permanently to the authoritative source with hundreds of millions of persistent identifiers. The Legal Entity Identifier creates unique company identities worldwide. And GS1 is extending the GS1 barcode in exactly this direction with the Digital Link: from a mute stripe code to a reference to the official product source.²⁸ The Master Data Card carries this logic over to business communication as a whole.

Technically, the timing has never been better. The AI industry is moving from the memorized web to real-time connections with reliable sources (retrieval/grounding). Open connection standards such as the Model Context Protocol make authoritative data sources directly connectable to AI systems. Provenance standards such as C2PA, supported across industries and still young in broad adoption, show that the industry has recognized the provenance question.²⁸ The Cards deliver into exactly this architecture: structured, source-bound, available in real time. **They are the missing authoritative source layer that these standards are waiting for.**

One clear commitment follows from this: plugilo supports the open protocols and standards of this new layer instead of developing competing protocols of its own. The Model Context Protocol (MCP) connects agents with tools and knowledge sources. The Agent2Agent protocol (A2A) connects pluggy with other AI agents. GS1 Digital Link connects globally recognized product identifiers with the official online information. And the commerce protocols of the large providers, the Agentic Commerce Protocol (ACP) from OpenAI and the Universal Commerce Protocol (UCP) from Google, connect AI-supported product search and advice with purchase and transaction (Chapter 15.2). Through these connections, approved owner knowledge reaches external agents, data sources, product identities, and commerce systems, and nobody has to learn plugilo-specific interfaces.

From this openness follows perhaps the most important step of the entire undertaking. As soon as a sufficiently large pool of certified owner data is assembled, it will be offered actively to the operators of the large AI models through open interfaces. Model operators receive answers with verifiable evidence, fewer incorrect statements, and falling liability risks; how this is measured is described by the reference case in Section 19.7. Owners gain control in every system in the world, and users see correct answers with a visible source. **The program in one sentence: replace the copy with the original, everywhere.**

Above the individual systems, an open meta level emerges: the source-bound reference layer above all AI systems and above the internet, one that everyone can point to. A later build-out stage (planned, funded by coming rounds) adds the plugilo meta-model to it. The distinction from Chapter 10.3 matters: the meta-model is not another implicit knowledge store, it is a trained mediation layer for matching and retrieval. An image for it: it is the table of contents, not the book. It knows which original belongs to which question, and it fetches it; the answer always comes from the source. For use today, none of this is needed. Every answer stays bound to the certified original Card, "Certified first" applies per statement and per scope (Chapter 8.1), and what the source does not cover is not asserted. The meta-model does not compete with the large language models; it sits one level above them and provides them with the open, source-bound reference. The language models remain the surface, and the meta level supplies the reference.

Nobody has to wait for that. AI answers can already be saved and organized as Cards today. Where a certified original Card exists, it takes the place of the generated text: "Certified first," including in dialogue with outside systems. Step by step, plugilo becomes the trustworthy memory of the AI era. Fleeting answers turn into verifiable, living knowledge objects with a direct line to the owner. Every owner connects his truth with customers, partners, and employees right away, and every trade partner gets more than data: a digital advisor that answers on the basis of that truth. **Connecting the large models is a build-out stage, not a precondition.** Everyone has it in his own hands.

A real-time original from the owner for every object, openly usable in every reputable AI system in the world. No manufacturer solves that alone. Alone, everyone remains an island; together we set the standard.

The same sovereignty applies to the demand side. The demand profile of the user is planned as an open, exportable standard: the user releases it selectively to third parties and systems and can take it with him in full at any time. That is portability instead of platform lock-in.

PART 3 · THE SOLUTION

The tools for it: the original for each object, the architecture of three layers, the services, the realistic limits of real time and control, the trust architecture, and the comparison with today's approaches. Every partner can use them, none has to, because participation in the Alliance requires no service. Anyone looking for an overview only should read Chapter 10 (the services) and Section 12.1 (how they work together).

Part 3 at a glance: For every object there is a living original with visible provenance. Three layers bring it to every point of decision and make it usable. The services are an offer, not a precondition. Real time is enforced where it can be enforced, and named honestly where it cannot. Trust is preserved across partner boundaries. And the comparison with all of today's approaches shows: the building blocks exist, the connection is missing, and that connection is exactly the target architecture.

8. The original for each object: the Master Data Card

The original for each object is the building block everything else rests on, and it decides whether AI systems speak reliably or arbitrarily about a product.

8.1 The principle: the Card as a living original

The principle is simple: for every object there is a Master Data Card, the visible, certified Card with the owner's binding data, always current, with a visible sender and open to direct questions. Behind it, the Card carries the machine-readable original data set, the Master Data Record: open, structured, versioned and retrievable in real time, optimized first of all for the machines that will give the answers in future. In this paper we usually just say the original; the Card is its face, the record its memory. When the owner changes something, the change is re-indexed automatically and the Card is kept current on all connected surfaces, continuously (Chapter 11). **This is not a copy, it is a living connection between provider and person.**

One object, one ID: In the age of AI, data objects need a unique, verifiable identity. plugilo therefore pursues the stated goal of giving every data object, over time, one unique plugilo Card ID that ties together all existing identifiers of the same object. It is the same logic that makes DNS domains and DOI documents uniquely addressable worldwide. Treating all numbers of one object as equivalent is the founding idea that DCI started with in 1993 (Chapter 21).

The beginning: products. For today's focus, product data, this means something concrete: the Card ID connects the GTIN from GS1 with the article numbers of manufacturers, shops and ERP systems. Wherever the same product carries a number, that number is stored on the card. At the plug (Section 10.4), the right data is therefore always available, and it comes from the owner.

The outlook: objects of every kind. The same logic will apply over time to all data objects: to companies of every size with their address and company identifiers, from the OEM, meaning the manufacturer of the original product, to the Reseller to the restaurant or hotel. It applies equally to places, events, news, recipes and real estate, and later to personal identities as well. Initiatives such as World ID show how important unique, verifiable identities will become there. The target state is called Proof of Object: proof of a unique object together with truth

that someone stands behind. Standards identify, the owner authorizes, plugilo makes the object living and readable for AI.

Technically, every Card carries a continuous backbone for this, the Canonical Truth Backbone. It keeps four kinds of information cleanly apart, and in plain words they are: what is the object? What applies right now? What do others contribute? And what belongs to me as an End User? The first class is the object core (Canonical Object Core), source-bound and versioned; independent of any shop, media or advisory context it defines what the object is, and the owner stands behind it. The second is the authoritative live state (Live Authoritative State): recalls, lifecycle, warranty and certificate status, carried by the bodies responsible in each case. The third is attributed context (Attributed Context & Experience): prices, availability, tests and editorial assessment, always with a sender and always retractable. The fourth are the personal demand signals (Personal Demand & Runtime Signals), and they belong to the End User alone. The owner remains the source, the context supplements the core. **Context and signals must never silently overwrite the core.**

For the End User, the Card shows this order on three levels: the original core with the properties and statements the data owner is responsible for; the evidence, that is provenance, release, scope of testing and the validity period of a certificate; and the context of availability, prices and partner knowledge, each with a recognizable sender. Existing standards keep their role in this. GDSN synchronizes product master data through certified data pools, GS1 DQX assesses its quality against defined test specifications and seals verified records electronically, but it is not a platform for providing or distributing master data.²⁹ plugilo makes this distributed knowledge usable for advice as one connected object, permanently linked to its source. The core is not static for all that: the Card remains a living, versioned original.

The principle does not apply only to products. Company and partner data (accountable: the respective company), employees and responsibilities (the responsible internal unit), campaigns and terms (marketing and sales), news and publications (the publisher of the statement), policies and process rules (the releasing department): each of these objects is kept as a Card with the same structure, the same provenance and the same rules of use. The common bracket is the authorized knowledge object; stable master data, changing content, rules and events keep their own life cycles. In this way the same maintained original serves customer advice, internal automation and collaboration with partners and their agents (Section 4.4).

“Certified first” refers to the statement and its scope: The owner is authoritative for the official identity, the declared specifications, and the released object record. Authorities, certification bodies, and qualified testing organizations are authoritative for their own findings, within the scope tested and for version, market, and validity period. Part of this honesty is the status of the principle itself: “Certified first” is formulated as a binding target principle, it governs architecture and advice today and will be enforced more deeply in technical terms with every stage of expansion; it is not yet a fully guaranteed property of the system.

What happens in case of a conflict? If a competent external finding directly contradicts a statement by the owner, neither source silently overrides the other. The Card then opens a marked conflict note (Conflict Record), shows both statements with origin and status, and steers the advice by competence. On safety, compliance, recall, and legal questions, it follows the current, valid statement of the competent body; other contradictions stay visibly open until they are resolved. For authority over the origin of a statement is never proof of its factual quality.

In practice, the Card therefore separates three truths of an object. Identity truth answers what the object unambiguously is, evidenced by the owner and authoritative registries: the GTIN for trade items, the UDI for medical devices, the VIN for vehicles, the LEI for companies. Claim

and compliance truth answers which properties, approvals, and certificates apply, evidenced by owners, authorities, and testing bodies. Market and availability truth answers where, when, and at what price the object is available, evidenced by Resellers and distributors in real time. These truths are never mixed: a Reseller knows the price better than the manufacturer, but that gives it no authority over the specification.

In public, the network explains this architecture as three circles of responsibility, the audience-facing version of the same separation: the truth of the owners, the context of the partners and the real demand of the users. **This is where the central differentiation lies: cleanly separated, clearly labeled, never mixed.**

In practice, every card grows through three stages. Simple cards secure an object in lightweight form: a saved page, a link, a photo, a note. Enriched cards add structured depth: key facts, specifications, relationships, comparison attributes. The highest form is the Master Data Card (Section 8.3): provenance-aware, certifiable, owner-controlled from Level 2 upward, and therefore the visible unit of truth of an object.

The structure of every Card follows a fixed anatomy. Five base layers form the framework that is understood at once: cover, details, specifications, media, and connections. Depending on the object and the use case, four performance layers extend this framework: functions, AI advice, trust & provenance (mandatory on Master Data Cards, including the certificate), and the knowledge graph. Behind every meaningful Card sits a tenth, mandatory layer: the machine-readable layer for AI systems, interfaces, and Agentic Commerce. The order of the layers is fixed globally; missing layers are skipped, never reordered.

The separation of the two axes matters: the tier answers where and how a piece of information appears, while the four information classes of the backbone answer who owns it, how binding it is and how current it is. A tier is therefore not an authority label, and every material data element carries both: its place and its class.

Eleven mandatory features turn a Card into an original. If one of them is missing, the Card is a presentation; with all eleven it is an original (Table 2).

Feature	What it is for
Persistent Card ID	Unique object reference via GTIN, article number, domain, document or event ID; the address does not change with every update
Identified data owner	Clear area of responsibility: who stands behind the data?
Source references	Down to field or document level where required; distinguishes original from copy
Data quality level	Traceable criteria instead of a blanket grade
Timestamp, version, validity	Change history and validity status show whether a statement still applies
Countries, languages, target groups	Usage contexts in which the information applies
Releases, licenses, roles	Authorized domains and delegated roles control who may use the Card where
Mandatory presentation	Identity, source, quality and the way back to the owner may never be removed

Two readings	Machine-readable interfaces for AI, human-readable standard presentation for people
Connections	Partners, variants, accessories, compatible products, services, documents and Stacks
Actions	Save, follow, share, contact, support, inquiry and purchase directly on the Card


Table 2: Eleven mandatory features that turn a Card into an original; if one is missing, it remains a representation.

For the End User all of this condenses into the trust view of the Card. It answers eight questions at a glance, and each field states what the answer is based on; together they make the source attributable and the relationship lasting (Table 3 and Figure 6).

The End User's question	Field of the Card	Why it supports trust
Who is the owner?	Legally and organizationally responsible entity	Makes the source attributable
Where does the statement come from?	Source system, document or accountable entry	Distinguishes original from copy
How current is it?	Last change and synchronization status	Shows whether a statement can still apply
Which version applies?	Card version and change history	Allows traceability
What has been checked?	Owner release, rule-based check, external certification, shown separately	Prevents vague blanket promises
Who may display the Card where?	Partner release (Section 13.4)	Makes the touchpoint officially verifiable
What does the advice say, and where does it end?	pluggy with sources, limits and escalation path	Links information to responsibility
How is the connection kept?	Save, follow, contact and consent	Keeps the connection beyond the moment

Table 3: The trust view of a Card answers eight questions; each field states what the answer is based on.

Anatomy of a Master Data Card: how the End User sees the original



Schematic representation; the product and the owner are fictitious.

Figure 6: How the End User sees the original: sender, seal, version, three levels and the trust view on one Card.

And nobody starts from zero: more than six million prepared data records on products, companies and content form the initial fill (the group's holdings; company's own statement, as of September 2026), supplied and maintained through CF1 (Section 10.6). They contain three decades of the group's data expertise, bundled into a source-bound structure. Every owner takes over the profile that already exists and improves it step by step (the five trust levels: Section 8.3).

Breadth is also supported by a licensed data foundation. Level 1 Cards can be filled from licensed, CF1-harmonized reference data of established data suppliers; over time, further databases and suppliers will extend this foundation, from GS1 GDSN through commercial product content networks such as NIQ Brandbank or Icecat to industry catalogs. The same rules apply to every licensed source: clear labeling of origin, a clean chain of rights, and an invitation to the owner to take over the original. This data is useful, scalable and readable for AI, but clearly labeled: a Level 1 Card is a Master Data Card with the status of a Structured Reference, not yet owner-confirmed Source Truth. Owner control begins at Level 2, Managed Source Truth is reached at Level 5, and Section 8.3 describes the full level logic.

8.2 Existing standards remain: the five levels of product data

Anyone looking at product data worldwide finds five levels that are often confused: identification (which product is this?), classification (which product group does it belong to?) and attributes (which properties does it have?). Then come data exchange (how is the data transmitted?) and standards with testing (by which rules, and who confirms compliance?). Established bodies exist for each level. GS1 identifies with GTIN and GLN and exchanges through GDSN. GS1 GPC, UNSPSC and ECLASS classify; ETIM and the IEC Common Data Dictionary standardize technical characteristics. Standards are set by ISO, DIN and the DKE, the joint electrotechnical standardization body of DIN and VDE. Conformity is confirmed by independent testing organizations such as TÜV, the VDE testing institute, DEKRA, SGS, Bureau Veritas, Intertek or UL Solutions. The interplay is a division of labor that works in practice:

standardization defines the rules, manufacturers develop against them, testers confirm compliance.

For plugilo, all these organizations are complementary building blocks, not competitors. Each standard answers part of the question. None of them defines a living, owner-controlled digital object with versioning, continuous updating, documents, media, AI knowledge, provenance, partner links, and user interactions. The Master Data Card takes on that role: it integrates and connects the existing standards instead of replacing them, and it makes their results visible, comparable, and machine-readable in one shared object for the first time. For this, plugilo is developing an open mapping framework that connects identifiers, classifications, and industry-specific data models step by step. Mapping all the standards named here in full is the target state, not today's status. The world of sources reaches further than any single database: from the owners' operational systems (ERP, PLM, PIM, DAM) to identity registries, exchange and content networks, government registers, testing laboratories, physical capture, and supply chain events. Each of these source types has its own authority and its own limits, and the mapping framework assigns them accordingly.

One principle is not negotiable: no owner has to rebuild its master data into a proprietary format. DIN and ISO standards, classifications such as ETIM, ECLASS and UNSPSC, GS1/GTIN identifiers and industry taxonomies are taken in through open interfaces, right down to the fields of the Digital Product Passport. Mapping the category and standards systems connects them. This is the equivalence expertise the group has worked with since 1993 (Section 21.1). This mapping is created together with the associations and expert bodies of the Alliance, so that nobody has to give up their own ordering system.

In addition, plugilo maintains its own AI-ready object classification, fully mapped onto the existing standards. It replaces no standard. It makes objects precisely addressable for advice, comparison and retrieval, including where standards do not (yet) apply.

8.3 The trust model: three questions, three answers

The trust model answers three questions separately.

Why should a person, a company or an AI system trust the data on a card? Three questions answer that separately from one another: who is accountable for the data, how good is it, and what has been proven externally? These questions are never merged into a single grade or a blanket seal:

The trust model: three questions, three separate answers

Who is responsible for the data? Five trust levels (Level 1–5)	How good is the data? Data Quality Profile, measurable	What is proven externally? Evidence Passport
Level 1 · Structured Reference Level 2 · Owner Connected Level 3 · Owner Sourced Level 4 · Owner Approved & Reviewed Level 5 · Managed Source Truth	Completeness Currency Consistency Traceability Coverage	Audits and certificates Standards and mappings Compliance documents Organization, scope, validity Conflicts stay visible
Blue seal from Level 2, star for Level 5	Shown separately from the level	Evidence stays with the issuing body

A higher level means more owner responsibility, not automatically better quality. That is why the three answers stay separate.

Figure 7: Trust is answered in three separate answers, not one blanket score: who is responsible, how good, what is verified.

1. **Who is accountable for the data?** Five trust levels show it (Master Data Card Level 1–5). Level 1 is the automatically prepared reference ("Structured Reference"), Level 2 the official connection of the owner, Level 3 the direct link to his original sources, Level 4 the human review. Level 5 is the permanently managed source ("Managed Source Truth"), where CF1 takes over the maintenance on behalf of the owner. The blue plugilo seal appears from Level 2 upward, as soon as the owner has officially adopted his object; Level 5 additionally carries a star and is awarded only when all maintenance, rights and quality conditions are permanently in place. The criteria for awarding, reviewing and withdrawing it are documented and disclosed in the Partner Portal (Figure 8); the final design of the seal is currently being defined.
2. **How good is the data? Measurable:** Separately from the level, every Card carries a measurable quality profile (Data Quality Profile): completeness, currency, consistency, traceability, and coverage, by category, source, and market. There is deliberately no single overall score; it would be misleading. Owners see precisely which improvement has the largest effect on how they appear in AI.
3. **What is externally verified?** Tests and certificates from the standards and testing organizations in Section 8.2 are carried along as independent evidence, in the Standards, Evidence & Certification Passport. Every entry names the organization, the scope tested, the underlying standard, the date, and the validity, machine-readable and, where available, digitally signed. Every piece of evidence stays attributed to its organization; logos and seals appear only where authorized.

Two rules protect the meaning of the model. A higher level does not automatically mean complete data, and an external certificate does not automatically raise the level. Each dimension answers its own question, and that is why it is possible to say precisely what has been checked and what has not. This makes the card more than a data set. It is a traceable digital proof of trust for people and AI systems. The binding to sources makes the basis of every answer traceable; a guarantee that every generated answer is free of errors does not lie in it. What counts are the controlled knowledge base and its correct use. In the long run, the review model can also be opened to qualified or accredited third parties. Behind the model sits a field-level source architecture for this, because a universal database with all real data does not exist. Every single entry is taken from the source responsible for exactly that field, that market, that version and that point in time, and it carries a controlled provenance class: from the owner original through the official filing, the external test attestation and physical capture to labeled AI derivation. Two distinctions are non-negotiable here. Reported is not verified. And a trust level applies per field family, market and version, not across a whole product.

Five trust levels: who maintains the data, and what does the End User see?




Figure 8: Five levels, one principle: the higher the level, the more responsibility the owner takes on; the blue seal exists only at Level 5.

The same clear separation applies to the four information classes of the backbone from Section 8.1, which are never mixed. For partners from the data world and the measurement world, two rules follow from this. Licensed content enters the network as a clearly labeled reference layer, attributed and never silently absorbed. And measurement and market data stay in a separate market and insight layer, which protects the neutrality of the measurement partner as much as the control of the owners.

Looking ahead, a new dimension of testing is emerging within it, AI-ready Product Truth: confirmation that a product truth is uniquely identified, cleanly structured, citable, and fit for AI with a clear basis for use. No such module exists today as an established standard. It will be developed with qualified partners.

One original per object, five visible trust levels, external evidence with a named sender, and the start is a release, not a project. The plugilo Card: the digital original. All product knowledge. Released by the owner.

9. The architecture: Master Data, Connect, Application

Three layers show how verified original data becomes advice in every place, and why each of them has to remain usable on its own.

plugilo is not another content platform, not a marketplace, and not a generic chatbot. It is an open data, structure, and distribution infrastructure for trustworthy information. It consists of three connected layers, each of which can be used on its own:

- **A) Master Data Layer:** the owner-controlled, structured, AI-friendly base of originals, with Master Data Records, Cards and Stacks, mapping, profiles, and the Partner Portal.
- **B) Connect Layer:** take along, follow, and connect. It brings the truth to every place on the internet, through plug-it, the Browser Extension and the Follow Service, and, as a later build-out stage, through the Connect Service and the Business Graph (planned).

- **C) Application Layer:** the services for people and systems: the plugilo Wallet as the user's personal knowledge and demand wallet, pluggy, and the Demand Trainer. The open Service Store follows as a later stage.

Three layers: secure the truth, distribute it, make it usable




Figure 9: Each layer can be used and connected on its own; the Master Data Layer carries the whole.

Aggregated, privacy-compliant user signals flow back to the affected owners as Demand Intelligence (beta). What matters here is that each of the three layers stays open and can be connected on its own. **The plugilo applications are an offer, not a precondition.**

One layer secures the truth, one carries it to every place, one makes it usable. Each stands on its own. Together they are the infrastructure.

10. The services

Here are the services with which the infrastructure works day to day, so that it becomes visible what a partner can actually put to use from the first day.

All services at a glance, by layer and status




Figure 10: The productive core is live; everything beyond it is marked as an expansion stage with a clear status.

Figure 10 shows all services by layer and status. Four of them form the productive core and build on one another: the object, the End User's place, the advisor, the distribution; in the Partner Portal everything converges. What is planned is set out with status in Chapter 17.

10.1 Cards & Stacks: the object and its context

The Master Data Card from Chapter 8 is the smallest unit of the network, and it can be questioned directly: the customer asks in the Card, and the answer comes from the data behind it. Cards are created everywhere, from shop content, websites and catalogs, and anyone can create them, from a product to an editor's news item or a chef's recipe. Automatically generated and user-created Cards are preliminary stages open to takeover, invitations to the owner: when he takes over his record, the Card becomes the certified original, uniquely linked and unique, and all other representations of the same object are assigned to it. The Master Data Card is shown first, while shop texts, variants and additional offers remain traceable references. **One genuine Card per object, all others cleanly behind it.** "Certified first" applies throughout, visible in the five trust levels from Section 8.3. AI-assisted condensation keeps every Card concise: essence first, depth behind it.

- **Stacks as curated context:** Users, advisors, editorial teams and organizations bundle Cards into topic worlds, rankings, performance comparisons, projects, assortments or saved lists, ready to follow, to share and to use together, with discussion directly at the content, with family, colleagues, specialist advisors or AI. An explicit selection is an unambiguous signal: explicit intent rather than inferred attention, the basis of Demand Intelligence and of value creation from curation (Section 14.3). Stacks can be nested and kept current rule-based from sources, searches and databases, through logic rather than manual upkeep, fed from the binding original data: "all five-star hotels in Berlin", "the new products of the month", "the most-read reports of the week". Whole Stacks are taken along like a single Card, the new products of a brand, the best-of list of a magazine, dynamic shop lists such as top sellers and pre-filtered selection lists. The network accompanies the selection phase, not just the purchase.

- **Standard Stacks, pre-structured along life:** So that nobody starts with an empty collection, the Wallet is pre-structured along the major areas of life, from everyday matters through food, living and mobility to friendships, media and finances, and in business mode from products through marketing and sales to projects.
- **Business and personal profiles:** Employee Cards belong to the company domain, tied to position and role; personal Cards belong to the person alone. This separation is the basis of company-centric networking in the Connect Layer.

10.2 The plugilo Wallet: the End User's knowledge and demand

- **The plugilo Wallet, the personal knowledge and demand wallet (beta):** Free in its basic form and neither a saved-items list nor a bookmark store, but the personal counterpart to the Master Data Card. Here the End User brings together everything that occupies him: interests, research, comparisons, projects, possessions and plans, as Cards on products, companies, news, content and events, connected live to the master data, curated in Stacks and networked through follows in the Trusted AI Advisory Network. This creates a structured decision memory that stays usable for the long run, and pluggy connects both sides: Trusted Knowledge × Real Demand = Superior Advisory. Payment wallets manage money, identity wallets the ID document, the plugilo Wallet knowledge and demand; the stated goal is to connect up to 80 percent of all online shops in one application (target figure, Section 10.4).
- **Who the Wallet belongs to (beta):** to the End User, with an export guarantee, granular approvals and revocation at any time. Whoever shares nothing stays anonymous, and the demand profile is intended to become an open, exportable standard (Chapter 7). In everyday use it works like a loyalty card: recognized everywhere in the network, and pluggy never starts from zero.
- **How it is used (beta):** as a mobile app and as a web portal on the desktop, with the same content, today in German and English, further languages to follow. pluggy becomes the intelligent surface of the Wallet (Chapter 10.3). On the basis of the Trusted Data, transactions directly from the Wallet are to become possible in future (Chapter 17).
- **Retention is built in, not added afterward (beta):** Today every saved-items list demands a registration of its own; the Wallet replaces this scattering. A saved Card stays alive. The Follow Service keeps it current, price changes, availability, recalls and updates come to the End User, and Stacks turn them into lasting projects. From this come three engagement levels and three reasons to return. Active End Users save, compare and ask. Passive ones stay connected through follows and notifications. Inactive ones come back through a trigger, through a price or availability alert, a shared Stack or the next purchase decision. The longer someone uses plugilo, the more valuable their Wallet becomes.

If Coca-Cola knew who drinks Coke, it would own the largest end-customer database in the world. We make that possible: not through observation, but through voluntary connection. The End User takes the original into the Wallet, follows the brand and releases what he wants to share, and for the first time the brand knows the people who actually use its products.

- **Demand Trainer, Wallet and plug-it as a learning system (beta):** Together, the Wallet and plug-it are the Demand Trainer, not a product but the learning layer above the Wallet and

the plug. What End Users save, share and organize shows the system real demand. Chapter 10.5 describes how the signals flow back.

10.3 pluggy: the AI advisor

- **pluggy, the trustworthy AI advisor:** Not a generic chatbot, but a digital employee that answers from sources on the principle of "Certified first" (Chapter 6.1). Alongside the certified owner data, he uses vetted registered sources and, for coverage, currency and local answers, live connected search and map services (including Google Search and Google Maps). He receives his knowledge as Skills: the operator's own Skill and, with consent on both sides, the released original Skills of his partners (on the model: Chapter 16). This is where the division of roles of the Alliance sits: pluggy advises, plugilo makes pluggy trustworthy ("pluggy advises. plugilo makes pluggy trustworthy."). He is only one front end among many, because the same advisory intelligence works as a service inside external interfaces and systems. He does not replace specialist advice; he gives it back knowledge, time and quality, where the shortage of skilled staff hurts. This colleague knows every product, speaks every language and never tires. His tasks follow the questions of the End Users: understanding (what can the product do?), deciding (does it suit me, is it compatible with what I already have, which ingredients or allergens does it contain?) and applying (how do I use and care for it properly?). Answers come directly at the product, in the online shop, at the GTIN scan in the store or through the Browser Extension on supported product pages. Whoever takes a Card along takes him along too. Your best brand ambassador sits inside every plugilo Card.
- **A knowledge graph instead of a black box:** The foundation is the Master Data Layer as a trustworthy knowledge graph: every object a unique node with source, status and approval, every relationship explicit, changeable, correctable in real time, with clear ownership. Language models work implicitly and statistically, cannot show provenance and are not accountable for their knowledge. They do not replace the graph, they need it: they supply the language, the graph supplies the truth, and the Card is the bridge.
- **pluggy learns from three sources, each of which stays incomplete on its own:** The first are the data owners with their original knowledge about properties, applications and evidence, the factual foundation. The second are Resellers and publishers with their context: which assortment is available, which services exist and what matters in this particular place. The third is the End User with his need: what he is looking for, which requirements and constraints he has and which questions remain open. An example makes the division of labor tangible. The manufacturer knows what is in the soup. The Reseller knows which variant is on the shelf and what is on offer today. And the customer knows that his daughter cannot tolerate nuts. Good advice brings brand knowledge, Reseller context and customer need together. Technically, pluggy answers by retrieval at runtime: the sources are fetched for each question, not trained into model weights, and targeted fine-tuning remains reserved for selected tasks. "Learning" therefore means building, selecting and updating the knowledge and context base here, not continuous training through every use. Released demand signals sharpen the advice, but they never change the owner's product properties.
- **The Mirror, the AI visibility analysis (beta):** As an entry point, plugilo uses GEO tools (Generative Engine Optimization) to analyze how the large AI systems talk about a brand and its products today: which sources they draw on, what is outdated or wrong, what is missing entirely. The result is the Mirror of a company's own AI representation, the diagnosis for what it needs to do, and the simplest first step into the network.

Beyond answering, pluggy becomes the personal knowledge manager of the End User. He says in natural language what he intends to do, pluggy understands the intent, researches,

creates Cards, sorts them into Stacks, adds tags, keeps the collection current. "I am planning a family trip to Japan" turns into a project Stack with destinations, accommodation, flights, activities and documents, ready to continue and to share. Three rules apply without exception: every answer carries its source, every automatic action is visible and reversible, the user keeps control. A fourth is demanded by the market itself: pluggy is always recognizable as an AI agent, discloses its knowledge base, names uncertainty and gaps in knowledge, and hands over to a human or to the responsible body in safety-critical, legal, medical, or unclear cases; 72 percent of customers want to recognize when they are talking to an AI (Chapter 1). The path there: a product roadmap in five stages, today's chat, effortless collecting, the self-organizing Wallet, proactive prompts, commissioned actions (separate internal paper; status discipline per Chapter 17).

10.4 plug-it and the distribution services

The Connect Layer follows a simple principle: one button, the plug, does the same thing everywhere. It takes the Card along, one click, and the binding object sits in your own collection, in sync. Everything else happens on the Card, visible only where it makes sense: following (subscribing to updates) and connecting (mutually confirmed relationships, only on person and company Cards). **For the customer it is one button, for the network it is three depths of relationship: taking along, following, connecting.**


- **plug-it, the connection at the place of decision:** As a widget or plugin, it connects websites, shops, portals, and media to the verified data without an IT project. Around 80 percent of all online shops can be reached through a handful of widely used shop systems (estimate based on the market shares of the leading shop systems); every operator activates the service with a switch.
- **Universal saving, the way into the Wallet:** plug-it standardizes the saving scattered across shops and sites (Section 10.2), the way PayPal once standardized paying: products, recipes, travel, real estate, vehicles, any object from any source, taken into the plugilo Wallet with one click of the plug, whole Stacks included (examples in Section 10.1). Over time, the certified original layers itself automatically over the copies, and purchase routes and sources are preserved.
- **Imprint standard:** The plug sits in the imprint of every company website. One click, and the owner-maintained company data moves into your own system; if it changes, the Card updates automatically, on request all the way into every connected database. After product and company data, personal data is to follow, in close cooperation with government bodies, certified and officially recognized.
- **Share button:** In future, the plug will sit next to the familiar social icons. One click turns any web page into a Card in the user's own collection, with no third-party feed, no distraction, and no data flowing to third parties.
- **The plug on the product (QR/NFC):** Through QR codes and NFC, the plug moves onto packaging, devices and type plates; the coming EU product passport requires exactly this data carrier on the product (Chapter 18). A technician scans the plug on the machine and within seconds has the verified manual, the training video, the line to the manufacturer and the spare part order, in his language, always current.
- **plug-it in the Store (in development, concept phase):** The same plug a Reseller uses in his online shop or on a simple homepage also works in his physical store. The End User scans the electronic shelf label, a QR code, an NFC tag or the barcode at the shelf, and the physical product lands as a Master Data Card in his Wallet: with the certified manufacturer data for which there is no room on the label, with videos, accessories and the direct line to the brand. A web view works without an installed app. pluggy advises at

the shelf, and the sales staff assemble a Stack from the same source and hand it to the customer's smartphone instead of a printed note. The Card carries its place of origin: if the customer buys from the Wallet at home days later, the store remains recognizable as the introducing point of contact. The store thus becomes as measurable as a web shop, and the customer no longer leaves it anonymously (Section 14.2).

- **Browser Extension (beta):** It makes practically every suitable web page connectable, including where the plug is not (yet) built in. plugilo recognizes brands, companies, products and identifiers in the page text, the small plug appears automatically in front of company, brand and product names, and one click opens the certified original data, from the search engine to the marketplace. The framework is clear: suitable, unobjectionable pages, safe-site filtering, full control by the user, not indiscriminately everywhere. The original data therefore reaches the customer wherever he reads, regardless of who owns the page. Brands simply recommend plugilo to their customers, and all further services come in one.
- **Follow Service (beta):** users follow companies, products, people and topic worlds; blogs, feeds and catalogs are turned automatically into Cards from certified original sources. This is the communication channel that has been missing: mass newsletters are losing attention and measurability, organic reach is eroding, and the algorithmic feeds and mass-produced AI content of the large networks distract (Section 6.4). Here, following works without spam, without an opaque algorithm, without platform dependency: the customer decides what he sees, and every partner invites its own customers.

plug-it in the Store: from the shelf into the Wallet, and the store stays recognizable

The same plug as in the Reseller's online shop and on his homepage, here at the shelf (status: in development).



Billing as everywhere: the store is a domain, the data owner pays the Touchpoint, the store receives the Reward.
Attribution of a later purchase from home is part of the ongoing concept work.

Figure 11: From the shelf into the Wallet: the plug works in the store as it does in the shop; the Card takes its place of origin along, and the store remains recognizable as the point of contact even for a later purchase from home.

The Card is the object. The Wallet is the End User's place. pluggy is the advisor. The plug connects it all to every place on the internet.

10.5 The Partner Portal: control center and return channel

- **The Partner Portal as control center:** Open to every provider, from the sole trader to the global corporation. A company's own data profile is created in minutes, imported with AI support from websites, catalogs and data sources, reviewed and approved. This is where Records, Cards, Stacks and campaigns are maintained and published, plug-it is

configured, and rights, partners, roles and the company domain are managed. Dashboards show where data is working and what demand is emerging. Data maintenance turns into market management, with the owner in control at all times.

The Portal is more than administration: the sales, self-service, onboarding and expansion platform, the one place where partners start, grow and book and configure all services, including pluggy with its sources, focus areas and places of use. **Getting started requires no major IT project** and no rebuilding of existing systems: it is enough to name the sources, website, catalogs, press and service areas. The infrastructure handles the rest: the Cards are created automatically and stay connected to their sources, and are therefore as current as the source. In perspective the Portal becomes a lean control and exception interface, while the daily work takes place in the customer's AI environment (Section 10.7).

The Portal is at the same time the return channel. **This data is the real treasure of the network.** Because what counts here is not anonymous page views but Cards, every interaction produces a structured, explainable event instead of a click signal:

- **Demand and interest:** what is saved, followed, and organized into Stacks
- **Comparison behavior:** which products are set against each other
- **Decision criteria:** which properties are asked for
- **Advisory signals:** which questions reach pluggy and which recommendations are accepted
- **Transitions:** when a user moves to the Reseller, the offer or the service
- **Behavior after the purchase:** support, documents, accessories, and follow-up needs

These signals flow back to the owners in aggregated, privacy-compliant form. They answer questions that classic market research clarifies only with a delay: what do people care about, which alternatives do they consider, where does a decision break down? They are provided through the Partner Portal as paid services, and the personal demand profile stays with the user. Chapter 16 describes how the packages are structured.

10.6 The CF1 Data Services: professional data work for all partners

Behind the infrastructure stands a data factory. CF1 (ContentFactory1), a wholly owned subsidiary of DCI AG with around 200 employees, brings together three decades of experience with product and manufacturer data at industrial scale. Its services are open to all partners of the Alliance and carry the licensed data foundation of the Level 1 Cards (Section 8.1).

- **Data creation and structuring:** websites, catalogs, data sheets and feeds become structured, source-bound data records, multilingual and standardized.
- **Harmonization and mapping:** identifiers, classifications and characteristics are mapped onto common standards, the basis for the CF1-harmonized reference data and the open Mapping Framework (Section 8.2).
- **Ongoing maintenance and quality assurance:** data holdings stay current, checked and documented, from the Guided Launch to permanently managed Source Truth within the business model (Chapter 16).
- **Content production:** descriptions, translations, media and advisory-ready content that make Cards complete and pluggy able to answer.

Important for the trust architecture: **CF1 is the group's data services company, not an external auditor.** It creates, structures and maintains data on behalf of the owners. External

conformity attestations remain with the independent testing and standards organizations (Section 8.3); CF1 makes sure the data can meet their requirements in the first place.

10.7 plugilo in your AI: the original where the work is done (target picture)

Chapter 4 described it: what used to be software becomes a capability you describe in your own AI environment. plugilo draws the consequence for itself. Instead of teaching companies yet another portal, they should be able to use their authorized original knowledge where their work will be done in future: in ChatGPT, Claude or another supported AI environment. A working day then looks like this: a marketing director starts in her AI with a short report, three product details need release, two partners are not yet activated, one campaign generates an unusually high number of contacts. She asks for details, delegates data maintenance to a colleague, has a partner information prepared and releases the checked change. The AI obtains the data via plugilo; plugilo controls which changes are permissible and records version, release and use. A colleague with a different AI sees the same released Cards and the same task status, within his rights.

The principle behind it: one service core, several accesses. Leading CRM, content, project, payment and commerce platforms already open their data and actions to external AI clients via the Model Context Protocol; rights, data model and evidence remain in the business system.³⁰ In the same way plugilo intends to provide Cards, relationships and controlled services through an open access, while original knowledge, rights, releases, delegations and traceable use remain anchored in the network. Three kinds of work are envisaged: interactive requests, recurring jobs such as a weekly report or a maintenance reminder, and event-driven flows, for example when a published Card change informs the connected partners. The realistic promise is “no daily portal duty”, not “no set-up”: connecting the account, choosing the business unit and authorizing third-party systems remain one-time work. The Partner Portal thus gradually becomes a lean control and exception interface, and Solution Partners take over set-up, mapping and process operation as a concrete service. Status: target picture in concept phase (Chapter 17).

11. Real time, governance and rights: what “always current” really means

A promise like “always current” is worth only as much as the technical conditions behind it, and this chapter names them, including the limits.

The claim that the original is current everywhere is the core of this paper, and for that reason it has to be precise. On plugilo-controlled surfaces, real time can be enforced technically through live retrieval or live rendering. In external systems, currency depends on their integration mode: cache, index, query cadence, event processing and contractual obligation. A Master Data Record changed at 9 a.m. is current at 9 a.m. in the Wallet and on every plug-it surface; in a third-party vector index it is current only once that index has processed the event. Anyone who conceals that promises something he cannot deliver. Table 4 and Figure 12 therefore show four guarantee levels, and every statement in this paper about currency can be assigned to one of them.

Level	Environment	Defensible statement
Level A	plugilo-controlled surfaces: Wallet, Partner Portal, Card rendering	Live source and presentation are technically enforced and audited
Level B	Contractually integrated partners with plug-it or API	Currency, mandatory fields, rights and reporting are agreed contractually and technically

Level C	External AI systems with a voluntary connector, for example via MCP	The current original is offered; actual use and presentation are decided by the host
Level D	Open internet, independent language models without integration	plugilo improves findability, reference and counter-source but does not control the answer

Table 4: Four guarantee levels: what is enforced on plugilo's own surfaces can be agreed with partners and, with external AI systems, only offered.

Four guarantee levels: enforced, agreed, offered, not controlled

Level A	enforced	plugilo-controlled surfaces: Wallet, Partner Portal, Card rendering Live source and presentation are technically enforced and audited
Level B	agreed	Contractually integrated partners with plug-it or API Currency, mandatory fields, rights and reporting agreed contractually and technically
Level C	offered	External AI systems with a voluntary connector, e.g. via MCP The current original is offered; the host decides on use and presentation
Level D	not controlled	Open internet, independent language models without integration plugilo improves findability, reference and counter-source but does not control the answer

In this paper, "live" means levels A and B. "Target picture" means levels C and D.

Figure 12: "Real time" applies only where it is enforced or agreed; with external systems the original is offered, not forced.

The levels determine the language of this paper. "Real time" is claimed only for Level A and Level B. For Level C the rule is: the current original is offered, and actual use and presentation are decided by the host. For Level D: plugilo improves findability, reference and counter-source, but does not control the answer. Language models that are not integrated cannot be forced to use or display the original; the path there runs through the open connections from Chapter 7 and through the weight of the Alliance, not through assertions. **What this paper calls the "target picture" is Level C and Level D. What it calls "live" is Level A and Level B.**

11.1 Ten technical minimum requirements for a living original

For Levels A and B to deliver what they promise, ten minimum requirements apply. Identity stays persistent, because Card ID, object ID and canonical address do not change with every edit. Every relevant change is versioned and receives a version, a time, a source, a responsible party and a status. Fields such as valid from, valid until, expired, revoked and replaced by control validity across the lifecycle. The Card is rendered live wherever possible or retrieved fresh at the moment of use. Webhooks, event streams or delta feeds report changes and recalls. For caches, lifetime and invalidation are defined bindingly and measurably. The presentation contract protects the mandatory fields source, owner, quality, version and return path from being removed. Permissions authorize use by domain, region, target group, partner status and purpose. An audit trail logs system, domain, Card, version, time and action traceably. And if the connection is missing, the fallback applies: the last valid version is marked as such or the display is stopped. These ten points are the standard against which every plug-it integration and every partner agreement is measured.

For acting agents an eleventh requirement is added, the action contract per process: reading, recommending and acting are separate permissions. A price information does not authorize an agent to grant a discount; budget limits, exceptions, human approvals and stop conditions are checked technically, and source and version used remain traceably linked to every executed action. Copies must not develop uncontrolled authority: caches, backups and historical

records remain useful; what matters is reference to the source, update rule, expiry and revocation.

11.2 Revocation and recall: what control means in practice

Owner control covers provision, maintenance, responsibility, release, delegation and recall. A revocation removes the official status and ends further live use within the released scope; a recall, a safety notice for example, is reported as an event to all connected surfaces and overwrites the older version there as the authoritative live state (Chapter 8.1). Historical events remain traceable within the permitted scope, because an audit trail without a past would be no audit trail. And where a system is not integrated, Level D applies: revocation governs future use in the network, and no one can guarantee the retroactive deletion of external copies (Chapter 21.3). Delegation is clear in both directions: the owner can delegate maintenance and releases to distributors, agencies or national subsidiaries, and can withdraw that delegation just as easily.

Real time is not a claim but an integration level. On our own surfaces it is enforced, with partners it is agreed, with external systems it is offered. That honesty is what makes the promise reliable.


12. How it all works together: the loop and the overall comparison

At the end everything comes together, because only in the interplay of the parts does the loop arise that connects data, advice and demand, and only then can the comparison with today's approaches be drawn in full.

12.1 The loop closes

The loop closes: the original goes out, trust and demand come back

forward: the original with provenance, currency and advice



Four properties are preserved all the way: provenance, currency, advice, and the link to the owner. That makes trust reproducible, and the owner learns what demand arises.

Figure 13: The loop closes: the original goes out with provenance, currency and advice; trust and consented demand come back to the owner.

The original at every point of contact: on the shelf, in the online shop, in the AI. The customer takes it along and stays connected to brand and Reseller. And what he does there arrives as a signal with whoever maintains the data.

Card and pluggy carry the knowledge to the places where it is needed: to the AI systems and agents, to the Point of Sale in the shop, in the store and directly at the product, and to the Point of Interest, that is websites, media and service. The goal is an open connection through the available interfaces; that does not mean that every AI system is connected today, but that the knowledge base is organized so that suitable systems and partners can use it in a controlled way. This closes the circle back to Chapter 5.2: the architecture meets each of the five requirements. Owner control in the Master Data Layer and the Partner Portal; visible provenance through "Certified first"; structured, versioned, multilingual Records including the product passport fields; distribution to every place of decision through the Connect Layer; open interfaces for all systems. And a Charter that also applies to the initiator.

In doing so, the network generates a class of signals that classic measurement cannot observe: consented, explainable decision events before, at, and after the purchase, the signal layer. Every event carries its context, that is product, category, market, contact point, and time; every link in the chain produces an explainable event instead of an anonymous click (Table 5). The volume of these signals grows with usage, not with field work, and signals from a clean, trustworthy environment are categorically worth more than signals from noisy, algorithmically driven environments. The governance is settled: signals are used only with consent, aggregated and privacy-compliant; the personal demand profile stays with the End User (Chapter 8.1).

Phase of the decision journey	Signals	What measurement partners read from them
Before the purchase	Saving, following, Stacks, comparison sets, questions to pluggy, accepted and rejected recommendations	Real demand situations, shortlist dynamics, decision criteria per category, brand consideration in real time
At the transition	Transitions from the Card to Reseller, offer, and point of sale	Channel choice, offer resonance, the path to the verified purchase
After the purchase	Manuals, accessories, warranty, service, and follow-up needs, opened in the Card	Usage phase, service needs, repurchase and accessory potential

Table 5: The signal layer makes visible the decision path that panels and clicks cannot see.

12.2 The overall functional comparison: ten criteria, one pattern

Chapter 3 showed where today's methods end. Now that the solution has been described, the comparison can be made in full, along ten criteria that make the five requirements from Chapter 5.2 concrete: owner-controlled source, product semantics, live update, AI usability, provenance, visible original, presentation control, partner release, lasting owner relationship and demand loop. What is assessed are categories of solution, not companies. A dash means that the criterion is not part of the primary purpose, one dot means partial or implementation-dependent coverage, and two dots mean strong coverage within the defined target profile (Table 6 and Table 7).

A look at today's methods reveals a pattern. PIM, MDM, and DAM create a strong internal source, but no binding presentation in third-party AI systems. GDSN synchronizes standardized trade data worldwide, but delivers no advisory object and no interface surface. Marketplace feeds come with clear obligations and reach, but with their own formats and rules per platform. Schema.org is broadly compatible, but obliges no one to adopt anything. GS1 Digital Link gives the product a persistent reference, but no complete knowledge and advisory model. An open API is flexible and potentially live, but the recipient decides. RAG and vector stores ground a specific application, but create exactly there another copy with its own index and ranking. MCP makes resources and tools accessible in a standardized way, yet the AI host decides on connection, selection, consent, and interface.³¹ C2PA, Verifiable Credentials, and PROV are strong provenance building blocks without product knowledge and without distribution. The Digital Product Passport brings the transparency required by regulation, but is not intended as an advisory layer for AI. **All of them are building blocks. The gap lies in their connection.**

Solution	Owner source	Product semantics	Live update	AI usability	Provenance
PIM, PXM	●●	●●	●	●	●
GDSN	●●	●●	●●	–	●
Schema.org	●	●●	●	●	●
GS1 Digital Link	●	●	●●	–	●
RAG, vector store	●	●	●	●●	●
API, MCP	●	●	●●	●●	●
C2PA, VC, PROV	–	–	●	–	●●
plugilo target architecture	●●	●●	●●	●●	●●

Table 6: Source, semantics, currency, AI use, provenance: strong building blocks, each within its purpose.

Solution	Visible original	Presentation control	Partner release	Owner relationship	Demand loop
PIM, PXM	–	–	●	–	●
GDSN	–	–	●	–	–
Schema.org	–	–	–	●	–
GS1 Digital Link	●	–	–	●	–
RAG, vector store	–	–	●	–	●
API, MCP	–	–	●	●	●
C2PA, VC, PROV	●	–	●	●	–

plugilo target architecture	●●	●●	●●	●●	●●
-----------------------------	----	----	----	----	----

Table 7: Visible original, presentation, partner release, owner relationship, demand loop: this is where the open gap lies, and this is where the target architecture begins.

The result is clear, and it should be read honestly. Existing standards and platforms are strong building blocks, yet no category meets all ten criteria at once, and for all of them the gap lies in the same half: visible original, presentation control, partner release, owner relationship and demand loop. The plugilo target architecture is the only one in this matrix designed precisely for that combination. The values here describe the target architecture, not an independently audited market leadership; they have to be proven through product functions, tests and reference integrations, and Chapter 19.7 names the reference case for that. **Under the target profile defined here, plugilo is the most complete solution architecture, because it does not merely distribute data but keeps the controlled original visible, usable, updatable and connected to the owner.**

All building blocks exist. What is missing is the connection: a layer that holds source, currency, visible original, rights, relationship and return channel together. That layer is the target architecture.

13. The trust architecture: why trust becomes scalable

To close the solution part, the question that stands above everything else: what makes a piece of information trustworthy, how does the release granted by the brand itself become a trust signal, and how is that trust preserved when knowledge travels through Resellers, publishers and AI systems to the End User?

Trust does not come from correct information alone. It comes about when people can recognize where a piece of information originates, who stands behind it, how current it is, and whom they can turn to with questions. That is exactly what the infrastructure of this paper is built on, and this chapter names the architecture behind it. The need has arrived in the market: at ECR iN MOTiON 2026, Paul Vetter, General Manager Henkel Consumer Brands, set out three theses for the coming years in his keynote "Future? Ready! Ready together", in essence: data and information are the new gold, trust becomes more important than data alone, and success comes from strong partnerships; the publicly documented core message of his keynote is that future readiness needs collaboration, shared responsibility, and the exchange of knowledge across company and community boundaries.³² That applies to every brand and should be read as a market trend. Anja Stolz, CMO of R+V Versicherung, puts it in the formula that AI scales business models, but not trust.³³ The data belongs to the owners. The trust comes about when the original is recognizable as the original. And the partnership is the Alliance.

13.1 What trust is: three preconditions

Trust research has agreed for decades: people trust a source when they attribute three qualities to it. Competence (does the brand know what it is talking about?), integrity (are the statements genuine and traceable?) and benevolence (is someone there for me when I need help?). This model by Mayer, Davis and Schoorman is the most cited in the discipline.³⁴ Trust matters most when someone has to decide although he cannot check everything himself; Niklas Luhmann therefore describes trust as a mechanism for the reduction of complexity.³⁵

Applied to the purchase decision: the End User does not have to compare hundreds of sources if he can identify the original clearly and reach the party responsible for it. Table 8 shows how the infrastructure delivers all three preconditions technically.

Prerequisite	The End User's question	How the infrastructure delivers it
Competence	Does the brand know what it is talking about?	Current original data, documents and expertise of the data owner in the Master Data Card
Integrity	Are the statements genuine and verifiable?	Provenance, release, version history and a clearly stated audit status; the blue seal as the visible anchor
Benevolence	Is someone there for me when I need help?	Direct dialogue via pluggy, support and the lasting connection to the brand in the plugilo Wallet

Table 8: The three preconditions of trust and how the infrastructure delivers them.

13.2 Seven levels on which trust is built

- **Trust begins with the original.** The End User must be able to recognize who owns the information, whether the statement really comes from that source, when it was last updated, whether it still applies, and whether it was only released by the owner or additionally reviewed independently. The Master Data Card turns anonymous information into an attributable statement: not "somewhere on the internet it says", but "this information comes from the responsible data owner".
- **Trust needs currency.** A piece of information may once have been correct and still be wrong today; a static copy is therefore not enough. The Card stays connected to its source, and changes made by the owner become visible at the connected points of contact. That is the difference between a copy and a digital original. Research on supply chains shows the same pattern: transparency and traceability strengthen trust in products, and the accuracy and traceability of the information weigh more heavily than its sheer volume.³⁶
- **Trust needs someone accountable.** People trust information more readily when someone visibly stands behind it. The connection to the owner is therefore more than a source reference; it gives the End User a responsible address. Here a brand stands behind the answer, shows its data, explains it and remains reachable. The End User no longer experiences the brand as the anonymous sender of an advertising message, but as an approachable partner who does not leave him alone with questions and problems.
- **Trust arises in dialogue.** The possibility of asking a question changes the relationship: passive content becomes active advice. A meta-analysis of 144 studies with 57,952 participants shows that digital interactivity strengthens brand loyalty and purchase intention through trust and emotional attachment.³⁷ pluggy makes the knowledge of the brand approachable at any time, at the Point of Sale and at the Point of Interest, before the purchase, during use and with later problems: a consumer hotline in your pocket that knows the entire released knowledge of the brand.
- **Trust needs separate proofs.** Owner release, rule-based data quality checking and independent certification are three different proofs, and they must never merge into one vague overall seal (Chapter 8.3). Data quality itself follows a recognized benchmark: ISO/IEC 25012 describes fifteen characteristics, among them accuracy, completeness,

consistency, currency and accessibility; the Card does not hide the quality status in an overall grade but explains it by dimension and source.³⁸

- **Trust needs authorized partners.** The End User rarely meets the original at the brand itself, but at the Reseller, in the marketplace, at the publisher. The brand's release turns any data recipient into an officially connected point of contact; identity, scope of release, country, domain, location, product area and validity become visible (Section 13.4).
- **Trust needs a network with the same rules for everyone.** Uniform rules and verifiable roles preserve the meaning of the Card across many partners and countries. Reach is created without losing the identity of the source; that is the difference between a seal one brand awards and a standard a network carries (Section 13.5).

Trust needs a recognizable original. Current data. Provable provenance. A responsible owner. And a direct connection for every question.

13.3 The blue seal as an anchor of trust

A seal only works over time if its meaning is unambiguous and verifiable. Research on digital trust seals shows that it is not the symbol that creates trust, but the institution behind it, its review and its traceable statement.³⁹ The blue seal of the Master Data Card (from Level 2 upward, Chapter 8.3) therefore tells the End User exactly three things. The original source is confirmed. The Card is connected directly to the data owner. Provenance, quality and updating are traceable. The star at Level 5 (Managed Source Truth) adds a fourth statement: the source is managed permanently. Just as important is what the seal does not claim, namely that all content is independently certified. The owner release is proof of provenance and responsibility, the data quality check is an additional proof, and independent certifications from DQX, TÜV, VDE or other bodies are shown separately in the Evidence Passport. **This separation is exactly what makes the seal credible.**

Mark	Confirms	Explicitly does not confirm
Verified Original (blue seal)	Identity of the owner, provenance, release, version and active connection	No blanket independent certification of all content
Data Quality Checked	Check against named rules, with scope and date	No product safety or performance certification unless stated
Independent Certification	Check by a named independent body for a named subject	No statement outside the stated scope of the audit
Authorized Card Point (planned)	Release of the specific partner for a defined use	No general endorsement of all of the partner's services

Table 9: Four marks, four clearly separated statements: every mark says what it confirms and just as clearly what it does not.

The separation also has a legal side. The mark "Verified Original" confirms the identity of the owner, provenance, release, version and active connection, expressly not a blanket independent certification of all content. A data quality check confirms an examination against named rules, with scope and date, not product safety. An independent certification applies only to the named subject. And the status of an authorized partner confirms the release for a defined use, not a general endorsement of all of the partner's services. Precedents show that

this precision is what makes a mark valuable: Content Credentials lead from a visible symbol to verifiable provenance details, and authenticity programs in retail gain their value from a clear process, a named examiner and a traceable result.⁴⁰

13.4 The brand's release as a trust signal: the Authorized Card Point

When a brand determines which Resellers, publishers and points of sale may use its Master Data Cards, the release itself becomes a trust signal. The partner then shows not just any product description, but an original provided by the brand and released for exactly this use. The target picture for this is called the Authorized Card Point (planned): a verified Reseller or publisher that the brand has released for specific Cards at named Points of Sale or Points of Interest. The release is object-specific, verifiable and revocable, and it does not concern ownership of publicly available facts but the use of the live-connected Master Data Card, of the brand seal, of the advisory access and of the official relationship functions. That boundary is what keeps the model clean under competition law: the EU rules on vertical agreements recognize selective distribution systems under conditions and permit quality criteria, but they set limits on restrictions of online distribution; plugilo therefore does not try to monopolize product facts, but protects and grants the original, the mark, the advice, the partner status and the relationship.⁴¹ Brand releases are reviewed under competition, distribution, trademark and data protection law before rollout.

The process is laid out in five steps. The partner registers his company via a business domain and documents responsible persons, roles, locations and the domains in use. He applies for the release for a brand, for specific assortments and for named places of use. The brand or a delegated administrator checks identity, suitability, rights, presentation quality and compliance. The brand grants a release defined by territory, subject matter and time, and the system provides the Card, pluggy and the partner mark. plugilo logs use and status, and the brand can renew, restrict, pause or revoke releases. For the partner this is not a technical integration but a visible status: a publicly verifiable profile, an embeddable badge linking to the audit status, and a display of the authorizing brands and areas of use. This status arises from criteria and cannot be bought alone. Precedents have existed for a long time: authorized service partners receive original parts, tools and training only against verifiable qualification and ongoing review; official jewelers confirm the authenticity of a watch and activate the warranty; authenticity programs in online retail make the check visible through a badge, a code and a documented audit status.⁴⁰ **Trusted data does not bypass the channel. It makes every authorized partner a better and more trustworthy decision point.**

For Resellers and publishers, the release is a claim in both directions: they can request a release from the brand and thereby make demand visible in the Channel, and brands can invite preferred partners specifically. What runs today and what is an expansion stage is stated in Chapter 17: partner releases for Cards and Skills in the Network and in Channels (Chapter 16) run as beta, the public Authorized Card Point status with badge and audit profile is planned.

The release applies online and in the store alike. A Reseller authorized by the brand shows the same released original on his website, in his shop and via plug-it in the Store at the shelf (Section 10.4); the customer recognizes the same marking and the same connection to the brand in all three places.

13.5 Scalable trust: the relationship survives across partner boundaries

The most important strategic idea is this: plugilo does not simply scale information, it scales the trust relationship between owner, partner and End User. Normally the brand's control ends as soon as product information is copied and passed on through retail, media, publishers or AI systems. In the network, by contrast, four elements stay connected to one another (Figure 13

in Section 12.1): the data owner maintains the original and controls the releases. The real-time Master Data Card carries provenance, currency, review status and seal. Resellers and publishers bring proximity and reach and add their context. The End User decides at the Point of Sale and at the Point of Interest with the original, and through questions, following and renewed contact the circle closes at the owner.

Resellers and publishers bring this trust to where decisions are made; the brand supplies the original it is accountable for; plugilo makes sure that provenance, quality and the connection to the owner are not lost along the way. In this way trust in a single brand becomes a trust network across all partners and points of contact, reproducible worldwide, not as an advertising promise but as a traceable structure. **The brand is no longer just a sender. It becomes reachable at any time.**

PART 4 · THE VALUE

Now that the building blocks are known: what each group gains, why this position in the market has no counterpart as far as we can see, and how the business model lets every participant grow with his own success.

Part 4 at a glance: Every group gains immediately and on its own account. The position between data owners and AI systems is, as far as we can see, unoccupied. And what is paid for is usage and success, not ad space. Whoever contributes grows with the network.

14. What each group gains

This chapter shows group by group which benefit arises immediately, because nobody joins a shared infrastructure without an advantage of his own.

The network holds because every player has an immediate benefit of his own, and because these benefits reinforce each other. The same applies to every group: the infrastructure is technically in operation (status definitions: Chapter 17), and your own truth can be connected with your own customers and partners from the first day. The services referred to here, meaning Cards and Stacks, the plugilo Wallet, pluggy and plug-it, are described in detail in Part 3; Table 10 on the next page assigns them to the groups.

Group	What they gain	Ch.
Brands, manufacturers, OEMs	Control over their own truth in every AI answer, measurable demand, qualified contacts, pluggy as brand ambassador at every partner	14.1
Mid-sized companies and hidden champions	AI presence based on their own data, without a digital budget and without depending on a gatekeeper	14.1
Distribution and suppliers	Cross-brand advice across the full assortment, faster activation of new products, more valuable portals	14.2

Retail and specialist businesses	An always-available, brand-compliant advisor against the skilled-labor shortage; the customer returns to the business instead of an anonymous search	14.2
Publishers, media, associations, experts	Content stays findable and compensated as knowledge objects; curation becomes a product again; associations carry industry intelligence	14.3
End Users (Consumer, Prosumer, Business)	Answers with visible sources, a decision memory in the Wallet, and full control over their own demand profile	14.4
Business customers and procurement	Provenance-secured real-time data for Agentic Commerce, inquiries and projects traceable at the object	14.5
Government and the public sector	Digital sovereignty in practice, verifiable data quality for the economy	14.6
AI providers and technology partners	An authorized source layer that lowers error rates and reduces liability risks	14.7

Table 10: Every group gains immediately and on its own account; the sections of Chapter 14 provide the reasoning.

14.1 Brands, manufacturers, OEMs, mid-sized companies, and hidden champions

The Trusted AI Advisory Alliance pursues a long-term goal: original knowledge should remain under the control of its owners in the AI world as well. For manufacturers, however, the value does not begin in the future, because plugilo already creates new opportunities for sales, advice, and customer acquisition today. Product information that used to be content becomes an active advisory and sales channel. The value arises on three levels:

- **Today, the business value:** sell more, advise better, support the trade, and win new customers. pluggy turns existing company and product knowledge into active advice around the clock, first in the plugilo Wallet, on request on your own web presences.
- **Tomorrow, the network value:** equip distributors and Resellers with your own original knowledge, open up new Channels, measure demand and generate qualified contacts. That is exactly the scaling the business model provides through Cards and Skills in the Network and in booked Channels (Chapter 16).
- **Long term, the strategic value:** control your original knowledge, secure data sovereignty, influence your own AI presence, and help shape an open standard. The big mission remains fully intact; it simply comes after the immediate value.

Brands control their truth in AI advisory and thereby lay the foundation for the most reliable advisor for their own products and services. Verified data takes effect across the entire partner network, in specialist retail, in shops, on portals and in the media. **The brand determines which truth AI uses.** This protects the brand from false information in the AI era and keeps the information about products and services consistent across all sales channels and partner networks, including hybrid scenarios with distribution, specialist retail, online shops, and marketplaces. In doing so the brand strengthens its sales partners instead of bypassing them, up to transparent, performance-based support for the Resellers who advise on its products. The partner receives more than data. He receives a colleague, because pluggy works on the truth of the brand and stands beside every trade partner. How pluggy learns from brand knowledge, local context and customer need is described in Section 10.3. The operational effect shows immediately, in better advice at the point of sale, in measurable demand and in qualified contacts. Through the plugilo Wallet, every customer contact also

reaches beyond the first conversation. This creates the best brand ambassador a brand has ever had, and every customer can take him along: from the first question through the purchase to service and repurchase. A fleeting click becomes a lasting relationship.

For smaller and specialized providers, the same applies, only with a different lever. They gain what used to be reserved for large digital budgets: correct, findable AI presence based on their own data, without dependence on a single gatekeeper. The free entry lowers the barrier to zero, and visibility is decided again by performance and data quality, not by the size of the advertising budget.

Winning new customers. pluggy does not only serve existing customers and partners. Manufacturers additionally open up new audiences: in the Partner Portal they select relevant audiences and suitable Channels, such as trade media, publishers, communities, or influencers; on request, plugilo takes over their approach and activation (Section 19.6). As soon as a Channel is available, the approved manufacturer knowledge is used there to match specific customer needs. Touchpoints and Advisory Matches make this demand measurable and at the same time create a new revenue source for the channel operator.

14.2 Distribution, retail, and advice-intensive businesses

Distributors combine assortment depth with scalable, cross-brand advisory. That means better answers across all brand worlds, less manual searching through PDFs and portals, faster activation of new products, and a seamless transition to inquiry and quotation. Their portals become more valuable for manufacturers, Resellers, and business customers in the process, and distribution moves from data supplier to activation partner for the channel.

For Resellers, studios and specialist businesses, the system answers the most pressing problem in the industry: too few skilled staff for ever more complex product questions. An always available, brand-compliant digital advisor answers customer questions immediately, in the store, on the website and at every further point of contact. That improves the quality at the point of decision, generates qualified contacts and reduces research and support effort. Saved information leads the customer back to the business instead of into an anonymous search. Along the way, the company's own website maintains itself: with the original sources connected, product and information areas stay current automatically, maintained by the brands and delivered at the business. One basic principle is not negotiable here. **The system strengthens advice instead of bypassing it, and specialist retail remains the commercial path.** And through the plugilo Solution Partner Program, plugilo itself becomes a product for system houses, IT Resellers and agencies: they sell it to their business customers, roll it out, support it, and earn on an ongoing basis (Section 16.3).

For brick-and-mortar retail, plug-it in the Store solves the oldest problem of the sales floor: the anonymity of the store visitor. Customers see the product, do not find enough information at the shelf, leave the store without buying, and the retailer loses every contact and every attribution when the purchase is made online later (research offline, purchase online). With the plug at the shelf, the customer takes the original along, stays connected to brand and store, and for the first time the retailer sees which products are scanned often but not bought: the conversion of the shelf, as an aggregated signal without personal data. The sales staff advise from the same current source as the customer and hand the customer their advisory Stack. Advice-intensive assortments benefit first: kitchen and home appliances, consumer electronics, bathroom and DIY, furniture, bicycles, garden. The status is stated openly: in development (Chapter 17).

14.3 Publishers, media, event partners, associations, chambers of commerce, and subject experts

Media gain a new role as the trustworthy context of the AI era. Articles, reviews, guides and events remain findable, savable and connected to their source as structured knowledge objects, instead of disappearing after publication. **Reach turns into paid advisory context.** Publishers earn a share because they supply the access to the target audiences, and for advertising partners, interest and demand become reliably measurable. Curation itself becomes a product again: editorial rankings, topic worlds and recommendation Stacks are followable, permanently maintained works. Readers then no longer subscribe to issues alone, but to the selection expertise of the house. The plugilo Wallet creates the direct channel beyond newsletter and platform algorithms. And every curated list can be taken along as a whole: whoever takes a magazine's best-of list along carries the editorial selection with him, with every Card in it, permanently connected to the source.

Associations extend their role digitally and become carriers of industry intelligence: member news, awards, training, events and projects become structured and permanently usable. They organize Chapters, collective data standards and a shared voice toward the global platforms. In doing so they protect the interests of all market participants. Through their expertise, standards, classifications and industry taxonomies feed into the mapping, and the truth layer becomes a jointly supported industry standard rather than one player's isolated solution. The neutral association role and commercial activation stay cleanly separated, and services are compensated fairly and transparently.

For self-employed professionals, creators and subject experts, the basis shifts as well: anyone who works with recommendation and expertise gets verified facts as a foundation. In a world of increasingly machine-generated arbitrariness, demonstrable trustworthiness becomes the hardest currency, and provenance-secured data supplies the proof. Structured, shareable Cards and Stacks extend every recommendation beyond the single post or advisory moment.

14.4 Consumer End Users, prosumers, and Business End Users

Today the customer is himself the integrator of the market. He does the work nobody does for him: searching, comparing, finding things again, keeping track, spread across dozens of tabs, apps and screenshots. Finding something again is often harder than finding it in the first place, and what people look for is rarely the best specification but what fits their own situation. Nobody feels this more than the prosumers, the demanding, near-professional End Users, and the Business End Users who research, compare, and procure for a living. They are exactly whom the network addresses first. The network takes this work off his hands. It delivers clear answers with a visible source, savable performance comparisons, structured questions and a continuing connection to the provider and to specialist retail. Mobile app and web portal together become the plugilo Wallet, the personal knowledge and demand wallet, in which the originals come together as lasting, continuously updating assets. Scattered finds turn into an ordered decision memory, and the demand profile stays under his full control. The Wallet accompanies more than the purchase. It accompanies the whole life cycle, from planning through comparison, decision, ownership, service and warranty to repurchase. In perspective, the Wallet also becomes the place of the transaction. What is saved and compared today can in future be requested, ordered and bought directly from it, on the basis of the Trusted Data and secured through the planned blockchain anchoring (Chapter 17).

At the center is the End User with his demand, and across the whole network he is served like a king. He gets what he needs, at the right time and in the right place. This is the deliberate counter-position to information overload and fake news, to bot content, opaque algorithms

and the remaining side effects of the social media portals. **This is about truth with a source instead of noise.**

Daily use comes from discovery, not from an obligation to file things away, because nobody opens an app “to save something.” The Wallet is a calm, self-determined space for discovery made of marked original sources: new items from brands the user follows, recommendations from friends and experts, topic worlds from the media, with no flood of advertising and no outside algorithm. Whatever appeals lands in the user's own Stacks with one click. Discovery is the occasion, the ordered collection is the result, and together they train the personal demand profile along the way.

14.5 Business customers and Agentic Commerce

Under the term Agentic Commerce, a new pattern is establishing itself worldwide: AI agents research, compare, fill shopping carts, collect quotations and trigger orders, first assisted, increasingly on their own. This raises efficiency considerably. But it sharpens the data problem at issue here in the hardest way. **A person who reads an incorrect detail can pause. An agent orders it.** Errors then no longer scale in opinions but in orders, contracts and supply chains.

The value of the infrastructure shows most clearly here: agents need machine-readable, provenance-secured real-time data, and the Master Data Layer supplies it. Purchasing organizations record their requirements in the company's own demand profile, while positions and supplier relationships live in the Business Graph. Inquiries, quotations and projects remain traceable as shared Stacks from the first interest to the closing, discussed directly at the object and on the basis of binding data. **Anyone who wants to be found, compared and ordered correctly by agents in future needs his data in exactly this form.** With the planned blockchain anchoring, this becomes more than findability: provenance-secured data on which people and authorized agents will be able to trigger transactions safely, from the quotation to the binding order, without a break in the media chain (Chapter 17).

14.6 Government and the public sector

The public sector gains lived digital sovereignty instead of declarations of intent: an industry-backed infrastructure under European rules, with no need for subsidies, and with better information quality for citizens. Authorities and institutions can provide consumer guidance, official data, and educational and health information themselves as provenance-secured sources. In doing so they raise the quality of AI-supported information overall.

14.7 AI providers and technology partners

The AI providers themselves gain as well. Every serious system gets better when it can access authorized, structured, current sources: fewer wrong answers, lower liability risk, higher user trust. The interfaces are open. And the Alliance does not see itself as a counter-model to existing AI systems, but as the reliable source and reference layer above them, the one everyone points to.

Suppliers bring truth. Partners bring reach and advice. Associations bring trust and network. Customers bring demand. plugilo connects it all. — Brands bring trust. Partners bring reach. Users bring demand. plugilo is the network.

15. Market and position

This chapter sets out what the market delivers today and which layer above it is missing, because right now the large AI providers are building their own access routes for product data.

15.1 The market at a glance: who delivers what, and what is missing

The market for AI in commerce is moving, but structurally one-sided. Shopping assistants of the large platforms serve their own platform, not the market. Guided selling tools usually advise without a cross-owner source and without Demand Intelligence. Data enrichers deliver data but no network, and generic chatbots have neither owner control nor proof of provenance. plugilo occupies a different position that, as far as we can see, has so far been unoccupied: owner-controlled truth, advisory, distribution and real demand signals in one infrastructure layer between products, providers and customer decisions. Price comparison stays deliberately out of scope. **The network compares performance, fit and trustworthiness; it makes advice better, not prices lower.** The standards landscape and the trust model follow in Chapter 8.

A look at the structure helps to order the market. AI systems are moving from supporting the purchase decision to executing it: assistants build shortlists, agents compare and order, generated answers replace search results. Four questions decide who captures value in this commerce. Where does trustworthy product data come from? How is the new commerce measured? Where does the End User actually encounter this commerce? And who owns the interface, the signals, and the customer relationship? The data world of product data pools and content syndication answers the first question, the measurement world of panels, retail data, and market research the second. The third and the fourth question are open in both worlds. **Whoever operates the interface to the End User sits at the source of the signals and holds the customer relationship. That is exactly where plugilo stands.**

The line of development behind this reaches further than commerce. AI systems learn from ever more sources: the open internet, company data and, increasingly, the data of the physical world, from machines, vehicles and sensors that are in use by the million. The more sources there are, the more valuable the function becomes that separates the reliable from the noise. This is precisely the data gateway plugilo is building, as an open gateway from all data owners to all systems and AI models. Open means here: the owners control their data, the standards are disclosed, and every serious AI system can connect. **It is a gate that the owners open, not a guard who sells access.** Four functions make the gateway concrete. Source binding keeps every statement connected to its original. Rights and approvals govern use and distribution according to the will of the data owners. The structured query lets systems obtain defined views. And open distribution interfaces bring the data where it is needed. Every answer stays source-bound and rights-controlled. In the long-term target state, platforms, database operators and website operators obtain reliable original data automatically, because the original is simpler and better than any copy. The more AI systems prefer current, source-bound records, the more outdated copies without provenance lose relevance.

Provider group	What they deliver	What is missing
Shopping assistants of large platforms	Convenient selection within the platform	Serve the platform, not the market; no owner-controlled source
Guided-selling tools	Advisory journeys for individual vendors	No cross-vendor original source, no demand intelligence

PIM and product experience systems, data enrichers	Internal data maintenance and channel output	No network, no trust levels, no AI distribution
Product data pools and data quality (incl. GS1 GDSN/DQX)	Identification, master data exchange, verified data quality	No living, advisory-ready object with media, certificates, and partner relations
Product content and syndication networks (incl. NIQ Brandbank, Icecat)	Professional capture, images, pack and specification data, syndication to Resellers	Distribute content into commerce systems; no visible, market-wide advisory and demand layer with owner trust levels
Platforms for the Digital Product Passport	Compliance, sustainability, and lifecycle documentation	The passport is one layer in the Card, not the operational product object for advice and commerce
Industry data spaces (incl. Catena-X)	Trusted B2B data exchange within one industry	Industry focus; more likely an infrastructure partner than a competitor
Generic chatbots	Language interface	Neither owner control nor provenance
Product feeds of the AI providers	Direct access for Reseller data into AI answers	Start at the Reseller feed; no neutral, cross-model layer of originals and trust
Context governance (incl. Collibra, Atlan, Palantir)	Internal enterprise context: metadata, policies, lineage, ontologies, agent access	Operate inside the organization; no market-facing distribution across company boundaries
Enterprise AI platforms (incl. SAP Joule, Agentforce, Copilot)	Agents on internal processes, data, and workflows	No authorized external owner knowledge; End Users and distribution surfaces rather than competitors
Data and context infrastructure (incl. Databricks, Confluent, Neo4j)	Transport, consolidation, real-time streams, and graphs	Technical foundation without a market-facing truth and advisory layer
Market research and retail measurement (e.g., NIQ, Circana, Kantar)	Panels, retail and End User data, measurement, forecasts, increasingly measurement infrastructure for AI-driven commerce	No interface of their own in the hands of the End User; the decision path before and after the purchase remains unobserved


Table 11: Every building block of the market exists; nobody connects owner data, external evidence, and AI advice into one layer.

This overview shows two things. First, we know the market, because every one of these groups does its part, and many of them are possible partners. Second, they all work inside the same layer, as a platform, a tool, a data pool or a feed. plugilo does not compete with any of these groups for their layer, it connects them through the open meta layer from Chapter 7, which platforms, tools, data spaces and AI models can all draw on. **Whoever works inside one layer fights for market share. Whoever places the reference above them all makes every layer better.**

We do not compete for a layer. We are the open reference above all of them: the meta level that systems, platforms, and models can all point to.

Taken together, a strategic map of the environment emerges. Source systems such as PIM, PXM, and MDM platforms (incl. Akeneo, Salsify, Informatica, inriver, Pimcore) maintain the product and company data of the manufacturers. Product content and syndication networks (incl. NIQ Brandbank, Icecat, Syndigo, 1WorldSync) capture, enrich, and distribute product content at industrial scale, from images and pack data to technical specifications, to Resellers and commerce systems. Access, integration, and streaming providers (incl. Starburst, Boomi, Confluent) transport it. Data and AI infrastructure (incl. Databricks, Microsoft Fabric, AWS, Google Cloud) processes it. Governance platforms (incl. Collibra, Atlan, Palantir) organize internal enterprise context. GS1 Digital Link, MCP, A2A, ACP, UCP, and the Digital Product Passport secure identification and interoperability. Enterprise AI platforms (incl. SAP Joule, Salesforce Agentforce, Microsoft Copilot, Amazon Bedrock, Google Vertex AI, ServiceNow, Atlassian Rovo, Glean) are the agent surfaces of companies. And sovereign data spaces (Catena-X, Gaia-X, International Data Spaces) govern exchange between organizations.

The ecosystem map: occupied categories and the unoccupied layer



Every category is occupied and remains a partner. Unoccupied is the layer that makes released owner knowledge visible, mobile and advisory-ready across the market.

Figure 14: Every category of the ecosystem is occupied and remains a partner; only the layer above, which connects everything, is open.

Each of these categories is occupied, and many of their providers are natural source, infrastructure, or distribution partners. What remains unoccupied is the market-facing layer

that connects them: owner-controlled Source Truth becomes the portable Master Data Card, trusted AI advisory, market-wide distribution, reusable Cards and Stacks, and real demand intelligence. That individual providers are extending their offerings toward AI readiness and Agentic Commerce changes nothing about this, as Section 3.5 shows: they manage data inside the customer's organization, while plugilo makes approved owner knowledge visible, portable, and advisory-ready across company boundaries.

An honest assessment starts with the fact that the individual building blocks already exist. GS1 operates a globally established infrastructure with GTIN, GDSN and GPC, and checks master data with GS1 DQX. More than two million companies are connected to GS1, Verified by GS1 provides basic information on more than 300 million products, and more than one billion products carry a GS1 barcode. Digital Product Passport platforms document sustainability, origin and compliance. CEN/CENELEC published the first European DPP standards in 2026, and ISO/IEC and UNECE are working on further standards.⁴² Catena-X is building a data space for the automotive industry, a possible infrastructure partner rather than a competitor. PIM systems maintain product data, and AI shopping is growing. Classification, data pools, quality checks, product passports and AI product advice are not unique in themselves.

What is unique, based on our own market observation (as of September 2026), is the combination: an owner-controlled, versioned original for every object, with visible trust levels, measurable quality and external evidence. This original can be taken along everywhere as a living object, distributed model-independently to AI systems, partners and End Users, and activated in advisory, support and commerce. Based on our market observation too, no provider today combines all of this in one offering. The category that plugilo occupies with it is called Trusted Product Truth Infrastructure for AI: the independent trust and connection layer between data owners and AI systems. It cannot be defended by technology alone, but by the network behind it (Section 16.4).

15.2 Market validation: AI providers are opening their product feeds and need this layer

The strongest confirmation of the model comes from the market itself. OpenAI lets Resellers supply product information through specified feeds, expressly including through appointed service providers, and is expanding the Agentic Commerce Protocol into a connection layer for product search and discovery. Microsoft offers the same access through the Copilot Merchant Program and expressly describes trustworthy product truth as the foundation of Agentic Commerce. Google combines structured product data with Merchant Center feeds and, in January 2026, launched the Universal Commerce Protocol (UCP) together with Shopify as a further open commerce standard. OpenAI has broadened its initial checkout focus toward product search and discovery.⁴² At the same time, agents are taking over the conversation. OpenAI is expanding ChatGPT Ads with connections to HubSpot and Shopify, including a format that leads users directly to a brand's AI agent, and Google is testing AI Contribution, a payment for publishers whose content feeds into Gemini answers.⁸ So the world's largest AI providers are building their own access routes for product data right now. They need complete, current, structured data with usage rights, unambiguous identities and trust signals.

What all of them lack is an independent, cross-model layer of trust and original data in which the various sources of a product are ranked in a hierarchy anyone can follow. Today's programs start with the Reseller feed. plugilo starts with the data owner: the Reseller supplies the offer, the price and the availability, the manufacturer supplies the binding product truth. That separation is decisive, because Resellers describe the same product differently, incompletely or incorrectly. The Master Data Card is the shared reference, while Reseller offers and publisher content stay connected in the Connect Layer. **The principle is manufacturer truth instead of Reseller feed.** An honest assessment includes both sides. The

same programs are validation and disintermediation risk, because brands and Resellers can also supply feeds directly or through established service providers. plugilo therefore concentrates on what a single feed cannot deliver: the manufacturer-authorized layer of original data and trust, with evidence, conflict records and a demand return channel, across models and independent of platforms, rather than tied to a single program.

15.3 The Master Data Trust & Evidence Network: partner roles ahead

This architecture leads to the strategic next step: an open network of trust holders inside the Trusted AI Advisory Alliance, the Master Data Trust & Evidence Network. plugilo does not ask an external organization to certify the whole Card across the board; the Card is broader than the scope of most individual standards. Instead there are clearly defined roles. Standards partners (GS1, ECLASS, ETIM, IEC, DIN/DKE, bodies close to ISO) supply mappings, identifiers and classification models. Data quality partners (the group's own CF1, Section 21.1, and GS1 data quality services) check completeness, normalization and traceability. Testing and certification partners (the TÜV companies, the VDE testing institute, DEKRA, UL, SGS, Bureau Veritas, Intertek) confirm defined product, safety and compliance modules. Maintenance partners such as CF1 keep the source permanently up to date on behalf of the owner. And AI distribution partners, meaning models, agents and commerce platforms, consume owner-controlled, attributable product truth and return demand signals where that is permitted.

plugilo's role is defined precisely: the competitively neutral layer for orchestration, evidence and distribution between data owners, standards organizations, testing bodies, partners and AI systems. Neutral in competition, not in stewardship (Chapter 21). The first cooperation goal is modest and immediately feasible: show authorized certificates and attestations on the Card, validate them against original documents and registers, and transport them digitally. This is not about a new plugilo certification that replaces existing institutions. The joint Proof of Concept with GS1 Germany (Chapter 18) would be the first building block of this network.

This partner network reaches beyond testing and standards organizations. PIM, PXM and MDM platforms are the natural source systems: they manage the data, and plugilo makes the released version visible and mobile as a Master Data Card. The same holds for product content and syndication networks such as NIQ Brandbank or Icecat: they capture and distribute product content at industrial scale and are natural source and content partners of the truth layer. Enterprise AI platforms are the natural consumers: through open interfaces, plugilo can supply them with authorized external owner knowledge that their internal contexts do not hold, from SAP Joule to Glean; the map in Section 15.1 lists them in overview. **None of these platforms has to be replaced. Each one can become a distribution surface for the certified originals.** And context governance platforms such as Collibra or Atlan remain complementary: they organize internal company context, while plugilo makes owner-controlled truth usable across company boundaries.

One more role belongs in this network: the measurement partners. Measurement and market research companies such as NIQ, Circana or Kantar know more about End Users than almost any other organization. Their portfolio, however, is built for companies, with panels, measurement, dashboards, data feeds and forecasts. Panels capture the transaction, web analytics the click, and the decision process itself remains unobserved: what made it onto the shortlist, what was compared with what, which attributes decided the category, and what the customer does after the purchase. plugilo was designed from the start to work with the measurement world and not against it. plugilo operates the interface and generates the primary data; the measurement partners classify and measure the consented signals neutrally, connect them with market and retail data, and turn them into market intelligence. The fields complement each other with almost no overlap. Product catalogs and the matching of items are joined by living Master Data Cards. Preference measurement is joined by concrete

questions, saved objects and follows before the purchase. Purchase verification is joined by the decision path up to the verified purchase, and channel measurement by sessions, Matches and every measurable plug-it contact point. Two rules keep the roles clean: the measurement and market data of a partner never enters the owner-controlled object core (Chapter 8.3), and plugilo remains the competition-neutral interface and data layer. **Identity, measurement, and advisory from one source only emerge in the combination of standards partner, measurement partner, and plugilo; each party accelerates the other two.**

16. The business model: growing in the network

This chapter explains what is actually paid for in the network, so that every participant knows before joining how he grows with his own success.

16.1 The principle: you pay for demonstrable business value

Getting started carries no risk for anyone: a free entry package with a basic allowance that covers everyday use for smaller companies. Beyond that, everyone pays for demonstrable business value. Data owners pay for data services and maintenance, for the delivery of verified data into approved advisory and sales contexts, and for the customer contacts that result. Media and portals that contribute reach and context share in the value created. Specialist retail and sales partners can be supported by brands through transparent marketing budgets (Market Development Funds, MDF), paid out according to actual advisory performance: money flows to where advice happens.

Above all this sits a simple billing principle. What is paid for is the quality of the actual use, meaning the active, demonstrable engagement with a piece of content, graded by its depth, not the mere impression. That sets the model apart fundamentally from the impression logic of classic advertising: value arises where people or systems prepare a decision.

The architecture in Chapter 9 is at the same time the map of the business model. Revenue arises on all three layers: Master Data Services, enrichment and maintenance; distribution and networking in the Connect Layer; applications, services, and advice in the Application Layer. Added to this are the demand services from the Partner Portal (Chapter 10.5). Each layer can be used on its own and is economically viable on its own, and the Master Data Layer carries the whole as its foundation.

16.2 Two products, three roles: plugilo Cards and pluggy

The model fits in one sentence: **plugilo brings the original to every customer; data owners pay for contacts, operators earn along, and whoever wants to advise adds pluggy** (Figure 15; the services themselves were described in Part 3). It knows two products and three roles. A **data owner** is anyone who is accountable for original knowledge: manufacturers and brands as well as publishers, associations, service providers or retailers with their own offering. An **operator** is anyone who shows this knowledge on his website, in his shop, portal or store, or advises with it. Both at once is possible, and the third role, the **Solution Partner**, sells and supports plugilo at his business customers (Section 16.3). The first product is **plugilo Cards**: the owner's Master Data Cards appear in real time as the original on his own website, in shops, at partners, in trade portals and at publishers, with provenance and verified quality. For this the sources are connected once, automated and free of charge, guided by experts or permanently managed by CF1 (Section 10.6). On the owner's own pages Cards are free; at partners in the **Network** and in booked **Channels** at publishers, trade portals, communities and influencers, the owner pays for Touchpoints and Matches, never for the mere impression. The second product is **pluggy**, the Trusted AI Advisory Agent, and it is optional: it advises, recommends and sells on the operator's pages with the operator's knowledge, which leads

the advice as the master Skill, complemented by released original knowledge of other data owners as sub-Skills. What is paid is a monthly package per domain, measured in real advisory conversations, the Sessions, not in widgets, bots or clicks; Sessions of the package expire at the end of the month, reloaded ones do not, and a free package serves for testing. This yields three ways to use plugilo: show the Cards of others (costs nothing, brings Rewards), display your own Cards (free on your own pages, with reach at partners and in Channels), and add pluggy.

The business model: two products, three units, one balance, three roles

plugilo brings the original to every customer. Data owners pay for contacts, operators earn along. Whoever wants to advise adds pluggy.




Figure 15: Two products, three units, one rule: Cards carry the original, Touchpoints bring it to POS and POI, Sessions fund the advice, Matches create reach and revenue; operators share in the value through Rewards, the End User never pays.

16.3 Billing: Touchpoints, Sessions, Matches and Rewards

The logic behind it needs only two terms. Every pluggy has exactly one **master Skill**, which leads the advice and determines identity and basic context, and it can use any number of released **sub-Skills** that add fitting original knowledge. From this follow three units, one balance and one Reward. The **Touchpoint** arises when a released Card is actively opened at the point of sale or point of interest, by mouse-over, click or opening; page views and impressions do not count. The data owner pays for it: free on his own pages, at partners in the Network and in Channels according to the price list. The **Session** is a coherent advisory conversation between an End User and pluggy; several questions and answers belong to the same Session, page views and technical calls do not count. All Sessions of a pluggy are paid by the operator of the pluggy, the owner of the master Skill, in the monthly package, regardless of where that pluggy is deployed. The **Match** arises when pluggy recommends, in a Session, a released Card from a sub-Skill that fits the expressed need, in the owner's Network as well as in booked Channels; the respective data owner pays for it, in addition to the operator's Session, and per Session there is at most one Match per data owner. Touchpoints, Matches and Sessions are paid from a **plugilo balance** loaded in advance and consumed only for valid use. And the **Reward** is the operator's individually agreed share of the price of the Touchpoints and Matches on his site: for Touchpoints it goes to the operator of the domain, for Matches to the operator of the pluggy. For Resellers and distributors it can be arranged with the data owner as an MDF share; for publishers, communities and influencers it is a new,


demand-driven revenue source. Sessions fund the operation of your own pluggy; nobody shares in them. All payments run through plugilo. **The closer to the purchase decision, the more valuable the unit:** a Touchpoint shows interest, a Session is a conversation, a Match is a qualified hit on a need. Whoever has his Cards displayed beyond his own pages pays Touchpoints; whoever runs his pluggy pays the Sessions; whoever's knowledge is recommended in someone else's advice pays the Matches; whoever displays the Cards and Skills of others shares in the value.

Three rules protect the fairness of the model. A Card or a Skill appears only if both sides agree, the data owner and the operator; every release can be changed at any time, and the release is itself a trust signal, because customers see that the operator is authorized. A Match becomes billable only once a customer has expressed a real need and the original knowledge for it becomes visible; that is closer to a qualified lead than to classic advertising, and billing comes from the plugilo balance, which is consumed only for valid use. And several Cards of the same owner count as only one Match within a Session; only when Cards of different owners are recommended can one Match per owner arise. There is no billing by impressions, no hidden costs and no auction; the value of a Match depends on audience, topic, industry, segment and Channel, because specialized B2B audiences are more valuable than high-volume consumer environments. For future work from the customer's AI environment (Section 10.7) one more rule applies: an administrative query is not a Touchpoint, a comparison formulated by a general AI is not a Match, and a conversation about data maintenance is not a Session; there is no fourth unit and no new prices.

An example without prices makes the logic tangible. An End User opens, in a Reseller's online shop, the Master Data Card of a dishwasher that the manufacturer has released via plug-it: that is a Touchpoint at a partner, which the manufacturer pays from his plugilo balance, and the Reseller receives the agreed Reward for it; on the manufacturer's own website the same opening would be free. Then the End User asks the Reseller's pluggy for the right model. The conversation is led by the Reseller's master Skill; the Reseller pays this Session from his monthly package. The manufacturer has released his Skill as a sub-Skill, and pluggy recommends his Cards to fit the expressed need: that is a Match, which the manufacturer pays, regardless of how many of his Cards appear in this Session; the Reseller receives the Reward for that too. Had the same advice run on a booked trade portal, the publisher would have received the Rewards. **The End User never pays** (Figure 16).

One advisory conversation, three units: who pays what, who receives what

Example: an End User opens the Card of a dishwasher in a Reseller's online shop and then asks pluggy.



On a booked portal, the publisher would receive the Rewards. Several Cards of the same manufacturer in one session remain one Match. No prices: the paper names principles, not terms.

Figure 16: One conversation, three units: the manufacturer pays the Touchpoint and the Match, the Reseller the Session; the Reseller receives the Rewards, the End User never pays.

The store is a domain. For billing it makes no difference whether a released Card is opened in a Reseller's online shop, on his homepage or at his shelf: it is the same Touchpoint, paid by the data owner, for which the operator of the domain, here the store, receives the Reward; advice at the shelf runs in the same Session logic. There are no new units and no separate store fees. How the attribution of a later purchase from home is designed in detail is part of the ongoing concept work and will be published with the price model.

The third role is a commercial layer of its own: **the plugilo Solution Partner Program, Sell & Manage.** System houses, IT resellers and agencies not only use plugilo themselves, they sell it to their business customers, introduce it there and support it permanently. This happens in four steps: the Solution Partner runs his own Master Data Cards, plug-it and optionally his own pluggy. He resells plugilo, and every business customer receives his own Cards and his own domain, optionally with his own pluggy and package. He takes over set-up, data connection, training and ongoing support as his own services. And he shares recurrently in the ongoing plugilo revenue of the customers he has won, the Partner Revenue. plugilo provides platform, Partner Portal and billing; the Solution Partner owns the customer relationship and knows the need. **The Solution Partner's customer relationship remains his customer relationship:** plugilo strengthens it and does not replace it. Customer assignment, duration and level of the share, and managed-service roles are governed by the Solution Partner Agreement, whose terms are fixed bindingly before market launch; registration as a Solution Partner runs through the Partner Portal. The logic of the three units remains unchanged; the program adds the sales and support channel.

Free to start. You pay for what is used and what works: interest, advice, demand. And a Skill works only where both sides agree.

16.4 How every role grows with its own success

Every role follows the same principle, and the End User never pays (Figure 17): everyone runs his own Cards and optionally his own pluggy and makes it smarter with the Skills of his partners. What decides the billing is not the type of company but the role in the specific contact: the same Skill can be the master Skill in your own pluggy and a sub-Skill at a partner, and the same company can be data owner and operator at once. The manufacturer brings his original Cards and his knowledge to every point of sale of his trade partners and into booked Channels; he pays Touchpoints and Matches only once a Card is actually opened or recommended there, and steers this through balance, topics and releases. The distributor advises competently across his entire portfolio, because the released OEM Skills strengthen his own pluggy without him paying for that outside knowledge; the OEMs pay their Touchpoints and Matches and give him a share through Rewards. The Reseller brings the manufacturer's best advice directly to his own point of sale, online and in the store, keeps the customer in his own business and receives a Reward for every manufacturer Card opened on his domain. And the publisher earns money with relevant advice instead of ad banners: he runs his own Skill, activates booked Channels and receives a Reward for every Card opened and every valid Match. Each role scales exactly where its own success is created, and every Touchpoint and every Session can become a lasting connection: customers save Cards, follow brands and retailers and take advice and results along in the Wallet. That is the Trusted AI Advisory Network, the direct connection between data owner, operator and customer, independent of individual gatekeepers.

From this follows a strategic insight. A new form of advertising emerges, one that brings supply and demand together at the moment of need, without classic ad placements. plugilo

monetizes reach from day one, and every pluggy on a partner site brings new End Users into the network at the same time. **Growth and revenue come from the same mechanism; the business model therefore does not depend on a large user base at the start.**

Who contributes, who pays, who receives: the roles in the network

Role	Contributes	Uses	Pays	Receives
Manufacturer / OEM	Master Data Cards, original knowledge as a Skill	Cards at partners and in Channels, Skill releases, optional own pluggy	Touchpoints and Matches at partners and in Channels (own pages free), Sessions of own pluggy	The original at every POS and POI, advice, measurable demand, control over the truth
Distributor	Portfolio Skill, assortment depth	Own Cards, own pluggy plus OEM Skills as sub-Skills	Touchpoints of own Cards, Sessions of own pluggy	OEM knowledge without own maintenance, Rewards for Touchpoints and Matches (MDF)
Reseller	Offer Skill, point of sale, customer relationship	Manufacturer Cards via plug-it, own pluggy plus manufacturer and distributor Skills	Touchpoints of own Cards, Sessions of own pluggy	Best manufacturer advice, customer stays, Rewards for Touchpoints and Matches
Publisher / channel	Audience, content Skill	Provider Cards via plug-it, own pluggy plus booked provider Skills	Sessions of own pluggy	Rewards per Touchpoint and Match: revenue instead of banners
End User	Demand, questions, follows	plugilo Wallet (app and web), pluggy, Cards and Stacks	Nothing: core services free	Answers with sources, decision memory, control over the profile

It is not the type of company that decides, but the role of the Skill: master Skill in your own pluggy, sub-Skill at partners and in channels.

Figure 17: Every role brings something in and gains something; not the type of company decides but the role of the Skill; plus the End User, who never pays.

This model does not become defensible because plugilo initiated the standard. The lead has to come from execution: from the relationships with data owners and partners, the connected Master Data Cards, the mapping and maintenance expertise and the distribution through plug-it. Added to that are the Partner Portal with self-service onboarding, the Wallet with real demand signals, pluggy as the advisory application, operational data quality, plus reach, usage and network effects. The standard creates trust and adoption. The company wins through the best network and the best operations.

Execution calls for the same sobriety. plugilo builds its own platform for Cards, trust and distribution, but it does not reinvent base technology: lakehouse, ETL, vector, data catalog, integration and checkout functions come from established providers. plugilo connects them and concentrates its own capital on the layer nobody else occupies (Chapter 15.1).

The principles of pricing. The detailed price model is published on corporate.plugilo.com and continuously updated there; this Master Paper names the principles, not the terms. One rule stands above all: commercial logic must not distort the truth. Sponsored content is marked, paid information never overrides the certified core, and advice always stays anchored in the correct trust level. For orientation without figures: Cards on your own pages, the automated data connection and a free package for pluggy for testing are free of charge; guided and managed data connection by CF1 is available on quote. Costs arise only when released Cards are actively opened at partners or in Channels, your own pluggy runs in the monthly package, or your own knowledge generates Matches in someone else's advice. During the launch

phase, fixed entry prices apply; later, prices are formed dynamically by audience and location, with discounts for higher volumes by agreement. The terms in the respective offer are binding; they apply for the contractual term, and the model is developed further on the basis of real use.

We are in market launch, and the terms are developing along with the market. Binding terms come with the individual offer.

PART 5 · THE EXECUTION

What exists today, what follows as an expansion stage, how regulation and standards connect, how the network grows and how to take part.

Part 5 at a glance: The core is built, the expansion is ordered and openly labeled. Regulation and standards work in the same direction. The network grows through immediate individual value, clear formats, and effects that reinforce themselves. Entry is low-threshold and open to all.

17. Availability and expansion stages

What already runs today and what follows in stages belongs in the open, so that every partner knows what to rely on today and what is still ahead.

The core from Chapter 10 is in use. **The expansion stages that follow extend it, but they are not a precondition for its value.** They are identified as a roadmap, each with an unambiguous status: (beta) in use with partners and in continuous development · (in development) being built · (planned) later expansion, partly financed by coming rounds. "Beta" describes the technical operating status: the services run productively, and there is still much to develop; we deliberately do not call them "finished". Statements about partners, usage, and traction are the company's own statements; their figures are refined continuously and substantiated in personal conversation or in the data room. The logic follows the pattern of the iPhone and the App Store: first a clear, reliable core, later an open Service Store on top of it in the Partner Portal. Table 12 in this chapter sets out the current status of every individual service. What lies behind the expansion stages is described in the sections that follow.

- **Demand Intelligence, the return channel for owners (beta):** The return channel mechanics from Section 10.5 broaden out here. Owners gain visibility and continuous market knowledge, up to the voluntarily qualified contact. The personal demand profile stays with the user alone.
- **plug-it in the Store (in development):** The plug at the shelf (Section 10.4): scanning of electronic shelf label, QR, NFC or barcode, web view without app, advisory hand-over by the sales staff, store as the recognizable place of origin of the Card. In the concept phase; the connection follows the emerging standards for data carriers on the product (GS1 Digital Link, digital product passport, Chapter 18) and the systems of electronic shelf label providers.
- **plugilo in your AI (planned, concept phase):** Cards, quality status, reports and controlled actions directly from the customer's preferred AI environment (Section 10.7). Envisaged first are read-only queries and reports, then exactly one complete write process: delegate

data maintenance, submit a change, check, release, publish and prove the effect. Rights, releases and evidence remain server-side with plugilo; an administrative query is not a Touchpoint, an AI comparison is not a Match, a maintenance conversation is not a Session.

Service	What it delivers	For whom	Layer	Status
Cards & Stacks	The certified original per object, curated in context	Owners and users	Master Data	beta
Master Data Record	The machine-readable record behind every Card	AI systems and partners	Master Data	beta
Partner Portal	Control center for data, rights, and approvals; return channel for demand	Data owners	Master Data	beta
CF1 Data Services	Professional data creation, harmonization, and maintenance up to managed Source Truth	Data owners and partners	Master Data	beta
plugilo Wallet	The personal knowledge and demand wallet, as mobile app and web portal	Users	Application	beta
pluggy	The AI advisor that backs every answer with a source	Users, retail, media	Application	beta
The Mirror	Analysis of how AI systems talk about a brand today	Brands and vendors	Application	beta
plug-it	The plug that activates Cards at every point of decision	Websites, shops, portals	Connect	beta
Follow Service	Keeps the relationship between End User and object alive	Users and owners	Connect	beta
Demand Trainer	Wallet and plug-it as a learning system for real demand	Owners	Application	beta
Browser Extension	Makes practically every suitable web page connectable, with safe-site filtering	Users	Connect	beta
Demand Intelligence	Aggregated demand signals back to the owners	Data owners	Master Data	beta
Connect Service & Business Graph	Company-centric relationship layer with an open API	Companies	Connect	planned
Object-centric conversation	Exchange directly at the object instead of anonymous groups	Users and partners	Application	planned
Service Store	Standardized third-party services attached to Cards and Stacks	Partners and users	Application	planned
Blockchain anchoring	Verifiable provenance; later, transactions on Trusted Data	Owners and agents	Master Data	planned

Developer layer	Open APIs and components for custom applications	Partners and start-ups	all layers	planned
-----------------	--	------------------------	------------	---------

Table 12: What runs is called beta; what is coming is called planned. Nothing in between is blurred.

- **Object-centered conversation (planned):** The exchange does not happen in anonymous groups, but directly at the object: at the product, at the project, at the trip. Every Card and every Stack is also a conversation space, on request with a specialist advisor or an AI advisor in the same conversation, while owner updates flow in automatically. The established networks organize conversations around people. Here they organize themselves around the subject.
- **Connect Service, connecting with a relationship type (planned):** Mutually confirmed relationships link companies and people, each with a relationship type: for companies, customer, supplier, partner, reference, or agency; for employee Cards, purchasing, sales, or management. Private and business roles remain distinguishable. Between private individuals, connecting works as simply as in any other network.
- **Business Graph, company-centric instead of person-centric (planned):** Employee Cards belong to the company domain; if the person moves on, the domain owner reassigns the Card. Contacts built up by a position stay with the position and therefore with the company, including trade fair contacts and the customer and supplier network. In established business networks, an employee who leaves simply takes along the network built at the company's expense; here that is structurally impossible. Relationships become lasting company assets, usable through an open, owner-controlled API.
- **Service Store (planned):** Following the app store principle, third parties will connect standardized services directly to Cards and Stacks through an open interface: callback and advisory appointment, availability and quotation, demonstration and test drive, support, financing, logistics, appointment and project requests. Whoever has used one of these services once finds it again everywhere in the same form. That is the standard advantage from Chapter 5.1, applied to services. Frankly: the build-out follows with the coming financing rounds, and early partners help shape the interfaces.
- **Blockchain anchoring, from provability to transaction (planned):** a later cryptographic or blockchain-based anchoring can document the origin, versions and time stamps of released original data in a tamper-evident way. It proves that a Card matches the unaltered owner version. Whether the recorded data is true and authorized is still ensured by the trust model (Chapter 8.3). The target state reaches further: on this basis, transactions with the Trusted Data should become possible later, meaning binding orders, contracts, attestations and approvals, triggered directly from the Wallet or by authorized AI agents (Chapter 14.5). The blockchain remains invisible trust infrastructure: it does not define the user experience, it secures it. The plugilo Wallet is and remains a knowledge and demand wallet, not a payment wallet and not a crypto wallet.
- **Developer layer (planned):** Open APIs, components, and tools with which partners, system houses, and start-ups build on plugilo: maintain Cards programmatically, trigger events, embed their own services, run plugilo internally. In this way an ecosystem of its own grows around the truth layer.
- **Confirmed relationships and genuine editions (planned):** For the brand this is customer loyalty, not a collecting game. Original knowledge always stays freely distributable; only confirmed relationships, genuine editions, and earned status become scarce and valuable: a Connection Card as a personally activated connection to an object, a brand, and the partner who made the introduction, a Collector Edition tied to a real occasion, place, period, or print run, a Founding Member Card for early participation in a named

period. Safety, usage, and mandatory information is never made scarce, and plugilo uses no fake limits, countdowns, or hidden conditions; the European Commission counts such patterns among the problematic dark patterns.⁴³ The first benefit should be easy to reach, forgery hard, deeper relationships earned.

A roadmap is only as credible as its status labels. What runs is called beta; what is coming is called planned. Nothing in between is blurred here.


For the Wallet and pluggy there is an additional internal product roadmap in five stages (answering, collecting, organizing, understanding and reporting, acting). It details the services listed here from the user's perspective and changes no status information in this Master Paper.

18. Connection to regulation and standards

The Digital Product Passport and the established standards are coming anyway, and whoever organizes their data now handles both tasks with the same work.

The strongest tailwind for structured original data comes from regulation itself. The EU Ecodesign Regulation (ESPR) is introducing the Digital Product Passport step by step for a large share of physical products: machine-readable data on durability, repairability, materials, and conformity, accessible through a data carrier on the product. The battery passport comes first, mandatory from 18 February 2027 for EV/traction batteries, LMT batteries (for e-bikes and scooters, for example), and industrial batteries above 2 kWh. Further groups such as textiles and iron and steel follow the EU work plan; specific obligations take effect with the delegated acts and after transition periods.⁴⁴ Over the coming years, millions of companies will therefore have to build exactly what the Master Data Record already is. The infrastructure maps the product passport fields, and the data set that regulation demands is the same one that creates visibility in AI advisory. **Maintain it once, use it twice** (Figure 18).

Regulation is coming anyway: timeline for structured product data



The data set the regulation demands is the same one that creates visibility in AI advice. Maintain it once, use it twice.

Figure 18: Regulation does not ask whether a company organizes its data, but when; the battery passport is mandatory from 18 February 2027.

Tailwind also comes from AI regulation, because the EU AI Act raises the importance of provenance, rights, and transparent use. The Alliance takes the constructive route. Instead of forcing fees or shifting responsibility, the owners themselves control which information is used and how, through open interfaces and certified original data. That is fair, transparent use without expropriation and without forced licensing disputes. The NIST framework for trustworthy AI takes the same direction with validity, reliability, transparency, explainability, and privacy, and it notes that a system can meet these characteristics without users

automatically trusting it.⁴⁵ Trust has to become visible, and that is exactly what the Card delivers.

Regulation does not ask whether a company wants to organize its data. It only asks when. Whoever organizes early decides how they come across while doing it.

Outlook, the joint Proof of Concept with GS1: GS1 Germany and plugilo are working on a joint Proof of Concept. It examines how GTIN-identified, quality-checked owner data can be activated as living Master Data Cards, used through the AI advisor pluggy for trustworthy, source-bound advice, and distributed through open interfaces to retail, media, and AI systems. The basis is real, extensive owner data sets from leading brand manufacturers, individual records with well over 100 attributes, images, and documents. The success criterion is proof from end to end: the GTIN resolves unambiguously, the approved data is transferred completely and correctly into the Card, provenance and quality evidence are visible, and pluggy answers real user questions in a traceable, repeatable way that scales to further products and brands. In everyday use it would look like this: anyone who scans a product would receive not just the data set, but an advisor that explains it on the basis of the owner data; the path would lead from the GTIN straight into trustworthy advisory. A cooperation does not yet exist. The concrete design remains subject to joint evaluation; the target picture, however, would be a long-term, worldwide collaboration within the Trusted AI Advisory Alliance.

plugilo seeks the same connection to the sovereign data spaces and product passport initiatives. Catena-X, Gaia-X, and the International Data Spaces Association create frameworks for trusted data exchange between organizations, CIRPASS-2 prepares the Digital Product Passport in practice, and GS1 Digital Link connects the physical product identifier with the official online information. The division of labor is clear: data spaces exchange data between organizations, the product passport documents product and compliance, GS1 identifies. **plugilo turns the approved knowledge into visible, understandable, and usable advice for people and AI agents, connected rather than replaced.**

The roles are precisely defined. What would be activated is GTIN-identified, DQX-verified owner data (DQX, the data quality program of GS1 Germany, checks manufacturers' product master data and confirms verified records electronically). The approved, GS1/DQX-verified owner version always remains the authoritative upstream source, and plugilo holds no competing truth: the Card follows the source. The target picture would go beyond data delivery: GS1 as an active Master Partner for governance, standards, and access to the more than two million companies connected to it worldwide.

19. The growth plan: how the network realistically grows

Every network lives on critical mass, so the mechanics meant to achieve it are laid open here.

Every network initiative has to answer one central question: why should this particular network reach critical mass? **The answer is not a hope for viral accidents, but a fully engineered plan.** Its mechanics are disclosed here, because transparency about one's own method is part of the trust promise. The plan combines three levels: immediate individual value, targeted activation formats, and self-reinforcing network effects, international from the start.

19.1 Practice is starting: the mechanism shows first results

This paper does not describe a plan on the drawing board. The infrastructure is technically in operation (status definitions: Chapter 17), and partners from brands, retail, media, and associations already work with the model in different industries and countries (company's own statement; references are published continuously, key figures substantiated in personal conversation). Selected references from dozens of companies that already use plugilo or are currently setting it up include, in the IT world, Philips/AOC and the distributors api, Littlebit, and Siewert & Kau, and in home appliances, kitchens, and furniture, Miele Schweiz, Electrolux/AEG, and Liebherr. Trade associations and their companies organize their official product truth together, embed verified advisory into their own websites, and bring customers into the specialist conversation properly prepared. **The pattern is the same everywhere and can be transferred at will, industry by industry, country by country.** Further industry and national Chapters in Europe, North America, and Asia are being built, and the worldwide ICT OEM outreach is under way (Chapter 21). The Alliance is the core of this scaling path and will be expanded with its own infrastructure, its own team, and new investors. On the user side, plugilo is deliberately at the beginning. Over the past months, small tests ran with selected influencers and companies to compare paths and methods; the findings from them flow continuously into the solution. The focus is therefore on partnerships that bring the solution into the market in a direction-setting way: standards organizations such as GS1 Germany, associations such as Küche Schweiz, and strategically important brands with which proofs are running. The Alliance is designed to work across industries, from IT through home appliances, kitchens, and furniture to consumer goods. Operational key figures on activation, activity, and retention are presented to partners in working sessions, not in papers. Chapter 20 names the invited partners, Section 20.7 the entry point.

19.2 The foundation: individual value first, the chicken-and-egg problem is solved

Networks fail when their benefit only appears at scale. Here, participation pays from the first day, independent of the size of the network. Every company receives its prepared official data profile (Section 19.3), an individual analysis of how AI presents it today, control over its own presence, and free entry. **Whoever joins gains immediately, and the network is the bonus that grows every day.**

19.3 The activation formats: how we win companies

- **The Mirror (AI visibility analysis):** The most effective door opener is the truth itself. Every target company sees, individually and with evidence, how AI systems present it today (Section 10.3), and a picture of its own situation like that is hard to set aside. The urgency therefore comes from the company's own data situation, not from sales pressure: the analysis makes GEO measurable, the Cards make the presentation shapeable.
- **The prepared profile (take over instead of buy):** Before the first contact, the official data profile is already built. The invitation is not "buy a product" but "take over what belongs to you". Entry shrinks to a single step: review and approval.
- **Open letters and industry benchmarks:** Instead of writing to a hundred companies one by one, we put a single question publicly to an entire industry: who will be the first brand of your industry in the network? That first place exists only once, and whoever hesitates leaves it to the competitor. Associations and media partners co-sign such letters, flanked by an anonymized industry benchmark ("The most common AI errors across industry X") that gives the press its story without exposing individual companies. Sample texts are in the Partner Operations Guide.

- **Limited founding circles:** Each industry has only a few founding places: documented, transparent, binding. Founding partners shape the standards, carry the founding status permanently, and receive terms that exist only in this phase. The limit is not a gimmick, because a founding phase that never ends would not be credible, and credibility is the Alliance's capital.
- **Claim Days and founding dialogues:** Accessions are bundled into shared moments: one announced day per industry on which all prepared profiles are unlocked, and the founding dialogue as a flagship event with associations, policymakers, media, and the first founding partners. Every accession becomes a visible event instead of a quiet signature.
- **Personal outreach with substance:** The public formats are followed by the direct invitation: short, fact-based messages with the individual analysis result and a reserved founding place, extended by personal conversations for key partners. In the ICT sector, Michael Urban opens access to the decision-maker level (Chapter 21).

19.4 The network effects: why every participant brings more

- **Partner cascade:** Every brand won brings its network along: hundreds to thousands of retail partners, distributors, and publishers receive prepared profiles with the certified data of "their" brand. One participant becomes an entire partner network.
- **Reverse pull:** Activated Resellers and portals request the data of missing brands from their own channel ("We work with certified profiles, and there is none yet for brand Y"). The invitation to the next brand comes from its own sales channel rather than from the platform provider. There is hardly a more credible recommendation in the market.
- **Association multiplier:** One association sets an entire industry in motion: co-signature, industry Chapter, member activation, shared data standards. The federated Chapter model repeats this leverage internationally, and every country and every industry founds its own movement alongside.
- **Observer funnel:** The free tier turns interested parties into participants. Whoever follows analyses and developments first-hand stays able to decide and converts as soon as the occasion arises. Participation alone increases the reach and legitimacy of the network.
- **Quality spiral:** More owners create more verified data, better data improves the advice, better advice wins users and demand. These signals raise the value of participation for every further provider, the core of every successful standard.
- **Model operator step:** The pooled data itself becomes the lever: a real-time pool of certified original data that no serious operator can ignore, because it makes their answers source-bound and verifiable. Every new owner increases the value of this offering, and every additional AI system connected increases the value of participation for every owner: the network effect works in both directions.
- **End User channel:** The plugilo Wallet turns every partner into an ambassador: brands, shops, and publishers invite their own customers to follow and recommend plugilo along with it. Just a few of the largest brands in the world could carry the network to a large share of users (scenario, Chapter 6.1). The user side and the provider side grow together and reinforce each other.
- **User loop, the plug as the moment of acquisition:** Every plug wins users in exactly the second in which it is useful, in context and without advertising. Every shared Stack carries the network further, and every Card created by a user invites its owner to take it over. What is measured is what counts: activated users, return visits, Cards per user, referrals.

- **Topic worlds as starting capital:** Every industry begins with curated topic worlds, set up together with media and association partners: visible showcase examples that demonstrate the value, bring new Chapters to life, and solve the cold-start problem that otherwise makes networks fail.

19.5 The PayPal effect: one plug for all the databases in the world

The great growth stories of the internet follow documented patterns, and this plan builds on them. PayPal grew because it sat exactly where trade happened. As the payment button on the eBay marketplace, it grew with every transaction, because sending money to a recipient who was not yet registered triggered an invitation to open an account. Sign-up and referral bonuses of 5 to 10 US dollars each accelerated this usage-driven network effect. At the end of 2013, WhatsApp reached around 400 million monthly active users, by its own account without large marketing campaigns, carried by product benefit and personal recommendation. Dropbox, in turn, made referral part of the product. Both sides received additional storage, which according to co-founder Drew Houston raised registrations by 60 percent on a lasting basis, and every shared folder was an invitation as well.⁴⁶ The principle is the same in all three cases: **the service spreads in the moment of use, carried by those who already use it.**

This pattern is built into the network. The plugilo plug for all the world's data and websites appears in place of the heart symbol, beyond product data: news and articles, restaurants and holiday rentals, recipes, vehicles, and real estate. Every database and every platform can build it in, and the stated goal is the large portals, trading platforms, and established publishers, travel platforms such as Airbnb or Tripadvisor, and marketplaces such as eBay. We standardize saving the way PayPal once standardized paying.

The users themselves carry the network further as well. Advisory results, meaning Cards, performance comparisons, and Stacks, are passed on to family, friends, and colleagues with one click. **Every shared recommendation is an invitation, every adopted collection a new participant.** That is how results grow into a network. Brands recommend plugilo to their partners worldwide, partners invite the missing brands, and users bring their contacts along. Three flywheels mesh in this way.

One standard unified payment in e-commerce and accelerated online trade. A standard for data organization can trigger the same mass effect in the AI age.

The internal growth and activation mechanics, meaning principles, formats, and key figures, are documented in the separate Partner Operations Guide. That this pattern works is documented many times over in history (Chapter 5.1 and Section 6.5). Industries organize around shared trust infrastructure, even among competitors, when three conditions are met: a shared problem, reliable rules that are the same for everyone, and immediate benefit for the individual. All three are met here. The problem is obvious, the rules exist in the Charter and will be published with the first Chapters, and the individual benefit starts on the first day. The record pace of AI adoption and the first partnerships across industries and countries (Section 19.1) indicate that the mechanism is in place and showing first results (company's own statement, as of September 2026).

19.6 Five levers of scaling

Growth does not rest on one mechanism, but on five levers that reinforce each other. All five are built and in operation; the intensity at which they are driven determines the pace and the number of markets. The three flywheels from Section 19.5 belong to the fifth lever (Figure 19).

- **Lever 1, Alliance and standards partners: one standard, one AI strategy, one worldwide rollout.** The Trusted AI Advisory Alliance is the strategic umbrella; it stays market-neutral and open to all partners and is not the vehicle of a single company. plugilo delivers what a fast rollout requires: technology, Charter, activation formats, and the connection to brands, associations, publishers, and influencers. Standards organizations are a particularly strong lever within this: they create identity and measurable data quality, and plugilo makes this data usable as living Master Data Cards (Chapters 15.3 and 17). This lever does not require convincing companies one by one; existing global structures carry the rollout.
- **Lever 2, brands as the engine of the network.** A brand never comes alone. It brings its entire environment with it: products and original knowledge, employees, distributors, Resellers and shops, service partners, publishers and influencers, and existing and new customers. An OEM gives pluggy its knowledge once and makes it available to its entire partner network; every partner becomes a new advisory and contact point. At the same time, Resellers and shops start asking missing brands for their data and their pluggy. The pull works in both directions, and every brand won becomes a sales channel for plugilo. Monetization comes from advisory, Matches, Channels, services, and insights, not from selling personal data.
- **Lever 3, the plug as a universal standard for End Users.** You save once, from every shop and from every site, into the one Wallet. plug-it anchors saving where the heart symbol and the saved-items list sit today (Sections 10.2 and 10.4); where it is not yet integrated, the Browser Extension makes saving possible anyway. Every integration is also a growth surface, because every saved or shared item can bring in the next End User. This is the largest End User lever in B2C and in B2B, and it requires the largest build-out: ready-made plugins for the leading shop and content systems, partnerships with large shops and portals, languages and country logic.
- **Lever 4, publishers, influencers, and associations as compensated reach partners.** They already own audiences; plugilo does not have to build them, only connect them to the advisory network. plugilo sets up Channels by country, industry, and audience, the OEM selects audiences and Channels, and that selection becomes an activation request. A status system keeps the market transparent: red (not yet contacted or not available), yellow (in activation or negotiation), green (activated and bookable). Influencers build their own Stacks with recommendations and topic worlds, and associations activate their affiliated companies together. This lever links user growth directly to monetization: brands fund relevant advisory, publishers bring the audiences, and they receive a share of the value through Rewards.
- **Lever 5, the self-reinforcing network and data cycle.** Brands activate their original data, partners and Channels distribute Cards and bring pluggy into their environments, End Users save, share, follow, and ask. Real demand signals arise, the answers from pluggy become more relevant, the value for brands, Resellers, and publishers rises, and that brings in further participants and further original data. The invitation mechanisms from Section 19.4 are the gears of this cycle. The planned Business Graph (Chapter 17) is not yet built.

Five levers of scaling that reinforce one another




Figure 19: Five levers that accelerate each other; the intensity at which they are driven sets the pace and the number of markets.

The five levers can be driven at three intensities (Table 13). Prioritization follows user potential, speed, resource requirements, monetization, and strategic importance for the partners. **There is no “pilot”: the entry point is the operational mode, with shared goals and a shared roadmap. The focused start is an intensity level, not a test phase.**

Level	Content	Effort and effect
Focused Start	First brands, selected industries and countries	Lower effort, fast results
Accelerated Rollout	Several industries, countries, and Channel networks in parallel	Fast build-up of measurable user numbers
Global Standard Rollout	Parallel activation of brands, standards structures, shops, publishers, and End Users	Highest effort, greatest speed, a chance at standards leadership

Table 13: The focused start is an intensity level, not a test phase.

19.7 The reference case: what is proven end to end

The differentiation from Chapter 12.2 is only as robust as it becomes visible and measurable in the product. The reference case with which the Alliance shows this has eight steps, and each one can be verified on its own. A brand registers its domain and confirms the responsible administrator. At least one product family is created as Master Data Cards with sources and quality level. A master Skill trains the brand's own pluggable, and sub-Skills add support and partner context. The brand authorizes one Reseller and one publisher in the Partner Portal, and both domains integrate the Card and pluggable via plug-it or an equivalent live integration. An End User asks a question, receives advice with a visible Master Data Card and saves the result in the Wallet. The same proof is designed for the store as soon as plug-it in the Store is available: scan at the shelf, Card in the Wallet, store as the recognizable place of origin (Section 10.4); association partnerships in advice-intensive industries such as the one with Küche Schweiz are one of several bridges to it. The brand changes a relevant piece of information, and the updated version appears in all connected presentations. Dashboard and audit trail finally show delivery, version, domain, Advisory Session or Match, and the return channel. The whole chain is measured: whether partners actively request releases, whether End Users really open the trust view, which knowledge needs arise in sessions, whether use turns into a lasting connection (saves, follows, return after 30 and 90 days), how fast a change reaches all surfaces and how many divergent copies remain. The Proof of Concept

with GS1 Germany (Chapter 18) and the ongoing partnerships (Section 19.1) are first building blocks of this evidence.

Credibility also includes what this paper does not claim. That plugilo raises conversion by a certain percentage, lowers returns or support costs by a certain amount, that the blue seal is understood in the same way by all customer groups, or that authorized partners are automatically perceived as more trustworthy: such statements are permissible only after measurement. The paper also deliberately avoids three formulations. “Only plugilo can do this worldwide” is too absolute without a complete market review; what is defensible is that the target architecture is designed for the full combination of the defined criteria. “Controlled everywhere in real time” cannot be enforced in systems that are not integrated; what is defensible are the guarantee levels from Chapter 11. And “standardized worldwide” presupposes a formal standardization procedure; what is defensible is a platform standard that can be used globally and is compatible with the established standards. Measured effects, target pictures and hypotheses stay separate in this paper, and that is exactly what makes its statements quotable.

20. The call and the next steps

Who is needed now and what there is to do stands here as an open invitation to everyone who represents industries, the public, and capital.

The task described here is too large for individual companies and too important to leave to chance or to global platforms alone. This paper is therefore also addressed to the organizations that represent industries and the public interest.

20.1 To the associations and chambers of commerce

No actor can organize an industry faster and more credibly than its own association. We invite trade, industry, retail, and crafts associations as well as chambers of commerce:

- to join the Alliance and to represent the data sovereignty of the industry as a neutral voice toward members and the public;
- to found an industry Chapter of their own, with shared data standards, a founding circle, and member activation, following the example of the association partnerships already under way;
- to contribute their own formats such as member news, awards, training, and events as structured industry intelligence that remains usable over time;
- to secure access for small members: free entry guarantees that data sovereignty does not become a question of company size.

20.2 To business organizations and interest groups

The question of which sources AI systems draw on when they talk about the domestic economy belongs on the economic policy agenda. We invite business associations, associations of mid-sized companies, brand and media associations, and their representative offices in Berlin, Bern, Vienna, and Brussels, and equally their counterparts in every market in the world. Take up data truth and source sovereignty in your policy work, moderate the dialogue between data owners, AI providers, and policymakers, and carry the Alliance into these debates as a practical, industry-backed example of implementation. This paper may be used for that purpose in full or in excerpts.

20.3 To policymakers and public institutions

From policymakers we ask for neither subsidies nor a product endorsement, but for three contributions that only they can make:

1. **The dialogue:** the willingness to examine the Alliance with the responsible expert levels for AI, the data economy, and digital sovereignty, and to support founding dialogues with their expertise.
2. **The signal:** the public statement that trustworthy AI needs provenance-secured sources and accountable data owners: neutral toward technologies and vendors, yet setting the direction for the whole development.
3. **The participation:** the review of which consumer information, official data, and educational and health information can themselves be made available as authorized sources with verified provenance.

The dimension reaches beyond the economy. The same copy cascade that distorts product data affects news and public information; machine-generated content and deepfakes increasingly undermine trust in media and institutions. Verifiable provenance is therefore more than economic support. It is democratic infrastructure, because citizens can tell what comes from an accountable source and what does not.

Digital sovereignty does not arise from regulation alone, but from infrastructure that works according to its own verifiable rules. The Alliance is an industry-backed contribution to that, open to support at national, European, and international level. **What applies to Germany and Europe applies to every country: an economy deserves a truth layer that belongs to it.**

20.4 To investors and strategic partners

The next internet will not be built around search. It will be built around trusted decisions, and around the question of who owns the data behind them.

Building the worldwide truth layer is one of the great infrastructure opportunities of the AI age. Between the data of the global economy and the AI systems that advise on the basis of it lies a layer that no one has occupied so far. Added to that are network effects, recurring revenue models, and historical precedents that became global standards (Section 6.5). The ambition is deliberately chosen large: an infrastructure role on the scale of a search engine, as the basis for trustworthy information on an open meta level above all systems. A network with these goals is category infrastructure. It aims to connect up to 80 percent of online shops (target figure, Section 10.4), to establish the Trusted Data Wallet of the world, and to provide the source layer of the AI systems. In perspective, this reaches the transaction level, because on blockchain-anchored original data Trusted Data becomes transaction-capable (Chapter 17). The long-term target state reaches further still: the one Card ID per data object (Section 8.1) and with it the independent data hub, because the more systems retrieve the originals, the more copies lose their reason to exist. Whoever backs the coming rounds is not financing another language model, but the source-bound reference layer that all models point to.

In the end, whoever connects the existing standards prevails. plugilo claims this role as an integrator that competes with no camp: GS1, ECLASS, ETIM, and the testing organizations are brought in and strengthened, not replaced or bought up. Investors and strategic partners can follow as observers first and then contribute capital, market access, or both. **Whoever enters early helps shape the standard instead of following it later.** All details, meaning capital requirement, valuation framework, capital market perspective, and full documents, are

presented in person; this Master Paper is expressly not an offer and does not replace investment documents.

For large strategic partners, the calculation is simple. plugilo brings a finished, running system into every partnership: the prepared master data base with more than six million records, Cards and Stacks, pluggy, plug-it, the plugilo Wallet as a mobile app and web portal, Browser Extension, Partner Portal, Follow Service and Demand Trainer, plus the sessions, Matches, and Channels model, the Trusted AI Advisory Alliance with its Charter, and decades of experience with brands, product data, and Channels. Anyone who builds the same capabilities alone needs End User software, an advisory AI, a trust model, a distribution mechanism across brands and Resellers, and an alliance that sets standards. That costs years and considerable funds, and above all the time in which others set the standards. Scaling with plugilo, by contrast, starts on this foundation from day one: every dollar goes into reach and speed instead of into rebuilding what already exists, and plugilo contributes its development unit, so that interfaces and requirements of strategic partners receive priority in the roadmap. Conversely, plugilo is not looking for just any partner. They must know brands and Resellers worldwide, have data expertise, master the measurement of the new commerce, and bring the neutrality that keeps signals valuable. The joint work then flows into two areas: product development and technical scalability, and rollout and user generation. **What plugilo proposes goes beyond a customary partnership: one roadmap, one development program, one go-to-market.**

corporate.plugilo.com is the single reference address for all further communication: registration, documents, appointments, and contact.

20.5 To the operators of the AI systems and model providers

Your systems are only as good as the data they answer from, and the share of unreliable, copied, and synthetic content in the open web is growing (Chapter 4). plugilo offers you the connection to the other side: verified, owner-controlled, machine-readable original data with provenance, status, and licensing clarity, available in real time and citable with its source, following the principle "Certified first". **A single connection reaches many owners.** Instead of licensing and maintaining thousands of individual sources, one gateway connects your models to the pooled, certified knowledge of the data owners. That applies today to products, companies, and content, and in perspective it is meant to reach along the Digital Product Passport and the authorized live states into the physical world. We expressly invite the operators of the large systems to help shape this interface early. Information and contact at <https://corporate.plugilo.com>.

20.6 To company leadership: the commitment of headquarters

The change described in Chapter 4 cannot be delegated to a department. AI will change processes, software, communication and decisions in every area of your company, and the one foundation on which this succeeds or fails is your own, accountable knowledge. We want to shape this transition together with strong companies, and for that we need the binding commitment of their headquarters: a sponsor in executive management, accountable data and business teams, released resources and shared, measurable goals. Not as one more AI project, but as the decision which version of your knowledge applies and who maintains it. We invite the decision-makers and doers to implement the first end-to-end use case with us, from the released original to the answer and action at the customer, and to expand successful connections in the network. The metrics are on the table: share of source-bound statements, use of expired data, time until a correction is adopted, unauthorized accesses, rework and total cost per result. **Data sovereignty becomes effective when companies organize it**

together. Only together are we strong, and for that we need strong partners with full commitment.

20.7 Invitation and next steps

The public entry point is the same everywhere and deliberately low-threshold: “Get the Master Paper”, registration at corporate.plugilo.com. The observer journey follows from there: read, follow along, ask questions, and take part when it fits. The observer role is free of charge and non-binding. The Alliance was presented publicly at ECR iN MOTiON by GS1 Germany on September 16 and 17, 2026, at the World Conference Center Bonn, with plugilo as exhibitor at booth 42 and the talk “Master Data meets AI: How original knowledge becomes better advice” on September 17; the talk and the handout are available as a reference.³²

This paper may and should be passed on, to management teams, associations, editorial offices, partners, and political institutions. The best first test is simple: ask any AI system about your own products, services, or content and compare the answer with your official information. **In most organizations, the result is the most convincing part of this paper.**

The next steps are simple, and the first one has usually already happened, because with the registration through which this paper is provided, you are already an observer of the Alliance. Securing your own presence takes minutes and carries no risk with the free entry package. The founding circles for each industry are open to active partners as long as seats are available. Associations, interest groups, and political institutions are invited to the founding dialogue, the opening event in Berlin at which the Charter, references, and the first founding partners will be presented. We look forward to many taking part. Information, membership, and dates: <https://corporate.plugilo.com/contact/>

With strategic partners, we go straight into the operational mode (Section 19.6): define how the services interlock, and decide on a cooperation that can start immediately. The first milestone is a letter of intent on the shared direction and one concrete activity that starts right away, no further phase of getting acquainted. The standards are being set now; those who start together decide who sets them.

For manufacturers, the path can be described in six steps, and the value begins before any IT integration (Figure 20):

1. **Join.** Register the company and secure your presence.
2. **Present.** Make company, products, news, and knowledge visible on plugilo.
3. **Advise.** Use your own pluggy directly in the plugilo Wallet.
4. **Share.** Activate existing distributors and resellers with your own Skill.
5. **Grow.** Win new customers through publishers, communities, and further Channels.
6. **Integrate.** Only then, if needed, connect website, shop, databases, and further systems.

The path for manufacturers: value before any IT integration




Figure 20: The value begins before any IT integration; systems are connected only once it pays off.

The truth layer of AI is being built now, day by day, with or without the owners of the data. This paper is the invitation to be part of it: as an observer, a participant, or a co-designer. plugilo. The new Standard for Trusted AI Advisory. Plug it, love it, share it.

Michael Mohr Founder & CEO, plugilo / DCI · Initiator of the Trusted AI Advisory Alliance

The open letter and the invitation, communication, and sample texts are in the Partner Operations Guide, where the active implementation begins.

Legal notices, data protection, and disclaimer

Data protection and data sovereignty: plugilo works in compliance with the GDPR and operates its AI and data processing on infrastructure that follows European data protection rules. The principle of the Charter applies: the data belongs to whoever created it. Releases can be revoked at any time, user data is not sold to third parties, and personal data is processed only to the extent necessary for the services described. The binding versions of the privacy policy, imprint, terms of use, and further policies are available at <https://corporate.plugilo.com> and are kept up to date there.

Trademarks and third-party rights: The companies, brands, and standards named belong to their respective owners. They are named solely for factual context and establish no partnership or endorsement, unless such a relationship is expressly described as existing or planned.

Disclaimer: All figures, sources, and statements in this Master Paper have been checked to the best of our knowledge and supported by the sources cited. Even so, errors, transcription mistakes, or changes to the sources in the meantime cannot be ruled out; no warranty or liability is assumed for them. Only the respective original sources are binding. This Master Paper serves as information, establishes no legal claims, and is neither an offer nor investment, legal, or tax advice.

PART 6 · THE PEOPLE BEHIND IT

Who stands behind the Alliance: the origin, the name, the stewardship, the technology, and the team.

Part 6 at a glance: The mission is old, the architecture is new. It is carried entrepreneurially and made binding by verifiable rules. Its own infrastructure stays in its own hands. And the team is built so that it can connect camps.

21. Stewardship and team

Whether you trust an infrastructure depends on who stands behind it and how long they have been working on the same task.

21.1 The origin: from DCI to plugilo

The vision behind plugilo is neither short-term nor trend-driven. It is the continuation of a mission that began in 1993 with the founding of DCI, the Database for Commerce and Industry. The name said it all: connect demand and supply intelligently, and make market information structured, comparable, and usable. That was a world of bulletin boards, modems, and fax machines, where product data slept in catalogs and procurement was manual work. The underlying research work by Michael Mohr had already received the Philip Morris Research Award in 1990.

The founding design of DCI already worked with centrally maintained equivalence tables and the relations same, broader, and narrower. Those were exactly the principles that the field standardized much later. SKOS formalized semantic hierarchy and mapping relations in 2009, Google's Knowledge Graph followed in 2012, and "entity resolution" became an industry term.⁴⁷ The Master Data Record completes this idea: the one central reference that all systems point to.

In 2000 this vision led to a listed company and to a platform on which customers could state their needs and be served accordingly. The technology was ahead of its time. The principle worked. Then the Neuer Markt collapsed. **The platform was lost. The vision was not.** The company reinvented itself, through media, content, and operational discipline, and that persistence, that experience, and that infrastructure thinking are part of plugilo's DNA to this day. From 2014 the group built its development base in Hanoi, in 2016 plugilo Inc. was founded in the United States as a strategic vehicle for the AI age, and since 2024 the international scaling has been running with new investor backing. plugilo is thus an independent company that emerged from DCI. The group today comprises DCI AG, ContentFactory1 GmbH (CF1), and plugilo Inc., with teams in Europe, the United States, and Asia. DCI AG acts as a holding company with little operating business of its own and is listed on m:access, the mid-cap segment of the Munich Stock Exchange. The operating business is carried by CF1 with around 200 employees and plugilo with around 50 (company's own statements, as of September 2026). CF1 is a wholly owned subsidiary of DCI AG; its data services are described in Section 10.6.

21.2 The name: a plow and a plug

The name plugilo is deliberately chosen. In Esperanto, plugilo means plow. The plow prepares the ground, creates structure, and makes growth possible; without it, seed cannot realize its potential. In the same way, plugilo prepares the digital ground for better connections between demand and supply. At the same time the word plug resonates in it, as the standardized connection between people, data, and industry. plugilo is therefore symbol and program at once: create structure, enable connections, prepare growth. This becomes visible in the plug that accompanies the whole network as its sign, from plug-it on the Reseller's page to the punctuation mark in front of the brand name.

21.3 The stewardship: commercially backed, binding through rules

plugilo is the initiator of the Trusted AI Advisory Alliance and at the same time the commercial provider of this infrastructure. That is not a transitional state but a decision: plugilo remains a commercially backed operator, because only in that way can infrastructure of this scale be financed, operated, and developed further over the long run. The Alliance acts as a catalyst, and not only for plugilo: it accelerates the entire market for trustworthy data. The Charter applies to the initiator as well, and the open interfaces expressly allow other technology providers to build on it. A standard is worth more the more people use it. That was true for the GTIN, and it holds here in the same way.

Whoever has been working on the same task since 1993 does not have to chase a trend. That endurance is not nostalgia, but the reason to credit this infrastructure with staying power.

Binding force therefore does not come from the legal form but from verifiable rules: access and Charter rules, documented criteria for levels and seals that will be published with the first Chapters, and the same conditions for all participants. The strongest control lies with the owners themselves: they continue to determine which knowledge they release and where it may be used, and thus ensure that false copies no longer spread. Releases can be changed or ended in line with the contractually agreed rules; a revocation governs the permitted future use in the network, it cannot guarantee the retroactive deletion of all external copies in the open web, and defined immediate measures apply to rights violations, security problems, and impermissible uses. **Whoever breaks the rules loses the data on which the model rests.**

So that this commitment also holds beyond today's operator, four points are fixed. The Charter stays in force if ownership changes, and participants can end their releases on that occasion without a notice period. Owners receive an export and portability guarantee for their data in open formats. The specification of the truth layer, with Card ID, record structure, and interfaces, is published and remains reproducible. And as the first Chapters are built, the governance structures of the Alliance are expanded and published step by step: representation of the owners in the Chapters, documented access and Charter rules, and an appeals path against classification decisions.

The separation behind this is simple and permanent: **the standard is collective, the operation is entrepreneurial.** A say concerns the rules, meaning criteria, access, classification, and the Charter, not the prices and not the commercial services. This is how the GTIN has worked for more than fifty years: the standard belongs to everyone, the implementation is business.

This separation also settles what neutrality means in this paper and what it does not. Neutrality here means two things: plugilo competes with no data owner (competitively neutral) and is open to every reputable AI system (model-neutral). It expressly does not mean that the infrastructure belongs to no one. Anyone who promises more is promising something he cannot keep.

21.4 Sovereign technology: our own AI infrastructure

A paper about data sovereignty would not be credible if the infrastructure it describes ran on someone else's platforms. AI and data processing therefore run on dedicated infrastructure, provided by an investor with direct access to the ecosystem of leading AI hardware vendors (reference confidential; evidence in personal conversation). Approved data and the advice grounded in it do not sit in anonymous third-party clouds, but on controlled compute capacity under European data protection standards. **Sovereignty here is not a promise, it is architecture.** The same basis is intended to carry the open meta level and the planned plugilo meta model with the coming financing rounds (Chapter 7). We are currently preparing formal security and compliance certifications and putting the organizational and technical foundations in place; the respective status is evidenced in personal conversation.

21.5 The team today

plugilo is led by Michael Mohr, the founder and CEO, and by Michael Urban, who has worked alongside him for around three decades. Mohr founded DCI AG and has been structuring product and market data for more than three decades; he connects manufacturers, retail, and

media through digital business models. As an Industry Expert and investor with an entrepreneurial stake, Urban is responsible for the worldwide outreach to the OEMs and manufacturers of the ICT industry and is a driving force behind building the Alliance.⁴⁸

For a network that spans the world, one quality of the team is decisive: only someone who competes with no camp can connect camps that would never entrust their data to a competitor. That is exactly how the group is set up, from the data work of ContentFactory1 to the worldwide partner outreach. That is what we want to be measured against.

Appendix A: The theses of the Trusted AI Advisory Alliance

This appendix brings together the core theses of the Alliance in seven groups: the problem, ownership and trust, the advice, the market and the business model, the Alliance and the path, trust and the original, and systems and control. It is a living part of the Master Paper, because new theses and conclusions are added continuously and versioned. Press and partners may quote the theses freely, individually or in full, citing the source “plugilo Trusted AI Advisory Alliance, Master Paper”.

I. The problem

1. **Thesis 1.** The knowledge base of AI consists mostly of copies without provenance. Product information is adopted, abridged, translated, and recombined thousands of times over, usually without a source, a timestamp, or any marking of the original. In the open web there is no authority that tells the original from the transcript.
2. **Thesis 2.** Frequency is not correctness. AI systems weigh what they encounter, how often, and how consistently; when three hundred identical false claims stand against one manufacturer original, the error counts as confirmed consensus. An answer only becomes reliable through the verified source.
3. **Thesis 3.** Zero-Click moves the first customer contact into the AI answer. More and more information paths end before anyone reaches a company website, a Reseller, or a manufacturer. In future, visibility will not belong to whoever runs a well-maintained website, but to the company whose structured data feeds the answer.
4. **Thesis 4.** The answers of the AI systems already decide purchases, reputation, and liability. Whoever is not represented there with his original data is described from copies. The consequences are operational and strategic: poor advice, liability risks, lost demand, higher service and returns costs, and a loss of trust that affects everyone.

II. Ownership and trust

5. **Thesis 5.** The data belongs to whoever created it. Releases are granular and revocable at any time, and every use follows the rules of the data owner. Aggregation without this control only reproduces the problem that already exists.
6. **Thesis 6.** Every data object has a certified real-time original, permanently connected to its source instead of distributed as a copy. The Master Data Card is its visible form, the Master Data Record its machine-readable core. Regulation works in the same direction: the Digital Product Passport makes verifiable, machine-readable product truth mandatory.
7. **Thesis 7.** “Certified first” applies per statement and per scope. The data owner is authoritative for his original, the competent bodies are authoritative for their own findings. Contradictions are made visible as a marked Conflict Record, never silently overwritten.

- 8. Thesis 8.** Trust is measured, not asserted. Three dimensions stay separate: responsibility in Level 1 to 5, quality in the Data Quality Profile, external evidence in the Evidence Passport, integrated with the existing standards rather than in their place. The four information classes of every Card stay just as separate, because context and signals never overwrite the object core.

III. The advice

- 9. Thesis 9.** pluggy is not a chatbot, it is a digital employee. It answers from sources on the principle of "Certified first" and learns from three sources, each of which stays incomplete on its own: the brand knowledge of the data owner, the context of the Reseller or publisher, and the need of the End User. Every answer carries its source.
- 10. Thesis 10.** The plugilo Wallet is the loyalty card of the AI era. In it the End User opens up the certified originals as lasting assets, with an export guarantee and full sovereignty over his demand profile, and he is recognized everywhere in the network. In time, transactions on the basis of the Trusted Data become possible from it as well.
- 11. Thesis 11.** The network compares performance, not prices. Price comparison stays deliberately out of scope: what is compared is performance, fit, and trustworthiness, so that advice becomes better and not prices lower. Visibility is therefore decided again by performance and data quality.
- 12. Thesis 12.** The network generates a signal layer that classic measurement cannot observe: consented, explainable decision events before, at, and after the purchase. Every event carries its context of product, category, market, contact point, and time; the personal demand profile stays with the End User.

IV. The market and the business model

- 13. Thesis 13.** Reach is created from day one. Every pluggy on a partner site brings new End Users into the network, and growth and revenue come from the same mechanism. The business model therefore does not depend on a large user base at the start.
- 14. Thesis 14.** What is paid for is interest, advice and need, not ad space. Billed are the Touchpoint as an active opening of a Card, free on your own pages, the Session as a real advisory conversation, and the Match, which arises only when a need has been expressed; operators share in the value through Rewards. There is no billing by impressions, no auction and no hidden costs.
- 15. Thesis 15.** A Skill works only where both sides agree. Someone else's Skill becomes active in a pluggy only when the Skill owner and the operator consent, and any approval can be changed at any time. Control over how knowledge is played out therefore stays permanently with the data owners.
- 16. Thesis 16.** The data gateway between data owners and AI systems is still unclaimed. Standards organizations, measurement partners, and plugilo can claim it together, because their fields complement each other with almost no overlap: identity and verified quality, neutral measurement of the consented signals, living Master Data Cards, and the advice built on them.
- 31. Thesis 31.** If Coca-Cola knew who drinks Coke, it would own the largest end-customer database in the world. We make that possible: the Wallet connects brand and End User directly, voluntarily and under the End User's control; every Card, every follow and every release turns an anonymous consumer into a known relationship, without anyone being watched.

- 32. Thesis 32.** The original belongs where the work is done. plugilo does not build a new interface to learn, it brings Cards, rules and services into the AI the customer already uses; rights, releases and evidence remain in the network.

V. The Alliance and the path

- 17. Thesis 17.** What is emerging here is an open standard, not a closed service. Specification and interfaces are open, the operation is entrepreneurial, and existing standards are integrated and strengthened, not replaced or bought up. This is how the GTIN has worked for decades: the standard belongs to everyone, the implementation is business.
- 18. Thesis 18.** Only someone who competes with no camp can connect the camps. plugilo is competitively neutral toward data owners and model-neutral toward AI providers, but not neutral in who carries it: initiator and commercial operator. The Gateway stays a gate that the data owners open, source-bound, rights-controlled, and without a new gatekeeper.
- 19. Thesis 19.** The more data owners the Trusted AI Advisory Alliance represents, the better it can assert their interests worldwide. The leverage is concentrated: the 1,000 consumer brands with the widest reach in the world are likely to reach nearly all people in everyday life together, and around 50 active supporters could already carry the Alliance into the wider market (a working hypothesis and a scenario, not an adjusted measurement).
- 20. Thesis 20.** There is no pilot: the entry point is the operational mode, with shared goals and a shared roadmap. The goal is category infrastructure, the competitively neutral data hub through which owner-authorized originals reach all systems and copies lose their reason to exist. The alternative is that someone else does it.

VI. Trust and the original

- 21. Thesis 21.** Trust needs a recognizable original. People trust a piece of information when they can recognize where it comes from, how current it is, who takes responsibility for it, and whom they can turn to. The Master Data Card makes these four answers visible, and pluggy makes the brand approachable at any time; in this way trust becomes reproducible across Resellers, publishers, and points of contact.
- 22. Thesis 22.** Being in the AI is not enough. Being recognized as the original is what counts. Many become visible, old copies included; in the AI answer, what decides is whose data is recognizable as the original. No manufacturer solves that alone, it takes a shared rule for labeling, provenance, and release: alone, everyone remains an island; together we set the standard.

VII. Systems and control

- 23. Thesis 23.** AI creates new knowledge islands. The more companies build their own AI applications, the more often the same brand knowledge is imported, indexed and interpreted separately. The brand delivers the data N times and in the end controls not a single version.
- 24. Thesis 24.** Data access is not knowledge control. An interface or a feed makes data available; the recipient still decides whether, when and how to use it. An interface delivers data. A Master Data Card preserves the original.
- 25. Thesis 25.** Presentation is part of data quality. Correct data loses its effect when mandatory information is removed, wrongly prioritized or shown without a source. That is why owner, provenance, version and return path belong to the presentation contract of every Card.

- 26. Thesis 26.** Provenance is necessary but not sufficient. Origin and integrity have to be combined with quality checking, accountability and factual correctness; owner release, data quality checking and independent certification remain separate proofs.
- 27. Thesis 27.** Currency needs a live connection. Copies and indexes age. Real time applies only where it is enforced or agreed; with external systems the original is offered, not imposed. Honest guarantee levels are more credible than one promise for everyone.
- 28. Thesis 28.** The brand's release becomes a trust signal. Trusted data does not bypass the channel; it makes authorized Resellers and publishers better and more trustworthy decision points, and the End User can see that he is with the original.
- 29. Thesis 29.** The shelf is a point of contact like a website. Same source, same rules, same attribution: whoever takes the original along at the shelf stays connected to brand and store, and the store becomes as measurable as a web shop, without the customer losing his anonymity towards third parties.
- 30. Thesis 30.** Reliable automation needs a reliable knowledge foundation. AI takes over processes, software and decisions; the one task that remains is the decision which version applies. Whoever organizes his knowledge as an accountable original builds every AI on it; whoever does not automates his mistakes.

Appendix B: Key arguments per target group

This appendix collects the arguments that partners, press and sales ask about most often: what does each group gain in concrete terms, and how can that be said defensibly? Every statement is substantiated in the paper, and the chapter references lead to the derivation. The wordings may be used freely in conversations, offers and posts.

For brands, manufacturers and OEMs

- **Build once, use it many times under control.** Instead of running a separate knowledge project for every channel and every AI, one original emerges that is used across all systems (Chapter 3).
- **The original works where the decision is made.** Card and pluggy bring the original knowledge to the Point of Sale, to the Point of Interest and into the AI answer, before customers reach the website (Chapters 1 and 12).
- **Owner, provenance, quality and currency stay part of the presentation.** Old and contradictory copies lose against a connected Card with a recognizable version (Chapters 8 and 13).
- **The channel gets stronger instead of being bypassed.** Authorized Resellers and publishers are released in a targeted way and become verifiable advisory points; specialist retail remains the commercial path (Sections 14.2 and 13.4).
- **Content turns into advice, sales and service.** pluggy turns product knowledge into advice, recommendation, support and after-sales, around the clock and in every language (Section 10.3).
- **Demand becomes visible.** Questions, saves, follows, Stacks and Advisory Matches show which information customers and partners actually need (Section 10.5 and Chapter 12).
- **Less dependence on individual gatekeepers.** The original is distributed independently of model and platform; no single provider decides on visibility (Chapters 7 and 15).

Defensible wording: "Your product data is copied everywhere today. With plugilo you bring your recognizable original via authorized partners to every relevant Point of Sale and Point of Interest, and your customers take it with them."

For Resellers, distribution and specialist retail

- **The brand's current original instead of self-maintained copies.** Released original knowledge replaces uncertain adoptions and reduces manual content maintenance as soon as the source connection is active (Section 14.2).
- **Advice at the Point of Sale without building knowledge yourself.** pluggy supports product selection, compatibility, application and support exactly where specialist staff are missing (Section 10.3).
- **Your own release is visible and verifiable.** As an authorized partner, the business stands apart from unauthorized marketplace offers and purely price-driven providers (Section 13.4).
- **The customer stays connected.** When the End User takes a Card along through the business, that business remains visible as the introducing partner, and the relationship continues after the visit (Sections 10.2 and 13.5).
- **The store becomes as measurable as a web shop.** With plug-it in the Store, the customer takes the original along at the shelf, the store remains recognizable as the point of contact for a later purchase, and the conversion of the shelf becomes visible (Sections 10.4 and 14.2; in development).
- **Missing brands can be requested.** Resellers make demand in the channel visible and bring the originals of the next brands into their own channel (Section 19.4).
- **plugilo becomes a product of your own.** In the plugilo Solution Partner Program, system houses, IT Resellers and agencies sell, implement and support plugilo at their business customers and earn on a recurring basis (Section 16.3).

Defensible wording: "As an authorized partner you do not show just any product description, but the current and verifiable connection to the brand."

For publishers, media, associations and influencers

- **Structured original knowledge with clear provenance.** Editorial assessment and brand facts remain separate layers with their own sender; nothing is mixed (Section 8.1).
- **Content becomes savable, comparable and usable for the long term.** Tests, rankings and topic worlds live on as Cards and Stacks instead of disappearing after publication (Section 14.3).
- **Reach turns into compensated advisory context.** Publishers run their own Skills, open Channels and receive a share for every Card opened and every valid Match, without ad banners (Section 16.4).
- **Associations carry the industry intelligence.** Chapters, shared data standards and a common voice toward global platforms (Chapter 6 and Section 20.1).

Defensible wording: "You combine your editorial credibility with clearly marked original knowledge, without mixing the senders."

For End Users: consumers, prosumers and business End Users

- **Recognize what comes from the owner.** The trust view shows owner, provenance, currency, version, audit status and partner release (Section 8.1).

- **Ask questions, see the gaps.** pluggy answers with sources, names what it does not know and hands over to a human when needed (Section 10.3).
- **Take the original with you.** The result sits in the Wallet, stays current and does not have to be searched for again; the connection to brand and partner is preserved on request (Section 10.2).
- **Keep control.** Following, saving, personal details and ending the connection are decided by the End User; anyone who shares nothing stays anonymous (Sections 6.2 and 21.3).

Defensible wording: "Take the original with you. The Card stays current, answers questions and connects you with the brand and the partner."

For AI providers, standards organizations and measurement partners

- **One connection, many owners.** Instead of licensing thousands of individual sources, one gateway connects the models to pooled, certified owner knowledge, including provenance, status and licensing clarity (Chapter 7 and Section 20.5).
- **Better answers, less liability.** Source-bound, current original data reduces wrong statements and liability risks (Section 14.7).
- **Standards are strengthened, not replaced.** GS1, ECLASS, ETIM and the testing organizations supply identity, classification and evidence; the Card makes them visible and usable for advice (Sections 8.2 and 15.3).
- **Signals that panels do not see.** Measurement partners receive consented decision events before, at and after the purchase, neutrally classified (Sections 15.3 and 12.1).

Defensible wording: "Providers bring truth. Partners bring reach and context. End Users bring demand. plugilo connects it all."

The shortest version for every group: Trust begins with the recognizable original. plugilo brings it via authorized partners to every Point of Sale and Point of Interest and turns it into a lasting, voluntary connection.

Glossary

This document is called Master Paper throughout, in German as in English: the reference document of the Trusted AI Advisory Alliance. In-depth papers carry their own purpose-specific titles. The terms are firmly assigned: the Trusted AI Advisory Alliance (in German: die Allianz) is the association of the data owners with its Charter. The Trusted AI Advisory Network is the emerging network, this Master Paper is the reference document, and plugilo Solutions and Services are the commercial products and services of plugilo.

- **Match (Advisory Match):** Arises when pluggy recommends, in a Session, a released Card from a sub-Skill that fits the expressed need, in the Network as in Channels. The respective data owner pays for it from the plugilo balance, in addition to the operator's Session; per Session at most one Match per data owner.
- **Session (Advisory Session):** A coherent advisory conversation between End User and pluggy; several questions and answers belong to the same Session, page views and technical calls do not count. All Sessions of a pluggy are paid by its operator, the owner

of the master Skill, in the monthly package per domain; Sessions of the package expire at the end of the month, reloaded ones do not. Nobody shares in Sessions.

- **Authorized Card Point (planned):** Verified Reseller or publisher that the brand has released for specific Master Data Cards at named Points of Sale or Points of Interest; verifiable, purpose-bound, revocable (Section 13.4).
- **Operator:** Anyone who shows original knowledge on his website, in his shop, portal or store, or advises with it: retailers, shops, portals, publishers. Operators receive Rewards for Touchpoints and Matches on their site and pay the Sessions of their own pluggy. A company can be data owner and operator at once (Chapter 16).
- **Browser Extension:** Makes practically every suitable web page connectable, with safe-site filtering and under the user's full control. The plug appears as a new punctuation mark in front of brand and product names and opens the original data.
- **Business Graph:** See Connect Service & Business Graph; relationships stay bound to companies and functions when a contact person changes. Not yet built (Chapter 17).
- **Business Profile / Personal Profile:** Employee Cards belong to the company domain and stay with the position; personal Cards belong to the individual alone.
- **Canonical Truth Backbone:** The source-bound, versioned backbone of every Card. It separates the certified object core, the authorized live state, attributed context, and personal signals; context never overwrites the core.
- **Card ID:** The unique identifier of every Card in the target state of the data hub. It connects all existing numbers and standards of an object, so that every system can retrieve the correct, owner-authorized version automatically.
- **Card anatomy (5+4+1):** Five base layers (cover, details, specifications, media, connections), four performance layers (functions, AI advice, trust & provenance, knowledge graph), and one mandatory machine layer. The global order is fixed, and missing layers are skipped.
- **Card:** Structured, human- and machine-readable representation of an object, generated from the original data set, directly queryable, and embeddable anywhere.
- **CF1 (ContentFactory1):** The wholly owned subsidiary of DCI AG and the group's data services company: around 200 specialists for creating, structuring, harmonizing, and maintaining product and manufacturer data. Its services are open to all partners of the Alliance (Section 10.6).
- **Chapter / founding circle:** The national or industry-related carrier structure of the Alliance; founding circles are the limited starting formations per industry.
- **Connect Service & Business Graph:** Owner-controlled, company-centered relationship layer between companies, positions, and people, with an open API.
- **Connection Card / Collector Edition (planned):** Personally activated, non-transferable relationship of an End User to an object, a brand and a partner, or a genuinely limited Card edition tied to an occasion, a place, a period or a print run; the original knowledge stays separate from it and open (Chapter 17).
- **Presentation contract:** The mandatory fields that no integration may remove: source, owner, quality, version and the return path to the owner (Section 11.1).
- **Data Quality Profile:** Measurable quality profile of a Master Data Card: completeness, currency, consistency, traceability, and coverage, reported separately from the Level.
- **Data owner:** Anyone who is accountable for original knowledge: manufacturers and brands as well as publishers, associations, service providers or retailers with their own

offering. Data owners pay for Touchpoints and Matches beyond their own pages and determine which sources are used and where their Cards may appear (Chapter 16).

- **Data gateway:** Another name for the Gateway (see there); currently unclaimed (Chapter 4.2).
- **DCI AG:** The parent company of the group, founded in 1993, today a holding company with little operating business of its own, listed on m:access of the Munich Stock Exchange. The group includes CF1 and plugilo Inc. (Section 21.1).
- **Demand Trainer:** The learning layer above plugilo Wallet and plug-it, not a separate product. From what End Users save, share, and organize, the network learns real demand, as explicit intent instead of inferred attention. The customer is thus received like a regular customer everywhere in the network.
- **Digital Product Passport (DPP):** EU framework for machine-readable data accompanying products (ESPR); the Master Data Record maps the relevant requirements.
- **Engagement levels:** Active users save, compare, and ask; passive users stay connected through follows and notifications; inactive users return through a trigger (Section 10.2).
- **Follow Service:** Following instead of newsletters: content appears as Cards from labeled original sources of the real owners.
- **Five levers of scaling:** Alliance and standards partners; brands as the engine; the plug as a standard; compensated reach partners; the network and data cycle (Section 19.6).
- **Guarantee levels (A–D):** Four levels at which currency and presentation can be promised defensibly: enforced on plugilo-controlled surfaces, agreed with integrated partners, offered to external AI systems with a connector, not controllable in the open internet (Chapter 11).
- **Gateway:** The gate through which AI systems retrieve the approved originals: source-bound and rights-controlled, opened by the owners, open to every serious system (Chapter 7, Theses 16 and 18).
- **Intensity levels:** Focused Start, Accelerated Rollout, Global Standard Rollout: three levels at which the five levers are driven.
- **Master Data Card Level (1–5):** The governance dimension of the trust model: Level 1 “Structured Reference”, Level 2 “Owner Connected”, Level 3 “Owner Sourced”, Level 4 “Owner Approved & Reviewed”, Level 5 “Managed Source Truth”. From Level 2 upward the Card carries the blue plugilo seal, Level 5 additionally a star.
- **Master Data Card:** The visible and certified form of the Master Data Record, with a five-level maturity model and the principle “Certified first”.
- **Master Data Layer, Connect Layer, Application Layer:** The three layers of the infrastructure: the owner-controlled original base, the distribution to every point of decision, and the services for people and systems (Chapter 9). Every layer is open and can be connected on its own.
- **Master Data Record:** Owner-controlled, source-bound original data set in open, AI-friendly structures; existing schemas and standards remain usable.
- **Master Data Trust & Evidence Network:** The outlook on the open network of trust holders: standards, data quality, testing, maintenance, and AI distribution partners, with plugilo as the competitively neutral orchestration, evidence, and distribution layer (Chapter 15.3).
- **Master-Skill:** The leading Skill of a pluggy. It defines the identity, base context, and rules of the advice; every pluggy has exactly one master Skill, and its owner pays the sessions.

Not to be confused with the Master Data Card, the single official information unit within a Skill.

- **Measurement partners:** Measurement and market research companies that classify consented network signals neutrally and combine them with market and retail data into market intelligence. Their measurement and market data never enters the owner-controlled object core (Chapter 15.3).
- **Meta level / plugilo meta model:** The open, source-bound reference layer above all AI systems. The planned meta model implements it as a trained mediation layer for allocation and discovery; answers stay bound to the certified source, and no knowledge store of its own is created.
- **Demand Intelligence (beta):** The return channel of the network: aggregated, privacy-compliant demand signals from saving, comparing, questions, and transitions, made available to the owners concerned through the Partner Portal (Section 10.5). The personal demand profile stays with the user.
- **Network and Channels:** The two reaches of plugilo Cards. Network: the data owner's own pages (free there) and his existing partners. Channels: new audiences at publishers, trade portals, communities and influencers, booked via the Partner Portal; prices by audience (Chapter 16).
- **Neutrality:** Used in this Master Paper in two precisely defined senses. First, **competitively neutral:** plugilo competes with no data owner, neither with brands nor with retail, distribution, media, or associations. Second, **model-neutral:** the data is open to every serious AI system, with no exclusive tie. What is not meant is neutrality of stewardship: plugilo is the commercially backed operator of the infrastructure. Binding force comes from published rules, verifiable criteria, and the owners' release, which can be revoked at any time (Chapter 21).
- **Object-centered conversation:** Exchange directly at a Card or Stack instead of in anonymous groups; specialist advisors and AI can take part, and owner updates flow in automatically.
- **Open protocols (MCP, A2A, ACP, UCP, GS1 Digital Link):** The connection standards of the agent economy. Model Context Protocol connects agents with tools and knowledge sources, Agent2Agent connects the agents with each other. The OpenAI Agentic Commerce Protocol and the Google Universal Commerce Protocol lead from product search and advisory to commerce, GS1 Digital Link leads from the physical product identifier to the official online information. plugilo supports these standards instead of developing competing protocols of its own (Chapter 7).
- **Partner Portal:** The control center of the owners: profile creation, data maintenance, releases, rights, domain administration, partners, and analytics.
- **plug-it:** The plug at the point of action: widget or plugin for shops, portals, and media. It is the universal saving function, the sharing button, and the basis of the imprint standard.
- **plug-it in the Store (in development):** The plug at the shelf: scanning the price label, QR, NFC or barcode brings the physical product as a Master Data Card into the Wallet; the store remains recognizable as the Card's place of origin (Section 10.4).
- **pluggy:** The Trusted AI Advisory Agent, the network's trusted AI advisor. He answers according to the principle "Certified first", statement by statement and scope by scope: for every statement, the certified source that governs it is what counts (Chapter 8.1). Every answer is backed by a source ("pluggy advises. plugilo makes pluggy trustworthy."). Target state: the personal knowledge manager who fills, organizes, and keeps up to date the user's Wallet on their behalf.

- **plugilo balance:** The balance loaded in advance from which Touchpoints, Matches and Sessions are paid; it is consumed only for valid use and shown per unit in the Partner Portal.
- **plugilo Wallet:** The End User's personal knowledge and demand wallet, usable as a mobile app and as a web portal on the desktop, also called the Trusted Data Wallet. In it, the End User opens up the certified originals as lasting assets, curates them in Stacks, and turns scattered finds into a structured decision memory, with an export guarantee and granular, revocable approvals. In perspective, secured by the planned blockchain anchoring, transactions on the basis of Trusted Data are also intended to become possible; that does not make it a payment or crypto wallet.
- **Reward:** The operator's individually agreed share of the price of the Touchpoints and Matches on his site: for Touchpoints to the operator of the domain, for Matches to the operator of the pluggy. For Resellers and distributors in the Network also as an MDF share with the data owner; for publishers, communities and influencers a new, demand-driven revenue source. Nobody shares in Sessions.
- **ROPO (research offline, purchase online):** The customer gathers information in the store and buys online later; without attribution the store loses the contact. plug-it in the Store records the store as the Card's place of origin.
- **Service Store (planned):** The open connection for standardized third-party services to Cards and Stacks, following the app store principle in the Partner Portal (Chapter 17).
- **Signal layer:** The totality of the consented, explainable decision events before, at, and after the purchase, each with context (product, category, market, contact point, time). plugilo generates the primary data, measurement partners turn it into market intelligence (Chapter 12).
- **Skill:** A company's approved knowledge base from which pluggy advises: the connected sources, Cards, and Stacks. Depending on where it is used, the same Skill can be a master Skill or a sub-Skill; an external Skill becomes active only if owner and operator both agree, and any approval can be changed at any time.
- **Solution Partner Program (Sell & Manage):** The third role of the business model: system houses, IT resellers and agencies use plugilo themselves, sell it to their business customers, introduce it, support it and share recurrently in the ongoing revenue (Partner Revenue); the customer relationship stays with the Solution Partner, the terms are governed by the Solution Partner Agreement (Section 16.3).
- **Source Truth / Trusted Data:** Source Truth is the owner-controlled, provenance-secured, and versioned form of an object; Trusted Data is the collective term for data secured in this way in the network. Both terms describe responsibility and provenance, not infallibility (Chapter 8).
- **The Mirror (beta):** The AI visibility analysis for getting started. With GEO tools it shows how the large AI systems talk about a brand and its products today (Section 10.3).
- **Stack:** Curated collection of Cards: topic world, ranking, project, comparison, campaign, or product range. Followable, shareable, and portable as a whole.
- **Standard Stacks:** Pre-structured collections along the major areas of life and business: the immediately understandable starting point for every collection.
- **Standards, Evidence & Certification Passport (Evidence Passport for short):** Module of the Master Data Card for external evidence: standards, mappings, tests, certificates, and compliance documents, each with organization, standard, scope, date, and validity. Evidence remains attributed to the issuing organization.

- **Sub-Skills:** Additionally approved Skills that add matching original knowledge to the master Skill. A pluggy can use any number of sub-Skills, the master Skill always stays in the lead; the owners of the sub-Skills pay the resulting Matches.
- **Touchpoint:** The active opening of a released Card at the point of sale or point of interest, by mouse-over, click or opening; page views and impressions do not count. Free on the data owner's own pages; at partners in the Network and in Channels the data owner pays for it from the plugilo balance, and the operator of the domain receives a Reward.
- **plugilo Cards and pluggy (two products):** The two products of the business model. plugilo Cards: the owner's Master Data Cards appear in real time as the original on his own pages (free), at partners in the Network and in booked Channels (Touchpoints and Matches). pluggy: the operator's optional advisor, monthly package per domain by Sessions, free package for testing. Three ways to use it: show the Cards of others, display your own Cards, add pluggy (Chapter 16).
- **Trust chain:** The traceable connection from the data owner through the Card and the authorized partner to the End User and his voluntary relationship (Section 13.5).
- **Trusted AI Advisory Network:** The lasting connections that arise from Touchpoints and Sessions: customers save Cards, follow brands and retailers and take advice and results along in the Wallet; the direct connection between data owner, operator and customer, independent of individual gatekeepers (Section 16.4).
- **Trusted Product Truth Infrastructure for AI:** The category plugilo claims: the independent trust and connection layer between data owners and AI systems (Chapter 15).
- **Verified Original:** What the blue seal states: original source confirmed, directly connected to the data owner, provenance, release and updating traceable; no blanket independent certification of all content (Section 13.3).
- **Trust view:** The view of a Card that shows owner, provenance, currency, version, audit status, partner release, advice and relationship at a glance (Section 8.1).
- **Truth / data truth:** The programmatic guiding term of this Master Paper: the authorized, provenance-secured, versioned, and current version of the responsible party, sharpened as the counter-concept to the web of copies. Authority over provenance is no proof of factual quality. That is why the Data Quality Profile, the Evidence Passport, and Conflict Records complement the statement of the owner. Conflicts stay visible, and nothing is silently overwritten.
- **Truth layer / meta-level / reference layer:** Three names for the same thing from three angles: the layer of owner-controlled originals (truth layer), its position above all systems (meta-level), and its function as the point of reference that systems point to (reference layer). The Gateway is the access to this layer, and the signal layer is what flows back from its use.
- **Knowledge island / N-times problem:** The fragmentation that arises when every AI system imports, breaks up, indexes and displays the same brand knowledge separately; the brand delivers N times and controls no version (Chapter 3).

Sources and evidence

The following sources support the central figures and statements of this Master Paper. For notes on verification, warranty, and liability, see the section "Legal notices, data protection, and disclaimer".

¹ Studies of click behavior in web search (SparkToro/Datos, 2024) report that 58.5% of US and 59.7% of EU Google searches end without a click on an external website; the panel measurement also covers search refinements and abandoned sessions and excludes iOS, app, and voice searches. The effect of AI-generated answer boxes is documented separately by a Pew analysis (2025): when an AI summary is displayed, the click-through rate to external links drops markedly (8% versus 15%; only 1% clicks on sources cited in the summary); publishers report corresponding declines in reach. ² Adobe Digital Insights, 16 April 2026 (survey of more than 5,000 people in the United States: 39% had used AI for online shopping; traffic from AI sources to US retail websites in the first quarter of 2026 up 393% against the first quarter of 2025, a growth rate, not a share of all purchases; average machine readability score of product pages in US retail 66%, not an error rate). Adobe Consumer Report 2026 (4,000 respondents worldwide; 42% of those who use AI for research rely on AI often or always as their main source for advice, shopping, or problem solving; a subgroup, not a share of all buyers). ³ Pew Research Center, 22 July 2025 (browsing data from 900 US adults, March 2025): in Google searches with an AI summary, users clicked on a link cited in the summary in 1% of visits; for details see ¹. ⁴ Tow Center for Digital Journalism, Columbia University, "AI Search Has a Citation Problem", March 2025: test of eight AI search services with 1,600 source identification queries on 200 news articles from 20 media outlets; more than 60% of the source attributions were wrong, including references to syndicated or unauthorized copies; a limited news test, not a universal error rate. ⁵ Edelman, "2026 Trust Barometer Special Report: Brand Growth in an Insular World" (online survey of 17,688 people in 15 countries): 88% name brand trust as an important or critical purchase criterion, quality 89%, price/value 88%; 41% name unpaid third parties (customers, peers, experts) as the strongest influence on their brand trust, 20% the brand's own statement. Survey data describe attitudes and do not prove a causal effect of individual plugilo functions. ⁶ Salesforce, "State of the AI Connected Customer", 7th edition (15,015 consumers and 1,570 business customers in 18 countries): 42% credit companies with an ethical use of AI; 61% say that advances in AI make the trustworthiness of companies more important; 71% consider human validation of AI results important; 72% want to recognize when they are communicating with an AI agent; 71% are increasingly protecting their personal data, and 49% believe that companies use their data to the company's advantage. A company study with its own framing; the values apply to the stated reference groups. ⁷ PwC, "Voice of the Consumer 2025" (21,075 consumers in 28 countries, focused on food and shopping behavior): around 3.6 shopping channels used on average, around four among health- or climate-oriented respondents; 51% assign producers and manufacturers one of the three most important responsibilities for healthy eating. ⁸ Torsten Schwarz et al., "KI-Sichtbarkeit wird zur neuen Marketing-Währung" (AI visibility becomes the new marketing currency), marketing-BÖRSE, 21 September 2026: experiment by Erik Snoeijs (Vice President Corporate Brand Management, Bosch Rexroth) with contradictory AI verdicts depending on the question asked; observation by Burkhard Henn that AI overviews reduce direct traffic while the users who do click through convert better; funding round of Profound (180 million US dollars, valuation 1.8 billion US dollars); Google's test "AI Contribution"; OpenAI ChatGPT Ads with HubSpot and Shopify integration and an agent format (pilot at Wayfair, according to Digiday); note that press releases and customer statements carry more weight in LLM assessments than a company's own marketing. ⁹ Regulation (EU) No 1169/2011 (Food Information Regulation): ingredient and allergen labeling; outdated copies of an ingredient list are directly relevant to consumers. ¹⁰ Lee et al., "Deduplicating Training Data Makes Language Models Better", ACL 2022: duplicates in training corpora and their effects on language models. The ratio of 1:300 in the text is illustrative and does not claim a simple majority vote inside the model; AI search can additionally retrieve web sources without retraining the model weights. ¹¹ Foundational: Shumailov et al., "AI models collapse when trained on recursively generated data", Nature (2024): evidence under laboratory conditions that models measurably lose accuracy and diversity when recursively generated synthetic data progressively replaces real training data. Follow-up research (including Gerstgrasser et al., 2024) shows that accumulating authentic original data can limit this risk. ¹² Google Cloud, "Vertex AI APIs for search and RAG"; Microsoft Learn, "Retrieval augmented generation in Azure AI Search" (data access, chunking, indexing, ranking, permissions, updating, and citation tracking as requirements); OpenAI, API reference "Vector store file batches" (files are processed with a defined chunking strategy for File Search, and the knowledge base is tied to the application and the index). All three document application-specific ingestion and retrieval processes rather than model training. ¹³ Google Merchant Center, "Product data specification": platform-specific mandatory fields, formats, and currency requirements; wrong or missing data can lead to rejections or display errors. ¹⁴ Schema.org, "Product" (vocabulary for products, identifiers, properties, manufacturers, certifications); Google Search Central, "Introduction to Product structured data": structured data increases eligibility for rich presentations, while the concrete presentation is at the discretion of the respective Google experience and can change. ¹⁵ W3C, "Verifiable Credentials Data Model 2.0" (verifiers apply their own rules to the suitability

of claims); W3C, "PROV-O" (provenance model); C2PA, "Explainer" 1.4 and Content Credentials (contentcredentials.org): provenance makes origin and editing verifiable, but does not prove the factual truth of a piece of content.¹⁶ Public product pages (accessed 19 September 2026): Salsify (PIM, PXM, GDSN, syndication, Agentic Commerce), Syndigo (MDM, PIM, GDSN, syndication, retail network, AI surfaces), Akeneo (Product Cloud, Activation, AI Commerce), Icecat (catalog data, syndication, AI), 1WorldSync (GDSN, product content, commerce orchestration), Contentserv (PIM, PXM). On none of the pages reviewed is the end-to-end combination documented as one closed overall logic: a Card object visible to the End User and controlled by the owner, with visible provenance, partner release, the End User taking it along into the Wallet, and a direct return channel; provider statements are self-description and not an independent certification of performance. This assessment is reviewed regularly.¹⁷ A widely cited Gartner estimate (2021, based on research from 2020) puts the average annual cost of poor data quality per organization at around 12.9 million US dollars; methodology and transferability can only be verified publicly to a limited extent.¹⁸ Examples are the content licensing agreements between AI providers and publishers (including Associated Press and Axel Springer with OpenAI, 2023) and the increasing technical access control for AI crawlers by infrastructure providers (including Cloudflare, 2025; default block for newly onboarded domains, existing customers by opt-in, "Pay Per Crawl" initially as a beta).¹⁹ Stanford HAI, "AI Index 2026, Technical Performance", 2026: in the OSWorld benchmark for computer tasks, task success rose from around 12 to 66.3 percent. Benchmark values are not a success rate for business processes; around one third of tasks still fail.²⁰ McKinsey & Company, "The state of AI in 2026: On the road to ROI", 25 August 2026; survey of 1,719 participants in 97 countries (May to June 2026). 40% of respondents from companies with more than one billion US dollars in revenue report scaled AI agents (27% the year before), 22% at smaller companies; 32% did not buy at least one software product or function because of internal development with coding agents; 37% report a positive enterprise-wide EBIT contribution from AI. Self-reported. Supplementary practice signal: Anthropic, "Economic Index Report: Cadences", 26 June 2026, describes hours-long autonomous work through agentic tools; provider data from its own ecosystem, not readily transferable to all industries.²¹ Gartner, "Lack of AI-Ready Data Puts AI Projects at Risk", press release, 26 February 2025: in a survey of 1,203 data leaders (July 2024), 63% reported that their data management practices are not, or not certainly, suited for AI; Gartner predicts that through 2026, 60% of AI projects lacking AI-ready data will be abandoned. A forecast, not a measured failure rate.²² Gartner, "Over 40% of Agentic AI Projects Will Be Canceled by End of 2027", press release, 25 June 2025: analyst forecast that more than 40% of agentic AI projects will be canceled by the end of 2027 because of rising costs, unclear business value or insufficient risk controls.²³ NIST, "Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile" (NIST AI 600-1), 26 July 2024: covers, among other things, confabulation, information integrity and content provenance, and recommends documented provenance and checks before and during operation. W3C, "PROV-Overview", W3C Note, 30 April 2013: provenance described through entities, activities and responsible agents; provenance is an established principle, not a new invention. Microsoft Learn, "Retrieval Augmented Generation (RAG) in Azure AI Search" and "Document-level access control" (continuously maintained): grounding model answers in your own content at run time and checking permissions per query.²⁴ Gartner, "Top Predictions for Data and Analytics in 2026", 11 March 2026: forecast that by 2030 half of organizations will use AI agents to translate governance rules into machine-checkable data contracts. A forecast, not today's adoption level.²⁵ SWIFT: founded in 1973 by 239 banks from 15 countries as a cooperative, operational since 1977, today the global standard for financial messaging. GS1/barcode: industry and retail agreed on the Universal Product Code in 1973, first scan in 1974; today millions of companies use the GS1 standards. Visa: emerged from a cooperative owned by member banks (National BankAmericard Inc., 1970; renamed in 1976). Catena-X/Manufacturing-X: data space alliances carried by German industry with political support (Catena-X since 2021; Manufacturing-X announced in 2022, funded conceptually in 2023, first projects from 2024). GS1 baseline figures (around ten billion barcode scans per day; 118 local GS1 organizations and offices): GS1, "Why is GS1 leading the transition to 2D barcodes?", updated 8 January 2025, support.gs1.org.²⁶ Reach of large brand portfolios according to company statements: Procter & Gamble (around five billion people use the products; us.pg.com, April 2026); Unilever (3.7 billion people daily, brands in 190 countries; unilever.com); Hindustan Unilever (nine out of ten Indian households, more than nine million points of sale; annual report 2025/26); Coca-Cola system (2.2 billion beverage servings a day in more than 200 countries; servings are not distinct individuals); PepsiCo (more than one billion consumption occasions a day in more than 200 countries and territories); Nestlé (sales in 185 countries; Annual Review 2025). The figures apply to portfolios, not to individual brands, and are not adjusted for overlaps.²⁷ Australia as an example: the age restriction law for social networks adopted at the end of 2024 has obliged the platforms since December 2025 to take reasonable steps to ensure that users under the age of 16 do not hold accounts; comparable initiatives are under discussion in further

countries. In parallel, market studies have documented declining organic reach for years; classic mass newsletters are losing click measurability and attention, while curated formats grow.²⁸ Authoritative reference systems: Domain Name System (one authoritative source per zone, provided redundantly via authoritative name servers; caches with an expiry date, RFC 1034 ff.); the DOI system of the DOI Foundation (persistent identifiers that resolve permanently to the authoritative source; several hundred million registered DOIs); Legal Entity Identifier / GLEIF (since 2012/2014, globally unique, centrally resolvable company identifiers); GS1 Digital Link (extension of the barcode into a web reference to official product sources). Technology trends: Retrieval-Augmented Generation / grounding as the established architecture for tying AI answers to real-time sources; Model Context Protocol (open connection standard for AI systems and data sources, introduced in 2024, adopted across the industry since then and transferred into an independent foundation in 2025); C2PA / Content Credentials (provenance standard for digital content carried across industries; broad adoption still young).²⁹ GS1 Germany, "GS1 DQX Services" and "Funktionsweise GS1 DQX", version 2.7 (23 December 2025): GDSN synchronizes product master data through certified data pools; GS1 DQX assesses quality against defined test specifications, and after a successful check the electronic sealing follows; the data quality service is not a platform for providing or distributing master data.³⁰ Official product and developer documentation (as of 30 September 2026) documents the pattern: HubSpot (reading and changing CRM records from Claude, with user permissions and change log), Contentful (editing structured content via an MCP server), Atlassian (Rovo MCP for Jira and Confluence context), Asana (MCP access to the Work Graph), Stripe (MCP tools for the API), commercetools (managed MCP servers with separate user and machine access), Zapier (MCP connection to its app ecosystem, provider figures) and Shopify (Storefront Catalog MCP, migration to UCP). The providers serve as a comparison for the pattern "service core with external AI clients", not as partners of plugilo; what is documented are offered functions, not measurements for plugilo.³¹ Model Context Protocol, specification (as of 28 July 2026): MCP standardizes the exchange of resources, context, and tools between hosts, clients, and servers; the host decides on connection, selection, user consent, and interface.³² plugilo, talk and handout "Master Data meets AI: How original knowledge becomes better advice", ECR iN MOTiON 2026, Bonn, 17 September 2026 (handout as of 15 September 2026); GS1 Germany, press release "ECR iN MOTiON feiert Premiere in Bonn" (2026) on the keynote "Future? Ready! Ready together" by Paul Vetter, General Manager Henkel Consumer Brands; the three theses are paraphrased from the talk.³³ Anja Stolz, "KI skaliert Geschäft – aber kein Vertrauen" (AI scales business, but not trust), marketing-BÖRSE, 18 September 2026.³⁴ Mayer, Davis and Schoorman, "An Integrative Model of Organizational Trust", Academy of Management Review 20(3), 1995: trustworthiness as the attribution of ability (competence), integrity, and benevolence.³⁵ Niklas Luhmann, "Vertrauen: Ein Mechanismus der Reduktion sozialer Komplexität" (1968; in English "Trust and Power", Polity).³⁶ "How do transparency and traceability enhance consumer trust?", Journal of Consumer Marketing (Emerald, 2024/2025): transparency and traceability in supply chains strengthen trust in products; accuracy and traceability have a stronger effect than the volume of information.³⁷ "Owing to interactivity: a meta-analysis on consumer engagement", Journal of Research in Interactive Marketing 19(8), 2025: meta-analysis of 144 studies with 57,952 participants; digital interactivity has a positive effect on brand loyalty and purchase intention through trust and emotional attachment.³⁸ ISO/IEC 25012, "Data quality model" (15 data quality characteristics, edition confirmed in 2025); the standard itself is subject to a fee, and the official page describes purpose and status.³⁹ Study on digital trust seals, Electronic Commerce Research (Springer, 2020, DOI 10.1007/s10660-020-09415-2): it is not the symbol that creates trust, but the institution, the review, and the traceable statement behind it.⁴⁰ Apple, "Apple Authorised Service Provider Programme" (standards, certified technicians, regular audits; a service partner program, not a reseller program); Rolex, "Purchasing a Rolex" (Official Rolex Retailers confirm the authenticity of a watch and activate the warranty); eBay, "Authenticity Guarantee for Trading Cards" (check by a named partner, badge, QR code, documented audit status; only for defined categories and countries).⁴¹ European Commission, "Explanatory note on the VBER and Vertical Guidelines" (2022): vertical agreements, selective distribution, online restrictions, and the protection of authorized distribution systems. Not legal advice; concrete release programs require a case-by-case review.⁴² AI commerce programs and data quality: OpenAI enables merchants to provide product information through specified feeds (expressly also via commissioned service providers) and has been extending the Agentic Commerce Protocol since March 2026 into the connection layer for product search and discovery (including Salesforce, Stripe, and Shopify as integration paths); OpenAI has widened its initial focus on instant checkout in favor of search and discovery. Microsoft Copilot Merchant Program: provision of up-to-date product information; Microsoft describes trustworthy product truth as the foundation of Agentic Commerce. Google: structured product data in combination with Merchant Center feeds; in January 2026, Google launched the Universal Commerce Protocol (UCP) together with Shopify as a further open commerce standard;

for AI Overviews and AI Mode, Google states that it requires no special structured data beyond the usual Search requirements. GS1 DQX: automated and partly manual checking of product master data with electronic sealing, established in particular in the German food and near-food environment. DPP standardization: CEN/CENELEC JTC 24 published the European standards series EN 18216–18223 in spring 2026; the international committee ISO/IEC JTC 5 was established in March 2026 (work starting from September 2026); UNECE is developing a machine-readable DPP with Verifiable Credentials in the UN Transparency Protocol (pilot stage).⁴³ European Commission, press release IP/23/418 (January 2023), screening of online shops for manipulative practices: false countdown timers, hidden information, and steering interfaces as problematic dark patterns.⁴⁴ Digital Product Passport: EU Ecodesign Regulation for Sustainable Products (ESPR, Regulation (EU) 2024/1781, in force since 18 July 2024). It introduces the Digital Product Passport step by step for most physical products in the EU single market (excluding, among others, food and feed and medicinal products); the first work plan (2025) prioritizes textiles as well as iron and steel, among others; specific DPP obligations per product group only take effect with the delegated acts and after transition periods. Batteries: Regulation (EU) 2023/1542 requires the battery passport from 18 February 2027 for EV/traction batteries, LMT batteries (e.g. e-bikes, scooters), and industrial batteries above 2 kWh. GS1 Digital Link, among others, is positioning itself as the data carrier standard; EU preparatory projects (CIRPASS) have developed the technical foundations.⁴⁵ NIST, “Artificial Intelligence Risk Management Framework” (AI RMF 1.0): characteristics of trustworthy AI such as validity, reliability, safety, transparency, explainability, privacy, and fairness; a system can meet these characteristics without users automatically trusting it.⁴⁶ PayPal: Registration Statement Form S-1 (SEC, 2002), sections “Our History” and “Sales and Marketing”: transaction-based invitation, sign-up and referral bonuses of 5–10 US dollars each, account growth from 12,000 in early 2000 to around 10.6 million in September 2001. WhatsApp: “400 Million Stories”, company blog, 19 December 2013 (400 million monthly active users, 100 million of them within four months, without targeted advertising); Facebook, SEC Exhibit 99.1 of 19 February 2014 (more than 450 million users, more than one million new ones per day). Dropbox: Drew Houston, “Startup Lessons Learned” (2010): from 100,000 registered users in September 2008 to four million in January 2010, 35 percent of daily registrations through the two-sided referral program, a lasting plus 60 percent; confirmed in the Form S-1 (SEC, 2018, “Viral, bottom-up adoption”, with more than 500 million registered users at the time).⁴⁷ SKOS (Simple Knowledge Organization System): W3C recommendation from 2009 with semantic hierarchy relations (skos:broader, skos:narrower) within one concept scheme and separate mapping relations (skos:exactMatch, skos:broadMatch, skos:narrowMatch) for linking concept schemes. Google Knowledge Graph: introduced in 2012. Established classification and identifier standards that provide such mappings industry by industry are, among others, eCI@ss/ECLASS and ETIM (product classification), GS1 GPC (product categories), GTIN (articles), LEI (companies), and ORCID (people in research).⁴⁸ On the role of Michael Urban as Industry Expert and investor of DCI AG (stake in plugilo via DCI AG) — company announcements: corporate.pluginilo.com/industry-expert-michael-urbandci-ag-wins-investor-for-pluginilo-expansion and corporate.pluginilo.com/api-ceo-on-the-cooperation-with-pluginilo-the-thing-with-the-plugin-is-cool; external: it-daily.net (“DCI AG gewinnt Investor für pluginilo-Expansion”), register data: northdata.de (DCI Database for Commerce and Industry AG).

Closing note

All information is given to the best of our knowledge; editorial, research, and source status: September 2026. Figures, quotes, and time-dependent statements were checked as of that date and evidenced with the sources indicated. Information on partnerships, product status, and traction reflects the status as of September 2026; where marked, these are the company's own statements, whose evidence is provided in personal conversation or in the data room. Statements on the leadership team, the investor, and the infrastructure have been agreed with those involved and documented.

Figures and tables

Figure 1: The copy cascade: one source becomes hundreds of copies, and in the end the most frequent claim wins, not the correct one.

Figure 2: The same change, two outcomes: without the original the copy wins, with the Master Data Card the current ingredient list wins.

Figure 3: The N-times problem: the same original knowledge is imported, shortened, and interpreted anew in every AI system, and the brand sees none of these versions.

Figure 4: From answering to executing: agent capabilities are growing faster than companies' data foundations; the gap between them is the risk, and it closes only at the original.

Figure 5: Three levels of participation: observing costs nothing, and helping to shape is open to every data owner.

Figure 6: How the End User sees the original: sender, seal, version, three levels and the trust view on one Card.

Figure 7: Trust is answered in three separate answers, not one blanket score: who is responsible, how good, what is verified.

Figure 8: Five levels, one principle: the higher the level, the more responsibility the owner takes on; the blue seal exists only at Level 5.

Figure 9: Each layer can be used and connected on its own; the Master Data Layer carries the whole.

Figure 10: The productive core is live; everything beyond it is marked as an expansion stage with a clear status.

Figure 11: From the shelf into the Wallet: the plug works in the store as it does in the shop; the Card takes its place of origin along, and the store remains recognizable as the point of contact even for a later purchase from home.

Figure 12: "Real time" applies only where it is enforced or agreed; with external systems the original is offered, not forced.

Figure 13: The loop closes: the original goes out with provenance, currency and advice; trust and consented demand come back to the owner.

Figure 14: Every category of the ecosystem is occupied and remains a partner; only the layer above, which connects everything, is open.

Figure 15: Two products, three units, one rule: Cards carry the original, Touchpoints bring it to POS and POI, Sessions fund the advice, Matches create reach and revenue; operators share in the value through Rewards, the End User never pays.

Figure 16: One conversation, three units: the manufacturer pays the Touchpoint and the Match, the Reseller the Session; the Reseller receives the Rewards, the End User never pays.

Figure 17: Every role brings something in and gains something; not the type of company decides but the role of the Skill; plus the End User, who never pays.

Figure 18: Regulation does not ask whether a company organizes its data, but when; the battery passport is mandatory from 18 February 2027.

Figure 19: Five levers that accelerate each other; the intensity at which they are driven sets the pace and the number of markets.

Figure 20: The value begins before any IT integration; systems are connected only once it pays off.

Table 1: Wrong data hits every group differently, but every group concretely.

Table 2: Eleven mandatory features that turn a Card into an original; if one is missing, it remains a representation.

Table 3: The trust view of a Card answers eight questions; each field states what the answer is based on.

Table 4: Four guarantee levels: what is enforced on plugilo's own surfaces can be agreed with partners and, with external AI systems, only offered.

Table 5: The signal layer makes visible the decision path that panels and clicks cannot see.

Table 6: Source, semantics, currency, AI use, provenance: strong building blocks, each within its purpose.

Table 7: Visible original, presentation, partner release, owner relationship, demand loop: this is where the open gap lies, and this is where the target architecture begins.

Table 8: The three preconditions of trust and how the infrastructure delivers them.

Table 9: Four marks, four clearly separated statements: every mark says what it confirms and just as clearly what it does not.

Table 10: Every group gains immediately and on its own account; the sections of Chapter 14 provide the reasoning.

Table 11: Every building block of the market exists; nobody connects owner data, external evidence, and AI advice into one layer.

Table 12: What runs is called beta; what is coming is called planned. Nothing in between is blurred.

Table 13: The focused start is an intensity level, not a test phase.