



Building the connecting layer for trustworthy AI

17 June 2026

Executive summary

AI governance has reached a new stage. Regulation has moved from broad principles to specific obligations, international and regional standards for AI management systems are being adopted, and organizations across the AI value chain are investing in responsible AI to meet what their customers, partners, and the public expect. The foundations for trustworthy AI are now in place.

What remains is the practical means of acting on them: a common, recognized way for an organization to show that an AI system meets the obligations that apply to it, in a form others can rely on across the value chain and across borders. In more established fields this layer is familiar: a building is checked against building codes, an electrical product against published standards. AI is ready for the same.

The Appia Foundation is the international collaboration formed to build it. Its members develop specifications: the assessable criteria that let parties demonstrate conformity with their obligations, building on the foundational standards that bodies such as ISO/IEC develop rather than replacing them. As a Linux Foundation Joint Development Foundation (JDF) project, the work is built in the open by members from across sectors, the AI value chain, and around the world.

This paper explains why this layer is needed, what the Appia Foundation produces, how it is organized, and how to contribute. The trust it enables is meant to be held in common, by everyone who relies on AI.

The shift from principles to practice

A company begins using an AI tool that helps its recruiters screen job applications, so strong candidates reach a hiring manager sooner. Like almost every AI system in use today, it was not built in one place: a language model created by one company, adapted to assess candidates by another, connected to the company's hiring systems by a vendor, and configured by the company's own recruiting team for the way they work.

The benefit is real: faster shortlisting, fewer strong candidates overlooked, more time for recruiters to spend with people. But before any of that, a question must be answered, and it comes from several directions at once. The recruiters and hiring managers who rely on the tool need to trust that it is reliable. The applicants it screens want to know it is fair. The company's leaders, responsible for who it hires and for hiring within the law, need to be confident in it. Auditors and regulators ask for evidence of how it performs. And beyond any one company lies a wider question: whether this is a use of AI that society is prepared to accept.

Each party is asking a version of the same thing: can this be trusted?

Today, that answer is built from scratch each time, in a different form, for a different reader. Yet the groundwork now exists: there are real obligations to meet and real standards to build on, where a few years ago the picture was far thinner. What the ecosystem is now ready to build is the practical layer that turns those foundations into something an organization can show — a common, recognized way to demonstrate that a system meets the obligations that apply to it, in a form a hiring manager, an applicant, an auditor, or a member of the public can rely on.

In more established fields, this layer is so ordinary it goes unnoticed. A building meets building codes because it is checked against criteria everyone recognizes; an electrical product carries a mark because it was tested against published standards; food reaches the table through checks that hold from one country to the next. In none of these cases do standards create trust on their own. Trust comes from a shared, practical way of showing that something conforms to them.

AI is ready for the same kind of layer. It does not yet exist — and the conditions to build it are finally in place.

The connecting layer between standards and assessment

No single organization, sector, or country can build this layer on its own. A way of demonstrating trustworthy AI has value only if others recognize it: the next organization in the value chain, a customer in another sector, an assessment body in another country.

What has to be shared is not the threshold each party sets, but how a system is measured against it. A regulator may require more than a contract does, and one country may set a higher threshold than another; what they can hold in common is the way a system is measured, so a result carries the same meaning wherever it is read.

The Appia Foundation is the collaboration formed to build this shared layer: an international, openly governed effort bringing together participants from across the AI value chain and across jurisdictions. Together they develop specifications that translate foundational standards into assessable criteria — the criteria that let parties demonstrate conformity with the obligations they face, whether those come from law, regulation, a contract, a sector code, a voluntary commitment, or a supply-chain expectation.

Those obligations are where the work begins. Meeting them draws on the technical content of foundational standards from bodies such as ISO/IEC and CEN/CENELEC; Appia specifications carry that content into a form that can be assessed; and assessment evaluates a system against the resulting criteria, producing conformity others can recognize. The path is not always a single line — some standards can be assessed against directly, and some obligations shape the specifications without first passing through a standard.



The Foundation's own part in this "conformity chain" is deliberately narrow. It occupies the connecting layer between the standards that already exist and the assessment that relies on them: it builds on those standards rather than replacing them, takes the obligations that regulation and other sources set as given rather than defining them, and produces the criteria that assessment needs, so that assessment itself stays with those equipped to perform it.

The commitments that earn reliance

The parties who come to rely on these specifications will not have built them, and in most cases will never meet those who did. Reliance of that kind is earned in how the work is done — through a set of commitments that hold from the first specification to the last.

The Foundation builds on the institutions already at work. Its specifications start from foundational standards and stay traceable to them, so that as standards and regulation evolve, the specifications keep pace. This keeps the work aligned with the infrastructure the wider system already relies on, and lets the specifications carry the authority of their sources.

The Foundation's part is technical. Its specifications let a party show, with clear evidence, that an element meets defined criteria: a demonstration of conformity. Conformity is a technical result; compliance is the legal status of having met an obligation. The two are distinct, but they connect: the specifications translate foundational standards and regulatory frameworks into assessable form, so that a demonstration of conformity produces the kind of evidence applicable obligations call for. What that demonstration means for any particular obligation is decided by the framework behind it and the jurisdiction enforcing it, as in every established domain, where rule-makers set the threshold and rely on technical demonstration to show it has been met.

The specifications are modular and role-based. The AI value chain is made up of different actors carrying different responsibilities, so a party demonstrates only what relates to its role, the part of the system it is responsible for, and the context in which it operates. This keeps what a party must show in proportion to what it actually does.

The specifications will be publicly available and built for broad adoption. Their value lies in being widely shared; the Foundation develops them as shared infrastructure for the ecosystem it serves.

The three questions the work answers

Return to the company and its applicant-screening tool. Before anyone can establish that it is trustworthy, three questions have to be answered — the same three that arise for any AI system. Producing what it takes to answer them is the Foundation’s central work.

The first is what, exactly, is being examined. The tool is not a single object: there is the language model it draws on, the candidate-screening version built on that model, the connection into the company’s hiring systems, and the configuration the company itself applied. Each is a distinct thing that can be examined on its own, and a claim about one is not a claim about the others.

Appia specifications give every party a shared way to name what is being examined: a typology of the objects of conformity that make up an AI system, such as a component, a model, a system, or an integrated product or service. With it, everyone knows which part of the whole a given piece of evidence refers to.

The second question is who is responsible for what. The tool was built and refined by several parties, and trust in the whole depends on assigning responsibilities correctly among them. Appia specifications set out how responsibilities are divided across the value chain, so each party demonstrates conformity for the parts it is responsible for.

The third question is what good looks like, in terms specific enough to test. This is the substance of the matter: specifications that state what is required on a given dimension and how it is shown, such as how risk is managed across the life of an AI system.

Together, these three answers are the inputs conformity assessment is built from – modular, scalable, and interoperable – from which conformity can be assessed and assessment schemes developed.

The specifications themselves are organized into two layers. The Requirements and Guidance layer holds what is required, together with the guidance that goes with it. The Assessment Enablement layer holds how those requirements are evaluated: the conformity assessment materials and the typology of objects above. Each of the three answers draws on both layers, what is required and how it is evaluated, rather than sitting in one alone.

In the company’s case, these take concrete form: a specification for the underlying model, setting out what its builder should show about how the model was trained and tested; a specification for the deployment, setting out what the company should show about its own configuration and monitoring; and assessment materials defining how each is checked. These

are the Foundation's tangible outputs: the specific criteria each object of conformity must meet, and the means of checking them.

Two things let this body of specifications work across a value chain as varied as AI's. First, because the specifications are modular and role-based, a party draws only what its role calls for and composes it with the sector and regulatory requirements that already exist — so evidence produced once can serve more than once, rather than anyone starting over. Second, that evidence carries forward: the company does not re-establish what the model's builder already demonstrated, though it remains responsible for its own configuration and use. AI systems are rarely a simple line between two parties. Model developers, adapters, integrators, suppliers, and deployers relate in ways that shift from case to case. Building the specifications to compose and carry forward this way is what lets conformity hold across that complexity. This is the chain introduced earlier, seen from the inside.

Neutral ground, built by its members

Work like this needs a setting where organizations that compete with one another can build shared infrastructure on neutral ground, under rules agreed in advance.

The Appia Foundation operates under the Linux Foundation's Joint Development Foundation, a global, vendor-neutral home for open standards work, which provides exactly that: a convening authority answerable to no single company, government, or standards body; an intellectual property framework that lets participants from rival organizations contribute safely to common work; and the operational means to run a project of this scale. It also carries an established route to international standardization, as an approved PAS Submitter to ISO/IEC JTC 1, for specifications that are ready for it.

The specifications are written by the members, in working groups open to all of them. Initial workstreams include architecture, policy, mapping specifications to applicable obligations, and specific regulatory connection (including the EU AI Act). A layer meant to be recognized across the value chain and across borders has to be written by people who span it — and the parties who will come to rely on these specifications are the ones best placed to get them right.

The initial membership reflects that span. It runs across the AI value chain: organizations that put AI to work in automation, industry, energy, medical, networks, and payments (Ericsson, Mastercard, Mitsubishi Electric, Omron, Schneider Electric, Siemens); the model, platform, and compute providers they build on (Arm, Google, Microsoft, OpenAI); organizations that assess and certify AI systems and the governance tooling that supports them (Nemko, Naaia); and

insurers that rely on credible assessment data to underwrite AI systems (Armillar AI). The membership draws on Europe, North America, and Asia, with particular depth in the jurisdictions where the obligations this layer helps parties address are being written now.

The Foundation will also form an advisory board to bring academia, government, and civil society into the work. The members build the specifications; the advisory board is how perspectives beyond the membership, including the public interest, take part in shaping it.

The initial members are where the work begins, not the limit of it. The membership is meant to widen as more of those who depend on the result take part. What a member gains is a hand in shaping the criteria that will let AI be trusted — across the systems and markets they all work in — while those criteria are still taking form. The Foundation provides the means and the neutral ground; the substance comes from the members.

The invitation to take part

This layer is built by the organizations whose work depends on it, and used by far more who rely on what it produces. The most direct way to take part is to join the Foundation as a member and contribute to the working groups where the specifications are written. They are being drafted now, and what is written in the coming months will shape AI conformity for years. This is the time to help form it.

Others have a part too, and it grows as the work matures. Standards bodies and regulators, whose work the specifications build on, help keep the layer aligned with the standards and obligations it serves. Those who assess and assure AI systems, and those who put them to use, bring the practical experience that keeps the specifications grounded in how AI actually works. The advisory board brings academia, government, and civil society into the work, so perspectives beyond the membership have a place in it.

Wherever you sit, the next step is the same: get in touch, and the Foundation will help you find where you fit. The trust this work enables is meant to be held in common – by everyone who relies on AI, not just those who build it.

[Visit appiafoundation.org](https://appiafoundation.org) or email membership@appiafoundation.org to get started.