


BEYOND POLICY READINESS

GDPR Is a Data Programme Wearing a Legal Badge

What Two Years of Readiness Programmes Have Revealed About the Real Demand

Giovanni Leonardi

April 2018



GDPR Is a Data Programme Wearing a Legal Badge

Programme · No. 178

Originally written April 2018 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“The regulation's force — and, for many programme teams, its cruelty — is that it has encoded data management maturity into law.”

1. The Legal Green, the Data Red

The readiness programme board was green across its legal workstreams. Privacy notices had been rewritten. Consent mechanisms were being redesigned. A Data Protection Officer had been appointed — reluctantly, but appointed. The policy framework was, by any reasonable measure, comprehensive.

Then someone asked the question that has quietly derailed more GDPR programmes than any regulatory ambiguity: if a customer exercises their right to erasure next month, can we actually delete their data?

The room went quiet. Not because anyone doubted the legal right — Article 17 is unambiguous. But because nobody present could say, with confidence, which systems held that customer's personal data, which third parties had received copies, or what the technical sequence for verified deletion would look like across a landscape of forty-odd systems accumulated over two decades. The policy promised compliance within thirty days. The technology estate promised nothing at all.

I have watched this scene play out — with variations in the cast and the specific question — across readiness programmes in financial services, telecommunications, retail, and the public sector over the past eighteen months. The details change; the pattern does not. Organisations that handed GDPR to their legal function are discovering, with enforcement weeks away, that the regulation's true demands are not legal demands. They are data demands. And the programmes that treated this as a legal compliance exercise have, in too many cases, spent two years building the wrong thing.

2. Why Organisations Gave It to the Wrong Team

The misdiagnosis is understandable. GDPR arrived as legislation. It was drafted by lawyers, debated in legislative chambers, and published in the Official Journal of the European Union. Its language is dense with legal constructions — lawful bases for processing, legitimate interests, data protection impact assessments, supervisory

authorities with corrective powers. Any sensible executive, presented with Regulation (EU) 2016/679, would hand it to the General Counsel's office and ask for a readiness plan.

And Legal, to its credit, did what Legal does well. It interpreted the text. It mapped obligations against the old Data Protection Directive. It produced gap analyses. It drafted privacy notices that balanced regulatory precision against readability. It built consent frameworks and negotiated data processing agreements with third parties under Article 28. It designed training programmes so staff could recognise a data subject access request when one arrived.

All of this was necessary. None of it was sufficient. The two-year transition period was consumed by legal interpretation and policy production, and the question of whether the organisation's systems could actually deliver on those policies was deferred — sometimes consciously, more often through a quiet assumption that "IT will sort it out" once the legal framework was settled.

The assumption has proved expensive. In the programmes we have been close to, the cost of retrofitting data capability in the final six months has routinely exceeded the entire spend on legal readiness in the preceding eighteen. The ratio tells you where the real weight of this regulation lies.

3. The Real Demand Beneath the Legal Language

Strip away the legal apparatus and the regulation asks a set of questions that are, at root, questions about data:

- Where is personal data? Article 30 requires records of processing activities — which is a formal way of demanding a data inventory that most organisations have never built and some have actively avoided building, because the answer is uncomfortable.
- Can you produce it on request? The right of access under Article 15 requires an organisation to gather all personal data it holds on an individual and deliver it in a structured, commonly used format within one month. This presupposes a level of data lineage and retrieval capability that is simply absent in most enterprise architectures built over the past fifteen years.
- Can you delete it? The right to erasure under Article 17 requires verified deletion across every system and every third-party processor that holds a copy. In organisations with decades of system accumulation, shadow copies, and data duplication across operational and analytical environments, this is an engineering problem of genuine difficulty — not a policy problem.
- Can you demonstrate your basis for holding it? The accountability principle under Article 5(2) requires not merely compliance but demonstrable compliance — an auditable trail connecting every category of personal data to a lawful basis for processing. This is a data cataloguing and governance challenge, not a legal drafting exercise.

The regulation's force — and, for many programme teams, its cruelty — is that it has encoded data management maturity into law. Two decades of underinvestment in data architecture, data quality, and data lifecycle management have become, overnight, a

compliance exposure with sanctions attached.

4. What the Working Programmes Look Like

The readiness programmes that are approaching enforcement in genuinely good shape share a defining characteristic: they are, in their bones, data programmes. They began not with legal interpretation but with data discovery. Before anyone drafted a privacy notice, they mapped where personal data actually lived — not where the architecture diagrams said it should live, but where it actually was, including the spreadsheets on shared drives, the legacy databases that nobody wanted to touch, and the test environments still running with production data from a migration three years ago.

From that foundation, they built outward. The Article 30 records of processing became a natural output of the data inventory rather than a standalone compliance document produced from interviews. The technical capability for deletion was designed against the actual system landscape, tested with realistic data volumes, and timed to confirm that the one-month response window was achievable in practice, not merely in principle. Data subject access request processes were engineered as data retrieval workflows with exception handling for the systems that could not be automated in time.

In these programmes, legal expertise sits alongside the delivery team — advising on interpretation, validating the basis for processing, reviewing impact assessments — but it does not lead. The programme director is typically someone with data management or technology delivery experience. The workstreams are organised around data capabilities, not around legal articles. The legal framework wraps around the data reality, rather than the other way around.

The readiness programmes that are working are data programmes with legal advisors, not legal programmes with IT support. The distinction is not semantic — it determines what gets built, what gets tested, and what will actually function when the first data subject request arrives.

5. The Uncomfortable Position — and the Honest Response

For the organisations that took the legal-first path, the position with seven weeks to enforcement is uncomfortable but not beyond recovery. The policies, the notices, the consent frameworks — these are not wasted work. They are necessary components of a compliant posture. What is missing is the operational machinery to deliver on what those documents promise.

The temptation now is to attempt to close every gap simultaneously. We see programme teams trying to build a complete data inventory, redesign consent mechanisms, implement deletion capability across every system, retrain all staff, and renegotiate every third-party processing agreement in parallel, all against the May deadline. The result is predictable: nothing reaches genuine operational readiness, and the programme board

returns to the comforting fiction of percentage-complete metrics that measure activity rather than capability.

A more honest — and ultimately more defensible — approach is to triage. Three operational capabilities matter more than the rest in the opening months of enforcement:

1. Data subject request response. The ability to locate, retrieve, and where required delete an individual's personal data across the systems that hold it. This is the capability that individuals will exercise from day one, and failures will be visible to both the data subject and the regulator.
2. Records of processing. A defensible, operational record of what personal data is processed, on what basis, and where it flows. Not a perfect inventory — a credible one, maintained as a living document. The supervisory authorities have signalled clearly that a genuine, documented effort will carry more weight than a polished compliance pack that does not reflect reality on the ground.
3. Breach detection and response. The seventy-two-hour notification window under Article 33 demands that an organisation can detect a personal data breach, assess its scope and severity, and report to the supervisory authority — a chain of capability that depends on monitoring, escalation procedures, and rehearsed response, not on legal drafting.

Everything else — the refinement of consent mechanisms, the completion of every impact assessment, the renegotiation of every processor agreement — matters, but it is secondary. The enforcement posture that supervisory authorities across Europe are signalling, in these early months at least, favours organisations that can demonstrate genuine operational capability and a credible programme for continued improvement over those that present a complete policy library but cannot execute it.

6. The Broader Pattern

There is a point here that extends beyond GDPR, and it is worth stating plainly. The pattern of giving a data problem to a non-data function — because the problem arrived dressed in legal language, or regulatory language, or board-level language — is not unique to data protection. We see it in anti-money laundering programmes led by compliance functions that lack the data analytics capability to make transaction monitoring effective. We see it in regulatory reporting owned by finance teams that spend more time reconciling spreadsheets than engineering the data pipelines that would make reconciliation unnecessary.

In each case, the owning function does what it knows how to do. Policies are written. Frameworks are designed. Training is delivered. And the gap between what is committed on paper and what can be delivered in practice quietly widens until a deadline, an incident, or a regulator exposes it.

GDPR has forced this conversation into the open with unusual clarity. The right to erasure alone has done more to expose the true state of enterprise data management than a decade of data governance maturity assessments. If there is something valuable in the discomfort of these final weeks, it is this: the organisations now confronting the distance between their legal commitments and their data capabilities are seeing, many of

them for the first time, what their data landscape actually looks like — and what a serious data programme would need to address.

Whether they act on that clarity after May, or retreat into the comfort of periodic policy reviews, will determine whether this readiness effort becomes a genuine turning point in how the organisation manages its data — or merely the most expensive compliance exercise of the decade.