PORTFOLIO BOOK SERIES · No. 77

Perspective

BEYOND COMPLIANCE THEATRE

GDPR Did Not Create the Data Governance Crisis — It Exposed It

What the first year of compliance revealed about ownership, portfolios and machine-led decisions

Giovanni Leonardi

May 2019

GDPR Did Not Create the Data Governance Crisis — It Exposed It

Portfolio · No. 77
Originally written May 2019 | Re-edited and re-rendered for the WhereBeyond Edition,

"A data owner who cannot stop a release, fund remediation or accept risk is not an owner; they are a name in a register."

1. The morning after compliance

The request looked routine. A customer wanted a copy of her information and then asked for part of it to be erased. The privacy team opened a ticket, the service desk sent it to six application owners, and the search widened. A marketing extract sat outside the customer database. A risk file used a different identifier. Two spreadsheets had been emailed to an external fulfilment partner. The analytics team then disclosed that the customer's history was present in a model-development sample, although nobody could say whether the sample was still being used.

Seventeen systems, eleven working days and three senior escalations later, the request was answered. The organisation had met its statutory timetable. It had also exposed the truth: its data governance existed as policy, but not as an operating capability.

One year after the General Data Protection Regulation took effect, this is the lesson that matters. GDPR did not create the enterprise data problem. It supplied a demanding, public test of arrangements that had long been tolerated because their failure was dispersed across reconciliations, delayed reports, manual workarounds and arguments about whose number was correct.

The prevailing response has been to treat privacy as a legal and compliance programme. That was understandable for the immediate deadline. It is no longer sufficient. The durable response is to use GDPR as a reset of enterprise governance: to connect purpose, ownership, lineage, quality and decision rights to the portfolio choices that create and change data.

© WhereBeyond, 2024–

> *The regulation asks questions that only an operating model can answer: What information do we hold? Why do we hold it? Where did it come from? Who may use it? Who can correct or remove it? A policy can state the answers; only governed processes can keep them true.*

2. What the regulation actually tested

Before May 2018, many organisations could appear well governed while relying on three convenient ambiguities.

First, ownership meant sponsorship. A senior executive's name appeared beside a data domain, but the role carried no authority over project scope, operational controls or remediation budgets. A data owner who cannot stop a release, fund remediation or accept risk is not an owner; they are a name in a register.

Second, the system was treated as the unit of control. Application inventories recorded where information was processed, yet a customer's data moved through interfaces, extracts, document stores and suppliers. Rights belong to the person, not to the application boundary. The moment a subject access or erasure request crosses systems, a system-centred inventory becomes a map with the roads missing.

Third, data quality was separated from lawful use. Teams asked whether a record was accurate without asking whether it was still necessary, whether its source was understood, or whether it was being used for a compatible purpose. GDPR made those questions inseparable. Poor provenance is no longer merely an inconvenience for reporting; it weakens the organisation's ability to justify the processing itself.

The difference is visible in practice:

Paper governance	Operational governance
Names an accountable executive	Gives that executive defined decisions and escalation rights
Lists applications	Traces critical data across processes, interfaces, extracts and suppliers
Publishes retention rules	Applies deletion or anonymisation controls to live stores, archives and copies
Records consent or another basis	Connects the basis to each purpose and downstream use

© WhereBeyond, 2024–

Paper governance	Operational governance
Approves a project at a gate	Re-tests purpose, access and retention when the design changes
Logs quality issues	Assigns remediation funding and prevents recurrence at source

This is why cataloguing technology, by itself, cannot repair the problem. A catalogue may make assets visible, and that is useful. It does not decide whether an extract should exist, whether a disputed field may be reused, or who pays to correct an upstream control. Those are governance decisions. Technology records their consequences; it does not supply their legitimacy.

3. The portfolio is where ownership becomes real

Most broken data processes are manufactured by otherwise rational change. A programme adds a field to support a new service. A reporting project copies it into a warehouse. A digital initiative acquires a separate customer identifier to meet its timetable. A predictive model reuses the history for a purpose that was not considered when the data was collected. Each decision is defensible within its own project. Together they create an estate that no single project owns.

The pattern I have observed is that privacy teams are then asked to govern the accumulated result without influence over the investment decisions that produced it. They receive impact assessments late, after suppliers are selected and designs are politically committed. Their remaining choices are to approve with conditions that operations inherit, or to delay a launch and become known as the obstacle.

A composite portfolio review illustrates the mechanism. Fourteen active initiatives touched customer information. The initial register showed nine named data owners and sixty-three critical data elements. When the review asked owners to make three decisions—approve the purpose, accept the quality threshold and authorise the retention period—only twenty-one elements received all three decisions. Thirty-two had an owner who could advise but not commit funding. Ten had different owners in different programmes.

The response was not another policy. The portfolio office added a data decision to the investment process. Before design authority approval, every initiative had to show:

- the business purpose for each new or materially changed use of personal data;
- the accountable data owner and the decisions attached to that role;
- the source-to-use path, including transfers to processors and analytical copies;

© WhereBeyond, 2024–

- the quality, access and retention controls that would operate after launch;
- the funded owner of any remediation that the change created.

In the first six weeks, four projects changed their designs. One removed a duplicate customer store. Another replaced a free-text field with a controlled value because the information would feed a risk model. Two combined supplier assessments that had been proceeding separately. None of those outcomes came from a privacy lecture. They came from putting evidence and decision rights at the point where money and design were still movable.

Within one quarter, the number of elements with complete ownership decisions rose from twenty-one to fifty-four. The important number was not the increase itself; it was the remaining nine. They were no longer invisible. Each had a named exception, an expiry date and an executive willing to accept the consequence.

"Governance becomes real when an unresolved data question can change a design, move money or stop a release."

4. The strongest case for keeping privacy central

There is a serious opposing view. Privacy law is specialised, regulatory interpretation must be consistent, and distributed ownership can produce a patchwork of standards. If every business unit determines necessity, retention and acceptable risk for itself, the organisation may gain speed at the cost of control. A strong central privacy office, on this argument, is the safest response.

The diagnosis is right; the conclusion is incomplete. Central expertise is essential for interpreting obligations, setting minimum standards, handling difficult cases and maintaining an enterprise view. But centralising every decision confuses consistency with responsibility. A privacy office cannot know the operational purpose of thousands of fields, maintain every lineage path, or decide which product feature justifies which use. If it tries, it becomes a queue. Business teams wait for answers; specialists approve matters they cannot truly own; accountability migrates towards the people furthest from the process.

The better model is central standards with distributed, testable decisions.

- The privacy function defines legal interpretation, mandatory controls and escalation thresholds.

© WhereBeyond, 2024–

- Data owners decide purpose, quality, access and retention within those boundaries.
- Technology and operations implement controls and preserve evidence.
- The portfolio office makes unresolved data risk visible alongside cost, schedule and benefits.
- Internal assurance tests whether the operating evidence matches the declared position.

This is not a dilution of privacy authority. It is the only practical way to prevent privacy from becoming a parallel bureaucracy. The centre must own the rules and challenge; the business must own the uses and consequences.

5. Data ethics starts where compliance stops

The same governance weakness is becoming more visible in machine learning. In 2019, many analytical teams are moving from experiments towards operational models in credit, fraud, customer targeting, workforce planning and service prioritisation. The technical conversation is advancing faster than the management one.

A model can be statistically persuasive and still rest on data whose origin, purpose or meaning is poorly understood. A proxy variable can reproduce an historical pattern without anybody deliberately encoding a sensitive characteristic. An accurate model can become misleading as the process generating its inputs changes. A team may be able to reproduce the code while being unable to reconstruct which extract, exclusions and corrections produced the training data.

GDPR provides useful disciplines—purpose limitation, data minimisation, accuracy, transparency and accountability—but compliance alone cannot answer every ethical question. A use may have a lawful basis and still be unwise. A model may pass a threshold of accuracy and still allocate errors in a way that the organisation cannot defend. An explanation may be technically available yet unintelligible to the person affected.

The governance reset therefore needs to extend beyond a register of processing activities. For any model that materially influences a decision, leaders should be able to answer a small set of operational questions:

1. What decision is the model informing?
    a. Is it advisory, or does it trigger an action without human review?
    b. Who is accountable for the decision, not merely for the model?
2. What data created the result?

a. Can the team identify the source, extraction date, exclusions and transformations?

b. Was the information collected and prepared for a compatible purpose?

3. How can the result fail?

a. Which groups or cases bear false positives and false negatives?

b. What monitoring would reveal a change in performance or input quality?

4. What recourse exists?

a. Can an affected person obtain a meaningful explanation?

b. Can a responsible employee override, correct or escalate the outcome?

These questions do not require a new ethics committee for every analytical experiment. They require the existing governance system to recognise that data is not finished when it reaches the model. It remains subject to purpose, quality, accountability and challenge.

6. From annual compliance to continuous control

The temptation after a major regulatory deadline is to preserve the programme artefacts, reduce the team and move on. That would mistake evidence of mobilisation for evidence of control. Data maps decay when applications change. Retention rules fail when deletion is not built into operations. Impact assessments become ceremonial when they are completed after the consequential design choices. Consent records lose meaning when downstream uses are not connected to them.

A governance reset has a different rhythm. It converts one-off discovery into repeatable control:

- changes to purpose or data flow trigger a fresh decision;
- owners review exceptions and overdue remediation, not merely dashboards;
- operational teams test access, deletion and correction procedures;
- portfolio reviews examine cumulative exposure across initiatives;
- model reviews reconnect analytical performance to provenance and human consequence.

The immediate objection will be cost. Mapping flows, clarifying ownership and correcting source controls consume scarce capacity. Yet the comparison is not between governance and no governance. The existing alternative already has a price: repeated discovery exercises, manual reconciliations, delayed requests, duplicated stores, brittle interfaces, avoidable rework and decisions made on evidence nobody fully trusts.

© WhereBeyond, 2024–

The economic case improves when the work is organised around critical data and consequential uses rather than an attempt to govern everything equally. Start where failure affects customers, regulatory duties, financial decisions or strategic reporting. Make ownership executable there. Extend the discipline as the portfolio changes the estate.

GDPR's first year should therefore be remembered less as a successful compliance deadline than as an unusually clear diagnostic. It revealed that the enterprise cannot protect information it cannot trace, justify decisions it has not assigned, or govern models whose data it does not understand.

The organisations that learn this lesson will not treat privacy, quality and ethics as three programmes competing for attention. They will recognise them as different tests of the same management capability: the ability to make deliberate, evidenced decisions about data throughout its life.

That capability is not owned by the privacy office. It is not delivered by a catalogue. It is built when the portfolio, the operating model and accountable leaders finally govern data as seriously as they govern money.

© WhereBeyond, 2024–