



AI GOVERNANCE · SYSTEMIC RISK · INSTITUTIONAL RESILIENCE

Response to the UK Call for Evidence on *Data Regulation in the Age of AI and Other Data-Intensive Technologies*

Submitted to the Department for Science, Innovation and
Technology.

Themes 1, 3 and 5. Reproduced as submitted.

Respondent Essência A.I. Ltda, Rio de Janeiro, Brazil

Author André Afonso Silva, M.Sc.

Submitted 9 September 2026 · **Call closed** 9 September 2026

Published 11 September 2026

The record, unedited

NOTE ON THE RECORD

This document reproduces our response to the call for evidence exactly as it was submitted on 9 September 2026. Nothing has been corrected, tightened or improved after the fact. Where a sentence would now be written differently, it stands as filed.

That choice follows from the argument the response itself makes. A document revised after submission is no longer a record of what the Department received. The publication date is two days after the call closed, so that nothing here could have shaped another respondent's answer.

Any correction made after this date will appear as a dated note at the end of the document, never in the body.

The call for evidence invited responses across five thematic areas. We answered three. Theme 2, on data quality, accuracy and downstream impacts, and Theme 4, on transparency and rights in complex data environments, are left blank because we hold no evidence to offer on them.

The prompt questions set by the Department are reproduced before each response so that the answers can be read against what was asked. That material is Crown copyright, reproduced under the Open Government Licence v3.0.

RESPONDENT DETAILS AS FILED

Responding as	On behalf of an organisation
Organisation	Essência A.I. Ltda
Size	Sole trader (no employees)
Sector	Professional, Scientific, Technical – AI governance research, institutional design and advisory
Role in relation to AI	Operating AI or data-intensive systems – AI governance research, institutional design, and human oversight
Stages of activity	Evaluation/testing; deployment/adoption; monitoring – governance design, human oversight, model-output verification and decision-control architecture
Attribution	Content may be published and attributed to the organisation or to the respondent
Follow-up contact	Agreed
Evidence attached	<i>Temporal Sovereignty and Systemic Augmented Intelligence (SAI)</i> – SSRN working paper, DOI 10.2139/ssrn.6266298

Accessing and using data

PROMPT QUESTIONS SET BY THE DEPARTMENT

1. What are the challenges you face, or can foresee emerging, when accessing and using personal and non-personal data when developing or using AI and data-intensive technologies?
2. Regarding data protection, how do you assess and justify the lawful use of personal data in these contexts?
3. How appropriate are the data protection framework's definitions/obligations outlined above and are they likely to remain fit for purpose as data-intensive technologies continue to evolve?

1 What are the challenges you face, or can foresee emerging, when accessing and using personal and non-personal data when developing or using AI and data-intensive technologies?

Essência A.I. is a AI governance advisory firm based in Rio de Janeiro, Brazil. Our work concerns institutional capacity, systemic risk and the architecture of data governance in organizations adopting AI. Alongside advisory work, we maintain a dated public analytical record across Medium, Substack and LinkedIn, and have published a working paper on SSRN. The sources relevant to this response are listed below so that the Department can verify their dates and content independently. Our practice operates primarily under Brazil's LGPD rather than the UK GDPR. Both regimes distinguish a controller from an operator or processor acting on its behalf, but their statutory definitions and treatment of downstream processing are not identical. Under the UK GDPR, a processor may engage another processor, commonly referred to as a sub-processor, only with the controller's prior specific or general written authorization, subject to the applicable contractual requirements. We therefore offer the observations below as structural evidence, not as claims that the two regimes are interchangeable.

In the limited set of engagements from which we can draw direct evidence, access itself has seldom been the binding constraint. In those engagements, the recurring difficulty has been evidential rather than technical, because clients cannot produce an independent record showing that the access and subsequent use remained within the authorized processing perimeter in a form that would survive inspection by someone who did not write it.

One engagement illustrates the problem in its most ordinary form and we describe it without identifying the organization. The client is a large enterprise. The organization had an enterprise AI service covered by documented data-protection terms and an internal approval process. A senior employee in a market-facing function used that contracted assistant during working hours. In the evening, at home, the employee continued working and used a different commercial assistant on a personal consumer account.

The example is not primarily about employee misconduct. We have no evidence of malicious intent. The issue is observability, not motive. The employee was doing more work, in their own time, with a tool they found effective.

The instructive part is what happened to the organization's records. The compliance record was accurate as to the approved processing perimeter. It showed contracted processing, within the assessed perimeter, under negotiated terms. The evening processing was not under-recorded. It is outside what that record is capable of recording. An audit confined to the approved processing records could find the contract, assessment and approval while still missing the separate consumer-account use. The absence of a finding could therefore be read as evidence of a condition, when the instrument that produced the absence cannot detect the condition at all. This is the difficulty we consider most consequential for the Department's purposes.

Three features make this outcome structurally easy to reproduce. First, in this case, separation between the approved organizational environment and the individual's personal environment functioned as a de facto control, although it was not designed as one. Second, individuals can access capable commercial AI assistants outside the organizational procurement perimeter. Third, the processing through the personal account was governed by terms accepted by the individual rather than by terms negotiated by the organization, which may have had no prior visibility into those terms.

The evening use fell outside the organization's approved AI-service perimeter. Where personal data are processed through such a second service, the legal characterization depends on the actual purposes, instructions, and terms governing the provider's processing; a provider processing for its own purposes may be acting as a controller in its own right. The narrower point from the observed case is that the organization's approved records did not describe the separate processing relationship.

There is a second and, in our assessment, more serious version of the same structural problem, and it involves no individual deviation at all. A controller delegates to a processor. The processor may engage a sub-processor under the controller's prior general written authorization, subject to the applicable notification, objection, and contractual requirements. An external service may be technically integrated before the organization has completed the classification, authorization, and documentation required for its use. Even where those requirements are formally satisfied, legal authorization does not by itself give continuous operational visibility into how the downstream chain changes between review events. Nothing here requires anyone to have broken a rule. The question is whether the controller can reconstruct not only who was contractually permitted to process the data, but which services actually did so in the relevant operation.

We think the distinction between these two versions matters for policy. In the first case someone stepped off a path that existed. That is a problem of discipline and culture. In the second case, no actor necessarily departs from the authorized path; opacity can arise within the authorized chain itself. Deviation and structural opacity require different responses. The former can often be addressed through controls and enforcement; the latter requires mapping and, where necessary, redesign of the dependency chain. An institution that cannot reconstruct that chain cannot demonstrate that it governed the processing, while its accountability obligations continue regardless.

The challenge we foresee is the same problem under agentic execution. In an agentic system, the technical dependency graph may change between procurement or governance reviews as external services are invoked in the course of execution. That does not itself create a lawful processor or sub-processor relationship. It creates a governance requirement: Where a newly invoked external service receives or processes personal data, the organization may need to assess its role and the resulting authorization, contractual and documentation requirements, while technical change may occur faster than the institutional process that performs those tasks. Where dependency mapping is periodic, the controller's recorded understanding of the chain may therefore lag behind the chain actually used in execution.

PRIMARY SOURCES. UK GDPR, Article 28, on processor engagement of another processor under the controller's prior specific or general written authorisation. ICO guidance on controllers and processors.

PRIOR ANALYSIS BY THE RESPONDENT. The core delegation-chain analysis was developed before this call for evidence opened on 15 July 2026. Later entries extended it to infrastructure and supply-chain dependencies. *The Chain No One Designed: Why Unmapped Delegation Is Ungoverned Delegation.* Medium, 27 May 2026; Substack, 28 May 2026. *The Locus That Was Designed Away.* Medium, 3 June 2026; Substack, 4 June 2026. *The Layer That Does Not Appear on the Invoice.* Medium, 12 August 2026. *Two Questions, One Supply Chain.* Medium, 26 August 2026. medium.com/@andreglobal and substack.com/@andreglobal

2 Regarding data protection, how do you assess and justify the lawful use of personal data in these contexts?

In our practice, the assessment begins by specifying the intended purpose and mapping the distinct processing operations before deployment. We then identify the actors involved, the data required for each operation, the applicable legal basis under the relevant jurisdiction, downstream processing relationships, and the controls needed before the architecture is approved. Where a proposed design cannot be justified within those constraints, the design is changed or abandoned.

For UK GDPR purposes, we understand that the equivalent analysis requires an appropriate lawful basis for each distinct processing operation, together with any additional conditions applicable to special category data.

That sequencing is deliberate. In our practice, data-protection assessment is intended to constrain design and to run alongside development, rather than merely document a completed system. Where a DPIA is required under the UK GDPR, this is also consistent with ICO guidance that the assessment should begin early, before processing, and run alongside planning and development. We regard assessment conducted only after deployment as evidentially weaker because material design choices and sunk costs already exist.

We are direct with clients about what that sequence cannot reach. The organization described above completed its approved pre-deployment assessment, yet that assessment did not capture processing outside the assessed perimeter.

This is a structural problem in how compliance is evidenced and it applies to our own work as much as to anyone else's. Under the UK GDPR accountability principle, the controller must be able to demonstrate compliance and to document the lawful basis relied on for each processing purpose and the justification for that choice. Such accountability records may be created or retained under the controller's own governance processes. That creates a separate question about what the record can establish independently of the party whose compliance it describes. One test can be applied to any such record. Could this document exist, exactly as it is, if the assessment it describes had not happened that way? If it could, it is testimony. We use "testimony" and "evidence" here as an analytical distinction about independence from the party being assessed, not as categories of UK evidence law. Testimony can be reliable and valuable; the distinction concerns what a record can establish independently of the party's own account. This is not an allegation of bad faith. It holds for entirely honest records.

Under the weighting rule we propose, a contemporaneous record generated in the ordinary course of operation and not under the assessed party's sole control would carry greater independent weight than a narrative reconstructed after the event. This is a proposed assurance principle rather than a new obligation, and not a statement of UK data-protection law. Guidance on evidencing lawful basis could distinguish such contemporaneous records from records reconstructed after the fact. In our engagements, we have encountered accountability records that are primarily reconstructive. We draw no conclusion about prevalence from that.

PRIMARY SOURCES. ICO guidance on AI and data protection, on identifying each distinct processing operation and an appropriate lawful basis for each. ICO guidance on data protection impact assessments, on beginning the assessment before processing and running it alongside development. ICO guidance on the accountability principle and on documenting the lawful basis relied on.

PRIOR ANALYSIS BY THE RESPONDENT. The evidence and testimony distinction was developed in the following, published after this call for evidence opened. *The Page Does Not Keep the Difference*. Medium, 22 July 2026. *What the Instrument Cannot Read*. Substack, 23 July 2026. *The Label Is Cheap, The Review Is Not*. Medium, 19 August 2026. *A Stack of Things People Said About Themselves*. Substack, 27 August 2026. *The Records Nobody Required*. Substack, 3 September 2026. medium.com/@andreglobal and substack.com/@andreglobal

3 How appropriate are the data protection framework's definitions/obligations outlined above and are they likely to remain fit for purpose as data-intensive technologies continue to evolve?

We read the obligations named in this theme as data minimization, purpose limitation and the selection of a lawful basis. We consider all three to be structurally sound. The difficulty is temporal and it runs in two directions.

Taking the forward direction first, we do not read these obligations as legally one-off. Current ICO guidance already requires organizations to assess compatibility where personal data are reused for a new purpose and to identify an appropriate lawful basis for the new processing. Our concern is operational rather than doctrinal. In continuous and agentic systems, data flows, invoked services, or processing steps may change between the institutional review events through which those obligations are reassessed. We have described this mismatch in our published work as cadence asymmetry: execution can change continuously while institutional judgment remains episodic. The obligations remain sound. The risk lies in the interval between a material change in the processing environment and the organization's ability to reassess and evidence continued compliance.

Data minimization is a distinct legal principle, but its application is purpose-relative. What counts as necessary can only be assessed against a specified purpose. If purpose specification becomes unstable, the benchmark against which minimization is assessed becomes less determinate with it. This does not mean that minimization automatically fails whenever a purpose changes. It means that the necessity assessment may need to be revisited when the purpose or processing operation changes materially. For that reason, we would treat purpose specification and data minimization as operationally linked even though they remain legally distinct.

Consent may present a sharper problem in dynamic systems because valid consent must remain specific and informed. Where future processing operations become less determinate, broad consent risks becoming insufficiently specific or informed, while increasingly granular consent may impose material operational friction. We do not claim that agentic processing makes valid consent impossible. Our narrower concern is whether consent alone can carry the governance burden where material processing operations are selected dynamically after the initial interaction with the data subject. We do not yet have a mechanism that resolves that tension without weakening either specificity or operational autonomy.

The second direction concerns long-lived confidentiality. We do not treat Harvest Now, Decrypt Later as a storage-limitation problem in itself. The two issues intersect, but they are not equivalent. Storage limitation concerns how long a controller should retain personal data. HNDL concerns quantum-vulnerable ciphertext that may already have been captured by an adversary and retained for later decryption. Deleting the controller's copy at a later date does not eliminate a copy that has already been captured elsewhere.

Under HNDL, the relevant exposure begins when quantum-vulnerable ciphertext is captured or intercepted, not when it is eventually decrypted. A future decryption event may require no new compromise of the controller's systems because the ciphertext is already in the adversary's possession. The governance question is therefore whether the confidentiality lifetime of the data exceeds the period over which the organization can reasonably rely on its existing cryptographic arrangements and complete migration to quantum-resistant alternatives.

This matters particularly for information expected to remain sensitive for long periods. Some health, genetic, and biometric data may fall into that category, but the relevant variable is confidentiality lifetime rather than legal category alone. A long retention period does not itself create HNDL exposure, and short retention does not remove exposure once ciphertext has already been captured.

This is a present planning issue rather than a prediction of when a cryptographically relevant quantum computer will become available. The UK National Cyber Security Centre recommends that organizations complete discovery, assessment and an initial migration plan by 2028, complete their highest-priority migration activities by 2031, and complete migration to PQC by 2035, prioritizing systems that process business and personally sensitive data or manage critical communications. Separately, on 11 August 2026, Google Cloud published a

detailed roadmap targeting store now, decrypt later mitigation by the end of 2027, and integrity, non-repudiation, and key management milestones by the end of 2028. On 7 April 2026, Cloudflare announced a 2029 target for full post-quantum security. We treat these dates as planning signals, not as predictions of when quantum decryption will become feasible.

The post-quantum work now underway at infrastructure providers protects data in transit across their networks. It does not reach personal data held at rest inside a controller's own systems, which is what retention decisions govern.

The change we suggest requires no new instrument. Where personal data are expected to remain sensitive for long periods and depend on quantum-vulnerable public-key cryptography, existing DPIA or security documentation could record the expected confidentiality lifetime of the data, the relevant cryptographic dependencies, the organization's PQC migration plan, supplier dependencies, and the events that would trigger reassessment. This connects existing PQC migration practice to data-protection accountability, supplementing rather than replacing storage-limitation and security analysis. It would require no new lawful basis, data-subject right, or separate register.

We close with the limit of our own evidence. We are a firm outside the United Kingdom, operating under a different statute. Our client evidence is drawn from a limited set of engagements and is necessarily anonymized, which limits what the Department can verify independently. Our analytical record is public and dated, but it is analysis rather than measurement. Where the Department needs quantitative evidence on how prevalent these patterns are across the UK economy, we cannot supply it, and we would rather mark that boundary than let the argument carry weight it has not earned.

PRIMARY SOURCES. ICO guidance on purpose limitation, on assessing compatibility where personal data are reused for a new purpose. NCSC, *Timelines for migration to post-quantum cryptography*, 20 March 2025; Google, post-quantum cryptography migration timeline, 25 March 2026; Cloudflare, post-quantum roadmap, 7 April 2026; Google Cloud, post-quantum cryptography roadmap, 11 August 2026.

PRIOR ANALYSIS BY THE RESPONDENT. The temporal argument was developed before this call for evidence opened on 15 July 2026: *Temporal Sovereignty and Systemic Augmented Intelligence*. Written 18 February 2026; posted to SSRN 1 April 2026. DOI 10.2139/ssrn.6266298. Attached; *When Institutional Time Breaks*. Substack, 8 January 2026; *When Encryption Fails, Institutions Are Exposed*. Medium, 11 February 2026; *Harvest Now, Decrypt Later and Temporal Sovereignty*. Substack, 12 February 2026; *The First Test of the Completed Framework: HNDL as Retrospective Auditability of Judgment*; Substack, 21 May 2026. medium.com/@andreglobal and substack.com/@andreglobal

Governing data use across organisations

PROMPT QUESTIONS SET BY THE DEPARTMENT

1. How do you collect and then govern data across supply chains or between organisations in practice, including allocating responsibility and enabling data sharing or access?
2. How do you approach automated decision-making, and will this approach remain fit-for-purpose as technology advances?
3. How effective and appropriate are the data protection framework's definitions in assigning responsibilities (e.g. controller/processor), and are they likely to remain fit-for-purpose over time as technology advances?

1 How do you collect and then govern data across supply chains or between organisations in practice, including allocating responsibility and enabling data sharing or access?

Our evidence is advisory rather than operator-side: we do not present ourselves here as a data-sharing intermediary or large-scale data processor. In client work, our first governance step is not a control review. It is a map. We ask the client to state, with documents, through which links each critical processing operation passes and what authorization or contractual basis each link relies on to engage the next service. In the limited set of engagements from which we can draw direct evidence, clients have generally been able to identify their direct processors. Visibility has been weaker for downstream processors and for whether those mappings remained current at the time of review.

That mapping exercise usually returns two findings and they are different problems requiring different remedies.

The first concerns the length of the chain. Processing may descend through providers the controller did not select directly, because an intermediate processor engaged them under the authorization structure governing downstream processing.

The second concerns the shape of the chain and it is the one we find least anticipated. A contractual vendor map can show several distinct application providers while obscuring shared dependencies further upstream. The UK Competition and Markets Authority's mapping of the foundation-model value chain identifies critical inputs including compute and data, foundation-model development and release, and downstream routes to market, and has highlighted concentration risks where a small number of incumbent firms occupy positions across multiple levels of that chain. An organization may therefore contract with several application vendors while remaining indirectly dependent on the same foundation-model or cloud infrastructure provider. That dependency may not be visible from the application-level vendor count alone. The CMA has made a related observation: where an algorithmic system is built and operated by many contributors, responsibility for a failure at any given point can become unclear.

This matters because substitutability is a material condition of an operationally credible exit plan. A contract may be legally terminable while the service remains difficult to replace without interruption, migration cost or loss of capability. In that case, the governance problem is not the absence of a contractual exit right; it is dependency despite the existence of one.

On records maintained to support auditability, we offer one instrument clients can apply directly. Take twelve months of governance minutes for any oversight arrangement over data use and count how many times that arrangement stopped something. Not concerns recorded. Stoppages. Then ask what each stoppage cost and who

absorbed it. Zero stoppages in twelve months does not establish that nothing needed stopping. Nor does it prove that the arrangement lacks authority to stop. It means that the operating record, by itself, has not demonstrated that the interruptive capacity was exercised. We use twelve months as a practical observation window, not as a regulatory threshold. Because the count is derived from ordinary-course governance records rather than reconstructed for an audit, we treat it as more probative than a retrospective narrative. It is not independent evidence merely because it is contemporaneous.

PRIMARY SOURCES. UK GDPR, Article 28, on processor engagement of another processor under the controller's prior specific or general written authorization. ICO guidance on controllers and processors. Competition and Markets Authority, *CMA AI strategic update*, 29 April 2024, and *AI Foundation Models: update paper and technical update report*, 11 April 2024. <https://www.gov.uk/government/publications/cma-ai-strategic-update/cma-ai-strategic-update>

PRIOR ANALYSIS BY THE RESPONDENT. The core delegation and supply chain analysis was developed before this call for evidence opened on 15 July 2026. Later entries extended it to infrastructure dependencies. *The Chain No One Designed* (Medium, 27 May 2026). *The Locus That Was Designed Away*. Medium, 3 June 2026; Substack, 4 June 2026. *The Layer That Does Not Appear on the Invoice* (Medium, 12 August 2026). *Two Questions, One Supply Chain* (Medium, 26 August 2026). medium.com/@andreglobal and substack.com/@andreglobal

2 How do you approach automated decision-making, and will this approach remain fit-for-purpose as technology advances?

Our approach starts from the current UK legal position rather than from a competing definition. Section 80 of the Data (Use and Access) Act 2025 repealed Article 22 of the UK GDPR and replaced it with Articles 22A to 22D. Article 22A(1)(a) treats a decision as based solely on automated processing where there is no meaningful human involvement in taking it.

We note that Parliament, in the same Act, used a related operational formulation for Part 4 of the Data Protection Act 2018, concerning intelligence services processing: a decision is based on entirely automated processing where the decision-making process does not include an opportunity for a human being to accept, reject or influence the decision. We do not import that definition into Article 22A. We use it as an adjacent statutory formulation that makes decision authority observable.

Our contribution is an operational test of that authority. We ask where, in the actual decision process, a person holds the power to decide differently from what the system proposes. Refusal power is not a complete legal definition of meaningful human involvement: a person could hold formal override authority and still fail to exercise meaningful judgment. But where no person has authority to reject or materially alter the system's recommendation, we do not see how the involvement can be substantive.

The distinction is therefore not machine versus human presence. It is between human presence and effective decision authority.

The test has two parts. Where is the point at which a person can decide differently from what the system proposes? Has that person ever, in the operating history of the system, decided differently? A negative answer to the second question does not prove that the authority is decorative. It means that operating history has not yet demonstrated the capacity. That distinction matters: a low refusal rate may reflect a system that performs well. Zero recorded departures are therefore a signal for examination, not proof of ineffective oversight.

Article 22A(2) requires controllers to consider, among other things, the extent to which a decision was reached by means of profiling when establishing whether human involvement was meaningful. That is a qualitative criterion. We have not identified in the current statutory provisions or ICO guidance a specific evidential test based on whether reviewers have actually rejected or materially altered automated recommendations over time.

Article 22D(1) gives the Secretary of State power to make regulations describing circumstances that are, or are not, to be taken to involve meaningful human involvement. We have not identified regulations made under that power as of 8 September 2026. Separately, the Data Protection Act 2018 (Code of Practice on Artificial Intelligence and Automated Decision-Making) Regulations 2026, in force since 12 May 2026, require the Information Commissioner to prepare a code of practice covering AI and automated decision-making under the UK GDPR and Part 3 of the Data Protection Act 2018. Part 4 intelligence-services processing is outside the scope of that code. We offer the refusal test as one candidate input to that developing framework. Its first limb, identifying where decision authority resides, can be answered through governance and role documentation. Its second, whether that authority has actually been exercised, can be examined through ordinary operating records.

Operational logs can record intervention events: who intervened, when, and what changed. They do not by themselves establish whether the reviewer exercised independent judgment. A record that someone reviewed a case is evidence that a review event occurred; it is not, without more, evidence that the review was substantive.

A recent enforcement action illustrates the same underlying evidential problem. In August 2026, the Dutch Data Protection Authority fined Uber after finding that certain driver-account deactivations were fully automated and involved no human assessment. Uber disputes the regulator's characterization and has said it will appeal. We cite the case as a cross-jurisdictional illustration that regulators may examine whether human involvement existed in fact rather than whether an organization states that review was available.

On fitness over time, our assessment is that the framework remains correctly aimed. The difficulty is evidential: as automated decision paths become more complex, organizations may find it harder to distinguish meaningful human judgment from routine acceptance by examining the existence of a review step alone.

PRIMARY SOURCES. Data (Use and Access) Act 2025, section 80; UK GDPR Articles 22A to 22D. The Data Protection Act 2018 (Code of Practice on Artificial Intelligence and Automated Decision-Making) Regulations 2026 (SI 2026/425), in force 12 May 2026. <https://www.legislation.gov.uk/ukxi/2026/425>. Explanatory Notes to the Act, Part 5, paragraphs 634 to 652. <https://www.legislation.gov.uk/ukpga/2025/18/notes/division/10/index.htm> Autoriteit Persoonsgegevens, "Uber fined nearly 825 million euros for automated driver blocking," 21 August 2026.

SECONDARY CORROBORATION. The Guardian, 21 August 2026. <https://www.theguardian.com/technology/2026/aug/21/netherlands-fines-uber-automated-driver-suspensions>

PRIOR ANALYSIS BY THE RESPONDENT. The refusal argument was developed before this call for evidence opened on 15 July 2026. Later entries extended it to the evidential question. *The Locus That Was Designed Away* (Substack, 4 June 2026). *What the Instrument Cannot Read* (Substack, 23 July 2026). *The Label Is Cheap, The Review Is Not* (Medium, 19 August 2026). medium.com/@andreglobal and substack.com/@andreglobal

3 How effective and appropriate are the data protection framework's definitions in assigning responsibilities, and are they likely to remain fit-for-purpose over time?

The controller and processor definitions remain useful at the organizational level. They distinguish the entity that determines the purposes and means of processing from an entity processing personal data on its behalf. Our concern begins one level below that classification.

Where the controller is an organization, the legal classification allocates accountability at entity level; it does not by itself identify where, inside that organization, authority for a particular consequential decision is operationally located. External accountability can therefore remain clear while internal decision authority becomes diffuse. We call the internal point at which a consequential choice can be attributed and interrupted the locus of attributable authority. This is a governance construct, not an additional legal role under the UK GDPR.

Chains can be built in which no such point survives. Each link receives an output, acts on it, and passes it forward. Each participant can say honestly that no independent judgment was exercised, because independent

judgment was never asked for. The chain works end to end, a decision is produced, and it dissolves when traced back. This pattern is not unique to AI. Distributed and multi-agent architectures may intensify it by splitting processing and execution across several agents or services. No single intermediate component need contain the full decision path. We treat this as a foreseeable governance risk, not as a claim that multi-agent systems inherently eliminate accountability.

Applied to such a structure, the controller classification remains legally correct while being operationally incomplete for a different question: where, inside the controller, is the human authority capable of accepting, rejecting or materially altering the consequential outcome?

We do not think a new category solves this and we would argue against creating one. Our suggestion is narrower and concerns evidence rather than definition. Where a controller relies on meaningful human involvement in determining that a significant decision is not solely automated, it could be expected to identify the role holding authority to reject or materially alter the automated recommendation, and to retain operational evidence that this authority is real. Actual instances of rejection or material modification would be probative, though not conclusive, evidence where they exist; their absence should trigger inquiry rather than automatically establish non-compliance.

We are aware of the objection and think it has force. A requirement to name a locus can be met on paper by naming someone with no real capacity, which reproduces the problem one layer up. Our answer is that a named role can be examined in a way an unnamed role cannot. The inquiry then becomes concrete: what authority does the role hold, what information can it access, what happens if it disagrees with the system, and does the operating record show that disagreement is genuinely possible?

PRIMARY SOURCES. UK GDPR Article 4(7) and Article 4(8), controller and processor definitions; UK GDPR Articles 22A to 22D, as inserted by section 80 of the Data (Use and Access) Act 2025. ICO guidance on controllers and processors.

PRIOR ANALYSIS BY THE RESPONDENT. The locus argument was published before this call for evidence opened on 15 July 2026. *The Locus That Was Designed Away* (Medium, 3 June 2026; Substack, 4 June 2026). *When No One Decides* (Medium, 8 April 2026). [mediu m.com/@andreglobal](https://medium.com/@andreglobal) and substack.com/@andreglobal

Effectiveness of data frameworks in regulating AI

PROMPT QUESTIONS SET BY THE DEPARTMENT

1. Overall, what is working well, where have you observed disproportionate or unintended barriers in practice, and where have risks to data protection arisen?
2. Are the existing data frameworks sufficiently adaptable to future developments in AI, and where might rapid technological advances give rise to future difficulties? Or are there any alternative approaches that may be more effective, both now and in the future?

1 Overall, what is working well, where have you observed disproportionate or unintended barriers in practice, and where have risks to data protection arisen?

Two features of the existing framework work well in our experience.

First, the controller and processor framework provides a workable allocation of responsibility where the purposes and means of processing can be identified. Second, purpose limitation requires an organization to be clear from the outset about why personal data are being collected and used. In our engagements, we have seen that exercise expose cases in which the intended purpose had not yet been stated with enough precision to support the proposed processing. The discovery itself is useful.

A foreseeable barrier is not prohibition but displacement. One response to uncertainty over the lawful basis or permissible use of training or evaluation data is to proceed with evaluation on proxy, synthetic, or other test data while the underlying data question remains unresolved.

The result can be an evaluation gap: the system may be tested under data conditions that differ materially from those it will encounter in operation. In such cases, legal uncertainty has not necessarily removed the underlying risk; it may have displaced part of it into model or system evaluation. That displacement may not be visible in the data-protection record through which the original data decision was assessed.

We consider this a potentially significant cost of unresolved legal uncertainty. Our response is not to relax the underlying obligations, but to reduce avoidable uncertainty earlier and to recognize that a deferred data question can reappear as an evaluation-validity problem downstream.

The clearest data-protection risk we have observed is the one described in Theme 1. A compliance record can be accurate as to the processing perimeter it was designed to document while remaining incapable of showing processing outside that perimeter. In that situation, absence of an adverse finding should not be treated as independent evidence that no off-perimeter processing occurred. The evidential weakness is greatest where the record is reconstructive, is generated under the assessed party's sole control, and lacks independent corroboration.

PRIMARY REGULATORY SOURCES. ICO guidance on controllers and processors. ICO guidance on purpose limitation.

PRIOR ANALYSIS BY THE RESPONDENT. The evidence-and-records argument was further developed during the period in which this call for evidence was open. *The Page Does Not Keep the Difference* (Medium, 22 July 2026). *A Stack of Things People Said About Themselves* (Substack, 27 August 2026). medium.com/@andreglobal and substack.com/@andreglobal

2 Are the existing data frameworks sufficiently adaptable to future developments in AI, and where might rapid technological advances give rise to future difficulties? Or are there any alternative approaches that may be more effective?

Our assessment is that the core principles remain adaptable, but some mechanisms through which those principles are operationalized will face greater strain as processing becomes more dynamic. Consent is one of the clearest examples.

Valid consent must remain specific and informed. In dynamic or agentic systems, future processing operations may become less determinate at the moment consent is obtained. Broader consent then risks losing specificity or informational value, while increasingly granular consent may introduce material operational friction. We do not claim that agentic processing makes valid consent impossible, nor that consent is the appropriate lawful basis for every such system. Our narrower concern is whether consent alone can carry the governance burden where material processing operations are selected dynamically after the initial interaction with the data subject. We do not yet have a mechanism that resolves that tension without weakening either specificity or operational autonomy.

On alternative approaches, we offer a caution rather than a proposal.

One family of alternative approaches relies heavily on declaration: disclosure of AI use, labeling obligations, or self-assessment against published requirements. Where compliance is evidenced primarily through a declaration produced by the regulated party, we apply the same analytical distinction used in Theme 1. Could the declaration exist in the same form if the process it describes had not occurred that way? Is any material part independently checkable by someone who did not produce it? A negative answer does not make the declaration false. It limits what the declaration can establish independently. We use "testimony" and "evidence" here as an analytical distinction about independence from the party being assessed, not as categories of UK evidence law. Repeated self-declaration does not, by itself, create independent evidence. We are not arguing against disclosure. A framework that relies on declaration in place of verification may report increasing formal compliance without independently establishing the underlying condition the declaration is intended to represent.

On the potential of AI to improve compliance, we see one application with unusual leverage.

Post-quantum migration illustrates a narrower opportunity. Algorithm availability is no longer the only preparatory problem: organizations must first discover where quantum-vulnerable cryptography is embedded across systems, services, software, hardware, data flows and supplier dependencies. NCSC guidance treats discovery and assessment as a core migration activity, notes that cryptographic services in older systems have often evolved in ways that make discovery harder, and expects large organizations to need two to three years for discovery, assessment and initial planning alone.

This creates, in our view, a candidate use case for AI-assisted compliance rather than an observed result from our engagements. AI-assisted code, configuration and dependency analysis could be tested as an additional discovery layer over large cryptographic estates, subject to independent validation against deterministic inventory and scanning tools. We do not claim that AI-based discovery is more accurate or more complete than existing automated methods.

A cryptographic inventory is also a precondition for connecting long-lived personal data to the cryptographic dependencies and PQC migration assumptions on which its confidentiality relies, as discussed in our response to Theme 1. NCSC guidance already asks organizations, as part of discovery, to record the data they hold together with its expected lifetime and its value to an adversary.

PRIMARY SOURCES. ICO guidance on valid consent. ICO guidance on AI and data protection. NCSC, *Timelines for migration to post-quantum cryptography*, 20 March 2025. <https://www.ncsc.gov.uk/guidance/pqc-migration-timelines>

PRIOR ANALYSIS BY THE RESPONDENT. *Temporal Sovereignty and Systemic Augmented Intelligence*. Written 18 February 2026; posted to SSRN 1 April 2026. DOI 10.2139/ssrn.6266298. Attached. *Harvest Now, Decrypt Later and Temporal Sovereignty* (Substack, 12 February 2026). *The Label Is Cheap, The Review Is Not* (Medium, 19 August 2026). *A Stack of Things People Said About Themselves* (Substack, 27 August 2026). medium.com/@andreglobal and substack.com/@andreglobal

A note on method

A note on method, offered because this submission argues for a distinction that should apply to itself.

This response was composed with the support of ChatGPT (OpenAI) and Claude (Anthropic), working from our own published analysis. The full cognitive signature is ours: intent, constraints, judgment under uncertainty, veto and responsibility. The errors are ours too. They drafted passages, argued against our structure where they saw a weakness, and checked the argument for internal contradiction. Facts were verified at primary source before entering the text, and where verification was not possible the text says so or the passage was removed.

We disclose this under a standing rule rather than by decision here. Since 10 August 2026 our published series has carried a method note in every edition and on 17 August 2026 we published the rule itself in full, with a date, in text held by readers. It is published weekly because the series argues that a record produced because someone asked is worth less than one produced in the normal course of work. Nobody asked us for that note and nobody asked us for this one.

We state the limit of what the disclosure establishes. This note is a statement we make about ourselves. Read from outside, it is testimony. What gives it any weight beyond assertion is that the rule predates this call for evidence, that applying it here carries a cost we would not choose, and that the rule sits in public where any reader who kept the earlier text can check whether we departed from it. Those are the same properties we would apply to any declaration regime, and we would not ask the Department to weigh ours differently.

We answered three of the five themes. We did not answer Theme 2 or Theme 4, not because we consider them less important, but because we have no evidence to offer on data quality, fairness measurement or the exercise of individual data rights. We would rather leave a theme blank than fill it.



PUBLICATION NOTE

Attachment. The evidence file submitted with Themes 1 and 5 was *Temporal Sovereignty and Systemic Augmented Intelligence (SAI)*, a working paper deposited on SSRN, DOI 10.2139/ssrn.6266298. It is not reproduced here and can be retrieved from SSRN directly.

Source of the prompt questions. The thematic prompt questions reproduced in this document were set by the Department for Science, Innovation and Technology and are published on GOV.UK. They contain public sector information licensed under the Open Government Licence v3.0.

Status. This document records a response to a public consultation. It does not describe any advisory, contractual or institutional relationship with the Department for Science, Innovation and Technology, the Department for Culture, Media and Sport, or any other UK public body.

Verification links. Autoriteit Persoonsgegevens decision, 21 August 2026:

<https://www.autoriteitpersoonsgegevens.nl/actueel/uber-krijgt-boete-van-bijna-825-miljoen-euro-voor-geautomatiseerd-blokkeren-van-chauffeurs>. The Guardian report, 21 August 2026:

<https://www.theguardian.com/technology/2026/aug/21/netherlands-fines-uber-automated-driver-suspensions>. These links are provided for the reader's convenience and form no part of the text as submitted.

Contact. andre@essencia.ia.br · essencia.ia.br