

Operationalizing Data Subject Rights in the AI Era: An SRE-Inspired, Socio-Technical Framework for Ethical Reliability

Prabal Pathak^{1,*}, Dr. Ranjit Rajak²

¹Principal PM Architect, Microsoft, Apex, NC, USA

²Assistant Professor, Department of Computer Science and Application, Dr. Hari Singh Gour Central University Sagar, MP, India

Abstract Data subject rights (access, rectification, erasure) are central to modern privacy law, but in AI-intensive enterprises they are often fulfilled via workflows that do not reliably cover derived artifacts such as features, embeddings, logs, and model-related outputs. This gap turns formally granted rights into inconsistently delivered outcomes and undermines trust. We propose a socio-technical framework that treats data subject rights request response (DSRR) as a reliability problem analogous to production services. The framework introduces a control plane that maintains canonical request state, scope decisions, and evidence, and a data plane that executes collection, deletion, and correction across heterogeneous data stores and AI pipelines using lineage and policy. We define service-level indicators and objectives for timeliness, coverage, and response accuracy, and show how error budgets and continuous improvement practices from Site Reliability Engineering can be applied to rights operations. A limited simulated evaluation illustrates phased adoption effects on AI-system coverage and timeliness while preserving accuracy. We also discuss AI-specific erasure challenges, including machine unlearning and partial erasure, and outline directions for empirical validation and governance.

Keywords Data subject rights, DSAR, AI governance, Privacy engineering, Site Reliability Engineering, Data lineage, Service levels, Autonomy, Accountability

1. Introduction

Over the last decade, privacy regulations such as the EU General Data Protection Regulation (GDPR), the California Consumer Privacy Act as amended by the CPRA (CCPA/CPRA), and the Indian Digital Personal Data Protection Act (DPDPA) have expanded individuals' rights to access, correct, delete, restrict, and obtain copies of their personal data. [1–3]

Meanwhile, enterprise systems have become AI-intensive and highly distributed: personal data is replicated and transformed across transactional backends and CRMs, clickstream/event logs, data lakes and warehouses, feature stores and vector databases, model training pipelines and registries, and online inference and monitoring stacks.

In this setting, a single data subject rights request (DSRR)—for example, “delete my data” or “show me what you know about me”—touches raw records as well as derived artifacts (features, embeddings, audit logs, and model-related outputs). Meeting the intent of the right therefore requires both correct scoping and reliable execution across a dependency

graph of systems.

However, many organisations still fulfil DSRRs through ticket-driven, non-scalable workflows (shared mailboxes, case tools, ad hoc queries, and per-system escalation). These processes may pass compliance checks at low volumes, but they degrade as the data stack grows and response deadlines tighten—especially when AI pipelines and telemetry are in scope.

This operational gap has ethical consequences: when infrastructure cannot reliably deliver outcomes, formally granted rights can fail to become lived rights in practice (Section 2.4).

We therefore frame DSRR fulfilment as a reliability property of AI and data infrastructure—analogue to availability or latency—and draw on Site Reliability Engineering (SRE) to operationalise the idea with explicit metrics, targets, and continuous improvement mechanisms. [8]

Our contributions are threefold. (1) We define “rights reliability” and motivate it as ethical infrastructure for AI-era DSRR. (2) We propose a socio-technical control-plane/data-plane architecture that integrates lineage for dynamic scoping and evidence. (3) We formalise service-level indicators and objectives for timeliness, coverage, and response accuracy, and illustrate their behaviour via a small simulated evaluation that models phased adoption. We also extend the

* Corresponding author:

prabalpathak@gmail.com (Prabal Pathak)

Received: Jan. 19, 2026; Accepted: Feb. 3, 2026; Published: Feb. 9, 2026

Published online at <http://journal.sapub.org/computer>

discussion to AI-specific erasure challenges, including machine unlearning and partial erasure of derived artifacts.

2. Background: Regulation, Practice and Ethics

2.1. Data Subject Rights and AI

GDPR and similar frameworks articulate a rich set of rights: access, rectification, erasure, restriction, portability and objection. Although their doctrinal details differ, they share a common ambition: individuals should not be entirely at the mercy of opaque data and AI systems. They should be able to see, contest and reshape how their data is used. However, these rights were largely drafted with traditional data processing in mind: well-bounded databases, clear controllers and processors, and relatively static pipelines. AI practice has moved towards continuous training cycles, large-scale logging and telemetry, complex feature engineering and global, multi-tenant cloud infrastructures.

2.2. DSAR Tooling and Operational Gaps

A commercial ecosystem of Data Subject Access Request (DSAR) and privacy management tools supports intake portals, case management, templated communications, and connectors to common systems of record. These platforms reduce coordination friction, but many assume a relatively static system inventory and provide limited integration with data lineage or AI-specific artifacts (feature stores, embeddings, model checkpoints, and inference logs). As a result, they often manage the paperwork of rights better than the infrastructure of rights. [4–7]

2.3. SRE and Reliability Beyond Availability

Site Reliability Engineering (SRE) emerged in large-scale Internet companies as a way to make reliability measurable and manageable. Rather than treating reliability as an abstract desideratum, SRE practices define service level indicators (SLIs), service level objectives (SLOs) and error budgets. Our contribution is to extend this idea to data subject rights: to treat the ability to fulfil DSRRs in AI systems as a reliability property with its own SLIs, SLOs and error budgets. [8]

2.4. Ethical Stakes: From Formal Rights to Lived Rights

From an ethical perspective, the decisive question is not whether rights exist in law, but whether people can meaningfully exercise them. In AI-era systems, weak scoping, incomplete system coverage, and delays can turn rights into symbolic assurances. We use the term formal rights to refer to doctrinal entitlements and lived rights to refer to the delivered outcomes shaped by infrastructure and practice. The remainder of the paper treats timeliness, coverage, and response accuracy as measurable levers that narrow (or widen) the gap between formal and lived rights. [13]

3. Conceptual Framework: DSRR Reliability as Ethical Infrastructure

We propose to understand the technical fulfilment of DSRRs as a form of ethical infrastructure: a set of systems, processes and metrics whose reliability directly affects people’s ability to exercise their rights. In many organisations, DSRR fulfilment is not treated as an explicit system property. It is handled through ad hoc processes: a privacy team receives requests, files tickets to engineering teams, and hopes that the relevant systems respond within regulatory deadlines. This approach may work when the number of systems and requests is small, but it does not scale gracefully. [11,12]

By contrast, treating DSRRs as a reliability property means asking what it means for an organisation to be reliable in respecting data subject rights, how that reliability can be measured over time and how infrastructure should be designed so that reliability improves rather than degrades as AI systems grow. We argue that this amounts to defining a new class of non-functional requirement for AI systems: rights reliability.

We operationalise rights reliability along three dimensions—timeliness, coverage, and response accuracy—because each directly affects whether a right is practically exercisable in AI-intensive environments. Section 5 provides formal definitions and example service level objectives, and Section 7 discusses ethical risks such as metric gaming and rights-washing.

4. Socio-Technical Architecture for AI-Era DSRR

To support reliable DSRR fulfilment in AI systems, we propose a two-plane architecture that explicitly connects legal rights, organisational processes and technical infrastructure. The control plane aggregates the capabilities needed to handle DSRRs end to end: an intake and API layer, case management and workflow, a policy engine, an orchestrator and an evidence store. The data plane encompasses the heterogeneous systems where personal and derived data reside: relational databases, data lakes and warehouses, SaaS platforms, feature stores and vector databases, model registries and observability platforms.

Figure 1 illustrates the proposed two-plane architecture for AI-era DSRR operations. The control plane encapsulates the logical capabilities required to handle DSRRs end to end: a DSRR intake and API layer acting as the single entry point for portals and programmatic calls, case management and workflow that coordinate privacy, legal and engineering stakeholders, a policy engine that evaluates regulatory and contractual requirements, an orchestrator that decomposes each DSRR into per-system tasks and an evidence store that captures immutable audit traces. The data plane groups the systems where personal and derived data actually reside, including connectors to databases and SaaS systems, data

lakes and warehouses, feature stores and vector databases, logs and telemetry platforms and ML models/model registries.

A key novelty of this architecture is the role of data lineage. Rather than maintaining static lists of systems that might contain personal data, we treat system inventory as a living graph, automatically derived from schema registries, ETL and streaming pipelines, feature stores and model registries, and deployment descriptors for inference services. [9,10]

Figure 2 shows a simplified data-lineage graph capturing how subject data propagates through an AI-driven ecosystem. Source systems such as CRM applications and transactional backends, together with web and app event streams, feed into a central data lake. From there, curated datasets populate a feature store, which in turn is used by training jobs that produce model artifacts registered in a model registry. The same features and models then power online inference

services. Both the data lake and online services emit monitoring and telemetry data that close the loop for observability. During DSRR processing, the control plane queries this graph to determine which nodes are in scope for a given subject, ensuring that both raw data and derived artifacts are considered in DSRR fulfilment.

Operationally, DSRR control planes should therefore treat erasure as a multi-layer obligation: (i) delete or de-identify raw records in source and analytical stores; (ii) invalidate or recompute downstream derived artifacts (features, embeddings, aggregates, caches); (iii) address model influence via retraining, machine unlearning, or bounded retention policies; and (iv) record evidence of what was erased, what was not, and why (e.g., legal exceptions, technical infeasibility, or disproportionate effort). This evidence is essential for accountability and for avoiding “partial compliance” that is invisible to the data subject.

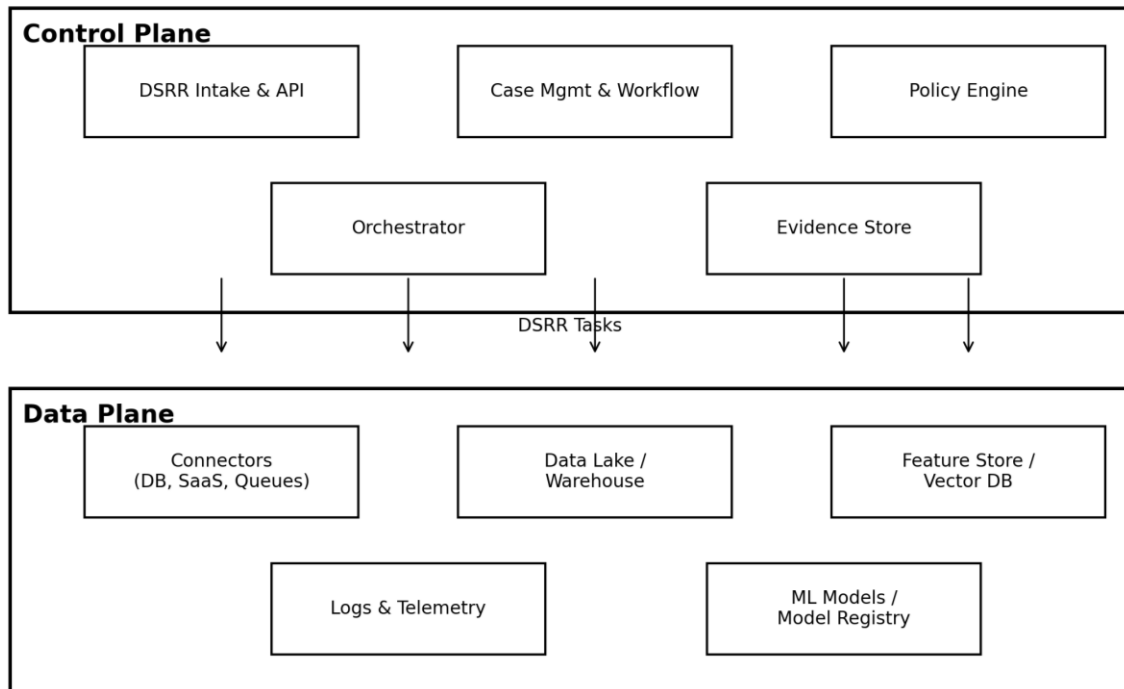


Figure 1. High-level DSRR control-plane and data-plane architecture

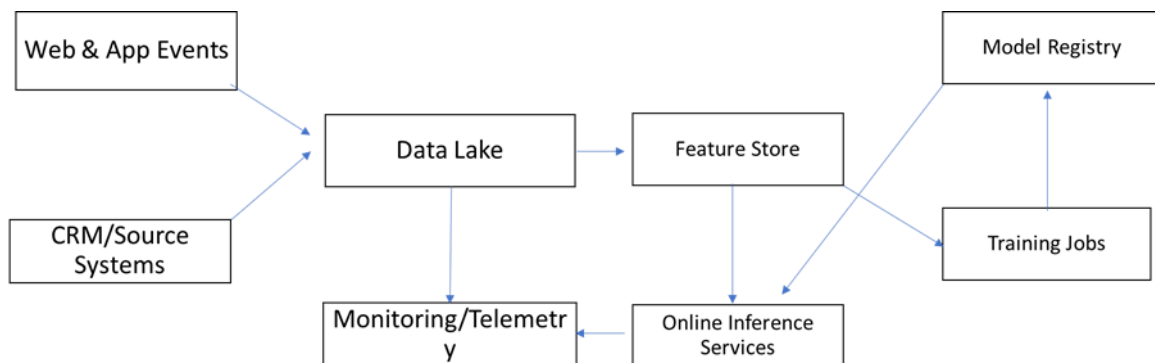


Figure 2. Simplified data-lineage graph linking source systems, data lake, feature store, model registry and online inference services

AI-era erasure is not equivalent to deleting rows in a database. When personal data contributes to learned parameters, the subject’s information can persist as model influence even after raw data is removed. Two classes of techniques are relevant. Machine unlearning aims to remove the effect of specific training points without full retraining (exactly or approximately), while partial erasure targets derived artifacts that may still carry personal information, such as feature store entries, vector embeddings, inference logs, caches, and monitoring traces. In practice, organisations often combine bounded retention (short log windows), versioned training datasets, and retraining triggers with approximate unlearning methods; the feasibility depends on model type, training regime, and audit requirements. [14–16]

4.1. AI-Specific Erasure: Machine Unlearning and Partial Erasure

AI-era erasure introduces a qualitative difference: the subject’s data may persist as influence rather than as an identifiable record. In addition to deleting rows in source systems, organisations must account for derived artifacts such as feature vectors, embeddings, materialised aggregates, model checkpoints, cached prompts/responses, and monitoring traces that can retain or reintroduce personal information.

Machine unlearning provides a spectrum of techniques to reduce or remove a subject’s influence on a trained model without always retraining from scratch. Operationally, organisations may choose among: (i) exact retraining from data-excluded corpora when feasible; (ii) approximate unlearning techniques that target specific examples or users while preserving model utility; and (iii) containment and mitigation controls, such as model version pinning, gated inference, and subject-level suppression rules that prevent the subject’s data from being reintroduced into future training or downstream features, even while legacy checkpoints are being phased out.

Because perfect erasure from learned parameters may not be technically or economically achievable in all settings, we treat partial erasure as an explicit operational state that must be declared, justified, and governed. The DSRR control plane should record: which artifacts were deleted (tables, logs, feature stores), which were regenerated (features, aggregates), which models were retrained or replaced (including model version identifiers and the selected erasure approach), and what residual risk remains (e.g., potential memorisation in legacy checkpoints or third-party replicas), along with the controls that reduce that residual risk.

From a rights-reliability perspective, partial erasure primarily impacts coverage (completeness of affected systems/artifacts) and response accuracy (truthfulness of the claim that erasure has been effected). Accordingly, the evidence ledger should attach verifiable signals—deletion receipts, pipeline run identifiers, model lineage links, and approved policy waivers—to each DSRR, enabling internal audit and, where appropriate, regulator-facing transparency.

Finally, we recommend a risk-tiered policy for AI-specific erasure. For higher-risk contexts (e.g., sensitive attributes, vulnerable populations, high memorisation risk), organisations should prioritise stronger guarantees (retraining or approved unlearning methods) and tighter error budgets. For lower-risk contexts, staged remediation with clear data-subject communication may be acceptable, but only if residual risk, scope, and remediation timelines are documented and monitored.

5. Formalizing Reliability Metrics for DSRR

To move from architectural principles to actionable governance, we define a set of service level indicators (SLIs) and associated service level objectives (SLOs) for DSRR operations. Consider an evaluation window with N DSRRs. For each request we observe its completion time, the set of in-scope systems, the subset that responded successfully and whether a confirmed issue was later identified.

We define three SLIs. The timeliness SLI measures the fraction of DSRRs completed within a target deadline T days. The coverage SLI measures, on average, what fraction of in-scope systems responded per request. The accuracy SLI measures the fraction of DSRRs that did not require rework due to incorrect or incomplete data. Each SLI has an ethical interpretation: timeliness relates to the idea that a right delayed is a right denied; coverage relates to equality and non-discrimination; and accuracy relates to truthfulness and harm avoidance.

6. Simulated Evaluation

Table 1. Example SLIs and SLO targets for DSRR reliability dimensions

Dimension	Example SLO Target	Metric Definition
Timeliness	99% within 20 days	Closed DSRR cases / all cases within 20 days.
Coverage	99.5% of in-scope systems	Systems that returned a response / in-scope systems per DSRR instance.
Accuracy	99.9% correct	Cases without confirmed issues / all closed cases.

To strengthen the framework with empirical grounding, we perform a limited simulated evaluation that exercises the proposed architecture and metrics over a synthetic but AI-realistic system topology. The simulator models (a) a lineage graph linking sources, analytical stores, feature/embedding layers, training pipelines, registries, and inference services; (b) DSRR arrivals with heterogeneous scopes and legal exceptions; and (c) per-system connector behaviour, including latency, partial failures, and occasional accuracy defects that require rework.

We evaluate four scenarios that mirror common adoption trajectories: Baseline (manual, ticket-driven coordination), Wave 1 (unified intake and case management), Wave 2 (lineage-driven scoping and explicit AI inventory integration), and Wave 3 (explicit SLIs/SLOs with error-budget-driven remediation of the highest-impact failure modes). For each scenario, we simulate $N=1,000$ requests over an evaluation window and report P90 completion time (timeliness), mean in-scope participation for AI systems (coverage), and the fraction of cases not requiring rework (accuracy).

Simulation parameters are chosen to be conservative rather than optimistic: connectors exhibit heavy-tailed latencies, a small probability of missed systems in manual scoping, and a non-zero rate of downstream inconsistencies (e.g., stale caches or delayed feature recomputation) that trigger rework. The purpose is not to claim industry benchmarks, but to test whether the framework’s mechanisms produce the expected directional effects.

Table 2. Normalised DSRR reliability metrics across adoption stages

Stage	P90 Timeliness (days)	Coverage (AI systems %)	Accuracy (%)
Baseline (manual)	45	80	98.5
Wave 1	38	82	98.7
Wave 2	30	96	99.4
Wave 3	22	99	99.9

Table 2 summarises the normalised outputs across the four scenarios, and Figure 3 plots the same metrics as curves. Relative to the baseline, Wave 1 improves timeliness primarily via coordination. Wave 2 yields the largest coverage

gain by reducing scoping omissions through lineage and AI inventory integration, and it improves accuracy by making derived-artifact handling explicit. Wave 3 applies SLOs and error budgets to prioritise fixes for the dominant contributors to missed deadlines and incomplete runs, driving further improvements in timeliness and accuracy.

Figure 3 plots the simulated metrics from Table 2. The timeliness curve shows P90 completion time in days (lower is better). Coverage shows the percentage of in-scope AI-related systems that successfully participate per request, and accuracy shows the percentage of closed DSRRs that do not require rework.

As a limited validation, these results demonstrate that the control-plane/data-plane separation plus lineage-driven scoping and SRE-style targets can jointly improve DSRR reliability properties that matter for rights delivery. The values should not be interpreted as compliance guarantees or cross-industry baselines; they reflect one plausible parameterisation of a complex operational space.

6.1. Limitations and Threats to Validity

This evaluation is intentionally modest. It is a simulation over a stylised topology rather than a longitudinal field deployment, and it abstracts away organisational factors such as staffing models, legal-review variability, and vendor dependencies. The main threat to validity is that real systems may exhibit different failure correlations (e.g., shared downstream bottlenecks) and different scoping ambiguity (e.g., identity resolution). Future work should therefore pair this framework with (i) pilot deployments in one or two production environments, (ii) open evaluation datasets for DSRR task completion traces, and (iii) comparative studies across organisations and sectors.

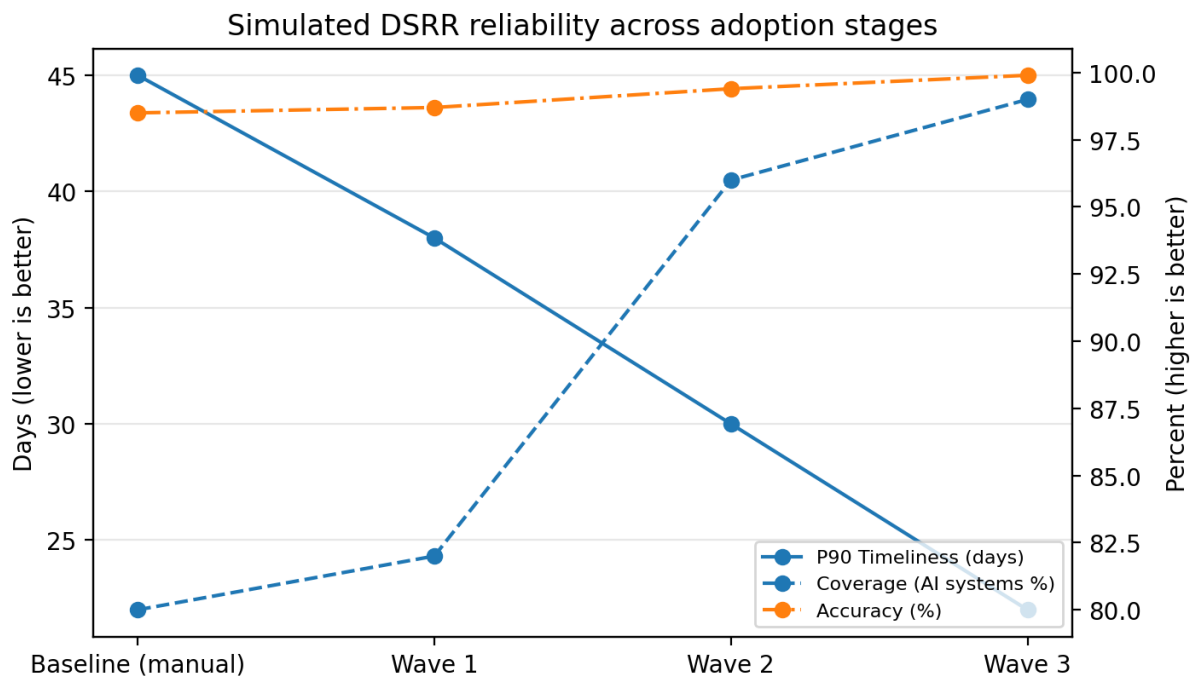


Figure 3. Simulated DSRR reliability metrics across adoption stages (timeliness, coverage, accuracy)

Table 3. Qualitative comparison with existing DSAR/DSRR approaches (representative commercial platforms anonymised as Product 1 and Product 2)

Approach	AI Artifact Coverage	Lineage-Driven Scoping	Explicit SLOs/Error Budgets	Control/Data Plane Separation
Commercial DSAR automation platform (Product 1)	Limited	No	No	Implicit
Commercial DSAR automation platform (Product 2)	Partial	Manual	Basic SLAs	Implicit
Provenance/compliance research (e.g., Ujcich et al.; Tan et al.)	Not AI-specific	Yes	No	N/A
Proposed framework	Explicit features/models	Yes (graph-based)	Yes (formalised)	Yes (architected)

Table 3 provides a concise comparison between representative DSAR automation platforms, prior provenance/compliance research, and the proposed framework. To avoid duplication, we use the table as the primary artefact and highlight only the key differentiators: explicit AI artifact coverage, operationalised lineage-driven scoping, and the use of SLIs/SLOs and error budgets as governance mechanisms.

7. Ethical Analysis and Discussion

The framework has normative implications beyond technical efficiency. By making DSRR performance measurable and reviewable—using the same operational discipline applied to production reliability—organisations create internal incentives to invest in rights delivery rather than treating it as an afterthought.

As discussed in Section 2.4, this supports a shift from formal rights (doctrinal entitlements) to lived rights (delivered outcomes). However, the shift is not automatic: metrics and dashboards can either surface gaps or be used to obscure them.

A primary risk is metricisation: teams may optimise for favourable numbers while undermining the underlying right (e.g., closing cases prematurely, narrowing scope definitions, or privileging easy-to-serve systems). Another is rights-washing, where internal SLO attainment is presented as ethical adequacy without meaningful external accountability. Mitigations include independent audits, explicit reporting of scope exceptions, and evidence-backed claims about derived artifacts and model influence (Section 4.1).

Finally, rights reliability reconfigures organisational decision-making. When error budgets constrain launches that increase DSRR debt, product and engineering teams become accountable for downstream rights impacts. To avoid concentrating power, oversight should include privacy, legal, and risk stakeholders, and—where feasible—public transparency reporting on aggregate reliability metrics.

8. Conclusions and Future Work

This article argues that in AI-intensive environments, the ability to honour data subject rights should be treated as a

reliability property of socio-technical systems, not merely as a legal obligation executed through manual workflows. We propose an SRE-inspired framing with explicit indicators, objectives, and error budgets, implemented via a control-plane/data-plane architecture integrated with data lineage and evidence.

A limited simulated evaluation illustrates the directional benefits of this framing: coordination improves timeliness, lineage-driven scoping improves coverage for AI systems, and SLO/error-budget governance improves both timeliness and accuracy by prioritising remediation of dominant failure modes.

Future work should focus on empirical validation in operational settings, including pilot deployments and cross-organisation studies. It should also deepen the integration with AI-era erasure techniques, including machine unlearning and partial erasure of embeddings, caches, and telemetry, and develop reporting standards that distinguish what was erased, what was retained under lawful exceptions, and what remains as model influence. [14–16]

Declarations

AI tool usage statement: The author used an AI-assisted writing tool to support language polishing and restructuring. All outputs were reviewed and edited by the author, who takes full responsibility for the content.

Competing interests: The author is employed by an organization that performs data subject rights request processing. The views expressed are the author’s own and do not necessarily represent those of the employer. No proprietary or confidential information was used. The author declares no competing interests beyond this employment relationship.

Data availability: No new datasets were created or analysed in this study.

REFERENCES

- [1] European Parliament and Council. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April

- 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation). Official Journal of the European Union L 119, 2016.
- [2] State of California. California Consumer Privacy Act of 2018 (CCPA), Cal. Civ. Code §1798.100 et seq., as amended by the California Privacy Rights Act (CPRA), 2018–2020.
- [3] Government of India. The Digital Personal Data Protection Act, 2023. Government of India Gazette, 2023.
- [4] Commercial DSAR automation platform (Product 1). “Data Subject Request Automation,” product documentation and solution overview, accessed 2025.
- [5] Commercial DSAR request management platform (Product 3). “Data Subject Request Management,” product page and technical overview, accessed 2025.
- [6] Commercial DSAR automation platform (Product 2). “Data Subject Access Request Automation,” solution description (G-Cloud and product documentation), accessed 2025.
- [7] Industry DSAR practice guide. “Data Subject Rights and DSAR: A Practical Guide,” online white paper, accessed 2025.
- [8] Beyer, B., Jones, C., Petoff, J., and Murphy, N. (eds.). Site Reliability Engineering: How Google Runs Production Systems. O’Reilly Media, 2016.
- [9] Ujcich, B. E., Nita-Rotaru, C., Bates, A., and others. “A Provenance Model for the European Union General Data Protection Regulation,” in Proceedings of provenance and compliance workshops, 2018.
- [10] Tan, W., and others. “Data Provenance and Compliance in Big Data Systems,” practice-oriented overview of provenance for regulatory compliance, 2019.
- [11] Stahl, B. C. “The Ethics of Data and Its Governance: A Discourse,” *Information*, vol. 16, no. 6, 2025.
- [12] Floridi, L., and Taddeo, M. “What is Data Ethics?,” *Philosophical Transactions of the Royal Society A*, vol. 374, 2016.
- [13] van Maanen, H. “AI Ethics, Ethics Washing, and the Need to Politicize Data Ethics,” open-access article, 2022.
- [14] Nguyen, T. T., Hua, T., and colleagues. “A Survey of Machine Unlearning,” *ACM Computing Surveys*, preprint 2025.
- [15] Liu, H., Yin, X., and colleagues. “A Survey on Machine Unlearning: Techniques and New Challenges,” *Journal of Information Security and Applications*, 2025.
- [16] European Data Protection Supervisor (EDPS). “TechSonar: Machine Unlearning and Data Protection,” EDPS TechSonar report, 2024.