
Compliance Debt: A Taxonomy for Deferred Governance in Regulated Systems

Benedikt Kolbeinsson
Regava
London, UK
benedikt@regava.com

Arinbjörn Kolbeinsson
Regava
London, UK
arinbjorn@regava.com

Abstract

Organisations routinely defer compliance work, accumulating costs that compound over time. We introduce *compliance debt* as the accumulated cost required to achieve regulatory alignment after deferral and present a taxonomy classifying it by source (regulatory lag, architectural, documentation, process), visibility (explicit vs implicit) and scope (component vs systemic). This framework provides a structured vocabulary for diagnosing governance gaps and supports conceptualisation and future measurement of continuous oversight mechanisms.

1 Introduction

Regulated organisations often postpone governance work until external pressure forces action. Financial institutions delay updating policies after new privacy regulations [1]. AI developers release models without documentation, intending to create it later [2]. Hospital systems implement digital tools before establishing audit trails [3]. This deferred work creates costs that compound as systems evolve and dependencies multiply.

This mirrors technical debt in software engineering, where expedient choices create future maintenance costs. However, the governance analogue lacks systematic treatment. Existing terms such as “compliance backlog” or “governance gap” log outstanding work but fail to convey compounding costs or provide classification frameworks [4–6]. The term *compliance debt* has appeared in requirements engineering [7]; we generalise this to the organisational governance level, offering a taxonomy that captures accumulation across regulatory, architectural, documentation and process dimensions.

We define **compliance debt** as the accumulated effort, cost and risk required to achieve full alignment with applicable obligations, measured from a state of known or *reasonably foreseeable* misalignment resulting from deferred governance activities, including gaps that would be detectable under a reasonable monitoring effort. Key characteristics include temporal accumulation (debt grows as gaps persist) and cumulative structure (deferring one task makes subsequent work harder). This differs from non-compliance: non-compliance describes a static state of misalignment, whereas compliance debt quantifies the accumulated remediation effort required to transition from that state to full alignment, including compounding costs from delayed action; and from technical debt by often being triggered by external regulatory change and *amplified* by internal design or process constraints [8].

Our contribution is a multi-dimensional taxonomy enabling organisations to classify and prioritise different forms of compliance debt. The framework is particularly relevant where regulatory obligations outpace organisational implementation [5, 9, 10], but generalises to any domain subject to oversight. See supplementary appendices for applications, future directions, and implementation guidance.

2 Taxonomy

This taxonomy synthesises compliance patterns observed across regulated sectors, informed by the technical debt framework [8], practitioner consultations in financial services, and documented regulatory adaptation challenges [5]. The three dimensions capture how compliance debt originates (source), whether it is recognised (visibility), and its organisational reach (scope).

① **By source.** Compliance debt originates from four sources, distinguished by trigger: one external (regulatory change) and three internal (system design, documentation practices, and operational processes). **Regulatory lag debt** arises when new or updated regulatory obligations take effect but organisational systems, processes or practices have not yet been adapted to meet them [5, 10]. When new data protection rules tighten consent requirements but an organisation’s practices remain unchanged for months, decisions made during this period (data collected or models trained) may require retrospective review. **Architectural debt** occurs when system design constrains future compliance. An AI model trained without provenance tracking [11, 12] cannot meet later transparency requirements without costly re-engineering. **Documentation debt** involves missing artefacts required for demonstrating compliance. Even with sound practices, absent documentation [13, 14] (model cards, risk assessments or audit trails) creates debt because proving alignment becomes difficult. **Process debt** reflects missing governance procedures such as review workflows, approval gates or accountability structures. An organisation may have compliant policies without processes ensuring consistent application, creating debt that becomes visible during audits when missing procedures lead to inconsistent decisions [15, 16].

② **By visibility.** Compliance debt varies in visibility. **Explicit debt** is recognised and tracked. Regulatory changes appear on risk registers, missing documentation is logged and known gaps are discussed in governance forums. Organisations consciously defer this work while monitoring backlogs, making trade-offs between priorities. Explicit debt is manageable because it can be prioritised and resourced systematically [17]. **Implicit debt** is unrecognised or underestimated. We include implicit debt where misalignment would be objectively detectable under reasonable monitoring, even if unrecognised internally. Teams may believe they are compliant while unaware that recent updates apply or that documentation standards have shifted. For example, an organisation may deploy infrastructure that violates new data localisation requirements because internal policies have not yet been updated to reflect the change. Implicit debt accumulates silently until audit or incident forces discovery, posing greater risk because organisations cannot manage what they do not measure. Compliance debt may also be recognised but underestimated, in which case the acknowledged portion constitutes explicit debt whilst the underestimated remainder is implicit.

③ **By scope.** Compliance debt also differs in scope. **Component-level debt** affects specific systems, models or departments. A single AI application lacking bias testing or one unit with outdated procedures exemplifies component-level debt. Remediation is localised without requiring organisation-wide coordination. However, component-level debt can proliferate across systems, eventually demanding broader intervention. **Systemic debt** spans multiple components or shared infrastructure. Organisation-wide absence of data lineage tracking creates debt across every data-dependent process. Cross-cutting concerns such as consent management or audit logging exemplify systemic debt. Remediation requires coordinated effort across teams, often involving restructuring or new tooling. Addressing systemic debt can also eliminate multiple component-level debts simultaneously, making the trade-off between targeted and systemic remediation a central governance decision [6].

3 Implications

The compliance debt framework provides a structured vocabulary for governance discussions, enabling teams to distinguish types of deferred work rather than using binary “compliant/non-compliant” framings. It suggests indicators for continuous oversight such as regulatory lag (time between obligation effective dates and implementation), documentation completeness (proportion of required artefacts present) and process maturity (coverage of defined workflows). These indicators can underpin private governance mechanisms [4, 18]. Appendices provide further detail on applications, future research, implementation guidance, governance regimes, and practical considerations.

References

- [1] Manuel Holler, Benjamin van Giffen, Seth Benzell, and Matthias Ehrat. The general data protection regulation in financial services industries: How do companies approach the implementation of the GDPR and what can we learn from their approaches? In *Proceedings of the 82nd Annual Business Researchers Conference (VHB 2020)*, Frankfurt, Germany, 2020.
- [2] Avinash Bhat, Austin Coursey, Grace Hu, Sixian Li, Nadia Nahar, Shurui Zhou, Christian Kästner, and Jin L.C. Guo. Aspirations and practice of ML model documentation: Moving the needle with nudging and traceability. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*, pages 1–17. ACM, 2023. doi: 10.1145/3544548.3581518.
- [3] Gina Rollins. Following the digital trail: Weak auditing functions spell trouble for an electronic record. *Journal of AHIMA*, 77(3), 2006.
- [4] Gillian K. Hadfield and Jack Clark. Regulatory markets: The future of AI governance. *arXiv preprint arXiv:2304.04914*, 2023.
- [5] Araz Taeihagh. Governance of generative AI. *Policy and Society*, 44(1):1–22, 2025. doi: 10.1093/polsoc/puaf001.
- [6] Jennifer Cobbe, Michael Veale, and Jatinder Singh. Understanding accountability in algorithmic supply chains. *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 2023. doi: 10.1145/3593013.3594073.
- [7] Bendra Ojameruaye and Rami Bahsoon. Systematic elaboration of compliance requirements using compliance debt and portfolio theory. In *Requirements Engineering: Foundation for Software Quality (REFSQ 2014)*, volume 8396 of *Lecture Notes in Computer Science*, pages 152–167. Springer, 2014. doi: 10.1007/978-3-319-05843-6_11.
- [8] Philippe Kruchten, Robert L. Nord, and Ipek Ozkaya. Technical debt: From metaphor to theory and practice. *IEEE Software*, 29(6):18–21, 2012. doi: 10.1109/MS.2012.167.
- [9] Oladapo Olatinsu. FinTech disruption and the compliance lag: Challenges in supervising non-traditional financial institutions. *International Journal of Science and Research Archive*, 10(2):1458–1472, 2023. doi: 10.30574/ijrsra.2023.10.2.0985.
- [10] Esmat Zaidan and Imad Antoine Ibrahim. AI governance in a complex and rapidly changing regulatory landscape: A global perspective. *Humanities and Social Sciences Communications*, 11, 2024. doi: 10.1038/s41599-024-03560-x.
- [11] Malte Herschel, Robert Diestelkämper, and Housseem Ben Lahmar. A survey on provenance: What for? what from? *The VLDB Journal*, 26(6):881–906, 2017. doi: 10.1007/s00778-017-0486-1.
- [12] NIST AI RMF 1.0. Artificial intelligence risk management framework. Standard, National Institute of Standards and Technology, Gaithersburg, MD, USA, 2023. URL <https://www.nist.gov/itl/ai-risk-management-framework>.
- [13] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. Model cards for model reporting. In *Proceedings of the 2019 Conference on Fairness, Accountability, and Transparency*, pages 220–229. ACM, 2019. doi: 10.1145/3287560.3287596.
- [14] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. Datasheets for datasets. *Communications of the ACM*, 64(12):86–92, 2021. doi: 10.1145/3458723.
- [15] NIST SP 800-53 Rev. 5. Security and privacy controls for information systems and organizations. Standard, National Institute of Standards and Technology, Gaithersburg, MD, USA, 2020.
- [16] ISO/IEC 27001:2022. Information security, cybersecurity and privacy protection — information security management systems — requirements. Standard, International Organization for Standardization, Geneva, Switzerland, 2022. ISO/IEC 27001:2022(E).

- [17] Heike Bockius and Nadine Gatzert. Organizational risk culture: A literature review on dimensions, assessment, value relevance and improvement levers. *European Management Journal*, 42(4):539–564, 2024. doi: 10.1016/j.emj.2023.02.002.
- [18] Anat Lior. Insuring AI: The role of insurance in artificial intelligence regulation. *Harvard Journal of Law & Technology*, 35(2), 2022.
- [19] Barbara Kiviat. The moral limits of predictive practices: The case of credit-based insurance scores. *American Sociological Review*, 84(6):1134–1158, 2019. doi: 10.1177/0003122419884917.
- [20] Martin Eling, Davide Nuessle, and Julian Staubli. The impact of artificial intelligence along the insurance value chain and on the insurability of risks. *The Geneva Papers on Risk and Insurance – Issues and Practice*, 47:205–241, 2022. doi: 10.1057/s41288-020-00201-7.

Appendix

A Potential Applications

The compliance debt framework can underpin various private governance mechanisms. Insurance markets could use debt indicators for premium differentiation [19, 20]. Certification schemes could grade organisations based on debt levels. Procurement frameworks could incorporate debt thresholds into vendor evaluation.

B Future Research

Future research should develop detailed measurement frameworks, case studies quantifying remediation costs and analyses of how compliance debt can support the private governance mechanisms increasingly proposed for AI and other regulated technologies. The framework contributes to emerging discussions on private governance [4, 6] and to broader debates on risk and compliance management [17] by providing a structured language for identifying and managing deferred oversight work, forming a foundation for continuous assurance and certification mechanisms.

C Implementation Guidance

The compliance debt framework serves different stakeholders in distinct ways:

Compliance teams can use the taxonomy to categorise and prioritise their backlog of deferred work. By distinguishing regulatory lag debt from architectural or process debt, teams can allocate resources appropriately: regulatory lag may require policy updates, whilst architectural debt demands engineering involvement.

Risk managers can incorporate debt indicators into enterprise risk frameworks. Tracking explicit versus implicit debt helps identify where monitoring gaps may leave the organisation exposed to unrecognised compliance failures.

Auditors and assurance providers can use the framework to structure compliance assessments. Rather than binary pass/fail evaluations, auditors can characterise the nature and severity of compliance gaps, providing more actionable findings.

Executive leadership can use debt metrics to inform resource allocation and strategic planning. Understanding whether debt is component-level or systemic helps determine whether targeted remediation or broader transformation is required.

Regulators and policymakers may find the framework useful for understanding implementation challenges. Patterns of architectural or process debt across a sector may indicate that regulatory requirements need clearer technical guidance.

D Application Across Governance Regimes

The compliance debt framework applies differently across governance contexts:

Hard law and mandatory regulation. In contexts with binding legal obligations and active enforcement (e.g., financial services regulation, data protection law), compliance debt carries direct legal and financial risk. Organisations in these sectors already maintain internal compliance registers, and the taxonomy provides a structured approach to existing practice. The primary value is systematic prioritisation and resource allocation.

Soft law and voluntary standards. Industry codes of conduct, voluntary certification schemes, and best practice guidelines create compliance expectations without direct legal enforcement. Here, compliance debt may be more tolerable but still carries reputational and market access implications. The framework helps organisations decide which voluntary commitments to prioritise.

Emerging regulatory domains. In areas like AI governance, where formal regulation is still developing but governance frameworks and ethical guidelines proliferate, organisations face uncertainty about which standards will become mandatory. The taxonomy helps track alignment with emerging expectations, reducing future regulatory lag debt when formal requirements crystallise.

Internal governance policies. Organisations often set internal standards that exceed regulatory minimums. Debt against internal policies (e.g., data ethics guidelines, security standards) may not carry regulatory risk but affects operational consistency and organisational culture. The visibility dimension is particularly relevant here, as implicit debt against internal standards can undermine governance credibility.

E Practical Considerations

One potential objection is that organisations would avoid quantifying compliance debt due to regulatory exposure risk. However, this misunderstands compliance practice in regulated sectors. Organisations face ongoing obligations to achieve and maintain compliance. Identifying and quantifying gaps is not optional but a necessary first step in remediation. Compliance teams routinely maintain internal registers of regulatory gaps precisely to address them before regulators discover violations. The taxonomy provides a structured framework for this existing practice, enabling systematic prioritisation and resource allocation. Internal debt quantification serves risk management rather than creating avoidable exposure.