Journal of Information Systems Engineering and Management
2026, 11(2s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

The Role of SRE Governance in Ensuring Ethical and Transparent Digital Service Delivery

Krishnarjun Senthilvelan
Independent Researcher, USA

ARTICLE INFO

Received: 27 Feb 2026

Accepted: 02 March 2026

ABSTRACT

Digital services used in financial transactions, healthcare provision, education or civic engagement represent critical infrastructure. This means that site reliability engineering must also adhere to ethical, institutional and social dimensions that transcend technical service level objectives and related metrics. While SRE's typical definition of service reliability involves operational metrics of uptime, latency, and errors, there are equity and digital opportunity dimensions to reliability. This article reconsiders SRE as a governance discipline grounded in structures of accountability. The Multi-Layer SRE Governance Model is a novel governance model composed of five inter-related layers. It comprises: normative reliability commitments (SLOs) as socio-technical reliability obligations; policies and controls forming a domain of decision rights and monitoring authority; operational transparency protocols to build trust in incident communication; architectural equity design to provide fairness in distributed systems design; and algorithmic accountability frameworks to govern automated remediation systems. With roots in multi-disciplinary service level management, distributed systems architecture, fairness in resource allocation, and algorithmic accountability, the model articulates how assumed distributions of acceptable failure, prioritized recovery, and differential impact on users are reflected in reliability engineering practice. The governance model advances actionable implementation recommendations to support embedding transparency, auditability and equity in cloud-native and DevOps. In this way, we position SRE governance as an emergent cross disciplinary field at the intersection of distributed systems engineering, digital infrastructure ethics and institutional accountability to support service equity.

Keywords: Site Reliability Engineering Governance, Ethical Digital Service Delivery, Equitable Cloud Architecture, Automated Remediation Accountability, Reliability Transparency Frameworks

1. Introduction to Ethical Dimensions of Site Reliability Engineering

Originally seen as underlay, digital systems now serve as infrastructures in realizing values of economic participation, access to healthcare, opportunity in education, and civic engagement. Like electricity grids or transportation infrastructures, their ethical considerations and social implications grow in importance. This is especially true for decisions of system reliability. Site reliability engineering, the discipline of working to maintain service availability and performance at scale, must therefore go beyond the technical optimization of services, towards governance frameworks that include transparency, accountability, and equity.

Classic metrics for SRE practice assess service operational characteristics like availability, latency, error budgets, and incident response times. While important, these metrics are not the only means of articulating the ethical framework governing reliability commitments. Service-level objectives are not only service targets, but also social contracts that guide who has reliable access to a service, whose outages take priority to respond to, and how to share the effects of failure. The decisions reliability engineers make about which groups get access and opportunities for health care platforms, financial transaction systems, digital educational resources, and civic engagement platforms directly affect issues of justice, equity, and institutional accountability.

Journal of Information Systems Engineering and Management
2026, 11(2s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

The July 2024 global CrowdStrike outage has since acted as a socio-technical case study of reliability incidents propagating through tightly coupled, globally interconnected digital infrastructures, with subsequent multi-sectoral service degradation (e.g. in healthcare, telecommunications, aviation, financial services) across millions of endpoints, that were temporarily common-cause failures due to a bad security update disseminated by centralized digital distribution channels, leading to the cancellation of surgeries, flight disruptions, transaction suspensions, and public service outages. An analysis by industry and government described the cause as a systemic reliability failure caused by architectural coupling and central control mechanisms which increased the single points of failure into multiple interconnected systems.

The incident was prominent not just for its impact on operations, but also for the disproportionate effect on some users, in particular healthcare services, emergency services, and vulnerable members of the community. Though the economic and demographic costs have not yet been fully registered, the case illustrates how reliability performance failures have social and institutional impacts that can vary across different affected groups, and thus shows the need to include equity criteria and frameworks, prioritization tools, and transparency measures into the governance setup for reliability engineering-decision required systems [1][2].

Governance of an SRE program is the policies, accountability, decision rights, and transparency practices that govern how reliability objectives should be defined, enforced, governed, and communicated to stakeholders. Governance practices also include governance assessments of who has decision rights over the reliability behavior of the system. What are the review and audit processes for reliability decisions? What processes are in place to understand and measure differential impacts across populations? What transparency obligations apply to service-level commitments? Policy-as-code tools, structured change management processes, equity-aware postmortems, and incident reporting guidelines are all part of the reliability governance package that envelopes SRE practice and situates it in ethical, organizational, and socio-political terms, rather than technical operational terms [1][2].

This article conceptualizes SRE governance as a new interdisciplinary field intersecting distributed systems engineering, the ethics of digital infrastructures and the governance of institutions, proposing a design framework for realizing principles of transparency, equity and accountability in contemporary cloud-native infrastructure and DevOps. The framework does so by making the practicalities of governance explicit at different levels and across different governance mechanisms, positioning reliability engineering as a governed socio-technical practice aimed at ethical digital service delivery.

2. Governance Frameworks for Transparent Service Level Management

Service level objectives have historically been an internal operational goal but are now used to codify ethical commitments, as well as introduce transparency into the provision of digital services. Governance frameworks that support the negotiation, declaration, and management of reliability targets can help build trust between service providers and users. Research into service level management shows that there is a wide variation in the transparency of the practice and the degree of user understanding and fairness [3][4].

A number of traits of cloud computing services have been defined and grouped in eight top-level attributes, namely security and privacy, finance, performance, agility, usability, accountability, assurance, and management. Siomko et al. developed a hierarchical approach to determine the overall service quality multi-attribute global inference of quality (MAGIQ) considering 64 primary performance indicators. In assigning weights to the elements of the service, a rank order centroids approach has been applied, with weights of 0.611, 0.278 and 0.111 for security, performance, and finance, respectively. In a number of service evaluation contexts, a finding has been that organizations incorporating user-preferred features into a service selection model reach better alignment [3].

Journal of Information Systems Engineering and Management
2026, 11(2s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

This transparency includes the communication of faults, how the service can be recovered, and how the user can seek redress. Machine-readable SLA formats such as WSLA (Web Service Level Agreement) and WS-Agreement enables SLA monitoring and management. SLA specification standardization research seeks to address this challenge. Theoretically, WSLA allows for precise specifications of the parameters, metrics and obligations for each service level [4]. In practice, a complete set of SLA parameters is required to define the measurable attributes of a service, such as availability (99.9% uptime), latency (two-second response time), capacity (maximum number of concurrent requests), etc. It is critical to provide verifiable commitments that allow compliance to be assessed objectively [4].

Governance processes are essential to determine whether users experience the reliability promises in practice and across various contexts. The SLA life cycle consists of the six phases of SLA template and service development, negotiations, preparations, execution, assessments, and terminations [4]. QoS assessments, including QoS audits, user satisfaction assessments and changing QoS requirements, can easily identify gaps between promised and achieved QoS. SLA monitoring architectures are designed to ease such measurements and are generally placed on both sides of the measurement target. Assessment subsystems calculate the service level parameters, and the evaluation functions check if the service meets the agreed terms, helping to ensure consistent quality for various user groups.

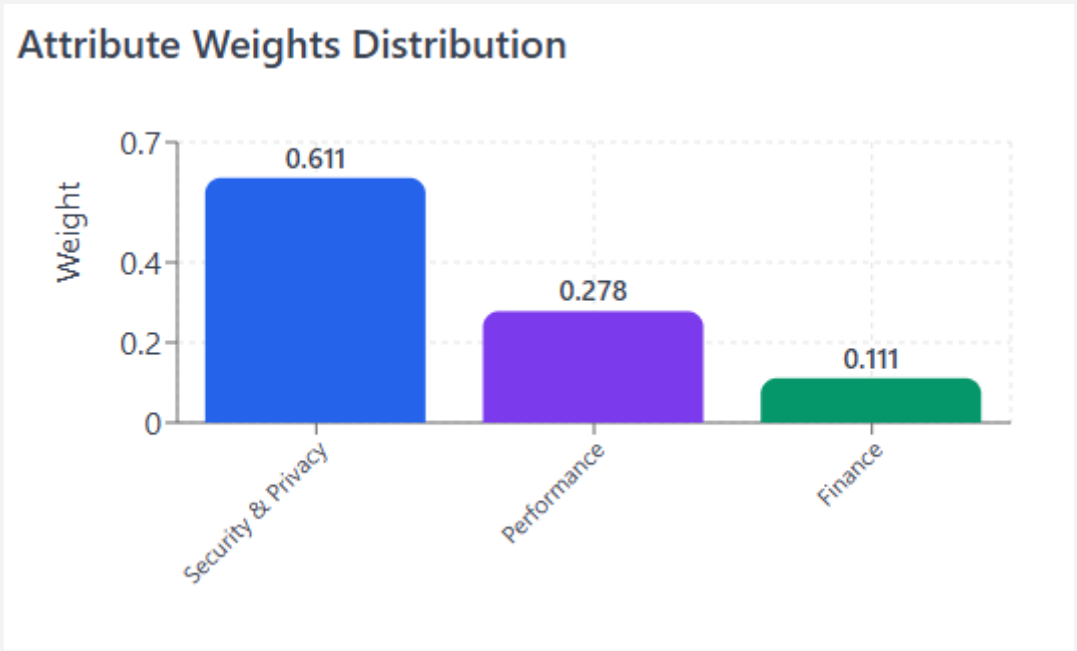

Fig 1: Service Level Objective Attribute Prioritization Using Rank Order Centroids [3, 4]

3. Incident Management as an Ethical Practice

Incident management is one area where reliability engineering and moral obligations to users and other stakeholders meet. The methods, expectations, and guidelines for how an organization detects, responds to, communicates about, and learns from outages can signal its commitment to transparent accountability and equitable treatment of affected populations. Empirical studies of incident management procedures show variation in procedures and their impact on users' trust and perceptions of fairness [5][6].

Service recovery attempts on social media also have an important effect on third parties, as transparency around how complaints are addressed affects their reactions to the situation. High process transparency, or knowledge of what happens between a company and a complaining customer, also increases customers' forgiveness (Mprocess=5.55, SD=1.11 vs. Conversely, participants

Journal of Information Systems Engineering and Management
2026, 11(2s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

who were exposed to the process The results showed that the variable led to higher e-satisfaction compared to another group (t(211)=3.68, p<.001, d=.505), increased e-loyalty (mean process standard deviation vs. mean result standard deviation t(211)=7.84, p<.001), and a lower intention to spread negative electronic word-of-mouth (SD=.91). Trust among consumers was found to be lower when only the results were revealed compared to when both the process and outcomes were shared (t(211)=-4. It was found that consumer forgiveness played a role in how transparency affected e-loyalty and the intention to share negative online comments, showing that both factors work together. These quantitative results suggest transparent incident communication is not merely a reputation management exercise but rather a trust-building practice that acknowledges both user impact and organizational learning.

New methods are being developed in the ethics of incident detection that include fairness evaluations as part of looking back at past incidents. Upon carefully reviewing 30 articles on hospital incident reporting and complaint systems, organizations has discovered that many harms, particularly those affecting marginalized groups, remain unreported. An automated method to find safety issues revealed that voluntary reporting systems set up by providers underreported negative events impacting Black patients and patients with limited English skills. Organizations that performed equity-guided assessments of adverse events identified bias and social determinants of health as causes of adverse events. This procedure helped them determine what adjustments to implement, such as providing situational bias training and including patient-advocacy organizations in root cause analysis meetings [6]. This process ultimately resulted in policy changes that prevented these adverse outcomes from recurring [6].

Ethics of communication during the crisis can include the timing of communication, the medium of communication, and the type of information communicated. For example, one study found that the transparency process and hybrid recovery (psychological acknowledgment + concrete compensation) lead to considerably more customer forgiveness (Mhybrid=4.99, SD=1.21) than using concrete (Mtangible=4.15, SD=1.57) or psychological recovery (psychological=4.39, SD=1.58) alone [5]. Interactions between transparency and service recovery were meaningful for customer forgiveness (F(1,315)=4.60, p=.011, η^2p=.028) and e-loyalty (F(1,315)=5.71, p=.004, η^2p=.035) [5]. Thus, being transparent when dealing with incidents is an ethical obligation. This transparency must include procedural transparency and appropriate redress that sufficiently addresses the emotional and material losses sustained by affected users [5][6].

Communication Dimension	Process Transparency Approach	Result-Only Transparency Approach	Impact on User Trust
Timing of Disclosure	Early notification with ongoing updates throughout resolution	Communication only after incident resolution	Influences perception of organizational responsiveness
Information Depth	Detailed explanation of what happened, why it occurred, and remediation steps	Summary of final outcome without procedural details	Affects user understanding and confidence
Recovery Communication	Combination of acknowledgment and concrete compensation measures	Single-dimension recovery (either tangible or psychological)	Shapes forgiveness and loyalty patterns
Stakeholder Engagement	Inclusive communication involving affected parties in root cause analysis	Internal-only investigation with limited external communication	Determines level of collaborative problem-solving
Learning	Public sharing of lessons	Minimal disclosure of	Influences long-term

Documentation	learned and preventive measures	organizational learning	trust building
Equity Considerations	Assessment of differential impacts across user populations	Aggregate impact reporting without segmentation	Affects fairness perceptions among diverse groups

Table 1: Quantitative Impact of Transparency on Service Recovery and Incident Management [5, 6]

4. Designing Equitable Reliability Architectures

Architectural decisions in distributed systems strongly influence the expectations of reliability by different types of users as the architecture implicitly defines the types of failures and recovery events that are tolerable. Architectural reliability requires consideration of the distribution of failures and recovery across affected populations. Instead of focusing on aggregate measures of availability, such as uptime, equitable reliability requires careful consideration of user-impact distributions. Research into architectural patterns can help determine how equity considerations can be integrated into technical design decisions [7][8].

Resource allocation policies in multi-tenant systems are another equity-related architecture. For example, various implementations of DRF (dominant resource fairness) in clouds have compared the customary CPU-memory DRF to the multi-resource 3D DRF, which includes disk, memory, and CPU resources. In terms of CPU, memory, and disk perishability, the 3D DRF achieved 87.50%, 94.44%, and 100% system utilizations, respectively, compared to the 79% and 83% CPU and memory utilizations achieved by the CPU-mem DRF [7]. By stopping jobs that need less CPU and memory but a lot of disk I/O from using too many resources, the better DRF algorithm increased average resource use by 15% and cut the time to finish jobs by 59%. These results illustrate how architectural resource allocation mechanisms encode fairness decisions, which affect performance and fair resource allocation patterns within multitenant systems.

Another example of an architectural decision affecting infrastructure equity is the design of failover mechanisms and deployment strategies in cloud computing environments. Different levels of failure independence can be achieved based on replica placement within the infrastructure, directly impacting both reliability guarantees and equitable service delivery across user populations [8]. The architectural topology of cloud data centers—spanning rack-level components, cluster configurations, and geographic distribution—creates natural failure boundaries that must be considered when designing resilient systems.

Failover mechanism selection represents a critical architectural choice with significant equity implications. Research on fault tolerance in cloud environments identifies several replication strategies with varying characteristics [8]. Active replication simultaneously invokes all replicas to process the same request at the same time, ensuring all replicas maintain identical system states unless designed for asynchronous operation. This approach enables continuous service delivery even with single replica failure but requires substantial resource investment. Passive replication designates one processing unit as the primary replica to handle requests while backup replicas only preserve system state during normal execution, taking over only when primary failure occurs. Semi-active replication executes all instructions on both primary and backup replicas but suppresses backup output, allowing immediate resumption when primary failure happens [8].

The deployment context of these replication mechanisms significantly affects their efficacy in providing equitable reliability. Multiple deployment scenarios demonstrate varying levels of failure independence [8]:

- Multiple machines within the same cluster: Replicas placed on hosts connected by a top-of-rack (ToR) switch benefit from low latency and high bandwidth but obtain very limited failure independence. A single switch or power distribution failure may result in complete application outage, as both replicas cannot communicate to complete fault tolerance protocols. Analysis using Markov models demonstrates that semi-active replication achieves

Journal of Information Systems Engineering and Management
2026, 11(2s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

0.9871 availability, semi-passive replication achieves 0.9826 availability, and passive replication achieves 0.9542 availability when deployed within the same cluster [8].

- Multiple clusters within a data center: Replicas distributed across different clusters connected via ToR switches and aggregation switches achieve moderate failure independence. This configuration avoids outages from single power distribution or switch failures. Availability improves to 0.9913 for semi-active, 0.9840 for semi-passive, and 0.9723 for passive replication schemes [8].
- Multiple data centers: Geographic distribution of replicas across data centers connected via switches, aggregation switches, and access routers provides the highest level of failure independence despite drawbacks in latency and bandwidth. Single power failures have minimal effect on application availability. This deployment achieves 0.9985 availability for semi-active, 0.9912 for semi-passive, and 0.9766 for passive replication [8].

These availability differentials demonstrate how architectural choices regarding failover mechanisms and deployment topology directly encode equity decisions. Organizations must balance resource costs, performance requirements, and failure independence levels when determining appropriate architectural configurations. Hot standby configurations with geographic redundancy provide maximum protection but require significant infrastructure investment, while cold standby systems within single clusters minimize costs but create concentrated failure risk that may disproportionately affect user populations during outages [8].

Equity has been incorporated in commercial production settings as an architectural configuration for DRF-based dependency risk assessment. In the context of real cloud applications, a telecommunication data analysis application could achieve 95.33% CPU, 91.66% memory, and 92.00% I/O using equity-aware DRF compared with 79.41% CPU, 85.00% memory, and 81.64% I/O using default DRF-based configurations [7]. This new fairness algorithm led to a standard deviation of only 8% in terms of I/O variance in resource consumption, versus 26% in the default algorithms [7]. These design choices show how considerations of fairness in resource consumption can be represented in concrete technical design when a system is planned and implemented to achieve fair resource consumption among tenant applications and workloads [7][8].

Resource Allocation Approach	CPU Utilization	Memory Utilization	I/O Utilization	Equity Impact
CPU-Memory DRF	Moderate	Moderate	Lower	Standard variance
3D DRF (Multi-resource)	High	High	Maximum	Improved distribution
Equity-Aware DRF	Highest	High	High	Minimal variance

Table 2: Failover Mechanism Characteristics [7, 8]

5. Automated Remediation and Algorithmic Accountability

Automated incident response systems enable excellent reliability outcomes, but they also introduce new governance challenges of explainability and accountability as incidents are triaged through self-healing, machine learning-based anomaly detection, and algorithmic decision-making components. New governance frameworks are needed to ensure these systems are fair and safe to operate in SRE practice.

Fairness testing of automated systems has been shown to contribute to managing discrimination in the outputs. In one study, Aequitas generated up to 70% of discriminatory inputs in the total inputs for testing of machine learning classifiers. Of six state-of-the-art machine learning algorithms, including fairness-optimized classifiers, all exhibited important fairness violations. In benchmark tests, Aequitas was on average 9.6x and 20.4x better than random testing strategies and up to 83.27%

Journal of Information Systems Engineering and Management
2026, 11(2s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

faster than standard testing strategies in terms of speed, with peak performance gains of 96.62% depending on model architecture [9]. Organizations that check their algorithms for fairness can use these automated testing methods to help fix any unfair behavior.

The second principle of governance for automated reliability systems is auditability, or the ability to access information about decision-making after the fact. Digital Object Identifiers (DOIs) enable traceability from technical artifacts to institutionally accountable entities when applied to algorithms [10]. This can be done via semantic metadata structures to ease versioning and auditing of algorithmic traces so that any problems in an algorithm's behavior can be attributed to the version of the algorithm. A three-level classification system (abstract algorithm logic, code-level reference implementation, and AI systems with trained models) would include all systems that make automated decisions.

These automated retraining capabilities have the potential to yield meaningful fairness benefits, re-training on discriminatory inputs identified through the system. Aequitas reduces discriminatory input usage by 43.2% (max: 94.36%) on average by adding only 22.92% additional data points to a training set [9]. These improvements were computationally inexpensive, with the enormous majority of classifiers retraining within one hour. As such, they are suitable for use in high-accountability production environments.

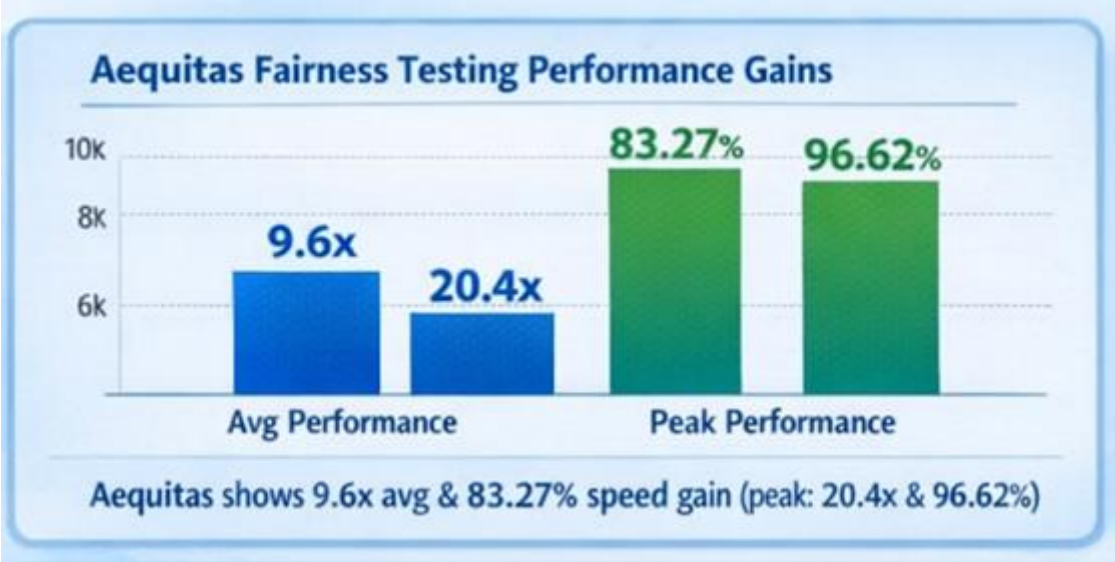

Fig 2: Comparative Performance Analysis of Aequitas Fairness Testing Framework [9, 10]

6. Public Accountability Through Reliability Transparency

Transparency of reliability performance represents an emerging mechanism for accountability in digital service provision. Beyond accountability functions, transparency mechanisms may foster competition among service providers and enable more informed user decision-making, thereby incentivizing improvements in reliability practices. Contemporary transparency measures span a spectrum from rudimentary uptime metrics to comprehensive scorecards incorporating multidimensional performance indicators and normative benchmarks. The efficacy of disclosure practices depends upon the interplay among transparency strategies, user comprehension capabilities, and organizational behavioral responses.

Research on information signaling within service transparency contexts has demonstrated that structured presentation of service attributes positively influences user acceptance and decision-making processes, particularly when information is systematically organized [11]. For organizations implementing transparent disclosure protocols, empirical evidence indicates substantial alignment

Journal of Information Systems Engineering and Management
2026, 11(2s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

between user perceptions and actual service properties, suggesting that well-designed disclosure mechanisms effectively facilitate user understanding across diverse service domains. The effectiveness of disclosure protocols correlates more strongly with the accuracy and relevance of disclosed attributes than with the volume of information provided [11].

Standardized frameworks for reliability reporting offer potential improvements in cross-provider comparability and information comprehensibility. Implementation research has established that compliance-driven standardization can simultaneously advance accountability objectives and operational excellence [12]. Financial institutions have employed automated security and compliance verification mechanisms within continuous delivery pipelines to maintain regulatory adherence while enabling innovation velocity. Similarly, healthcare organizations have adopted standardized incident response and monitoring procedures that preserve compliance obligations while ensuring service continuity. These implementations demonstrate how structured transparency frameworks can streamline and normalize compliance processes across organizational contexts [12].

Governance frameworks must reconcile transparency obligations with operational security requirements. Organizations implementing standardized cross-cloud governance processes have reported enhanced regulatory compliance alongside reduced operational complexity, facilitating more uniform approaches to data governance [12]. A principal advantage of these frameworks lies in their capacity to embed regulatory compliance within operational workflows rather than treating accountability as an externally imposed constraint. Through automation tools that enforce standardized configurations and compliance policies across heterogeneous platforms, organizations can maintain robust security postures while delivering transparent services [12]. This evidence supports the proposition that transparency and security objectives need not constitute competing priorities; rather, they can function as mutually reinforcing elements when appropriate standardization and verification mechanisms are implemented [11][12].

7. A Formal Model of Ethical SRE Governance

To consolidate the ethical dimensions discussed throughout this paper, this study proposes a structured model of Ethical SRE Governance composed of five interrelated components: Governance Scope, Accountability Actors, Fairness Validation, Transparency Reporting, and Adaptive Oversight.

Governance Scope defines the boundaries within which reliability decisions are made. This includes SLO definitions, error budget policies, architectural failover strategies, and automated remediation thresholds [4][8]. In this model, service-level objectives are treated not only as performance metrics but as institutional commitments that determine acceptable levels of service risk across user populations [1][3]. Formalizing governance scope ensures that reliability configurations align with compliance obligations, equity considerations, and organizational accountability [12].

Accountability Actors specify who holds decision rights and oversight authority over reliability behavior [1]. These actors typically include SRE teams (operational execution), product or business owners (risk prioritization), and governance or compliance bodies (regulatory alignment). Clear role definition ensures traceability of architectural changes, SLO adjustments, and automated system deployments [10][12].

The Fairness Validation Layer integrates equity assessment into reliability engineering practices. Architectural decisions—such as resource allocation policies, failover mechanisms, and traffic routing—should be evaluated for distributional impact across user groups [7][8]. Automated remediation and anomaly detection systems require periodic fairness testing and auditability controls to prevent the amplification of structural inequities [9][10]. This layer ensures that efficiency optimization does not unintentionally disadvantage vulnerable populations [6][7].

The Transparency Reporting Layer transforms internal reliability governance into public accountability. Structured incident communication, uptime reporting, and documented lessons learned help reinforce trust [5][6]. Transparency must balance disclosure with security, but it should

provide sufficient procedural visibility to demonstrate fairness, responsiveness, and institutional learning [11][12].

Finally, the Feedback and Adaptive Oversight Loop ensures continuous refinement. Post-incident reviews, QoS audits, fairness testing outcomes, and user feedback are incorporated into policy updates and architectural recalibration [4][5][6]. Governance maturity is reflected in the organization's ability to adapt reliability objectives based on changing societal expectations and regulatory environments [1][2].

To support implementation, Ethical SRE Governance may be assessed using measurable indicators, such as: documented equity impact assessments for SLOs [3][4], defined accountability matrices for reliability decisions [10], frequency of fairness testing in automated systems [9], timeliness of incident disclosure [5][11], and inclusion of differential impact analysis in post-incident reviews [6]. Organizations may adopt this model incrementally, beginning with formal SLO governance documentation [4][12], progressing toward structured transparency practices [11], and ultimately embedding fairness validation and adaptive oversight into automated reliability systems [8][9][10].

This model positions reliability engineering not merely as an operational function, but as a governed socio-technical practice aligned with ethical digital service delivery [1][2].

Conclusion

The article sees SRE governance (SREG) as a next step in maturity for digital services, determining whether users have the same access rights to money, health, education and civic services. This article introduces a Multi-Layer SRE Governance Model that maps reliability obligations across user populations with institutional obligations across this wide range of impacts. Unlike traditional IT governance frameworks, which primarily focus on compliance alignment, enterprise risk management, and high-level policy oversight, SRE governance operates at the level of real-time operational reliability where technical configurations immediately shape the continuity and distribution of essential services. In this sense, reliability decisions are not abstract governance artifacts but direct determinants of infrastructural access across diverse user communities. Service-level objectives can define acceptable failure distributions and priority levels for recovery, which may reflect equity-related priorities, and such governance mechanisms can be implemented through service-level management, inclusive incident response, equity-aware architecture, algorithmic accountability, and transparency reporting. Architectural decisions, including resource allocation, failover decisions, and topologies, must consider fairness. Automated remediation systems should be validated for fairness and be auditable to prevent structural inequities.What is needed for future academic research is empirical evidence from longitudinal studies, maturity assessments and equity reliability tools, governance frameworks for new technology, and regulatory frameworks of critical infrastructure operators. Comparative studies across sectors may further clarify how reliability governance influences institutional trust, resilience outcomes, and distributive service equity. Ethical SRE governance therefore represents not merely an extension of operational practice, but a structural redefinition of reliability engineering—from technical optimization toward accountable infrastructure stewardship. As digital systems increasingly function as critical societal infrastructure, the governance of reliability becomes inseparable from questions of equity, transparency, and institutional responsibility. Designing reliability as a governed socio-technical practice is thus a prerequisite for building resilient and trusted digital infrastructures in an inclusive society.

References

[1] David Tilson et al., "Research Commentary---Digital Infrastructures: The Missing IS Research Agenda," Information Systems Research, Volume 21, Issue 4, 2010. [Online]. Available: https://dl.acm.org/doi/10.1287/isre.1100.0318

[2] Rajiv D. Banker et al., "R&D Versus Acquisitions: Role of Diversification in the Choice of

Innovation Strategy by Information Technology Firms," Journal of Management Information Systems, Volume 28, 2011. [Online]. Available: https://www.tandfonline.com/doi/abs/10.2753/MIS0742-1222280205
[3] Farrukh Nadeem, "A Unified Framework for User-Preferred Multi-Level Ranking of Cloud Computing Services Based on Usability and Quality of Service Evaluation," IEEE Access, 2020. [Online]. Available: https://ieeexplore.ieee.org/document/9208735
[4] Philip Bianco et al., "Service Level Agreements in Service-Oriented Architecture Environments," Software Engineering Institute, 2008. [Online]. Available: https://apps.dtic.mil/sti/tr/pdf/ADA528751.pdf
[5] Jie CAI and Yoonseo Park, "Managing Service Recovery via Social Media: The Impact of Transparency and Service Recovery Type in the Distribution of Feedback," Journal of Distribution Science 22-1, 2024. [Online]. Available: https://koreascience.kr/article/JAKO202402443615612.pdf
[6] Joanne Goldman et al., "Integrating equity into incident reporting and patient concerns systems: A critical interpretive synthesis," 2022. [Online]. Available: https://qualitysafety.bmj.com/content/qhc/35/1/64.full.pdf
[7] Jia Ru et al., "Providing Fairer Resource Allocation for Multi-tenant Cloud-based Systems," 7th IEEE International Conference on Cloud Computing Technology and Science, 2015. [Online]. Available: https://nzjohng.github.io/publications/papers/cloudcom2015.pdf
[8] Ravi Jhawar and Vincenzo Piuri, "Fault Tolerance and Resilience in Cloud Computing Environments," Computer and Information Security Handbook (Third Edition), 2017. [Online]. Available: https://www.sciencedirect.com/science/chapter/edited-volume/abs/pii/B9780128038437000090
[9] Sakshi Udeshi et al., "Automated directed fairness testing," ACM Digital Library, 2018. [Online]. Available: https://dl.acm.org/doi/pdf/10.1145/3238147.3238165
[10] Juliana C. Braga et al., "Algorithmic Identity Based On Metaparameters: A Path To Reliability, Auditability, And Traceability," arXiv:2601.16234v1, 2026. [Online]. Available: https://arxiv.org/pdf/2601.16234
[11] DAVID ZAR et al., "Information Disclosure for Increasing User Satisfaction From a Shared Ride," IEEE Access, 2023. [Online]. Available: https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=10086493
[12] Michael Armbrust et al., "A View of Cloud Computing," Communications of the ACM, 2010. [Online]. Available: https://dl.acm.org/doi/pdf/10.1145/1721654.1721672