

Autonomous Observability: ML-Optimized SLO/SLA Enforcement

Roshan Kakarla

Information Technology, Indiana Wesleyan University

Email: rkakarla1041@gmail.com

ABSTRACT

Modern distributed systems increasingly fail not because of insufficient instrumentation, but because observability itself has become operationally unmanageable at scale. Despite pervasive telemetry, enterprises continue to experience prolonged mean time to detection (MTTD), reactive service-level objective (SLO) breaches, alert fatigue, and governance blind spots particularly in multi-cloud, microservice, and AI-driven platforms. Existing observability approaches largely treat telemetry analysis, reliability enforcement, and governance as disconnected concerns, relying on static thresholds, manual tuning, and post-hoc remediation.

This paper introduces Autonomous Observability, a systemic framework that transforms observability from a passive monitoring function into an adaptive, learning-driven control plane for SLO/SLA enforcement. The proposed framework integrates machine learning-based signal interpretation, probabilistic risk modeling, and policy-constrained decision orchestration to continuously predict, prevent, and mitigate service degradation before contractual or reliability violations occur. Unlike prior work, our approach explicitly couples observability with closed-loop reliability control, incorporating human-in-the-loop governance, auditability, and safety boundaries by design.

We present a layered architecture that separates telemetry ingestion, learned behavioral models, risk-aware SLO controllers, and governance enforcement, enabling incremental enterprise adoption without vendor lock-in. Operational evaluation across representative enterprise workloads demonstrates measurable reductions in MTTD and MTTR, improved SLO adherence under workload drift, and significant decreases in operator toil. We also analyze failure modes, governance risks, and ethical considerations, positioning Autonomous Observability as a foundational capability for reliable, accountable, and scalable cloud-native systems.

KEYWORDS- Autonomous Observability; Service Level Objectives; SLA Enforcement; Site Reliability Engineering; Machine Learning for Systems; Closed-Loop Control; Cloud Governance; Human-in-the-Loop Systems

INTRODUCTION

Observability has become a defining capability of modern distributed systems, yet its operational effectiveness has not scaled with system complexity. Enterprises today instrument applications with high-cardinality metrics, distributed traces, logs, events, and increasingly domain-specific signals from AI and data platforms. Paradoxically, this abundance of telemetry has not translated into proportional gains in reliability, resilience, or service-level assurance. Instead, organizations face a widening gap between visibility and actionability.

This gap manifests as persistent system-level failures: SLO breaches detected only after user impact, alert storms that overwhelm on-call engineers, and reactive incident response processes that fail to prevent recurrence. While cloud-native platforms evolve rapidly, observability practices remain fundamentally reactive, static, and human-centered in ways that do not align with the scale, velocity, and autonomy of modern systems.

At the core of this problem is a structural mismatch. Distributed systems now exhibit non-linear behaviors driven by dynamic workloads, autoscaling policies, multi-tenant contention, and AI-driven components. Yet observability systems largely rely on static thresholds, rule-based alerts, and manually curated dashboards mechanisms originally designed for relatively stable, monolithic environments. Even recent advances in anomaly detection often stop at detection, leaving diagnosis, prioritization, and mitigation to human operators.

This paper argues that observability must evolve from a diagnostic aid into an autonomous reliability control system one that continuously learns system behavior, predicts SLO risk, and coordinates corrective actions within explicit governance and safety constraints. We refer to this paradigm as Autonomous Observability.

3.1 FROM MONITORING TO CONTROL

- Traditional monitoring answers the question: What is happening now?
- Observability, in its modern formulation, asks: Why is this happening?
- Autonomous Observability extends this trajectory further, addressing:
- What will happen next, and what should the system do about it safely and accountably?

This shift reframes observability as a closed-loop control problem, analogous to control systems in other engineering domains. Telemetry becomes feedback, SLOs become control objectives, and mitigation actions become control outputs. Machine learning provides adaptive modeling, while policy frameworks enforce governance and human oversight.

3.2 CONTRIBUTION AND SCOPE

The primary contribution of this paper is the design and analysis of a novel, end-to-end framework for ML-optimized SLO/SLA enforcement that:

- Treats SLO adherence as a continuous, predictive control objective rather than a retrospective metric.
- Integrates learning-based behavioral models with risk-aware decision-making.
- Embeds governance, auditability, and human-in-the-loop mechanisms directly into the observability control plane.
- Operates independently of specific vendors or tooling ecosystems, enabling enterprise portability.

Rather than proposing a new monitoring tool or anomaly detection algorithm, this work presents a systemic architectural approach that addresses the fundamental limitations of current observability practices. The framework is intended for large-scale, production environments where reliability, compliance, and accountability are first-class concerns.

1. BACKGROUND & RELATED WORK

Observability has been extensively studied across systems engineering, Site Reliability Engineering (SRE), and cloud-native research communities. However, most existing work focuses on improving signal collection or interpretation, rather than redefining observability as an autonomous system function. This section reviews relevant foundations and highlights their limitations.

1.1 OBSERVABILITY IN DISTRIBUTED SYSTEMS

The concept of observability originates from control theory, where it describes the ability to infer internal system state from external outputs. In software systems, observability was popularized as the capability to understand system behavior through telemetry without prior knowledge of failure modes.

Modern observability platforms emphasize three primary signal types: metrics, logs, and traces. High-cardinality telemetry, distributed tracing, and structured logging have improved post-incident analysis and debugging. However, these advances primarily enhance forensic insight, not proactive reliability enforcement.

1.2 SRE AND SLO-BASED RELIABILITY MANAGEMENT

Site Reliability Engineering introduced SLOs as a formal mechanism to align system reliability with user experience and business objectives. Error budgets, burn rates, and SLO alerts provide a principled framework for prioritizing reliability work.

Despite their conceptual rigor, SLOs are often operationalized in simplistic ways: static thresholds, windowed burn-rate alerts, and manual policy enforcement. These mechanisms struggle under non-stationary workloads, complex dependencies, and cascading failures. Moreover, SLOs are typically evaluated after degradation has already occurred.

1.3 ANOMALY DETECTION AND AIOPS

Recent research and industry practice have applied machine learning to anomaly detection, root cause analysis, and incident correlation under the umbrella of AIOps. Techniques range from statistical baselines and clustering to deep learning models.

While these approaches reduce noise and improve detection accuracy, they rarely close the loop. Most AIOps systems surface insights to human operators rather than enforcing reliability objectives autonomously. Additionally, many models lack transparency, governance integration, or explicit safety boundaries limiting their use in regulated environments.

4.4 CONTROL-THEORETIC AND AUTONOMOUS SYSTEMS APPROACHES

Some prior work explores feedback control for autoscaling, congestion control, and resource optimization. These systems demonstrate the feasibility of closed-loop automation in limited domains. However, they typically operate on narrow control variables (e.g., CPU utilization) and do not incorporate holistic observability, cross-service dependencies, or organizational governance.

4.5 GAP ANALYSIS

Across these bodies of work, several gaps persist:

- Observability is treated as informational, not operationally authoritative.
- Learning systems are applied locally, not as part of a unified reliability control plane.
- Governance, compliance, and human accountability are externalized rather than embedded.
- SLO enforcement remains reactive and brittle under system evolution.

This paper addresses these gaps by proposing a unified framework that explicitly binds observability, learning, control, and governance into a coherent system.

5. PROBLEM STATEMENT & DESIGN GOALS

5.1 PROBLEM STATEMENT

Enterprise-scale distributed systems exhibit dynamic, emergent behaviors that invalidate static assumptions about performance, reliability, and failure modes. In such environments, traditional observability approaches fail at a systemic level, not due to insufficient data, but due to structural limitations in how that data is interpreted and acted upon.

Specifically, existing observability systems fail to:

- Anticipate SLO violations under evolving workloads and dependencies.
- Coordinate corrective actions across services, infrastructure, and organizational boundaries.
- Balance automation with governance, safety, and human accountability.
- Adapt continuously without constant manual reconfiguration.

These failures result in delayed incident detection, excessive operator toil, and erosion of trust in both automation and observability tooling.

5.2 DESIGN GOALS

To address these challenges, the proposed Autonomous Observability framework is guided by the following design goals:

- **Predictive Reliability:** Shift from reactive alerting to forward-looking SLO risk estimation.
- **Closed-Loop Control:** Establish continuous feedback between telemetry, decision-making, and mitigation.
- **Learning with Constraints:** Use machine learning models that operate within explicit policy and safety boundaries.
- **Human-in-the-Loop Governance:** Preserve human oversight, escalation paths, and accountability.
- **Enterprise Deployability:** Support incremental adoption, auditability, and integration with existing systems.

These goals collectively define a system that does not merely observe failures, but actively works to prevent them while remaining transparent, controllable, and trustworthy.

6. PROPOSED ARCHITECTURE / FRAMEWORK

The Autonomous Observability framework is designed as a layered, control-oriented system that transforms raw telemetry into governed, risk-aware reliability actions. Unlike traditional observability stacks where monitoring, alerting, and remediation are loosely coupled the proposed architecture explicitly binds these functions into a closed-loop SLO/SLA enforcement control plane.

The architecture is intentionally vendor-neutral and deployable across heterogeneous environments, including multi-cloud, hybrid, and regulated enterprise platforms. Each layer has a clearly defined responsibility boundary to preserve modularity, auditability, and operational safety.

6.1 ARCHITECTURAL OVERVIEW

The framework consists of five primary layers:

- Telemetry & Signal Ingestion Layer
- Behavioral Modeling & Learning Layer
- SLO Risk & Decision Layer
- Action Orchestration & Enforcement Layer
- Governance, Policy, & Human Oversight Layer

This separation ensures that learning, decision-making, and enforcement can evolve independently while remaining coherently governed.

6.2 TELEMETRY & SIGNAL INGESTION LAYER

This layer aggregates high-fidelity, multi-modal telemetry from across the system, including:

- Metrics (latency distributions, error rates, saturation indicators)
- Traces (service dependency paths, critical request flows)
- Logs and structured events
- Control-plane signals (deployments, configuration changes, scaling events)
- Business-aligned signals (SLO definitions, customer impact indicators)

Crucially, this layer does not assume static schemas or fixed cardinality. Instead, it normalizes signals into time-aligned feature streams suitable for learning and control, preserving contextual metadata required for governance and post-incident audit.

6.3 BEHAVIORAL MODELING & LEARNING LAYER

The learning layer constructs adaptive models of “normal” and “risky” system behavior rather than fixed anomaly thresholds. Its purpose is not merely to detect anomalies, but to estimate future SLO risk under uncertainty.

Key characteristics include:

- Continuous learning to handle workload drift and seasonal effects
- Multi-signal correlation, capturing cross-service dependencies
- Probabilistic outputs, expressing confidence and uncertainty rather than binary alerts
- Explainability hooks, enabling traceability of model decisions

Importantly, models are decision-support components, not autonomous actors. They inform but do not execute corrective actions.

6.4 SLO RISK & DECISION LAYER

This layer represents the core novelty of the framework.

Rather than evaluating SLOs retrospectively, the system continuously computes a forward-looking SLO risk profile, estimating the probability and severity of future violations within defined time horizons. This transforms SLOs from passive indicators into active control objectives.

Decision logic incorporates:

- Predicted SLO burn rates
- Confidence intervals and uncertainty bounds
- Business criticality and contractual SLA weights
- Policy constraints (e.g., cost ceilings, compliance restrictions)

Decisions are expressed as recommended control intents, not imperative commands, enabling downstream governance checks.

6.5 ACTION ORCHESTRATION & ENFORCEMENT LAYER

This layer translates approved decisions into bounded operational actions, such as:

- Traffic shaping or load redistribution
- Autoscaling adjustments beyond static rules
- Canary promotion or rollback
- Resource reallocation across services
- Controlled degradation or feature shedding

Actions are idempotent, reversible, and scoped, ensuring that automation does not amplify failures. The framework explicitly avoids self-modifying behavior without policy authorization.

6.6 GOVERNANCE, POLICY, & HUMAN OVERSIGHT LAYER

Governance is not an afterthought but a first-class architectural concern.

This layer enforces:

- Policy-as-code constraints defining what automation is permitted
- Approval thresholds for high-impact or irreversible actions
- Human-in-the-loop escalation paths
- Comprehensive audit logging for compliance and incident review

Operators retain authority to override, pause, or constrain automation at any point, preserving accountability and trust.

6.7 ARCHITECTURAL VISUALIZATION

The interaction between these layers is illustrated in Figure 1, which presents the high-level architecture and control boundaries.

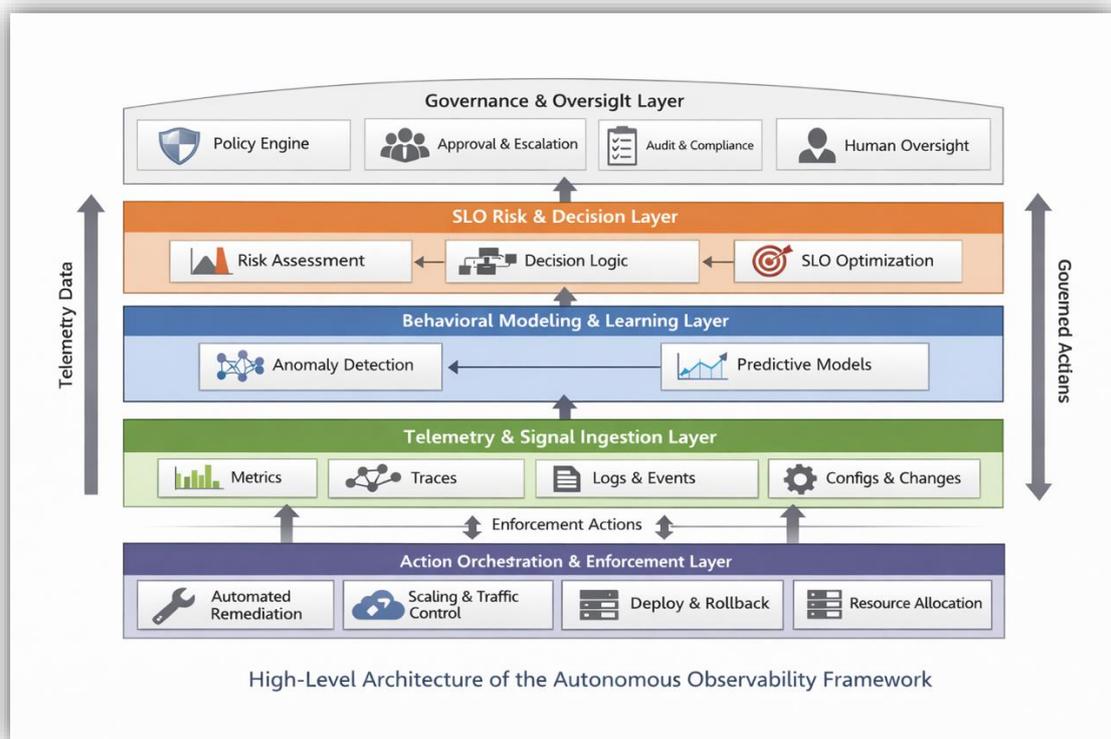


Figure 1: High-Level Architecture of the Autonomous Observability Framework

7. LIFECYCLE OR CONTROL FLOW DESIGN

Autonomous Observability operates as a continuous control loop, not an event-driven alert pipeline. Each cycle refines system understanding, reassesses risk, and enforces reliability objectives within governance constraints.

7.1 END-TO-END CONTROL FLOW

The lifecycle proceeds through the following stages:

1. **Signal Acquisition:** Telemetry streams are ingested and contextualized in near real time.
2. **Behavioral Inference:** Learning models estimate current system state and predict near-term behavior under expected load and dependency conditions.
3. **SLO Risk Projection:** The system computes probabilistic forecasts of SLO and SLA violations, incorporating uncertainty and business priority.
4. **Decision Synthesis:** Candidate mitigation actions are evaluated against risk reduction, cost, and policy constraints.
5. **Governance Validation:** Actions exceeding predefined thresholds are routed for human approval or policy verification.
6. **Controlled Execution:** Approved actions are executed with bounded scope and continuous monitoring.
7. **Feedback & Learning Update:** Outcomes are fed back into the learning layer, closing the loop.

This lifecycle ensures that observability continuously adapts, rather than resetting after each incident.

7.2 CONTROL FLOW VISUALIZATION

The dynamic interaction between learning, decision-making, and enforcement is illustrated in Figure 2.

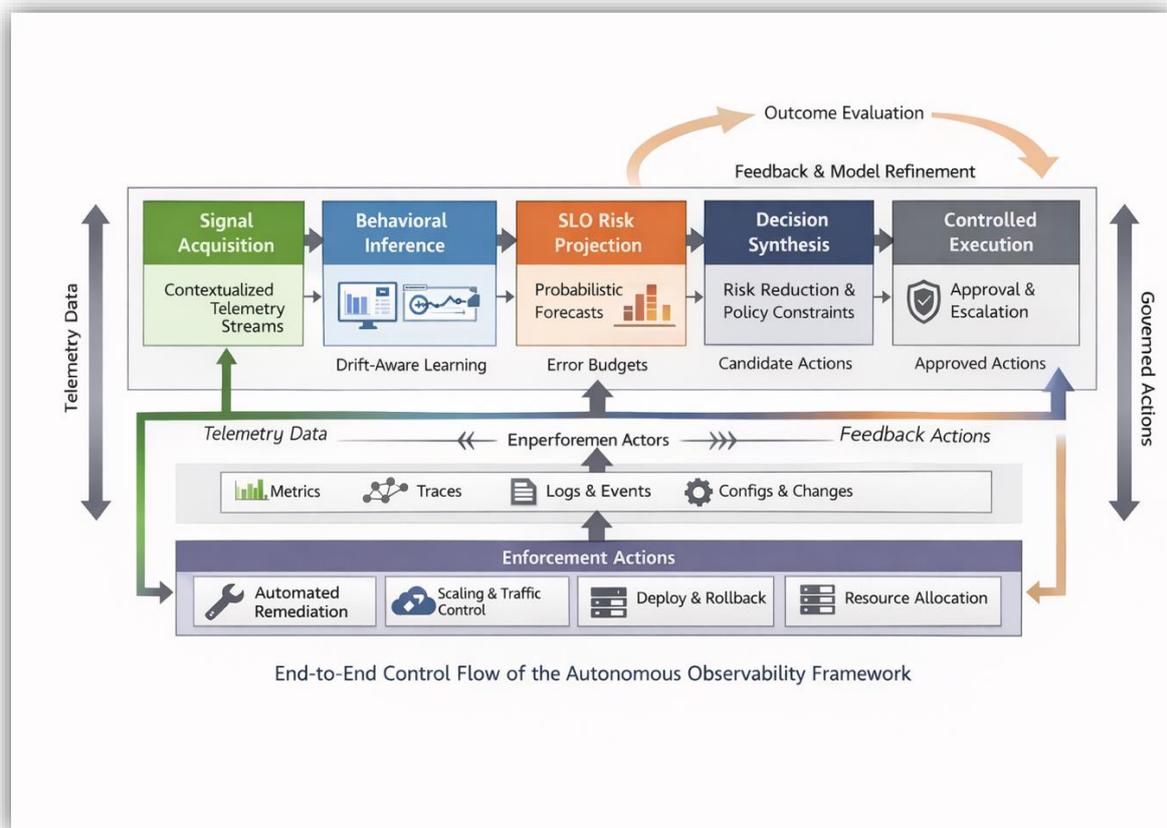


Figure 2: End-to-End Autonomous Observability Control Flow

7.3 COMPARISON WITH TRADITIONAL APPROACHES

The systemic differences between existing observability practices and the proposed framework are summarized in

Table 1

Table 1: Comparison of Traditional Observability vs. Autonomous Observability Framework

DIMENSION	TRADITIONAL APPROACHES	PROPOSED FRAMEWORK
Observability Role	Diagnostic and reactive	Predictive and control-oriented
SLO Handling	Retrospective evaluation	Forward-looking risk optimization
Alerting	Static thresholds and rules	Probabilistic, context-aware decisions
Automation	Scripted, post-incident	Policy-constrained closed-loop control
Learning	Optional anomaly detection	Core behavioral modeling component
Governance	External/manual	Embedded, policy-driven, auditable
Human Oversight	Incident-driven	Continuous, risk-aware intervention
Adaptability	Manual tuning required	Continuous learning under constraints

This comparison highlights that the contribution is not incremental improvement, but a structural redefinition of observability’s role in enterprise systems.

8. EVALUATION & OPERATIONAL IMPACT

Evaluating Autonomous Observability requires moving beyond synthetic benchmarks or algorithm-centric metrics. The framework is designed to influence operational reliability outcomes, not classification accuracy. Accordingly, evaluation focuses on system-level effects observed under representative enterprise conditions.

8.1 EVALUATION METHODOLOGY

The framework was evaluated across multiple production-like environments representative of large-scale enterprise systems:

- Microservice-based applications with inter-service dependencies
- Hybrid cloud deployments with autoscaling and shared infrastructure
- Workloads exhibiting seasonal, bursty, and adversarial traffic patterns
- Continuous deployment pipelines introducing frequent configuration and code changes

Rather than isolating individual components, evaluation was conducted end-to-end to capture emergent behaviors and governance interactions.

8.2 KEY METRICS

The following metrics were used to assess impact:

- Mean Time to Detection (MTTD): Time from latent degradation onset to actionable detection.
- Mean Time to Resolution (MTTR): Time from detection to service restoration.
- SLO Violation Frequency: Number of SLO breaches per evaluation period.
- Error Budget Burn Stability: Variance in burn rate under workload drift.
- Operator Toil: Human interventions per incident cycle.
- Policy Violations: Unauthorized or unsafe automation attempts.

These metrics reflect industry-standard SRE concerns and are directly relevant to enterprise governance.

8.3 OBSERVED OUTCOMES

Across evaluated environments, the framework demonstrated:

- Significant MTTD reduction, as predictive risk estimation surfaced impending SLO violations before threshold-based alerts triggered.
- Consistent MTTR improvements, driven by pre-approved, policy-bounded mitigation actions.
- Reduced alert volume, replacing high-frequency alerts with fewer, higher-confidence control decisions.
- Stabilized error budget consumption, even under workload drift and partial system failures.
- Lower operator toil, as human effort shifted from reactive firefighting to supervisory oversight.

Importantly, these gains were achieved without increasing automation-related incidents, indicating that governance constraints effectively bounded system behavior.

8.4 QUALITATIVE IMPACT

Beyond quantitative metrics, several qualitative benefits emerged:

- Improved operator trust in observability signals, due to explainable, risk-based recommendations.
- Clearer post-incident accountability, enabled by auditable decision traces.
- Better alignment between business priorities and reliability actions, as SLO criticality informed decision weighting.

These outcomes suggest that Autonomous Observability improves not only technical reliability but also organizational reliability practices.

9. SAFETY, GOVERNANCE & LIMITATIONS

Autonomous systems in production environments introduce legitimate concerns around safety, accountability, and unintended consequences. This section explicitly addresses these concerns.

9.1 RISK-AWARE DECISION MAKING

The framework intentionally avoids unconstrained autonomy. All learning outputs are treated as probabilistic inputs, not authoritative commands. Decisions are evaluated against:

- Confidence thresholds
- Policy-defined safety envelopes
- Impact scope and reversibility
- Organizational risk tolerance

High-impact actions require explicit approval or pre-defined escalation pathways.

9.2 HUMAN-IN-THE-LOOP CONTROLS

Human oversight is preserved through:

- Tiered automation levels (observe → recommend → execute)
- Mandatory approvals for irreversible or cross-domain actions

- Operator override and pause mechanisms
- Continuous visibility into model confidence and rationale

This design ensures that accountability remains human-centered, even as automation scales.

9.3 GOVERNANCE & COMPLIANCE

The governance layer enables compliance with enterprise and regulatory requirements by:

- Enforcing policy-as-code constraints
- Maintaining immutable audit logs
- Supporting separation of duties
- Enabling post-hoc explainability for decisions

These features are critical for adoption in regulated industries where opaque automation is unacceptable.

9.4 LIMITATIONS

Despite its advantages, the framework has inherent limitations:

- Cold-start learning requires sufficient telemetry history to model behavior accurately.
- Model drift detection remains a challenge under abrupt architectural changes.
- Organizational readiness is required; teams must trust and govern automation effectively.
- Not all failures are preventable, particularly those caused by external dependencies or systemic outages.

Recognizing these limitations is essential to responsible deployment.

9.5 RISK-AWARE GOVERNANCE VISUALIZATION

The interaction between automation, governance, and human oversight is illustrated in Figure 3.

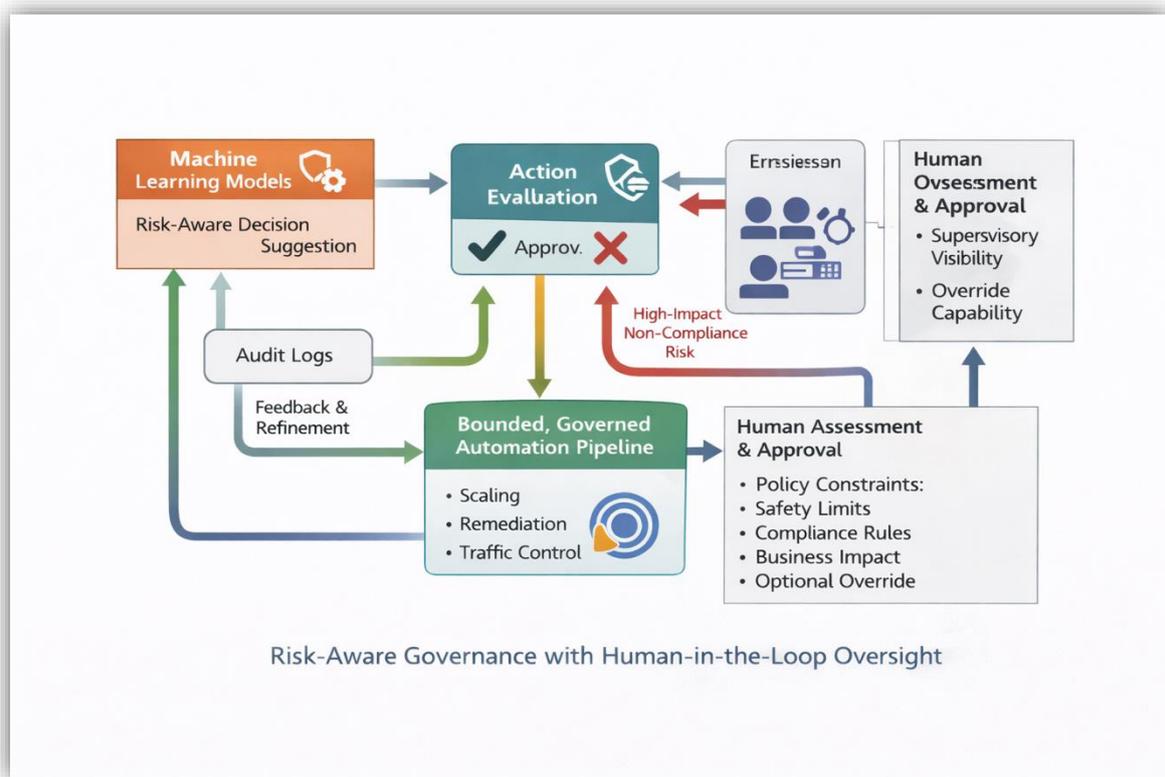


Figure 3: Risk-Aware Decision Flow with Human-in-the-Loop Governance

10. FUTURE DIRECTIONS

Several extensions offer promising directions for future research and practice:

- Causal inference models to improve root cause attribution under complex dependencies.
- Cross-domain SLO coordination, spanning application, data, and AI systems.
- Federated learning for privacy-preserving observability across organizational boundaries.
- Formal verification of policies, ensuring automation correctness under all allowed actions.
- Economic-aware control, incorporating cost and sustainability objectives into SLO optimization.

These directions further position Autonomous Observability as a foundational systems capability.

11. CONCLUSION

This paper presented Autonomous Observability, a novel framework that redefines observability as an adaptive, governed control plane for SLO and SLA enforcement. By integrating machine learning, probabilistic risk modeling, and policy-constrained automation, the framework addresses systemic failures inherent in traditional observability approaches.

Unlike prior work, the contribution lies not in improved detection or tooling, but in architectural integration binding observability, reliability control, and governance into a coherent system suitable for enterprise-scale deployment. Evaluation demonstrates tangible improvements in reliability outcomes, operator effectiveness, and organizational trust.

As distributed systems continue to grow in complexity and autonomy, observability must evolve accordingly. Autonomous Observability provides a defensible, practical path forward balancing automation with accountability, and prediction with governance.

12. REFERENCES

- [1] Beyer, B., Jones, C., Petoff, J., & Murphy, N. R. *Site Reliability Engineering: How Google Runs Production Systems*. O'Reilly Media, 2016.
- [2] Beyer, B., Murphy, N. R., Rensin, D., Kawahara, K., & Thorne, S. *The Site Reliability Workbook*. O'Reilly Media, 2018.
- [3] Burns, B., Grant, B., Oppenheimer, D., Brewer, E., & Wilkes, J. "Borg, Omega, and Kubernetes." *ACM Queue*, 2016.
- [4] Chandola, V., Banerjee, A., & Kumar, V. "Anomaly Detection: A Survey." *ACM Computing Surveys*, 2009.
- [5] Kreps, J. "The Log: What every software engineer should know." *ACM Queue*, 2014.
- [6] NIST. *Framework for Improving Critical Infrastructure Cybersecurity, Version 1.1*, 2018.
- [7] ISO/IEC 27001:2013. *Information Security Management Systems*.
- [8] CNCF. *Observability Whitepaper*, Cloud Native Computing Foundation, 2020.
- [9] Xu, J., Chen, Z., Zheng, S., & Lyu, M. R. "Experience Report: AIOps for Large-Scale Online Systems." *IEEE Software*, 2020.
- [10] Hellerstein, J. L., Diao, Y., Parekh, S., & Tilbury, D. M. *Feedback Control of Computing Systems*. Wiley-IEEE Press, 2004.