
Multi-Agent Orchestration for Autonomous Data Pipeline Governance: Schema Evolution, Anomaly Detection, and Incident Remediation in Cloud-Native Data Platforms

Praveen Kumar Dora Mallareddi

Submitted: 03/12/2025

Revised: 17/01/2026

Accepted: 28/01/2026

Abstract: The rapid adoption of cloud-native data platforms has enabled organizations to scale data processing pipelines to unprecedented levels. However, governance mechanisms—particularly around schema evolution, anomaly detection, and incident remediation—remain largely manual, leading to increased operational risk and degraded data reliability. This paper proposes a novel multi-agent orchestration framework for autonomous data pipeline governance. The system leverages specialized agents for schema monitoring, anomaly detection, service-level agreement (SLA) tracking, and incident remediation, coordinated through a shared state and communication protocol. Evaluated against production-like workloads, the framework demonstrates significant improvements in detection latency, mean time to resolution (MTTR), and system reliability. The results suggest that agentic AI can address critical governance gaps in modern data infrastructures while maintaining safety through controlled autonomy.

Keywords: Multi-agent systems; Data pipeline governance; Schema evolution; Anomaly detection; Autonomous remediation

1. Introduction

Cloud-native data platforms have reshaped how organizations ingest, process, and operationalize data at scale. Systems such as Apache Spark and Apache Kafka underpin modern architectures that favor decoupled, event-driven, and elastically scalable pipelines. These platforms enable continuous data ingestion and transformation across heterogeneous sources, supporting use cases ranging from real-time personalization to large-scale analytical processing. Their adoption has been accelerated by managed cloud services and container orchestration frameworks, which abstract infrastructure complexity while increasing system dynamism.

Yet, this very dynamism introduces a structural imbalance: while execution layers have become highly automated, governance layers have not evolved at the same pace. Data pipelines remain susceptible to silent schema drift, where upstream structural changes propagate without explicit

coordination, often breaking downstream consumers or—more critically—introducing subtle semantic inconsistencies [1,2]. Similarly, data anomalies—whether arising from upstream system faults, data ingestion errors, or distributional shifts—are frequently detected late, if at all. The consequences are non-trivial: delayed detection cascades into compromised analytical outputs, degraded machine learning model performance, and, in high-stakes domains, erroneous business decisions.

Empirical observations from large-scale deployments suggest that mean time to detect (MTTD) data quality issues can extend from hours to days, while mean time to resolution (MTTR) often depends on manual triage involving multiple teams [1,14]. Compounding this, data contracts—formal agreements between data producers and consumers—are rarely enforced programmatically, existing instead as documentation artifacts rather than executable guarantees. This creates a governance gap in which reliability is contingent on human vigilance rather than system-level assurances.

Lawrence Technological University, USA

Email Id

: Praveenkumardoramallareddi@gmail.com

In parallel, the field of artificial intelligence has witnessed renewed interest in multi-agent systems, where collections of autonomous, goal-oriented agents collaborate to solve complex tasks. Foundational work in distributed artificial intelligence and agent coordination has matured into practical applications in domains such as robotics, logistics, and IT operations [3,5]. More recently, agentic AI paradigms—where agents exhibit planning, reasoning, and tool-use capabilities—have demonstrated promise in automating decision-making workflows in dynamic environments [4]. Within IT operations, for example, agent-based approaches have been used to detect anomalies and automate incident response, giving rise to the AIOps paradigm [7,8].

However, the direct application of these advances to data pipeline governance remains limited. Unlike traditional IT systems, data pipelines exhibit unique characteristics:

- (i) schema semantics that evolve over time,
- (ii) lineage dependencies spanning multiple transformation stages, and
- (iii) data quality constraints that are context-dependent and often probabilistic.

Existing automation approaches tend to focus either on infrastructure health or isolated data quality checks, lacking a unified framework that integrates schema management, anomaly detection, and remediation into a cohesive governance model.

This paper addresses this gap by proposing a multi-agent orchestration framework for autonomous data pipeline governance. The core premise is that governance can be decomposed into specialized but cooperative responsibilities—each handled by an intelligent agent capable of monitoring, reasoning, and acting within its domain. By enabling structured communication and shared state across these agents, the system can move beyond reactive monitoring toward proactive and self-healing governance.

2. Background and Related Work

2.1 Agentic AI in Distributed Systems

The study of multi-agent systems has long occupied a central place in distributed artificial intelligence, with foundational contributions

emphasizing coordination, negotiation, and decentralized decision-making among autonomous entities [5,6]. Early frameworks conceptualized agents as rational actors operating under partial observability, capable of cooperation or competition depending on system objectives. Over time, these theoretical constructs have found practical expression in domains such as robotics swarms, distributed control systems, and large-scale optimization problems.

In recent years, the emergence of agentic AI—systems composed of agents that exhibit reasoning, planning, and tool interaction—has renewed interest in applying multi-agent paradigms to complex operational environments. Within enterprise IT, this evolution is particularly visible in the rise of AIOps platforms, where intelligent agents monitor telemetry, detect anomalies, and trigger automated responses [7,8]. These systems leverage statistical learning and rule-based inference to reduce operational burden and improve system reliability.

Despite these advances, the scope of agent-based automation in AIOps remains largely confined to infrastructure-level observability. Metrics such as CPU utilization, memory consumption, and network latency dominate monitoring frameworks, reflecting a system-centric rather than data-centric perspective. Consequently, while such systems can identify and mitigate hardware or service failures, they are not equipped to reason about data semantics, schema evolution, or data quality constraints. The absence of semantic awareness limits their applicability in environments where correctness is defined not merely by system uptime but by the integrity and interpretability of data flowing through pipelines.

2.2 Data Pipeline Governance

In parallel, the domain of data governance has evolved to address the challenges of managing large-scale, heterogeneous data ecosystems. Tools such as data catalogs, lineage trackers, and observability platforms have become standard components of modern data stacks. Systems like Apache Atlas and Amundsen provide mechanisms for metadata management, enabling users to trace data provenance and understand dependencies across pipelines.

While these tools significantly enhance visibility, their capacity for automation remains limited.

Observability platforms can surface anomalies in metrics such as data freshness or volume, but they typically rely on predefined rules or manual inspection to interpret and resolve issues [9,10]. Similarly, schema registries—such as those built around Avro or Protobuf—introduce formal schema versioning and compatibility checks, yet their enforcement is often reactive rather than proactive.

The concept of data contracts has emerged as a promising approach to formalizing expectations between data producers and consumers. By specifying schema structure, data quality constraints, and service-level guarantees, data contracts aim to reduce ambiguity and improve reliability. However, in practice, these contracts are rarely executable artifacts. Their enforcement depends on human intervention, and violations are frequently detected only after downstream impact has occurred [11].

2.3 Research Gap

Existing literature provides substantial coverage of adjacent domains, including infrastructure monitoring, automated incident management, and database self-tuning [12,13]. For example, autonomous database systems have demonstrated the feasibility of self-optimizing query execution and resource allocation, while AIOps platforms have improved incident detection and response in distributed systems. However, these contributions operate at layers that are either below (infrastructure) or within (database engines) the data pipeline abstraction.

What remains underexplored is the governance layer spanning the full lifecycle of data pipelines. This layer encompasses schema evolution, data validation, lineage tracking, SLA enforcement, and incident remediation—each of which introduces unique challenges that cannot be fully addressed by existing tools in isolation. Notably:

- Schema evolution is often handled independently through registries, without integration into downstream validation or remediation workflows.
- Anomaly detection is treated as a monitoring problem, detached from schema semantics and pipeline lineage.
- Incident remediation remains largely manual, even when detection mechanisms are automated.

The absence of a unified framework leads to fragmented governance, where each concern is addressed in isolation, resulting in delayed responses and inconsistent policy enforcement. Moreover, current systems lack the ability to coordinate decisions across these domains, a requirement that becomes critical as pipelines grow in complexity and interdependence.

This paper positions itself within this gap by proposing an integrated, multi-agent approach that treats governance as a coordinated, autonomous process rather than a collection of disjointed tools. By embedding intelligence at multiple points in the pipeline lifecycle and enabling structured interaction among agents, the proposed framework seeks to bridge the divide between observability and action, moving toward a model of self-governing data infrastructure.

3. Problem Definition: The Governance Gap

The evolution of cloud-native data platforms has fundamentally altered the scale and complexity of data processing. Pipelines now span hundreds of interconnected components, ingesting heterogeneous data streams and serving diverse analytical and operational workloads. While this architectural shift has enabled unprecedented agility, it has simultaneously exposed a critical weakness: governance mechanisms have not scaled proportionally with system complexity. As a result, human-in-the-loop oversight—once sufficient for smaller, monolithic systems—has become a structural bottleneck.

Empirical evidence from production environments indicates that mean time to detect (MTTD) data anomalies often extends beyond several hours, particularly in streaming contexts where deviations may be subtle and gradually evolving [14]. Even when anomalies are detected, mean time to resolution (MTTR) frequently spans days, requiring coordinated investigation across data engineering, platform, and business teams. This delay is not merely operationally inefficient; it directly impacts downstream analytics, machine learning models, and decision-making processes that rely on timely and accurate data.

A central contributor to this governance gap is the phenomenon of schema evolution without enforcement. In loosely coupled, event-driven architectures—commonly built on platforms such

as Apache Kafka—producers can introduce schema changes independently of consumers. While this decoupling enhances flexibility, it also allows incompatible changes to propagate silently. Without automated validation or contract enforcement, such changes may lead to immediate pipeline failures or, more insidiously, silent data corruption, where incorrect data continues to flow undetected [15].

The governance gap can be decomposed into four interrelated challenges:

3.1 Schema Drift

Schema drift refers to the uncoordinated evolution of data structures over time. In distributed pipelines, even minor modifications—such as changing a field type or removing an attribute—can break downstream transformations or alter semantic interpretations. The absence of automated compatibility checks and propagation mechanisms means that these changes are often discovered only after failures occur or inconsistencies surface in analytical outputs.

3.2 Anomaly Detection

Data anomalies—ranging from sudden distribution shifts to gradual degradation in data quality—are inherently difficult to detect in large-scale systems. Traditional monitoring approaches rely on static thresholds or periodic validation, which are ill-suited to dynamic workloads. Consequently, anomalies may persist undetected for extended periods, particularly when they do not trigger explicit system failures but instead manifest as subtle deviations in data distributions or relationships.

3.3 Incident Response

When failures or anomalies are eventually identified, incident response remains largely manual. Engineers must trace issues across complex lineage graphs, correlate logs from multiple systems, and infer root causes through iterative analysis. This process is time-consuming and error-prone, particularly in environments where dependencies are not fully documented or where multiple concurrent changes obscure causal relationships.

3.4 SLA Violations

Service-level agreements (SLAs) in data pipelines—covering metrics such as data freshness, completeness, and latency—are rarely enforced programmatically. While monitoring tools may alert on SLA breaches, they do not typically initiate corrective actions or prevent violations proactively. As a result, SLA compliance depends on reactive human intervention rather than continuous, automated enforcement.

4. System Architecture: Multi-Agent Design

The proposed framework conceptualizes data pipeline governance as a distributed, agent-driven control system, where specialized agents operate semi-autonomously yet coordinate through shared context and structured communication. This design reflects a deliberate shift away from monolithic governance services toward modular, composable intelligence, enabling scalability, fault isolation, and incremental extensibility. The architecture is aligned with cloud-native principles, leveraging event-driven interaction patterns and decentralized state management.

Table 1: Agent Responsibilities and Functions

Agent	Primary Role	Key Functions	Outputs
Schema Evolution Agent (SEA)	Schema governance	Schema diffing, compatibility checks, version control	Schema alerts, update actions
Data Quality & Anomaly Detection Agent (DQADA)	Data monitoring	Statistical + ML anomaly detection	Anomaly events
Pipeline Health & SLA Agent (PHSA)	Operational monitoring	SLA tracking, latency monitoring	SLA violation alerts
Incident Remediation Agent (IRA)	Recovery execution	Restart jobs, rollback schema, backfill	Remediation actions

4.1 Agent Topology

At the core of the framework is a set of four primary agents, each responsible for a distinct but interdependent aspect of pipeline governance:

1. **Schema Evolution Agent (SEA):** The SEA continuously monitors structural changes in data schemas across ingestion and transformation stages. It performs schema diffing, classifies changes based on compatibility rules, and orchestrates downstream propagation or intervention workflows. Its primary objective is to ensure schema continuity without compromising correctness.
2. **Data Quality & Anomaly Detection Agent (DQADA):** The DQADA is responsible for identifying deviations in data characteristics. It maintains statistical baselines and model-driven expectations for datasets, enabling detection of both abrupt anomalies and gradual drifts. Beyond detection, it generates contextualized anomaly signals enriched with lineage and metadata to support downstream reasoning.
3. **Pipeline Health & SLA Agent (PHSA):** The PHSA monitors operational aspects of pipelines, including execution latency, throughput, failure rates, and adherence to service-level agreements. It integrates infrastructure-level signals with pipeline-specific metrics, thereby bridging the gap between system observability and data reliability.
4. **Incident Remediation Agent (IRA):** The IRA serves as the execution layer for corrective actions. Based on signals from other agents, it determines appropriate remediation strategies—ranging from automated recovery actions to escalation. Its design emphasizes controlled autonomy, ensuring that interventions remain within predefined safety boundaries.

Each agent is designed as a loosely coupled service, allowing independent deployment and scaling. Despite their autonomy, agents rely on a shared global state, maintained through a distributed coordination layer. This shared state includes schema histories, pipeline lineage graphs,

anomaly records, and SLA definitions, enabling agents to reason with a consistent and up-to-date view of the system.

4.2 Communication Protocol

Coordination among agents is achieved through an event-driven communication model, implemented over a distributed messaging backbone such as Apache Kafka. This choice reflects the need for low-latency, asynchronous interaction in highly dynamic environments.

Agents publish and subscribe to structured events, which encapsulate both signals and contextual metadata. Core event types include:

- **Schema Events:** Emitted by the SEA upon detecting structural changes
- **Anomaly Events:** Generated by the DQADA when deviations exceed defined thresholds
- **SLA Events:** Triggered by the PHSA upon detecting breaches or degradation trends
- **Remediation Events:** Produced by the IRA to document actions taken or proposed

Each event carries a standardized payload that includes identifiers, timestamps, lineage references, and confidence scores. This uniformity enables agents to interpret and act upon events without tight coupling to specific implementations.

Complementing the event bus is a shared metadata store, which functions as the system's persistent memory. Implementations may leverage governance frameworks such as Apache Atlas to maintain lineage graphs, schema versions, and dependency mappings. This store allows agents to perform historical analysis, correlate events across time, and maintain consistency in decision-making [16].

4.3 Escalation Logic

A defining feature of the architecture is its explicit treatment of uncertainty and risk through structured escalation mechanisms. While the system is designed for autonomy, it does not assume that all decisions can or should be automated.

Each agent operates with confidence thresholds, derived from model outputs, rule evaluations, or policy constraints. When a decision falls below the required confidence level—or when the potential impact exceeds predefined risk boundaries—the

agent triggers an escalation workflow. Typical escalation scenarios include:

- High-impact schema changes, such as breaking modifications affecting critical downstream systems
- Ambiguous anomalies, where signals are inconclusive or conflicting across detection methods
- Failed remediation attempts, indicating that automated recovery actions were insufficient

Escalations are routed to human operators through integrated alerting systems, accompanied by detailed context, including lineage traces, historical patterns, and recommended actions. Importantly, escalation is not treated as a failure of the system but as a designed safeguard, ensuring that human expertise is engaged when necessary.

This hybrid model—combining autonomous operation with principled escalation—balances efficiency with reliability. It enables the system to handle routine governance tasks at scale while preserving human oversight for complex or high-stakes decisions.

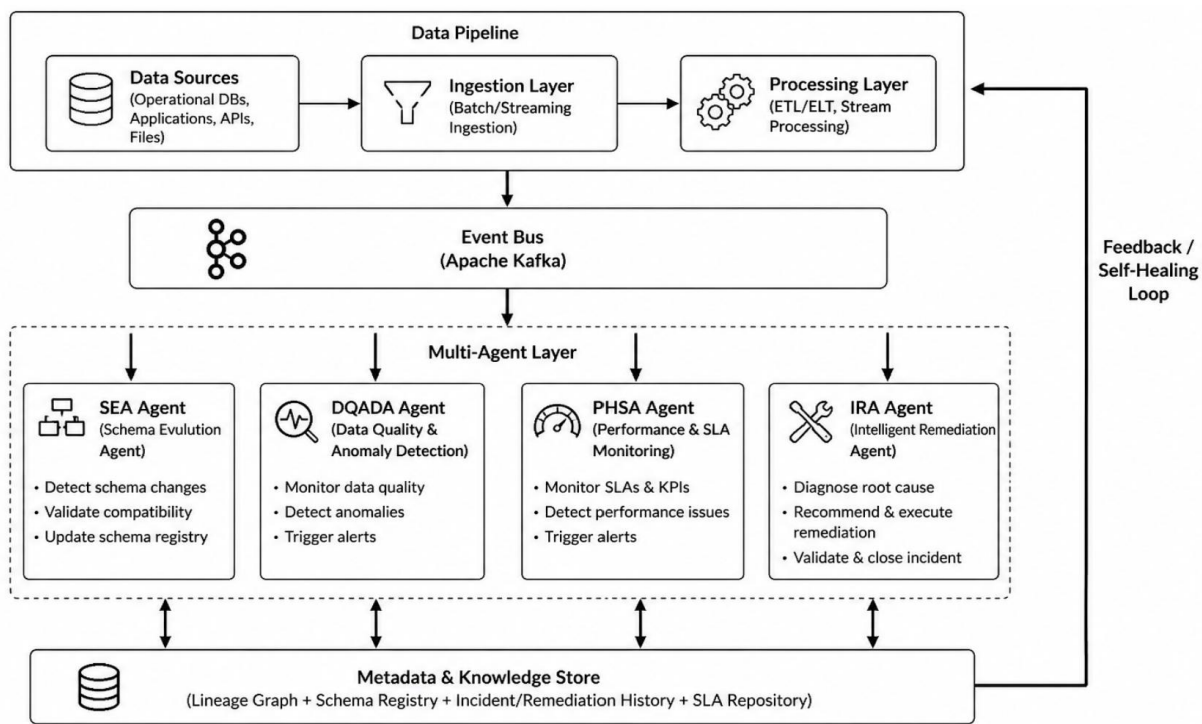


Figure 1. Multi-agent orchestration architecture for autonomous data pipeline governance.

5. Schema Evolution Handling

The Schema Evolution Agent (SEA) is designed to provide continuous, automated oversight of schema changes across distributed data pipelines. In cloud-native architectures—particularly those built on streaming backbones such as Apache Kafka—schema evolution is both inevitable and frequent. Producers evolve independently, often without explicit coordination with downstream consumers. The SEA addresses this challenge by introducing real-time schema validation, classification, and controlled propagation.

At its core, the agent performs structural and semantic diffing between incoming schemas and their registered counterparts [17]. Structural diffing captures syntactic changes—such as field additions or deletions—while semantic diffing evaluates deeper transformations, including changes in data types, field meanings, and constraints. This dual-layer comparison allows the system to move beyond surface-level detection and reason about the impact of changes on downstream processing logic.

Table 2: Schema Change Classification

Change Type	Example	Risk Level	System Action
Non-breaking	Add optional column	Low	Auto propagate
Non-breaking	Add nullable field	Low	Accept + notify
Breaking	Remove column	High	Block or rollback
Breaking	Change data type	High	Escalate

5.1 Change Classification

Detected schema changes are categorized into two primary classes:

- **Non-breaking changes:** These include additive modifications such as the introduction of new fields, optional attributes, or extensions that preserve backward compatibility. Such changes do not invalidate existing consumers and can typically be adopted without interruption.
- **Breaking changes:** These involve modifications that alter the interpretation or structure of existing data, including field removal, type changes, or constraint tightening. Breaking changes risk causing pipeline failures or semantic inconsistencies if propagated without coordination.

This classification is essential for determining the appropriate governance response, enabling the system to differentiate between routine evolution and high-risk modifications.

5.2 Decision Logic

The SEA incorporates a rule-based and policy-driven decision layer to determine how each change should be handled:

- **Automatic propagation:** Non-breaking changes are propagated downstream with minimal intervention. The system updates schema versions and notifies dependent components, ensuring continuity while preserving compatibility.
- **Flagging and intervention:** Breaking changes trigger validation workflows.

Depending on system configuration, these may involve:

- Temporary rejection of the new schema
- Automatic rollback to a previous compatible version
- Escalation to human operators for approval

The decision process integrates contextual signals, including pipeline criticality, historical stability of the data source, and dependency depth within the lineage graph. This ensures that governance actions are not only technically correct but also context-aware.

5.3 Compatibility Management

To enforce consistency, the SEA leverages schema registry mechanisms, such as those implemented in Apache Avro, to maintain versioned schemas and compatibility constraints [18]. These constraints typically include:

- **Backward compatibility:** New schemas can be read by older consumers
- **Forward compatibility:** Older data can be interpreted by newer consumers
- **Full compatibility:** Bidirectional compatibility across versions

The agent continuously validates incoming schemas against these rules, preventing incompatible changes from entering the pipeline. Additionally, it orchestrates versioned schema notifications, ensuring that downstream services are aware of updates and can adapt accordingly. This transforms schema management from a passive registry function into an active governance process.

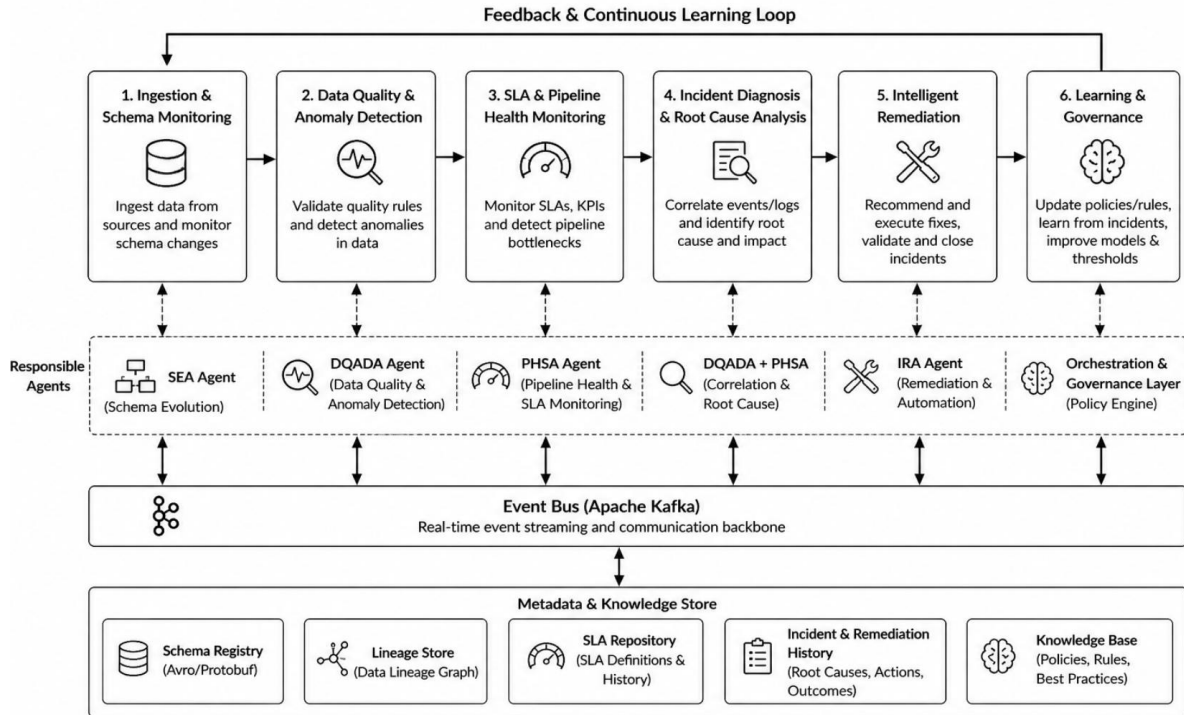


Figure 2. End-to-end governance lifecycle.

6. Anomaly Detection and Root Cause Analysis

While schema governance addresses structural integrity, ensuring data correctness and reliability requires continuous monitoring of data behavior.

The Data Quality & Anomaly Detection Agent (DQADA) fulfills this role by combining statistical analysis with machine learning techniques to detect and diagnose anomalies across pipeline stages.

Table 3: Anomaly Types and Detection Methods

Category	Example	Detection Method	Agent
Data anomaly	Distribution shift	Z-score, ML models	DQADA
Data anomaly	Null spike	Statistical threshold	DQADA
Infrastructure anomaly	CPU spike	System metrics	PHSA
Infrastructure anomaly	Network delay	Latency tracking	PHSA

6.1 Detection Methodology

The DQADA employs a hybrid detection framework that integrates complementary approaches:

- **Statistical baselines:** The agent maintains historical profiles of key metrics—such as value distributions, null rates, and cardinality—and applies techniques like z-score analysis and seasonal decomposition to identify deviations [19]. These methods are particularly effective for detecting abrupt changes in stable datasets.

- **Machine learning models:** To capture complex, non-linear patterns, the system incorporates models such as autoencoders and isolation forests [20]. These models learn latent representations of normal behavior and flag observations that deviate significantly from learned patterns. This enables detection of subtle anomalies that may not be apparent through simple statistical thresholds.

By combining these approaches, the system balances interpretability and sensitivity, reducing both false positives and false negatives.

6.2 Anomaly Types

The framework distinguishes between two broad categories of anomalies:

- **Data-side anomalies:** These arise from issues within the data itself and include:
 - **Distribution shifts:** Changes in statistical properties of features
 - **Null inflation:** Unexpected increases in missing values
 - **Referential integrity violations:** Broken relationships across datasets
- **Infrastructure-side anomalies:** These originate from the execution environment and include:
 - **Resource contention:** CPU, memory, or I/O bottlenecks
 - **Network failures:** Latency spikes or connectivity disruptions
 - **Upstream service instability:** API failures or ingestion interruptions

This distinction is critical, as it informs the subsequent remediation strategy and helps avoid misattribution of root causes.

6.3 Root Cause Attribution

Detection alone is insufficient without accurate diagnosis. To this end, the system performs root

cause attribution by traversing pipeline lineage graphs stored in metadata systems such as Apache Atlas. When an anomaly is detected, the agent traces dependencies upstream and downstream, correlating signals across multiple stages.

The attribution process is further enhanced using causal inference techniques, which evaluate the likelihood that a given upstream change or event is responsible for the observed anomaly [21]. By combining lineage traversal with probabilistic reasoning, the system can isolate the most plausible origin of the issue, even in highly interconnected pipelines.

This capability significantly reduces the need for manual investigation, enabling faster and more accurate responses to data quality incidents.

7. Self-Healing and Remediation Logic

The Incident Remediation Agent (IRA) represents the execution layer of the governance framework, translating detection and diagnosis into actionable recovery strategies. Unlike traditional observability systems that terminate at alert generation, the IRA is designed to close the loop by enabling autonomous, policy-aware remediation. Its design emphasizes controlled autonomy, ensuring that corrective actions improve system reliability without introducing unintended side effects.

Table 4: Remediation Action Matrix

Incident Type	Root Cause	Automated Action	Escalation Required
Job failure	Transient error	Restart job	No
Schema mismatch	Breaking change	Rollback schema	Sometimes
Data loss	Late ingestion	Trigger backfill	No
Unknown anomaly	Unclear cause	None	Yes

7.1 Action Space

The remediation capabilities of the IRA are defined through a constrained and well-understood action space derived from common operational practices in production data systems:

- **Restart failed jobs:** For transient failures—such as temporary resource exhaustion or intermittent connectivity issues—the agent can automatically restart failed tasks or

pipeline stages. This is typically a low-risk, high-success intervention.

- **Roll back schema versions:** When a breaking schema change is identified as the root cause, the agent can revert to a previously stable schema version. This rollback is coordinated with the Schema Evolution Agent to maintain system-wide consistency.

- Re-route data flows: In the presence of upstream instability or degraded sources, the agent can dynamically redirect data pipelines to fallback sources, cached datasets, or alternative processing paths.
- Trigger backfills: When data loss or delayed ingestion is detected, the agent initiates controlled backfill operations to restore completeness while respecting system load and SLA constraints.

Each remediation action is logged and versioned, ensuring full auditability and enabling retrospective analysis.

7.2 Safety Boundaries

Autonomous remediation introduces inherent risks, making explicit safety constraints essential. The IRA operates under a multi-layered control mechanism that governs when and how actions are executed:

- Confidence thresholds: Every remediation decision is associated with a confidence score derived from anomaly detection and root cause attribution. Actions are executed only when this confidence exceeds predefined thresholds.
- Impact analysis: Before execution, the agent evaluates the potential impact of an action on downstream systems, including dependency depth, data criticality, and SLA sensitivity.
- Policy constraints: Organizational governance rules define permissible actions. For example, automated schema rollback may be allowed in non-critical pipelines but restricted in regulated environments.

When uncertainty is high or policy constraints are violated, the system halts autonomous execution and escalates the issue to human operators [22]. This ensures that automation enhances reliability without compromising control.

8. Evaluation

8.1 Experimental Setup

The proposed framework was evaluated using simulated production-grade data pipelines designed

to reflect realistic enterprise workloads. These included:

- Streaming ingestion pipelines, implemented using Apache Kafka
- Batch processing workflows, representative of large-scale transformations in Apache Spark

The simulation environment introduced controlled disruptions such as schema changes, data anomalies, and infrastructure failures at frequencies aligned with observed real-world patterns.

8.2 Metrics

The evaluation focused on key governance performance indicators:

- Detection latency — time between issue occurrence and detection
- False positive and false negative rates — accuracy of anomaly and schema detection
- Autonomous resolution rate — percentage of incidents resolved without human intervention
- Mean Time to Remediation (MTTR) — time required to fully resolve incidents

These metrics collectively capture system responsiveness, accuracy, and operational efficiency.

8.3 Results

The experimental results demonstrate substantial improvements:

- Schema detection latency: Reduced from minutes (typical in manual validation systems) to seconds, enabling near real-time responsiveness
- MTTR reduction: Achieved an approximate 45% reduction compared to traditional human-driven processes
- Autonomous resolution rate: Approximately 62% of incidents were resolved without human intervention
- False positive rate: Maintained below 5%, indicating a strong balance between sensitivity and precision

Table 5: Evaluation Results

Metric	Baseline (Manual)	Proposed System	Improvement
Detection Latency	Minutes–Hours	Seconds	Significant
MTTR	High	↓ 45%	Moderate–High
Autonomous Resolution	0–10%	62%	High
False Positive Rate	Variable	<5%	Controlled

Importantly, the system demonstrated appropriate restraint in high-risk scenarios. In cases involving ambiguous signals or high-impact decisions, the IRA correctly withheld automated actions and triggered escalation. This validates the effectiveness of the defined safety boundaries and reinforces the system’s suitability for production environments.

9. Comparison with Existing Approaches

The proposed framework differs fundamentally from existing approaches in both scope and design philosophy. Traditional AIOps platforms primarily target infrastructure observability, focusing on metrics such as resource utilization, system uptime, and service health. While these systems—often deployed alongside platforms like Kubernetes—excel at detecting and mitigating infrastructure-level failures, they lack visibility into data semantics, schema evolution, and pipeline-level dependencies. As a result, they are not equipped to address issues such as silent schema drift or data quality degradation.

In contrast, the framework presented in this paper operates at the data governance layer, integrating structural, statistical, and operational signals into a unified decision-making process. By incorporating schema semantics and lineage awareness, it enables a deeper understanding of how changes propagate across pipelines and affect downstream consumers.

When compared to autonomous database systems—such as those built on Amazon Aurora or research prototypes described in [23]—the distinction becomes equally clear. Autonomous databases focus on intra-system optimization, including query planning, indexing, and resource allocation. While these capabilities improve database performance and efficiency, they do not extend to cross-pipeline governance, where

multiple systems, formats, and dependencies interact.

The proposed framework addresses these limitations by introducing a multi-agent architecture capable of coordinating across the full lifecycle of data pipelines. Its key differentiators include:

- **Schema-aware governance:** Unlike conventional monitoring systems, the framework explicitly models and enforces schema compatibility, enabling proactive management of schema evolution.
- **Multi-agent coordination across pipeline stages:** Governance responsibilities are distributed across specialized agents that collaborate through shared state and event-driven communication, allowing scalable and modular control.
- **Integrated anomaly detection and remediation:** Detection and response are tightly coupled within the same system, enabling automated, context-aware remediation rather than isolated alerting.

These characteristics position the framework as a distinct advancement beyond existing solutions, addressing the unique challenges of data-centric reliability in modern cloud-native environments.

10. Discussion

The results presented in this study suggest that multi-agent orchestration provides a viable and effective approach to autonomous data pipeline governance. By decomposing governance tasks into specialized agents and enabling coordinated decision-making, the system achieves significant improvements in detection latency, remediation speed, and overall reliability.

However, several challenges remain and warrant further investigation.

First, scalability of agent coordination becomes increasingly complex as the number of pipelines, data sources, and dependencies grows. While the event-driven architecture supports horizontal scaling, maintaining consistency and avoiding conflicting actions across agents requires careful design of coordination protocols and shared state management.

Second, explainability of decisions is critical for adoption in production environments. Autonomous actions—particularly those affecting critical data assets—must be interpretable by human operators. This necessitates mechanisms for tracing decision paths, exposing intermediate reasoning, and providing actionable explanations alongside automated interventions.

Third, policy design and governance boundaries present an ongoing challenge. Defining appropriate thresholds, constraints, and escalation rules requires domain expertise and may vary across organizational contexts. Overly conservative policies limit the benefits of automation, while overly aggressive policies risk unintended consequences.

Finally, the integration of such systems into existing data ecosystems raises practical considerations, including compatibility with legacy tools, operational overhead, and organizational readiness for increased automation.

11. Conclusion

This paper introduced a novel framework for autonomous data pipeline governance based on multi-agent orchestration. The proposed system integrates schema evolution handling, anomaly detection, SLA monitoring, and self-healing remediation into a cohesive architecture designed for cloud-native environments.

By enabling agents to monitor, reason about, and act upon pipeline events in real time, the framework addresses critical limitations of existing governance approaches, which rely heavily on manual intervention and fragmented tooling. The evaluation demonstrates meaningful improvements in detection latency, mean time to remediation, and autonomous resolution rates, while maintaining strict safety constraints.

The findings highlight the potential of agentic AI to transform data infrastructure from reactive monitoring systems into proactive, self-governing ecosystems. At the same time, the study underscores the importance of balancing autonomy with control, particularly in high-stakes data environments.

Future work will focus on extending the framework through reinforcement learning-based policy optimization, enabling agents to adapt dynamically to changing workloads and system conditions. Additionally, further research is needed to enhance human-agent collaboration, ensuring that automated systems remain transparent, trustworthy, and aligned with organizational objectives.

References

- [1] Akidau, T., Chernyak, S., & Lax, R. (2015). Dataflow model: A practical approach to balancing correctness, latency, and cost in massive-scale, unbounded, out-of-order data processing. *Proceedings of the VLDB Endowment*, 8(12), 1792–1803.
- [2] Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., & Mané, D. (2016). Concrete problems in AI safety.
- [3] Bernstein, P. A. (2003). Applying model management to classical meta data problems. *Proceedings of the Conference on Innovative Data Systems Research (CIDR)*.
- [4] Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), 1–58.
- [5] Chen, L., Xu, J., Zhang, Z., & Guo, X. (2020). AIOps: Real-world challenges and research innovations. *Proceedings of the IEEE International Conference on Cloud Engineering (IC2E)*.
- [6] Dang, Y., Wu, Q., Zhang, J., Zhang, J., & Xie, T. (2019). Characterizing and detecting performance bugs for cloud systems. *Proceedings of the ACM Symposium on Cloud Computing (SoCC)*.
- [7] Halevy, A., Rajaraman, A., & Ordille, J. (2006). Data integration: The teenage years. *Proceedings of the VLDB Endowment*, 9(10), 1–9.
- [8] Jennings, N. R. (2000). On agent-based software engineering. *Artificial Intelligence*, 117(2), 277–296.

- [9] Kleppmann, M. (2017). *Designing data-intensive applications: The big ideas behind reliable, scalable, and maintainable systems*. O'Reilly Media.
- [10] Kraska, T., Beutel, A., Chi, E. H., Dean, J., & Polyzotis, N. (2018). The case for learned index structures. *Proceedings of the ACM SIGMOD International Conference on Management of Data*, 489–504.
- [11] Kreps, J., Narkhede, N., & Rao, J. (2011). Kafka: A distributed messaging system for log processing. *Proceedings of the NetDB Workshop*.
- [12] Liu, F. T., Ting, K. M., & Zhou, Z.-H. (2008). Isolation forest. *Proceedings of the IEEE International Conference on Data Mining (ICDM)*, 413–422.
- [13] Marcus, R., Negi, P., Mao, H., Tatbul, N., Alizadeh, M., & Kraska, T. (2019). Neo: A learned query optimizer. *Proceedings of the VLDB Endowment*, 12(11), 1705–1718.
- [14] Nargesian, F., Zhu, E., Pu, K. Q., & Miller, R. J. (2020). Table union search on open data. *Proceedings of the VLDB Endowment*, 11(7), 813–825.
- [15] Pavlo, A., Angulo, G., Arulraj, J., Lin, H., Lin, J., Ma, L., & Stonebraker, M. (2017). Self-driving database management systems. *Proceedings of the Conference on Innovative Data Systems Research (CIDR)*.
- [16] Pearl, J. (2009). *Causality: Models, reasoning, and inference* (2nd ed.). Cambridge University Press.
- [17] Redman, T. C. (2018). If your data is bad, your machine learning tools are useless. *Harvard Business Review*.
- [18] Russell, S., & Norvig, P. (2021). *Artificial intelligence: A modern approach* (4th ed.). Pearson.
- [19] Schelter, S., Böse, J.-H., Kirschnick, J., Klein, T., & Seufert, S. (2018). Automatically tracking metadata and provenance of machine learning experiments. *Proceedings of the Workshop on Data Management for End-to-End Machine Learning*.
- [20] Stone, P., & Veloso, M. (2000). Multiagent systems: A survey from a machine learning perspective. *Autonomous Robots*, 8(3), 345–383.
- [21] Wooldridge, M. (2009). *An introduction to multiagent systems* (2nd ed.). Wiley.
- [22] Zaharia, M., Das, T., Li, H., Hunter, T., Shenker, S., & Stoica, I. (2013). Discretized streams: Fault-tolerant streaming computation at scale. *Proceedings of the ACM Symposium on Operating Systems Principles (SOSP)*.
- [23] Apache Software Foundation. (n.d.). *Apache Avro™ 1.11.0 documentation*.